

ASTRO® 25

INTEGRATED VOICE AND DATA

Enterprise OS Software Version 18.2 User Guide

System Release AN2024.HS, AN2024.1, 2022.HS, 2022.1, 2021.1, 2020.HS,
2020.1, 2019.x, 7.18

MARCH 2025

© 2025 Motorola Solutions, Inc. All Rights Reserved.



8971760H62-T

Intellectual Property and Regulatory Notices

Copyrights

The Motorola Solutions products described in this document may include copyrighted Motorola Solutions computer programs. Laws in the United States and other countries preserve for Motorola Solutions certain exclusive rights for copyrighted computer programs. Accordingly, any copyrighted Motorola Solutions computer programs contained in the Motorola Solutions products described in this document may not be copied or reproduced in any manner without the express written permission of Motorola Solutions.

No part of this document may be reproduced, transmitted, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, without the prior written permission of Motorola Solutions, Inc.

Trademarks

MOTOROLA, MOTO, MOTOROLA SOLUTIONS, and the Stylized M Logo are trademarks or registered trademarks of Motorola Trademark Holdings, LLC and are used under license. All other trademarks are the property of their respective owners.

License Rights

The purchase of Motorola Solutions products shall not be deemed to grant either directly or by implication, estoppel or otherwise, any license under the copyrights, patents or patent applications of Motorola Solutions, except for the normal nonexclusive, royalty-free license to use that arises by operation of law in the sale of a product.

Open Source Content

This product may contain Open Source software used under license. Refer to the product installation media for full Open Source Legal Notices and Attribution content.

European Union (EU) and United Kingdom (UK) Waste of Electrical and Electronic Equipment (WEEE) Directive



The European Union's WEEE directive and the UK's WEEE regulation require that products sold into EU countries and the UK must have the crossed-out wheeled bin label on the product (or the package in some cases). As defined by the WEEE directive, this crossed-out wheeled bin label means that customers and end users in EU and UK countries should not dispose of electronic and electrical equipment or accessories in household waste.

Customers or end users in EU and UK countries should contact their local equipment supplier representative or service center for information about the waste collection system in their country.

Disclaimer

Please note that certain features, facilities, and capabilities described in this document may not be applicable to or licensed for use on a specific system, or may be dependent upon the characteristics of a specific mobile subscriber unit or configuration of certain parameters. Please refer to your Motorola Solutions contact for further information.

© 2025 Motorola Solutions, Inc. All Rights Reserved

Contact Us

The Centralized Managed Support Operations (CMSO) is the primary contact for technical support included in your organization's service agreement with Motorola Solutions. To enable faster response time to customer issues, Motorola Solutions provides support from multiple countries around the world.

Service agreement customers should be sure to call the CMSO in all situations listed under Customer Responsibilities in their agreement, such as:

- To confirm troubleshooting results and analysis before taking action

Your organization received support phone numbers and other contact information appropriate for your geographic region and service agreement. Use that contact information for the most efficient response. However, if needed, you can also find general support contact information on the Motorola Solutions website, by following these steps:

1. Enter motorolasolutions.com in your browser.
2. Ensure that your organization's country or region is displayed on the page. Clicking or tapping the name of the region provides a way to change it.
3. Select "Support" on the motorolasolutions.com page.

Comments

Send questions and comments regarding user documentation to documentation@motorolasolutions.com.

Provide the following information when reporting a documentation error:

- The document title and part number
- The page number or title of the section with the error
- A description of the error

Motorola Solutions offers various courses designed to assist in learning about the system. For information, go to <https://learning.motorolasolutions.com> to view the current course offerings and technology paths.

Document History

Part Number	Description	Date
8971760H62-A	Original release of the <i>Enterprise OS Software Version 18.2 User Guide</i> .	October 2012
8971760H62-B – 8971760H62-P	General revision of the <i>Enterprise OS Software Version 18.2 User Guide</i> .	April 2024
8971760H62-R	Revised for the AN2024.HS and AN2024.1 system releases.	September 2024
8971760H62-T	Updated sections: ● IPQoS Priority Queueing on page 194 ● Configuring IPQoS Based on Priority Queueing for IPIP and IPSec Interfaces on page 195 New section: ● Configuring IPQoS Based on Priority Queueing for Physical and VLAN Interfaces on page 197	March 2025

Contents

Intellectual Property and Regulatory Notices.....	2
Contact Us.....	3
Document History.....	4
List of Figures.....	27
List of Tables.....	33
List of Processes.....	37
List of Procedures.....	38
About Enterprise OS Software Version 18.2 User Guide.....	46
Related Information.....	46
Command Line Interface Symbol Conventions.....	47
Text Conventions.....	49
Chapter 1: Basic Ports and Paths Configuration.....	50
1.1 Path and Port Numbering.....	50
1.1.1 Multiple Paths to a Wide Area Port Configuration.....	51
1.1.2 S2500 Port and Path Numbering.....	51
1.1.3 S6000 Path-to-Port Mappings (FlexWAN Modules in Both Slots).....	54
1.1.4 S6000 Path-to-Port Mappings (UltraWAN Modules in Both Slots).....	55
1.1.5 S6000 Path-to-Port Mappings (T3/E3 Modules in Both Slots).....	57
1.1.6 S6000 Path-to-Port Mappings (FlexWan Module in Slot A and UltraWAN II Module in Slot B).....	58
1.1.7 S6000 Path-to-Port Mappings (UltraWAN II Module in Slot A and FlexWAN Module in Slot B).....	59
1.1.8 S6000 Path-to-Port Mappings (FlexWAN Module in Slot A and T3/E3 Module in Slot B).....	60
1.1.9 S6000 Path-to-Port Mappings (T3/E3 Module in Slot A and FlexWAN Module in Slot B).....	61
1.1.10 S6000 Path-to-Port Mappings (UltraWAN II Module in Slot A and T3/E3 Module in Slot B).....	62
1.1.11 S6000 Path-to-Port Mappings (T3/E3 Module in Slot A and UltraWAN II Module in Slot B).....	63
1.1.12 S6000 Path-to-Port Mappings (12-Port T1/E1 II Modules in Both Slots).....	64
1.1.13 GGM 8000 Port and Path Numbering.....	65
1.1.14 MCG 8000 Port and Path Numbering.....	69
1.2 Path, Port, and Connector Syntax.....	71
1.3 Configuring Local Area Interfaces.....	73
1.4 Configuring Wide Area Interfaces.....	74
Chapter 2: Virtual Ports and Multiple Instance Lists Configuration.....	77
2.1 Virtual Ports Over Frame Relay.....	78
2.2 Parent Ports for Frame Relay.....	79
2.3 Virtual Ports Over PPP.....	80
2.4 Virtual Ports for Ethernet WAN Links.....	80
2.4.1 Configuring Ethernet Site Links by Using IPIP with BFD.....	80

2.4.2 Configuring Ethernet Interzone Links by Using IPIP with BGP.....	83
2.5 Virtual Ports for IPv6 Ethernet WAN Links.....	86
2.5.1 Configuring Ethernet Site Links by Using IPIIPv6 with BFD.....	86
2.5.2 Configuring Ethernet Interzone Links by Using IPIIPv6 with BGP.....	89
2.6 Configuring Frame Relay and PPP Virtual Ports.....	93
2.7 Configuring Virtual Ports for 802.1Q Virtual LANs.....	94
2.8 Multiple Instance Lists.....	96
Chapter 3: Autostartup Configuration.....	98
3.1 Autostartup Phase 1.....	98
3.1.1 Determining the Progress of the Establishment of the Physical and Data Links.....	99
3.2 Autostartup Phase 2.....	99
3.3 Preparing Autostartup.....	100
3.4 Configuring the Central Node and the BootP/TFTP Servers on the Central Site Network.....	101
3.5 Autostarting the Central Site Node.....	104
3.6 Sample Autostartup Configurations.....	104
3.6.1 Setting Up Configuration with Frame Relay WAN.....	104
3.6.2 Setting Up Configuration with PP WAN.....	106
Chapter 4: ASCII Text Configuration File.....	107
4.1 ASCII Text Configuration File Download.....	107
4.2 Configuration File Execution with LoadConfigs.....	107
4.3 Additional Configuration File Uses.....	108
4.4 StaticRP.cfg File.....	108
4.5 ASCII Boot and Capture Features.....	109
4.6 Configuring a Device by Using ASCII Boot.....	109
4.6.1 boot.cfg File Renaming.....	111
4.6.2 Macros in the boot.cfg File.....	111
4.6.3 Include Command in boot.cfg Files.....	113
4.6.4 Include Command Limitations.....	114
4.7 ASCII Boot Feature Limitations.....	114
4.8 ASCII Capture Feature.....	114
4.9 ASCII Capture Feature Limitations.....	115
4.10 Password Capturing Format.....	116
4.11 SNMP Command.....	116
4.12 Boot, Capture, and Log Files Contents Review.....	116
Chapter 5: IP Routing Configuration.....	119
5.1 Configuring Basic IP Routing over LAN Points and PPP Links.....	119
5.2 IP Routing for Wide Area Networks.....	121
5.3 Verifying the IP Routing Configuration.....	122
5.3.1 Checking the Forwarded Packets with the PING Command.....	123
5.3.2 Verifying the Overall Router Status.....	124

5.4 Customizing Routing Configuration.....	125
5.4.1 Configuring Multiple IP Networks/Subnets.....	126
5.4.2 Static Route Configuration.....	127
5.4.2.1 Static Route Override Feature.....	128
5.4.3 Configuring Packet Filtering.....	131
5.4.3.1 Packet Filtering Components and Protocol Reservation.....	131
5.4.3.2 IP Packet Filtering Examples.....	134
5.4.3.3 Setting Up the Filtering Options for IP Packet Filtering.....	136
5.4.4 RTP/UDP/IP Header Compression.....	137
5.4.5 Bidirectional Forwarding Detection (BFD) Configuration.....	138
5.5 Configuring RIP Routing Policies.....	140
5.6 Migration to a RIPv2 Network.....	142
5.7 RIPv2 Configuration for Networks with Variable Length Subnet Masks.....	143
5.7.1 Aggregate/Deaggregate Scheme.....	144
5.7.2 Range Table Mask Scheme.....	145
5.7.3 RIP V2 Compatibility.....	146
5.7.4 Authentication Block.....	147
5.7.5 RIPv2 Message Authentication.....	147
5.8 Configuring OSPF Routing Policies.....	147
5.8.1 Controlling the Reporting of Routes Learned from Other Sources and Advertised into the OSPF Domain.....	149
5.8.2 Configuring OSPF Route Aggregation.....	150
5.9 Configuring ISIS Routing Policy.....	150
5.10 Interautonomous System Routing Configuration with BGP.....	152
5.10.1 Configuring BGP Peers.....	152
5.10.2 Default Route Configuration.....	154
5.10.3 Configuring BGP Route Aggregation.....	154
5.10.4 Importing Routes from an IGP Domain to a BGP Domain.....	156
5.10.5 Route Import from a BGP Domain to an IGP Domain.....	157
5.10.6 Configuring Network Number Policies.....	158
5.10.7 Configuring AS-Path Permit or Deny Policies.....	159
5.10.7.1 Regular Expression Examples.....	160
5.10.7.2 Deny Filter Examples.....	160
5.10.7.3 Permit Filter Examples.....	161
5.10.8 Configuring AS-Path Weight Policies.....	161
5.10.8.1 Weight Filter Examples.....	162
5.10.8.2 Degree of Preference Calculations.....	162
5.10.9 MED/LOCAL-PREF Value Configuration.....	163
5.10.9.1 Configuring MED/LOCAL-PREF with a Network Policy.....	164
5.10.9.2 Configuring MED/LOCAL-PREF with an OSPF Export Policy.....	165

5.10.9.3 Configuring MED/LOCAL-PREF with an Interior Policy.....	165
5.10.10 BGP Convergence Optimization.....	166
5.11 IP Network Topology.....	166
5.12 Multipath Routing.....	168
5.13 Default Routes.....	171
5.14 Route Learning Within an Autonomous System.....	172
5.14.1 Route Learning with RIP.....	173
5.14.2 Network Reachability.....	173
5.14.3 Solution to Slow Convergence Problem with Split Horizon.....	173
5.14.4 Solution to Slow Convergence Problem with Poison Reverse.....	175
5.14.5 RIP Configuration Parameters.....	175
5.14.6 Different States of RIP-Learned Routes.....	176
5.14.7 Route Learning with OSPF.....	177
5.14.7.1 Different Functions of OSPF Devices.....	177
5.14.7.2 Area Border Router.....	178
5.14.7.3 Autonomous System Boundary Router.....	179
5.14.7.4 Routes Learning and Network Reachability.....	179
5.14.7.5 Network Costs Reduction with Demand Interface Circuits.....	181
5.14.7.6 Link State Advertisements.....	181
5.14.7.7 Metrics.....	182
5.14.7.8 OSPF Configuration Customization.....	182
5.15 Autonomous System Routing with BGP.....	183
5.15.1 External and Internal Peers.....	184
5.15.2 Peer-to-Peer Communication.....	185
5.15.3 Path Attributes.....	185
5.15.4 Path Selection.....	189
5.15.5 Policies.....	190
5.15.6 Route Aggregation.....	192
5.16 Integrated IS-IS (IISIS).....	192
5.17 Address Resolution.....	192
5.18 Other Global Routing Configurations.....	193
Chapter 6: Quality of Service (IPQoS).....	194
6.1 IPQoS Priority Queueing.....	194
6.1.1 Configuring IPQoS Based on Priority Queueing for IPIP and IPSec Interfaces.....	195
6.1.2 Configuring IPQoS Based on Priority Queueing for Physical and VLAN Interfaces.....	197
6.1.3 Interpreting Buffer Overrun Errors When IPQoS Priority Queueing Is Enabled.....	198
6.1.4 IPQoS Port Bundling.....	199
6.2 Configuring Rate Limiting with IPQoS Policy.....	199
6.2.1 Configuring Rate Limiting – Example 1.....	201
6.2.2 Configuring Rate Limiting – Example 2.....	201

6.3 Class-Based Queueing.....	201
6.3.1 CBQ Classes.....	203
6.3.2 Configuring IPQoS Based on CBQ.....	203
6.3.3 Configuring IP/TOS to CBQ Class Mapping.....	205
6.3.4 Class-Based Queueing Statistics View.....	206
6.4 Configuring IPQoS Based on Protocol Reservation.....	206
6.5 IEEE 802.1P Prioritization.....	208
6.5.1 Configuring 802.1P-to-TOS Mapping on an Ingress VLAN Port.....	208
6.5.2 Configuring TOS-to-802.1P Mapping on an Egress VLAN Port.....	209
6.6 Configuring TOS-to-IP Mapping.....	210
6.7 Defining a Default IPQoS Policy.....	210
Chapter 7: System IP (SYSIP) Configuration.....	212
7.1 System IP Configuration.....	212
7.2 Routing Issues.....	213
7.3 ARP Packets.....	214
7.4 NAT/Firewall Issues.....	214
7.5 VRRP Issues.....	214
Chapter 8: Resource Reservation Protocol Configuration.....	215
8.1 Configuring the RSVP.....	215
8.2 RSVP Proxy Sender and Receiver.....	216
8.2.1 Proxy Sender: Unicast Destination and One Sender Port.....	216
8.2.2 Proxy Receiver: Unicast Destination and One Sender Port.....	217
8.2.3 Proxy Sender: Multicast Destination with a Range of Sender Ports.....	217
8.2.4 Proxy Receiver: Multicast Destination with a Range of Sender Ports.....	217
8.2.5 Configuring RSVP with L2TP Tunnel.....	218
Chapter 9: Dynamic Host Configuration Protocol Configuration.....	220
9.1 Configuring DHCP on EOS Devices.....	220
9.2 DHCP Service Activation.....	221
9.3 Address Pools Configuration.....	221
9.4 DHCP Profiles.....	222
9.5 Default DHCP Profile (Profile P1) Contents Display.....	223
9.6 Changing DHCP Profile Configuration Values.....	223
9.7 Custom DHCP Profiles Configuration.....	225
9.8 Console Logging Feature.....	226
Chapter 10: IPv6 Routing Configuration.....	228
10.1 Configuring IPv6 for Local Area Networks.....	228
10.2 IPv6 Configuration for Wide Area Network Connectivity.....	229
10.2.1 Configuring IPv4 over IPv6 Tunnels.....	229
10.2.2 Configuring IPv6 over IPv4 Tunnels with an Encrypted Payload.....	230

10.2.3 Configuring IPv6 over IPv6 Tunnels.....	231
10.2.4 Configuring IPv6 over IPv6 Tunnels with an Encrypted Payload.....	231
10.3 Verifying the Configuration.....	232
10.3.1 Examining Network Devices.....	233
10.3.2 Retrieving IPv6 Statistics.....	233
10.3.3 Checking the Overall Status.....	233
10.4 Checking with PING6, TraceRoute6, and TELnet6.....	234
10.5 Customizing the IPv6 Routing Configuration.....	234
10.5.1 Configuring the Internal Port.....	234
10.5.2 Configuring Multiple IPv6 Subnets.....	235
10.5.3 Configuring Neighbor Discovery.....	235
10.5.4 Configuring Static Routes.....	235
10.5.5 Configuring RIPNG Routing Policies.....	236
10.5.6 Configuring RIPNG Route Aggregation.....	237
10.6 IPv6 Network Topology.....	238
10.7 Multipath Routing.....	238
10.8 Default Routes.....	239
10.8.1 Learning Routes with RIPNG.....	240
10.8.2 Network Reachability.....	240
10.8.3 Solution to the Slow Convergence Problem with Split Horizon.....	240
10.8.4 Solution to the Slow Convergence Problem with Poison Reverse.....	241
10.8.5 Different States of RIPNG-Learned Routes.....	241
Chapter 11: Network Address Translation (NAT) Configuration.....	243
11.1 NAT Guidelines.....	243
11.2 Configuring Network Address Translation.....	244
11.3 Address Mapping Examples.....	245
11.4 Basic NAT Operation.....	246
11.5 Traffic Directions.....	246
11.6 Address Mapping.....	248
11.7 Mapping the IPCP Address.....	249
11.8 TCP/UDP Port Mapping.....	251
11.9 Configuring Private Address Space.....	251
11.10 Configuring Load Sharing.....	254
11.11 Configuring Address Migration.....	255
11.12 Configuring Address Redirection.....	256
11.13 NAT Redundancy.....	257
11.14 Limitation on One-to-One Mapping.....	257
Chapter 12: IP Multicast Routing Configuration.....	258
12.1 Multicast Routing Terms.....	259
12.2 MBONE Connectivity with Multicasting.....	260

12.3 Multicast Addresses.....	260
12.4 Internet Group Management Protocol.....	261
12.5 Distance Vector Multicast Routing Protocol.....	262
12.6 Multicast Open Shortest Path First Protocol.....	263
12.6.1 Group Membership Learning.....	264
12.6.2 Shortest Path First Tree.....	264
12.6.3 Forwarding Cache.....	264
12.6.4 Interarea Multicasting.....	265
12.6.5 Interautonomous System Multicasting.....	266
12.7 Multicast Client.....	266
12.8 Multicast Group Limitation.....	267
12.8.1 IGMP (MIP) Multicast Group Limitation.....	267
12.8.2 PIM Multicast Group Limitation.....	268
12.9 Configuring IP Multicast Routing for Local Area Networks and Point-to-Point Links.....	268
12.10 Verifying the IP Multicast Routing Configuration Overall Status.....	270
12.11 Viewing IP Multicasting Statistics.....	271
12.12 DVMRP Configuration Troubleshooting.....	272
12.13 Multicast Routing Configuration Customization.....	273
12.13.1 Controlling Local Group Membership Queries.....	274
12.13.2 Configuring IGMP Proxy.....	274
12.13.3 Configuring IGMP Proxy with Source Address Translation.....	274
12.13.4 LAN QoS Configuration to Support Geographically Redundant IP Simulcast Prime Sites.....	276
12.13.5 Multicast Datagram Threshold Adjustment.....	277
12.13.6 Customization with DVMRP Protocol.....	277
12.13.6.1 Configuring a DVMRP Multicast Tunnel.....	277
12.13.6.2 Configuring DVMRP Scoping.....	278
12.13.6.3 Configuring DVMRP Multicasting over Frame Relay.....	280
12.13.6.4 Configuring a DVMRP Metric.....	281
12.13.6.5 Controlling the DVMRP Rate Limit for Multicast Traffic.....	281
12.13.6.6 Configuring DVMRP Route Aggregation.....	282
12.13.6.7 Controlling the Routing Table.....	283
12.13.6.8 Controlling the Forwarding Table.....	284
12.13.7 Customization with MOSPF Protocol.....	285
12.13.7.1 Configuring Interarea Multicasting.....	285
12.13.7.2 Configuring MOSPF Routing Policies.....	286
12.13.7.3 Configuring MOSPF Forwarding Policies.....	287
12.14 Configuring Multicast Border Router Functionality.....	288
Chapter 13: Protocol Independent Multicast-Sparse Mode (PIM) Configuration.....	290
13.1 Rendezvous Points.....	290
13.2 Bootstrap Router.....	290

13.3 Group Operations.....	291
13.4 Designated Router Election.....	291
13.5 Delayed PIM Hellos Disabling.....	293
13.6 GenID Option.....	293
13.7 Triggered Bootstrap Messages.....	294
13.8 Triggered Candidate RP Advertisements.....	294
13.9 Bootstrap Segmentation.....	295
13.10 Join/Prune Aggregation.....	296
13.11 Multiple Joins.....	296
13.12 Static RP Sets.....	297
13.13 Internal IP Address Support.....	298
13.14 Event Driven/Triggered PIM-SM.....	298
13.15 PIM Packet Formats.....	299
13.16 PIM SNMP Traps.....	299
13.17 Configuring PIM-SM.....	299
Chapter 14: Router Discovery Protocol Configuration.....	303
14.1 Setting Up RDP.....	304
14.2 Verifying the RDP Configuration.....	306
14.3 Troubleshooting the RDP Configuration.....	306
Chapter 15: Wide Area Networking Configuration with Frame Relay.....	308
15.1 Setting Up The Frame Relay Service.....	308
15.1.1 Configuring the Frame Relay Service.....	308
15.1.2 Configuring Congestion Control for NTTLMI Protocol Users.....	309
15.1.3 Configuring Congestion Control for Other LMI Protocol Users.....	310
15.1.4 Frame Relay Congestion Control Configuration Example.....	311
15.1.5 Configuring Frame Relay SVC Connections.....	311
15.1.6 Verifying the Frame Relay Configuration.....	313
15.2 Configuring Transparent Bridging over Frame Relay.....	313
15.3 Configuring IP over Frame Relay.....	314
15.4 Multilink Frame Relay Configuration.....	317
15.4.1 Adding an Aggregated Virtual Circuit.....	317
15.5 Priority Queuing and FRF.12 Fragmentation.....	318
15.5.1 Configuring Realtime Priority Queueing.....	318
15.5.2 Configuring FRF.12 Fragmentation.....	319
15.5.3 Priority Queue-to-TOS Mapping.....	320
15.6 FRIPHC Negotiation.....	320
15.6.1 Parameters Negotiated Between Peers.....	321
15.7 Configuring Disaster Recovery in a Frame Relay Environment.....	322
15.7.1 Configuring a Primary PVC.....	323
15.7.2 Configuring a Backup PVC.....	323

15.7.3 Configuring a Backup Link.....	324
15.8 Frame Relay Networks.....	324
15.8.1 PVC and SVC Connections.....	325
15.8.1.1 Establishing an SVC Connection.....	326
15.8.2 Frame Relay Topologies.....	327
15.8.3 Frame Relay Addresses.....	329
15.8.4 Local Management Interface (LMI) Protocol.....	330
15.8.5 CRTP Twice Algorithm.....	331
15.8.6 FRF.12 Fragmentation.....	332
15.8.7 Frame Relay Congestion Control.....	333
15.8.8 Disaster Recovery over Frame Relay.....	335
15.9 Frame Relay Queue-Depth Exceeded Traps and Variables.....	338
15.10 Frame Relay Autostartup.....	339
15.11 MFR Overview.....	339
Chapter 16: Wide Area Networking Configuration with PPP.....	341
16.1 Packet Size Negotiation.....	341
16.2 Serial Line Management.....	341
16.3 Serial Line Quality Maintenance.....	342
16.4 Authentication Methods.....	342
16.5 Load Sharing and Load Balancing.....	343
16.6 Point-to-Point Protocol Communication.....	344
16.6.1 Setting Up PAP.....	345
16.6.2 Setting Up CHAP.....	346
16.6.3 Verifying Your PAP or MPPE Configuration.....	347
16.6.4 Setting Up EAP.....	347
16.7 NCP Configuration.....	347
16.8 Activating LAPB to Reduce Noisy Lines.....	348
Chapter 17: Port Bandwidth Management.....	349
17.1 Dial-on-Demand.....	350
17.2 Bandwidth-on-Demand.....	350
17.3 System Bandwidth Disaster Recovery.....	351
17.4 System Bandwidth Path Configurations.....	351
17.5 Bandwidth Management Terms.....	351
17.6 Supported Communication Resources.....	352
17.7 Paths to Ports Associations.....	353
17.8 Resource Aggregation.....	355
17.9 Manual Bandwidth Management.....	355
17.10 Bandwidth Management Status and Statistics Displays.....	356
17.11 Configuring Wide Area Networking with Async PPP.....	357
17.11.1 Setting Up a Path for Async PPP and AT Dial.....	357

17.11.2 Initializing the Modem.....	358
17.11.3 Adding a Number to the Phone Number List.....	359
17.12 Modem Initialization Strings.....	359
17.12.1 Command Interaction.....	360
17.12.2 Signal Interaction.....	361
17.12.3 Online Operation.....	361
17.12.4 Initialization Strings Examples.....	362
17.13 Configuring WAN Resources.....	362
17.13.1 Configuring Dial-Up Lines with A Modem Or TA.....	363
17.13.2 Configuring Leased Lines.....	365
17.13.3 Configuring ASPC Leased Lines.....	366
17.14 Configuring System Bandwidth Management Mode (DOD).....	366
17.15 Configuring Bandwidth-on-Demand.....	367
17.16 Configuring the Dial List.....	367
17.16.1 Dial List Configuration Examples.....	368
17.17 Phone Number Operations.....	369
17.18 Paths to Ports Binding.....	369
17.18.1 Converting a Static Path to a Dynamic Path.....	370
17.18.2 Changing a Dynamic Path to a Static Path.....	370
17.18.3 Identifying Remote Sites with SCID and CLIP.....	370
17.19 Configuring the Path Preference List.....	371
17.20 Path Operations.....	373
17.21 Configuring Manual Bandwidth Management Mode.....	374
17.22 Verifying the Bandwidth Management Configuration.....	375
17.23 Troubleshooting the Bandwidth Management Configuration.....	376
17.24 Load Balancing over Multiple Dial-Up Links.....	376
17.25 Configuring Routing over DOD Links.....	377
17.26 Bandwidth Management Commands and Parameters.....	378
Chapter 18: Protocol Reservation Configuration.....	383
18.1 Protocol Reservation Procedural Overview.....	385
18.2 Protocol Reservation with Frame Relay Virtual Ports.....	387
18.3 Configuring Protocol Reservation for Bridged Traffic.....	387
18.4 Configuring Protocol Reservation for IP-Routed Packets.....	388
18.5 Protocol Reservation Allocating IP Protocol Types.....	389
18.6 Protocol Reservation on Virtual Ports.....	390
18.7 Protocol Reservation Functionality.....	391
Chapter 19: Data Compression Configuration.....	393
19.1 Supported Data Compression Types.....	393
19.2 Tinygram Compression.....	393
19.3 Link-Level Compression.....	394

19.4 Configuring Tinygram Compression.....	395
19.5 Configuring History-Based or Per-Packet Compression.....	395
19.6 Enabling LAPB for a PPP Link.....	396
19.7 Frame Relay Configuration Options.....	396
19.8 Verifying the Link-Level Compression Effectiveness.....	396
Chapter 20: Multiprotocol Data Prioritization.....	398
20.1 Data Prioritization Overview.....	398
20.1.1 Priority Assigned to Packets.....	399
20.1.2 Queue Types.....	399
20.1.3 Queue Arbitration Algorithm.....	401
20.2 Configuring Data Prioritization.....	402
20.3 Data Prioritization Parameters.....	403
Chapter 21: UDP/LAPB Packet Relay Configuration.....	405
21.1 Configuring UDP/LAPB Packet Relay on MNR S6000 Routers.....	405
Chapter 22: Cooperative WAN Routing Configuration.....	407
22.1 Configuring MNR S6000 Routers as CWR Peers.....	407
Chapter 23: Bridging Configuration.....	411
23.1 Transparent Bridging.....	411
23.2 Spanning Tree Algorithm.....	411
23.2.1 Spanning Tree Algorithm with Wide Area Bridges.....	416
23.2.2 Spanning Tree Protocol over PPP Configuration.....	417
23.2.3 Spanning Tree Parameters Modification.....	417
23.2.4 Topology Reconfiguration.....	418
23.3 Load-Sharing Algorithm.....	418
23.4 Routing Tables.....	418
23.5 Learning and Filtering.....	419
23.6 Configuring Basic Bridging.....	420
23.7 Verifying the Bridge Configuration.....	422
23.8 Getting Statistics.....	423
23.9 Troubleshooting the Bridge Configuration.....	424
23.10 Bridge Customization.....	425
23.10.1 Configuring Per-Port Transparent Bridging.....	425
23.10.2 Adding or Deleting Static Entries.....	426
23.10.3 Bridge Security Configuration.....	426
23.10.3.1 Configuring Source Explicit Forwarding.....	426
23.10.3.2 Configuring Source Explicit Blocking.....	427
23.10.3.3 Configuring Destination Explicit Forwarding.....	428
23.10.3.4 Configuring Destination Explicit Blocking.....	429
23.10.3.5 Combined Source and Destiny Security.....	430

23.10.3.6 Custom Filters Configuration.....	431
23.11 Optimizing Bridge Performance.....	432
Chapter 24: Mnemonic Filtering Configuration.....	433
24.1 Configuring Filters.....	433
24.2 Configuring a Filter Between Station Groups.....	434
24.3 Filter Parameters.....	436
24.4 Filtering Process.....	436
24.5 Built-in Bridge Masks.....	439
24.6 User-Defined Bridge Masks.....	440
24.7 Bridge Filtering Examples.....	441
24.7.1 Creating a Mnemonic Filter by Using the PROTOcolRsrv <TAG> Action Option.....	445
24.7.2 Reserving Bandwidth For Packets of Different Lengths.....	446
Chapter 25: User Datagram Protocol Broadcast Helper Configuration.....	447
25.1 Configuring UDP Broadcast Helper.....	448
25.2 Relaying BOOTP and DHCP Traffic.....	451
25.3 Verifying the UDP Broadcast Helper Configuration.....	453
25.4 UDP Broadcast Helper for BOOTP Customization.....	453
25.4.1 Limiting the Number of Hops.....	453
25.4.2 Determining the Booting Order.....	454
Chapter 26: Virtual Private Networks (VPNs) Configuration.....	456
26.1 Tunneling.....	456
26.1.1 Authentication Mechanisms.....	458
26.1.2 Multiprotocol Tunneling.....	459
26.1.3 Tunnel Peers.....	459
26.1.4 Compression.....	459
26.1.5 Network Address Translation.....	460
26.2 VPN Strategies.....	460
26.3 VPN Software Protocols.....	462
26.4 Remote Access for Individual Users (RAS VPNs).....	462
26.4.1 Configuring the Central Site of a Remote/Dialup VPN.....	464
26.4.2 Configuring the Central Site of a RAS VPN.....	464
26.5 VPN for a Remote Office (Site-to-Site VPNs).....	464
26.5.1 Configuring the Remote Office Route in a Site-to-Site VPN.....	465
26.5.2 Configuring the Central Site Device in a Site-to-Site VPN.....	466
Chapter 27: Remote Access Services Configuration.....	467
27.1 Configuring Remote Access with Tunnel Terminator.....	467
27.2 NCPs Configuration.....	470
27.3 RADIUS Framed Net Mask and Framed Route Attributes.....	470
27.4 Logging Messages.....	471

27.4.1 Specifying Log Content.....	472
Chapter 28: Security Features Configuration on EOS Devices.....	473
28.1 Protocol Authentication with Manual Security Policies.....	473
28.1.1 Configuring PIM Protocol Authentication.....	474
28.1.2 Changing the Protocol Authentication Keyset (Rekeying).....	475
28.1.3 Configuring PIM Authentication Rollback.....	476
28.1.4 Configuring X509 Certificates for the Router.....	477
28.1.4.1 Configuring X509 Certificates (Example).....	478
28.2 IP Traffic Protection with Dynamic Security Policies.....	479
28.2.1 Configuring a Dynamic-Key Security Policy for IPv4.....	479
28.2.2 Configuring a Dynamic Security Policy for IPv6.....	481
28.2.3 Customized Security Associations.....	483
28.2.4 Dynamic Key Configuration Examples.....	483
28.2.4.1 Configuring the GGSN and CBR.....	484
28.2.4.2 Setting Up the Dynamic Key Configuration with X509 Certificates.....	486
28.2.4.3 Setting Up a CWR Topology for Encryption.....	488
28.2.4.4 Configuring Routers in Fully Meshed Topology with a Dynamic Key.....	491
28.2.4.5 Creating an IPv6 IPsec Tunnel with an IKEv2 Dynamic Key.....	496
28.2.4.6 Creating a Hub and Spoke Topology.....	498
28.3 MD5 Authentication Configuration for OSPF and BGP.....	500
28.3.1 MD5 Authentication Control States.....	501
28.3.2 Configuring MD5 Authentication for OSPF.....	502
28.3.3 Configuring MD5 Authentication for BGP.....	503
28.3.4 MD5 Authentication Statistics.....	504
28.3.5 MD5 Authentication Traps.....	504
28.4 Secure Shell (SSH) Support.....	505
28.4.1 SSH Public Key Management.....	505
28.4.2 Starting an SSH Session (EOS SSH Client and EOS SSH Server).....	507
28.4.3 Starting and SSH Session (EOS SSH Client and Non-EOS SSH Server).....	508
28.4.4 Starting an SSH Session (Non-EOS SSH Client and EOS SSH Server).....	509
28.4.5 Stopping SSH Sessions.....	510
28.4.6 SCP Sessions.....	511
28.5 Password Protection.....	511
28.6 Encrypted Values for Password-Protected Parameters.....	512
28.7 Zeroization of Critical Security Parameters.....	512
28.8 Device Encryption Troubleshooting.....	512
28.8.1 CRYPTO MIB Trap Overview.....	512
28.8.2 Responses to Crypto MIB Traps.....	513
28.8.3 FRF.17, IPsec, and IKE Error Logging.....	516
28.9 CRYPTO Service Overview.....	519

28.9.1 Encapsulation Security Payload.....	520
28.10 Key Encryption Key Support.....	520
Chapter 29: IP Security Options.....	522
29.1 Configuring IP Security Parameters for End Systems.....	523
29.2 Configuring IP Security Options for Routers.....	524
29.2.1 Configuring Extended Security Option Labels.....	527
29.3 Verifying IP Security Options.....	528
29.4 ICMP Error Messages.....	528
29.5 Security Attacks Prevention on IP Routers.....	529
29.5.1 Secure Configuration Solutions.....	531
Chapter 30: L2Tunnel Connections Configuration.....	535
30.1 Configuring a PP Tunnel Terminator.....	535
30.2 Configuring a Router-to-Router Tunnel Initiator/Terminator.....	537
30.3 Configuring a PPTP Virtual Leased Line.....	539
30.4 Configuring Tunnel Switching.....	542
30.5 Tunnel Security.....	543
Chapter 31: Tunnel Route Shortcut (TRSC) Endpoints Configuration.....	545
31.1 NHRP.....	545
31.2 Configuring NHRP.....	546
31.2.1 Address and Routing Table Displays.....	549
Chapter 32: Internet Firewalls Configuration.....	551
32.1 Firewall Filter Types.....	552
32.1.1 Service-Independent Filters.....	552
32.1.2 Predefined (Service-Dependent) Filters.....	552
32.1.3 FTP Filters.....	553
32.1.4 TCP Stateful Filters.....	554
32.1.5 Generic Filters.....	554
32.2 Setting Up Internet Firewalls.....	555
32.2.1 Verifying the Firewall Configuration.....	558
32.2.2 Checking the Overall Status.....	559
32.2.3 Blocking Unwanted Traffic.....	560
32.2.4 User-Defined Services.....	561
32.2.5 IP Address Grouping.....	563
32.2.6 The RealPlayer Predefined Service.....	563
32.2.7 Enabling Security.....	566
32.3 Filters Management.....	566
32.3.1 Filter Definition with the ADD Filter Command.....	566
32.3.2 Filter Creation with an Off-line Editor.....	568
32.3.3 Filters Display and Deletion.....	569

32.3.4 Filters Assignment to Interfaces.....	569
32.3.5 Filters Activation and Deactivation.....	570
32.3.6 Firewall Filters Versus IP Filters.....	570
32.3.7 Execution Order for Traffic with Enabled Filters.....	570
32.4 System Logs Setup.....	571
Chapter 33: Access Control Service Configuration with RADIUS.....	574
33.1 Configuring RADIUS Authentication.....	574
33.2 RADIUS Parameters in the AC Service.....	575
Chapter 34: Lightweight Directory Access Protocol Configuration.....	576
34.1 Configuring LDAP for CCGW Clients.....	576
34.2 LDAP Working with CCGW Clients.....	577
34.3 Configuring LDAP for Firewall Clients.....	578
34.4 LDAP Working with Firewall Clients.....	579
34.4.1 Attributes and Object Classes.....	580
34.4.2 Adding Attributes.....	580
34.4.3 Creating an Object Class.....	581
34.4.4 Database Populating on the LDAP Server.....	581
34.4.5 Database Creation in LDAP Data Interchange Format.....	582
Chapter 35: GGSN Management Configuration.....	584
35.1 Basic GGSN Functionality.....	585
35.1.1 Data Flow.....	585
35.1.2 RADIUS Re-Authentication.....	588
35.1.3 RADIUS Authentication Server Switchover Algorithm.....	588
35.1.4 DHCP Server Switchover Algorithm.....	588
35.2 APN to Multiple IPIP Binding.....	589
35.3 HA GGSN Feature.....	590
35.4 Conventional Data Redundancy (HA and DSR) Feature.....	592
35.4.1 Additional Transport Network Redundancy.....	592
35.4.2 Mobile Agent Functionality.....	593
35.4.3 Support for Multi-Hop BFD Sessions Between FA and HA VRRP Addresses.....	594
35.4.4 Subscriber IP Address Allocation Through a DHCP Server on the Border Router.....	594
35.4.5 Replication Protocol.....	595
35.4.6 Conventional Channel Gateway Support for Conventional Data DSR.....	597
35.4.7 Support for Intersystem Data-Connected CEN.....	597
35.5 GGSN Charging.....	597
35.6 GGSN Dynamic DNS Update Feature.....	599
35.7 GGSN Hardware Redundancy.....	600
35.8 GGSN Configuration File Management.....	601
35.9 Configuring Basic GGSN Functionality on EOS Devices.....	601
35.10 Configuring Basic GGSN Parameters.....	602

35.11 Configuring APNs on the GGSN.....	602
35.11.1 Configuring APNs to Use Static Subscriber IP Address Allocation.....	604
35.11.2 Configuring APNs to Use Dynamic Subscriber IP Address Allocation with Local Server.....	607
35.11.3 Configuring APNs to Use Dynamic Subscriber IP Address Allocation with a DHCP Server.....	610
35.12 APN Configuration to Multiple IPIP Binding.....	614
35.13 Configuring Devices as HA GGSN Peers.....	614
35.13.1 Configuring GGSNs in the HA GGSN Sample Topology.....	616
35.14 Configuring Conventional Data Redundancy (HA and DSR) Feature.....	622
35.14.1 Configuring a GGSN to Support Conventional Data Redundancy (HA and DSR).....	622
35.14.1.1 Configuring the GGSN as a Backup VRRP Router.....	628
35.14.2 Configuring Border Routers to Support Conventional Data Redundancy (HA and DSR).....	629
35.14.2.1 Configure the Border Router as a Backup VRRP Router.....	635
35.15 Conventional Data Redundancy (HA and DSR) Configuration Example.....	636
35.15.1 Configuring the GGSNs in the Sample Conventional Data Redundancy (HA and DSR) Topology.....	643
35.15.2 Configuring Border Routers in the Sample Conventional Data Redundancy (HA and DSR) Topology.....	650
35.16 Configuring GGSN Charging.....	659
35.17 GGSN Configuration for Dynamic DNS.....	660
35.18 Configuring GGSN Hardware Redundancy for Dimetra Systems.....	661
35.18.1 GGSNs Preparation for Hardware Redundancy Configuration Through the EOS CLI.....	661
Chapter 36: Network Management Configuration.....	663
36.1 Simple Network Management Protocol.....	663
36.1.1 Configuring the SNMP Service for SNMPv1/v2c Management.....	663
36.1.2 Request Validation.....	664
36.1.3 Traps and Trap Profiles.....	664
36.1.3.1 Configuring Trap Profiles – Multiple Network Management Stations.....	667
36.1.3.2 Configuring Trap Profiles – Single Management Station.....	667
36.1.4 Set Operation.....	668
36.1.5 SNMPv2 GETBULK Support.....	668
36.1.6 Reports of Unconfigured Ports.....	668
36.2 SNMPv3 Support.....	668
36.2.1 SNMPv3 Configuration on EOS Devices.....	669
36.2.2 SNMPv3 Data Reset.....	669
36.2.3 Opening the SNMP V3 User Manager Menu.....	669
36.2.4 Predefined SNMPv3 USM Users Management on EOS Devices.....	670
36.2.4.1 Changing the MotoAdmin User Password.....	671
36.2.4.2 Changing the Privilege Level for the MotoMaster User.....	672
36.2.4.3 Inform Users.....	675
36.2.4.4 Creating MotoInformA (NoAuthNoPriv) Users Manually.....	676

36.2.4.5 Changing Credentials for Inform Users.....	677
36.2.5 General SNMPv3 USM Users Management on EOS Devices.....	679
36.2.5.1 Adding SNMPv3 USM Users.....	680
36.2.5.2 Deleting SNMPv3 USM Users.....	680
36.2.5.3 Changing Authentication Passphrases for SNMPv3 USM Users.....	681
36.2.5.4 Changing Encryption Passphrases for SNMPv3 USM Users.....	682
36.2.5.5 Viewing SNMPv3 USM Users.....	683
36.2.6 SNMPv3 Procedure Reference.....	684
36.3 Remote Network Monitoring Alarms.....	685
36.4 Network Maps.....	686
36.5 Statistics Retrieved Through SNMP.....	687
36.5.1 SNMPv2 GETBULK Support.....	687
36.5.2 Performance Statistics.....	687
36.5.3 Priority Queuing Statistics.....	690
36.5.4 PIM Statistics.....	690
36.5.5 Multicast IP (MIP) Statistics.....	691
36.5.6 BGP Statistics.....	691
36.6 Configuration Changes Logged with the AuditLog Service.....	691
36.7 Auditlog Files Viewing.....	693
36.8 Configuring Remote Logging of Syslog Messages.....	694
36.9 AuditLog Filters Management.....	695
36.10 Syslog Message Throttling.....	697
36.11 SNMP Event Notification Traps.....	698
36.11.1 Configuring SNMP for Notification Traps.....	698
36.12 Root and Admin Passwords Changes.....	699
36.12.1 Password Requirements.....	699
36.12.2 Password Complexity.....	700
36.12.3 Password Clearing.....	701
36.13 Executive Logon Timeout.....	702
36.14 System Remote Access.....	703
36.14.1 Displaying the Software Version with the REMote Command.....	704
36.14.2 Preventing Remote Access.....	705
36.14.3 Restricting Remote Access.....	705
36.14.4 Restricting Telnet Access.....	706
36.14.5 Restricting SSH Access.....	706
36.15 Resynchronization Feature for Encryption Devices.....	707
Chapter 37: Remote Polling Protocol Configuration.....	708
37.1 REMP Configuration.....	708
37.2 Target Group Priority.....	709
Chapter 38: Virtual Router Redundancy Protocol Configuration.....	710

38.1 Configuring the VRRP.....	710
38.1.1 Owner Router Enabling.....	711
38.1.2 Backup Routers Enabling.....	711
38.1.3 Priority Settings for Multiple Backup Routers.....	712
38.1.4 VRRP Enabling.....	714
38.1.5 VRRP Registration.....	714
38.1.6 Ping/Telnet Virtual Router IP (VIP).....	715
38.1.7 VRRP Disabling/Deleting.....	715
38.2 VRRP Customization.....	715
38.3 VRRP Scenarios.....	716
38.3.1 Configuring VRRP – Gateway to a WAN.....	716
38.3.2 Configuring VRRP – One Master Router Connecting Two LANs.....	716
38.3.3 Configuring VRR – Load Sharing with Redundancy.....	717
38.4 VRRP Traps Enabling and Disabling.....	718
Chapter 39: Scheduling and Event-Based Macro Execution.....	720
39.1 Schedules Creation.....	721
39.2 Scheduling WAN Connections.....	722
39.3 Event-Based Commands/Macros.....	722
39.3.1 Setting Up a Backup Port.....	723
39.3.2 Hanging Up a Port.....	723
39.3.3 Recovering from Port Loopback.....	723
39.3.4 VRRP Event-Triggered Macros Configuration.....	724
39.3.5 Configuring Bidirectional Forwarding Detection (BFD) with Event-Based Macros.....	724
Chapter 40: CCGW Functionality Configuration.....	726
40.1 Supported CCGW Features.....	726
40.2 CCGW Hardware Requirements.....	727
40.3 CCGW Port-to-Channel Mapping.....	729
40.4 CCGW LEDs.....	735
40.5 Configuring EOS Devices as a CCGW.....	735
40.6 CCGW Functionality.....	736
40.6.1 CCGW Identification.....	739
40.6.2 Conventional IV&D.....	739
40.6.3 IP Analog.....	740
40.6.4 IP MDC1200.....	740
40.6.5 IP Conventional.....	740
40.6.6 Mixed Mode.....	740
40.6.7 MDC 1200 Signaling.....	742
40.6.8 ACIM Link Channel Support.....	742
40.6.9 Advanced SECURENET Support.....	742
40.6.10 Relay Keying.....	743

40.6.11 CCGW Support for Conventional Data DSR.....	744
40.6.12 Conventional Talkgroups.....	744
40.6.13 Console Repeat.....	746
40.6.14 Main/Alternate Channel Status Awareness.....	746
40.6.15 Duplex Patch with Two-Party Conference.....	747
40.7 CCGW Troubleshooting.....	748
40.7.1 Inbound Emergency Call Error Detection Threshold Configuration.....	748
40.7.2 Inbound Call Statistics to Console Unknowns Comparison.....	749
40.7.2.1 Comparing the Inbound Call Unknown Start Count to Console Unknowns – Digital Channels.....	749
40.7.2.2 Comparing the Inbound Call Start Event Counts to Console Unknowns (MDC 1200 Channels).....	750
Chapter 41: Overload Protection Configuration.....	752
41.1 Configuring EOS Devices for Overload Protection.....	752
41.2 Overload Protection Statistics.....	753
Chapter 42: IP Link Converter (IPLC) Functionality Configuration.....	755
42.1 IPLC Feature Overview.....	755
42.1.1 Channel Type-to-CPort Type Mapping.....	762
42.1.2 V.24 and Analog Port Pairing.....	763
42.1.3 IPLC LEDs.....	764
42.1.4 IPLC Alarms.....	766
42.1.5 RemotePeerInfo Display.....	767
42.2 Configuring IPLC Functionality.....	769
42.2.1 Adding and Configuring Channels for IPLC Standalone Topology – IP-Digital (GGM 8000 with IPLC Functionality at Core).....	770
42.2.2 Adding and Configuring Channels for IPLC Standalone Topology – IP-Mixed.....	772
42.2.3 Adding and Configuring Channels for IPLC Paired Topology – IP-Mixed.....	775
42.2.4 Adding and Configuring Channels for IPLC Paired Topology – V24-Digital.....	779
42.2.5 Adding and Configuring Channels for IPLC Paired Topology – V24-Mixed.....	784
42.2.6 Creating a Configuration File with an IPLC Configuration Template.....	788
42.2.7 Downloading a Stored Configuration File to the GGM 8000.....	788
42.3 IPLC Configuration Templates.....	790
42.3.1 Configuring a Sample IPLC Standalone Topology – IP-Digital.....	791
42.3.2 Configuring a Sample IPLC Standalone Topology – IP-Mixed.....	792
42.3.3 Configuring a Sample IPLC Paired Topology – IP-Mixed (Core GGM 8000).....	793
42.3.4 Configuring a Sample IPLC Paired Topology – IP-Mixed (Site GGM 8000).....	797
42.3.5 Configuring a Sample IPLC Paired Topology – V24-Digital (Core GGM 8000).....	800
42.3.6 Configuring a Sample IPLC Paired Topology – V24-Digital (Site GGM 8000).....	801
42.3.7 Configuring a Sample IPLC Paired Topology – V24-Mixed (Core GGM 8000).....	802
42.3.8 Configuring a Sample IPLC Paired Topology – V24-Mixed (Site GGM 8000).....	803

42.3.9 Configuring a Sample IPLC Paired and Standalone Topology – Core GGM 8000 Supporting Multiple Channel Types.....	804
42.3.10 Configuring a Sample IPLC Paired Topology – V.24-Mixed (Site GGM 8000).....	806
42.3.11 Configuring a Sample IPLC Paired Topology – IP-Mixed (Site GGM 8000).....	807
Appendix A: Dial-Up Progress and Error Messages.....	809
A.1 Dial-up Progress and Error Messages for Modems.....	809
A.2 Dial-up Progress and Error Messages for EOS Devices.....	810
Appendix B: Loopback Testing.....	815
B.1 Dial-up Loopback Testing Using Modems.....	815
B.1.1 Performing a Local Loopback Test.....	817
B.1.2 Performing a Remote Loopback Test.....	819
B.2 Performing a Loopback Diagnostic Test for Built-In ISDN Ports.....	820
Appendix C: Internet Addressing.....	822
C.1 Internet Addresses and Classes.....	822
C.1.1 Dotted Decimal Notation.....	823
C.1.2 Addressing Rules.....	824
C.1.3 Sample Network Using the Class B Address Format.....	824
C.2 Subnet Addresses and Subnet Masks.....	825
C.2.1 Subnet Addressing.....	825
C.2.2 Subnet Masks.....	826
C.2.3 Establishing Subnets: Example 1.....	827
C.2.4 Establishing Subnets: Example 2.....	829
C.2.5 Establishing Subnets: Example 3.....	831
C.3 Variable Length Subnet Masks.....	833
Appendix D: Macro Features.....	834
D.1 Macro Conventions.....	834
D.2 Macros With Conditional Statements.....	834
D.2.1 Macro Variables.....	835
D.2.1.1 Macro Variable Types.....	836
D.2.1.2 Variables Comparison and Reassignment.....	839
D.2.1.3 Variable Substitutions.....	839
D.3 Control Structures.....	840
D.4 Keywords.....	841
D.5 Macro Caching and Shared Macros.....	842
D.6 Larger Macros.....	843
D.7 Macro Nesting.....	843
Appendix E: Supported MIBs and Traps.....	845
E.1 Supported Operations.....	845
E.2 Port Numbering Conventions in SNMP.....	845
E.3 MIBs Supported by EOS Devices.....	846

E.3.1 MNR NMA WAN Slot Object.....	847
E.3.2 MNR NMA LAN Path Object.....	849
E.3.3 MNR NMA Device Status Object.....	850
E.3.4 Ethernet WAN Link Object MIB.....	851
E.3.5 MNR Service ID MIB.....	854
E.3.6 MNR GGSN Fault Object.....	854
E.3.7 MNR Mobile IP Link Fault Object.....	857
E.3.8 MNR Mobile IP Service Fault Object.....	860
E.3.9 MNR DBReplicator Fault Object.....	862
E.3.10 MNR DHCP Service Fault Object.....	864
E.3.11 MNR CCGW Private MIBs.....	867
E.3.12 MNR NMA Private MIBs.....	868
E.3.13 Standard MIBs.....	869
E.3.14 SNMPv3 MIBs.....	869
E.3.15 MNR Private MIBs.....	870
E.4 Traps Supported on EOS Devices.....	871
E.4.1 EOS Non-Enterprise-Specific Traps.....	872
E.4.2 EOS Enterprise-Specific Traps.....	872
E.4.3 EOS Derived MIB-2 Enterprise-Specific Traps.....	876
E.4.4 Trap Delay Queueing and Event-Based Trap Suppression.....	877
E.4.5 SysName and Trap Event Timestamp in EOS Enterprise-Specific Traps.....	878
E.4.6 Support for SNMPv2 64-bit Counters.....	878
E.4.7 SNMPv3 and SNMPv1 Trap Generation.....	879
E.4.8 S6000 and GGM 8000 Supply System / Hardware Fault Monitoring Traps.....	879
E.4.9 CCGW Fault Traps.....	880
E.4.10 Trap Descriptions.....	888
Appendix F: Statistics Displays.....	967
F.1 ARP Service Statistics.....	967
F.2 BGP Service Statistics.....	968
F.3 BRIDGE Service Statistics.....	968
F.4 CRYPTO Service Statistics.....	969
F.5 DVMRP Service Statistics.....	970
F.6 FR Service Statistics.....	971
F.7 GGSN Service Statistics.....	971
F.7.1 GTP Statistics for the GGSN Service as a Whole.....	972
F.7.2 APN Statistics.....	972
F.7.3 RADIUS Statistics.....	973
F.7.4 GGSN Charging Statistics.....	974
F.8 IP Service Statistics.....	974
F.9 IPv6 Service Statistics.....	978

F.10 MIP Service Statistics.....	979
F.11 MOBILEIP Service Statistics.....	980
F.12 MOSPF Service Statistics.....	984
F.13 NAT Service Statistics.....	985
F.14 OSPF Service Statistics.....	986
F.15 PATH Service Statistics.....	986
F.16 PIM-SM Service Statistics.....	987
F.17 PORT Service Statistics.....	989
F.18 PPP Service Statistics.....	990
F.18.1 PPP Over Ethernet Statistics.....	991
F.19 RIPIP Service Statistics.....	991
F.20 RSVP Service Statistics.....	992
F.21 SNMP Service Statistics.....	992
F.22 STP Service Statistics.....	994
F.23 TCP Service Statistics.....	994
F.24 UDP Service Statistics.....	994
F.25 UDPHELP Service Statistics.....	995
Appendix G: Static Tables.....	996
Appendix H: Syslog Messages.....	997
H.1 SSH "Login Failed" Messages.....	1024
Appendix I: Regular Expressions.....	1025
I.1 Regular Expression Components.....	1025
I.2 AS Filter Examples.....	1027
I.3 GREP Command Examples.....	1027
Appendix J: Wide Area Network Setup Information.....	1030
J.1 Serial Port Utilization Percentage.....	1030
J.2 Serial Line Connectivity.....	1030
J.3 Network Clock Synchronization.....	1031
J.3.1 Source Quality Clocking.....	1031
Appendix K: Delay and Jitter Performance Management Tool.....	1032
K.1 Delay and Jitter Statistics View.....	1032
K.2 Delay and Jitter Performance Management Tool Functionality.....	1032
K.2.1 Latency (IPTD) Calculations.....	1033
K.2.2 Jitter (IPDV) Calculations.....	1033
K.2.3 Packet Loss Rate (IPLR) Calculations.....	1034
Appendix L: Upgrading EOS Software.....	1035
Appendix M: Downgrading EOS Software.....	1038
Appendix N: Glossary.....	1040

List of Figures

Figure 1: Relationship Between Connectors, Paths, and Ports.....	50
Figure 2: Possible Path-to-Port Assignments.....	51
Figure 3: S2500 Path-to-Port Mappings.....	52
Figure 4: S2500 Path-to-Port Mappings (V.24 Module in I/O Slot A).....	53
Figure 5: S6000 Path-to-Port Mappings (FlexWAN Modules in Both I/O Slots).....	54
Figure 6: S6000 Path-to-Port Mappings (UltraWAN Modules in Both I/O Slots).....	55
Figure 7: S6000 Path-to-Port Mappings (UltraWAN II Modules in Both I/O Slots).....	56
Figure 8: S6000 Path-to-Port Mappings (T3/E3 Modules in Both I/O Slots).....	57
Figure 9: S6000 Path-to-Port Mappings (FlexWAN Module in Slot A and UltraWAN II Module in Slot B).....	58
Figure 10: S6000 Path-to-Port Mappings (UltraWAN II Module in Slot A and FlexWAN Module in Slot B).....	59
Figure 11: S6000 Path-to-Port Mappings (FlexWAN Module in Slot A and T3/E3 Module in Slot B).....	60
Figure 12: S6000 Path-to-Port Mappings (T3/E3 Module in Slot A and FlexWAN Module in Slot B).....	61
Figure 13: S6000 Path-to-Port Mappings (UltraWAN II Module in Slot A and T3/E3 Module in Slot B).....	62
Figure 14: S6000 Path-to-Port Mappings (T3/E3 Module in Slot A and UltraWAN II Module in Slot B).....	63
Figure 15: S6000 Path-to-Port Mappings (12-Port T1/E1 Modules in Both I/O Slots).....	64
Figure 16: GGM 8000 Path-to-Port Pairings (Base Unit).....	65
Figure 17: GGM 8000 Path-to-Port Pairings (Expansion Module with Analog/V.24 Interface Kit).....	65
Figure 18: GGM 8000 Path-to-Port Pairing (FlexWAN Daughterboards in Expansion Module).....	66
Figure 19: GGM 8000 Path-to-Port Pairings (Low Density Enhanced Conventional Gateway Module).....	67
Figure 20: GGM 8000 Path-to-Port Pairings (High Density Enhanced Conventional Gateway Module).....	68
Figure 21: MCG 8000 Path-to-Port Pairings (Base Unit).....	69
Figure 22: MCG 8000 Port and Path Numbering (High Density Enhanced Conventional Gateway Module).....	70
Figure 23: Parent Port and Virtual Port.....	77
Figure 24: Virtual Ports on Both Sides of a Frame Relay Network.....	78
Figure 25: Ethernet WAN Links: Sample Single Site Link Configuration.....	81
Figure 26: Ethernet WAN Links: Sample Interzone Link Configuration.....	83
Figure 27: IPv6 Ethernet WAN Links: Sample Single Site Link Configuration.....	86
Figure 28: IPv6 Ethernet WAN Links: Sample Interzone Link Configuration.....	90
Figure 29: Multicasting Over VLANs.....	94
Figure 30: Autostartup Configuration for Frame Relay.....	101
Figure 31: Autostartup Configuration for PPP.....	101
Figure 32: Autostartup with Motorola Solutions Central Node and Frame Relay WAN.....	105
Figure 33: Autostartup with Central Node BootP Server and PPP WAN.....	106
Figure 34: Wide Area Device Configuration.....	123
Figure 35: Two IP Subnets Configured on the Same Ethernet.....	126
Figure 36: Routing Between Gateways.....	127

Figure 37: Adding a Route Statically in a Subnet Masked Environment.....	128
Figure 38: OSPF Meshed Network Example.....	130
Figure 39: IP Packet Filtering.....	133
Figure 40: TCP Header of a Compliant Packet.....	134
Figure 41: Migration to a RIPV2 Network.....	143
Figure 42: Route Aggregation/Deaggregation With RIPv2.....	144
Figure 43: Range Table Mask.....	146
Figure 44: OSPF Routing Policies.....	149
Figure 45: OSPF Route Aggregation.....	150
Figure 46: ISIS Routing Policies.....	151
Figure 47: BGP Network Topology.....	153
Figure 48: BGP Default Route.....	154
Figure 49: BGP Route Aggregation.....	155
Figure 50: Wide Area Router Connecting Four IP Networks.....	167
Figure 51: Typical IP Network Running OSPF.....	168
Figure 52: Recommended Autonomous System Topology.....	172
Figure 53: Route Advertisement Over Non-meshed Frame Relay Network.....	174
Figure 54: Route Advertisement Over Meshed Frame Relay Network.....	175
Figure 55: Different States of a RIP Route.....	176
Figure 56: Autonomous System With Multifunctional OSPF Devices.....	178
Figure 57: Backbone Area With One Isolated ABR.....	179
Figure 58: Forming Adjacencies on a Multiaccess Network.....	180
Figure 59: Internal and External Peers in an Autonomous System.....	184
Figure 60: Path Attributes.....	186
Figure 61: AS-PATH Attribute Example.....	186
Figure 62: NEXT-HOP Router.....	187
Figure 63: BGP Routing Metric.....	188
Figure 64: BGP Routing Selection and Policies.....	190
Figure 65: Sample CBQ Class Hierarchy.....	203
Figure 66: RSVP Proxy Sender with a Unicast Destination and One Sender Port.....	216
Figure 67: RSVP Proxy Receiver with a Unicast Destination and One Sender.....	217
Figure 68: RSVP for L2TP Tunnel.....	218
Figure 69: Different States of a RIPv2 Route.....	241
Figure 70: Outbound Traffic.....	247
Figure 71: Inbound or Load Sharing Traffic.....	247
Figure 72: Bidirectional Traffic.....	248
Figure 73: IPCP Address Mapping Network Diagram.....	250
Figure 74: Private Network Connecting to the Corporate Network.....	252
Figure 75: Load Sharing.....	254

Figure 76: Address Migration.....	255
Figure 77: Address Redirection.....	256
Figure 78: Basic IP Multicasting Configuration.....	269
Figure 79: Multicast Routing Topologies Troubleshooting.....	272
Figure 80: IGMP Proxy Configuration Example.....	275
Figure 81: Threshold Adjustment.....	277
Figure 82: Multicast Tunnel Configuration.....	278
Figure 83: DVMRP Scoping.....	279
Figure 84: Multicasting over Frame Relay.....	280
Figure 85: Controlling the Rate Limit.....	281
Figure 86: DVMRP Route Aggregation.....	282
Figure 87: Interarea Multicasting.....	286
Figure 88: MOSPF Destination Group Filtering.....	287
Figure 89: PIM-SM Sample Configuration.....	300
Figure 90: Routers Participating in RDP.....	304
Figure 91: Congestion Control Configuration Example.....	311
Figure 92: Routing Multiple Protocols over Frame Relay Network.....	314
Figure 93: Configuring IP over Frame Relay.....	315
Figure 94: Fully Redundant Network Among All Sites.....	322
Figure 95: Configuring Dual PVCs Over a Single Logical Port.....	324
Figure 96: SVC and PVC Connections over Frame Relay.....	326
Figure 97: Fully Meshed Frame Relay Topology.....	327
Figure 98: Non-meshed Frame Relay Topology.....	328
Figure 99: Transparent Bridging Unsuitable in some Non-meshed Topologies.....	328
Figure 100: Partially Meshed Frame Relay Topology.....	329
Figure 101: Frame Relay Addressing Example.....	330
Figure 102: Partially Redundant Network with Redundant Link at Site B.....	336
Figure 103: Partially Redundant Network with Redundant Link at Site A.....	336
Figure 104: Partially Redundant Network: Redundant Links at Site A and Site B.....	337
Figure 105: Fully Redundant Network Between Site A and Site B.....	337
Figure 106: Detecting Loopback of Packets with Magic Numbers.....	342
Figure 107: One Path Per Port Configuration.....	344
Figure 108: Two Paths Per Port Configuration.....	344
Figure 109: Use of Resources through the Virtual Pipe.....	349
Figure 110: Basic Bandwidth Management Port and Path Configurations.....	363
Figure 111: PPP MLP Load Balancing.....	376
Figure 112: IP Network Design Example: IP over a DOD Link.....	377
Figure 113: IP and Telnet Traffic Contention (No Protocol Reservation).....	384
Figure 114: IP and Telnet Traffic Contention (With Protocol Reservation).....	385

Figure 115: Hardware Configuration for Bridged Packets Example.....	387
Figure 116: Hardware Configuration for IP-Routed Packets Example.....	388
Figure 117: IP Protocols Allocation (UDP Included).....	390
Figure 118: IP Protocols Allocation (UDP Excluded).....	390
Figure 119: Protocol Reservation on Virtual Ports.....	391
Figure 120: Prioritizing and Forwarding Packets to Queues.....	400
Figure 121: Spanning Tree Structure.....	412
Figure 122: Network with a Loop.....	412
Figure 123: Root Bridge.....	414
Figure 124: Root Port.....	414
Figure 125: Selecting a Designated Bridge – Single Bridge.....	415
Figure 126: Selecting a Designated Bridge – Multiple Bridges Attached to Network.....	416
Figure 127: Two Wide Area Bridges Connected to Imaginary Backbone.....	417
Figure 128: Bridge Learning.....	419
Figure 129: Network Configuration After Host C is Moved.....	420
Figure 130: Effects of FireWall on Bridging and Routing.....	421
Figure 131: Source Explicit Forwarding.....	427
Figure 132: Source Explicit Blocking.....	428
Figure 133: Destination Explicit Forwarding.....	429
Figure 134: Destination Explicit Blocking.....	430
Figure 135: Combining Bridge Security Features.....	431
Figure 136: Network Showing Station Groups.....	435
Figure 137: Filter Selection Process.....	437
Figure 138: Filter Qualification and Action Processes.....	438
Figure 139: Discarding IP Packets.....	442
Figure 140: UDP Broadcast Helper Topology Example.....	447
Figure 141: UDP Broadcast Helper Configuration.....	449
Figure 142: UDP Broadcast Helper for BOOTP Configuration.....	451
Figure 143: Hops Limitation for BOOTPREQUEST Packets.....	453
Figure 144: Determining Which Clients are Booted First.....	454
Figure 145: Basic VPN Connection.....	457
Figure 146: Basic Tunnel Switching Components.....	458
Figure 147: Remote/Mobile Client Dial-up RAS VPN.....	460
Figure 148: Site-to-Site VPNs.....	461
Figure 149: ISP to Central Site Tunnel.....	463
Figure 150: Remote Workstation to Central Site Tunnel.....	463
Figure 151: Remote Office Tunnel.....	465
Figure 152: Distributed Remote Access.....	467
Figure 153: GGSN Topology: GGSN and CBR.....	484

Figure 154: GGSN Topology: GGSN and CBR (X509 Certificates).....	486
Figure 155: CWR Topology (FRF17).....	488
Figure 156: Fully Meshed Topology Between Three Routers (Dynamic Key).....	492
Figure 157: IPv4 in IPv6 Tunnel with IKEv2 (Dynamic Key).....	496
Figure 158: Hub and Spoke Topology Between Three Routers.....	499
Figure 159: IP Security Options Configuration Example.....	524
Figure 160: Spoofed Source IP Address Attack.....	529
Figure 161: CERT Recommended Filters.....	530
Figure 162: Alternative CERT Configuration.....	531
Figure 163: Noncontiguous IP Networks.....	531
Figure 164: Noncontiguous IP Networks (Alternative Topology).....	532
Figure 165: Subnets on the Internal Network.....	532
Figure 166: Multiple Contiguous IP Networks.....	533
Figure 167: Two-Router Configuration.....	533
Figure 168: Internet Firewall.....	534
Figure 169: PPTP/L2TP Tunnel Connections Between Router and Hub.....	535
Figure 170: Virtual Leased Line and Backup Connection.....	540
Figure 171: Configuring Tunnel Switching.....	542
Figure 172: Central Site VPN with Two Remote Sites.....	546
Figure 173: Firewall Setup Example.....	556
Figure 174: User-Defined Service Example.....	562
Figure 175: RealAudio Protocol Sequence.....	564
Figure 176: GGSN Functionality – Data Flow and IP Address Allocation	587
Figure 177: GGSN Charging System Architecture.....	598
Figure 178: Sample HA GGSN Topology.....	617
Figure 179: Sample Conventional Data Redundancy (HA and DSR) Topology with Backup VRRP Functionality Configured on the GGSNs in the Data and ZNM Subnets in Both Zones.....	638
Figure 180: Sample Conventional Data Redundancy (HA and DSR) Topology: IPIP Tunnels from GGSNs to BRs on Data Subnet.....	639
Figure 181: Sample Conventional Data Redundancy (HA and DSR) Topology: IPIP Tunnels from GGSNs to BRs on ZNM Subnet.....	640
Figure 182: Sample Conventional Data Redundancy (HA and DSR) Topology: Mobility Agents.....	641
Figure 183: Sample Conventional Data Redundancy (HA and DSR) Topology: Roamed Subscriber.....	642
Figure 184: Sample Conventional Data Redundancy (HA and DSR) Topology: DSR Failure Scenario.....	643
Figure 185: Accessing the System Remotely.....	704
Figure 186: VRRP Priorities.....	713
Figure 187: Owner Router with One Backup Router.....	716
Figure 188: Two VRIDs on One Router Connecting Two LANs.....	717
Figure 189: Load Sharing with Redundancy.....	718
Figure 190: MCG 8000 CCGW Port Numbering and Port-to-Channel Mapping.....	731

Figure 191: GGM 8000 Low Density CCGW Port Numbering and Port-to-Channel Mapping.....	732
Figure 192: GGM 8000 High Density CCGW Port Numbering and Port-to-Channel Mapping.....	733
Figure 193: MNR S2500 CCGW Port Numbering and Port-to-Channel Mapping (Analog).....	734
Figure 194: MNR S2500 CCGW Port Numbering and Port-to-Channel Mapping (Digital with V.24 in Slot A)..	734
Figure 195: MNR S2500 CCGW Port Numbering and Port-to-Channel Mapping (Digital with V.24 in Slot B)..	735
Figure 196: Standalone IPLC Topology: IP-Digital (GGM 8000 with IPLC functionality located at core).....	758
Figure 197: Standalone IPLC Topology: IP-Mixed.....	759
Figure 198: Paired IPLC Topology: IP-Mixed.....	760
Figure 199: Paired IPLC Topology: V24-Digital.....	761
Figure 200: Paired IPLC Topology: V24-Mixed.....	762
Figure 201: Low Density Enhanced Conventional Gateway Module Analog LEDs.....	764
Figure 202: High Density Enhanced Conventional Gateway Module Analog LEDs.....	764
Figure 203: Low Density Enhanced Conventional Gateway Module V.24 LEDs.....	765
Figure 204: High Density Enhanced Conventional Gateway Module V.24 LEDs.....	765
Figure 205: Typical Dial or Leased Connection.....	815
Figure 206: HSS V.35 and RS-232 Local Loopback Flowchart.....	817
Figure 207: HSS V.35 and RS-232 Remote Loopback Flowchart.....	819
Figure 208: ISDN Loopback Testing.....	820
Figure 209: Sample Network Using the Class B Address Format.....	824
Figure 210: Subnet Addressing.....	825

List of Tables

Table 1: S2500 Path/Port Syntax Rules.....	71
Table 2: S6000 Path/Port Syntax Rules.....	71
Table 3: S6000 UltraWAN Connector Syntax Rules.....	72
Table 4: GGM 8000 Path/Port Syntax Rules.....	72
Table 5: MCG 8000 Path and Port Syntax.....	73
Table 6: Number of Virtual Ports Supported per Hardware Platform.....	77
Table 7: Topologies that Require Virtual Ports.....	78
Table 8: Connectivity in Partially Meshed and Non-Meshed Topologies.....	79
Table 9: MIL Syntaxes for Instances and Their Meanings.....	96
Table 10: MIL Syntax for Instances and Group and Their Meaning.....	96
Table 11: RIP Routing Policies.....	140
Table 12: OSPF Routing Policies.....	148
Table 13: IISIS Routing Policies.....	151
Table 14: Route Precedence.....	168
Table 15: Routing Table Containing Multiple Paths.....	170
Table 16: Route Selection With and Without Load Splitting.....	171
Table 17: Configuring the EOS Device for RIP Updates Using RIPv2.....	175
Table 18: OSPF Configuration Parameters.....	182
Table 19: Information Added to the Route Table by System IP.....	212
Table 20: System IP with Same Subnet as an IP Address of an Active IP Interface.....	213
Table 21: System IP Different from All Active IP Interface IP Addresses.....	214
Table 22: RIPv2 Routing Policies.....	236
Table 23: Route Precedence.....	238
Table 24: Dynamic Many-to-One Mapping.....	248
Table 25: Mask Examples.....	249
Table 26: TCP/UDP Port Numbers and Services.....	251
Table 27: IGMP TOS Values by Packet Type.....	276
Table 28: Routing Table Items.....	284
Table 29: Forwarding Table Items.....	285
Table 30: Static RP Errors Reported by the EOS SHOW -PIM StaticRP Summary Command.....	297
Table 31: Priority Queue-to-TOS Mapping.....	320
Table 32: Values Used by EOS in IP Header Compression Negotiation.....	321
Table 33: Congestion Control Throughput Rate Reduction.....	334
Table 34: Covered and Uncovered Links – Partially Redundant Network with Redundant Link at Site B.....	336
Table 35: Covered and Uncovered Links – Partially Redundant Network with Redundant Link at Site A.....	336

Table 36: Covered and Uncovered Links – Partially Redundant Network: Redundant Links at Site A and Site B.....	337
Table 37: Path Configurations Available with System Bandwidth Management.....	351
Table 38: Communication Resources Supported by Enterprise OS Port Bandwidth Management Feature....	353
Table 39: Valid Port and Path Configurations.....	354
Table 40: Modem Initialization Strings.....	362
Table 41: Bandwidth Management Tasks and Commands	378
Table 42: Packet Types and Configuration Procedures for Protocol Reservation.....	385
Table 43: Traffic Type, Configured Bandwidth, and Normalized Bandwidth.....	392
Table 44: Packets with System-Assigned Priorities.....	398
Table 45: -PORT QueueInterLeave Parameter Values and 8-Slot Patterns.....	401
Table 46: Data Prioritization Parameters.....	404
Table 47: Physical T1/E1 Connectors and S6000 Logical Paths.....	408
Table 48: SHow BRidge CONFiguration Display Example.....	422
Table 49: Filter Service Parameters.....	436
Table 50: Built-in Bridge Masks.....	439
Table 51: User-Defined Bridge Masks and DataLink Locations	440
Table 52: Supported Service Name-to-UDP Port Mappings.....	448
Table 53: Well-Known Tunneling Ports.....	462
Table 54: MD5 Authentication States after Reboot.....	502
Table 55: Crypto MIB Trap Troubleshooting.....	513
Table 56: FRF.17 Syslog Messages and Reasons for Packet Discard.....	517
Table 57: IPSec Syslog Messages and Reasons for Packet Discard.....	517
Table 58: IKE Syslog Messages and Reasons for Packet Discard.....	518
Table 59: Description of the Fields in a Log.....	572
Table 60: APN to Multiple IPIP Tunnel Assignment Rules.....	590
Table 61: Predefined VACM Groups Supported on EOS Devices.....	684
Table 62: AuditLog Display Message Fields.....	692
Table 63: Allowable Channel IDs per CCGW Hardware Platform, Site Type, and Channel Type.....	727
Table 64: MCG 8000 Port-To-Channel Mapping.....	729
Table 65: CCGW LED Status Indicators (All Supported Channels In Service).....	735
Table 66: Channel Type-to-CPort Type Mapping (Standalone Channel at Core).....	762
Table 67: Channel Type-to-CPort Type Mapping (Standalone Channel at Site).....	763
Table 68: Channel Type to Upstream and Downstream CPort Type Mapping (Paired IPLC Topologies).....	763
Table 69: V.24 and Analog Port Pairing.....	763
Table 70: GGM 8000-to-MOSCAD NFM XC RTU Pinouts.....	766
Table 71: GGM 8000-to-MOSCAD SDM 3000 RTU Pinouts.....	766
Table 72: Modem Response Codes.....	810
Table 73: ISDN Dial Failure Cause Codes.....	810

Table 74: RS-232 Loopback Pin Assignments.....	818
Table 75: V.35 Loopback Pin Assignments.....	818
Table 76: Sample Subnets with Binary and Decimal Equivalents (Example 1).....	827
Table 77: Sample Subnets with Binary and Decimal Equivalents (Example 2).....	829
Table 78: Sample Subnets with Binary and Decimal Equivalents (Example 3).....	831
Table 79: Variables within Macros	835
Table 80: System/User Information Variables.....	837
Table 81: Comparison Operators.....	839
Table 82: Macro Cache Aging.....	842
Table 83: Mapping UI to SNMP Port Numbers (MNR S2500).....	845
Table 84: Mapping UI Port Numbers to SNMP Port Numbers (S6000 Routers).....	846
Table 85: MNR WAN Slot Object MIB States.....	848
Table 86: MNR LAN Path Object MIB States.....	849
Table 87: Device Status Object State Definitions.....	850
Table 88: Device Status Object State Change Trap Definition.....	850
Table 89: MNR Ethernet WAN Link Object MIB State Definitions.....	851
Table 90: MNR Ethernet WAN Interface Fault Object State Change Triggers.....	852
Table 91: Ethernet WAN Interface IPDV Alarms.....	854
Table 92: MNR GGSN Fault Object State Definitions.....	855
Table 93: MNR GGSN Fault Object Trap Definitions.....	855
Table 94: Mobile IP Link Fault Object State Definitions.....	857
Table 95: Mobile IP Link Fault Object Trap Definitions.....	858
Table 96: Mobile IP Service Fault Object State Definitions.....	860
Table 97: Mobile IP Service Fault Object Trap Definition.....	861
Table 98: DBReplicator Fault Object State Definitions.....	863
Table 99: DBReplicator Fault Object Trap Definition.....	863
Table 100: DHCP Service Fault Object State Definitions.....	864
Table 101: DHCP Service Fault Object Trap Definitions.....	865
Table 102: MNR CCGW Private MIBs.....	867
Table 103: MNR NMA Private MIBs.....	868
Table 104: Motorola Solutions Non-CCGW Private MIBs.....	870
Table 105: EOS Non-Enterprise-Specific Traps.....	872
Table 106: EOS Enterprise-Specific Traps.....	873
Table 107: EOS Derived MIB-2 Enterprise-Specific Traps.....	877
Table 108: CCGW Fault Traps.....	880
Table 109: CCGW Box State Definitions.....	886
Table 110: CCGW-to-Zone Controller State Definitions.....	886
Table 111: CCGW-to-Zone Controller Redundant Control Path State Definitions.....	886
Table 112: CCGW Conventional Channel State Definitions.....	887

Table 113: CCGW RNG Link State Definitions.....	887
Table 114: CCGW RNG Broadcast Link State Definitions.....	887
Table 115: Details about EOS Traps Supported on EOS Devices.....	888
Table 116: T1/E1 Line Statuses Represented by LineStatus Bit Positions.....	962
Table 117: Loopbacks Indicated by LoopbackType Variable.....	962
Table 118: Encryption-Specific Trap Variable Descriptions.....	963
Table 119: Number of Entries Allowed in Static Tables.....	996
Table 120: Syslog Messages.....	997
Table 121: Regular Expression Components.....	1025

List of Processes

Customizing Routing Configuration	125
Verifying the Configuration	232
Customizing the IPv6 Routing Configuration	234
Setting Up The Frame Relay Service	308
Configuring Disaster Recovery in a Frame Relay Environment	322
Configuring WAN Resources	362
Configuring the Central Site of a Remote/Dialup VPN	464
Configuring the Central Site of a RAS VPN	464
Configuring the Remote Office Route in a Site-to-Site VPN	465
Configuring the Central Site Device in a Site-to-Site VPN	466
Setting Up Internet Firewalls	555
Configuring Basic GGSN Functionality on EOS Devices	601
Configuring GGSNs in the HA GGSN Sample Topology	616
Configuring Conventional Data Redundancy (HA and DSR) Feature	622
Configuring a GGSN to Support Conventional Data Redundancy (HA and DSR)	622
Configuring the VRRP	710

List of Procedures

Configuring Local Area Interfaces	73
Configuring Wide Area Interfaces	74
Configuring Ethernet Site Links by Using IPIP with BFD	80
Configuring Ethernet Interzone Links by Using IPIP with BGP	83
Configuring Ethernet Site Links by Using IPIPv6 with BFD	86
Configuring Ethernet Interzone Links by Using IPIPv6 with BGP	89
Configuring Frame Relay and PPP Virtual Ports	93
Configuring Virtual Ports for 802.1Q Virtual LANs	94
Determining the Progress of the Establishment of the Physical and Data Links	99
Preparing Autostartup	100
Configuring the Central Node and the BootP/TFTP Servers on the Central Site Network	101
Autostarting the Central Site Node	104
Setting Up Configuration with Frame Relay WAN	104
Setting Up Configuration with PP WAN	106
Configuring a Device by Using ASCII Boot	109
Configuring Basic IP Routing over LAN Points and PPP Links	119
Verifying the IP Routing Configuration	122
Checking the Forwarded Packets with the PING Command	123
Verifying the Overall Router Status	124
Configuring Multiple IP Networks/Subnets	126
Configuring Packet Filtering	131
Setting Up the Filtering Options for IP Packet Filtering	136
Configuring RIP Routing Policies	140
Configuring OSPF Routing Policies	147
Controlling the Reporting of Routes Learned from Other Sources and Advertised into the OSPF Domain	149
Configuring OSPF Route Aggregation	150
Configuring IISIS Routing Policy	150
Configuring BGP Peers	152
Configuring BGP Route Aggregation	154
Importing Routes from an IGP Domain to a BGP Domain	156
Configuring Network Number Policies	158
Configuring AS-Path Permit or Deny Policies	159
Configuring AS-Path Weight Policies	161
Configuring MED/LOCAL-PREF with a Network Policy	164
Configuring MED/LOCAL-PREF with an OSPF Export Policy	165
Configuring MED/LOCAL-PREF with an Interior Policy	165

Configuring IPQoS Based on Priority Queueing for IPIP and IPSec Interfaces	195
Configuring IPQoS Based on Priority Queueing for Physical and VLAN Interfaces	197
Interpreting Buffer Overrun Errors When IPQoS Priority Queueing Is Enabled	198
Configuring Rate Limiting with IPQoS Policy	199
Configuring Rate Limiting – Example 1	201
Configuring Rate Limiting – Example 2	201
Configuring IPQoS Based on CBQ	203
Configuring IP/TOS to CBQ Class Mapping	205
Configuring IPQoS Based on Protocol Reservation	206
Configuring 802.1P-to-TOS Mapping on an Ingress VLAN Port	208
Configuring TOS-to-802.1P Mapping on an Egress VLAN Port	209
Configuring TOS-to-IPIP Mapping	210
Defining a Default IPQoS Policy	210
Configuring the RSVP	215
Configuring RSVP with L2TP Tunnel	218
Configuring DHCP on EOS Devices	220
Changing DHCP Profile Configuration Values	223
Configuring IPv6 for Local Area Networks	228
Configuring IPv4 over IPv6 Tunnels	229
Configuring IPv6 over IPv4 Tunnels with an Encrypted Payload	230
Configuring IPv6 over IPv6 Tunnels	231
Configuring IPv6 over IPv6 Tunnels with an Encrypted Payload	231
Examining Network Devices	233
Retrieving IPv6 Statistics	233
Checking the Overall Status	233
Checking with PING6, TraceRoute6, and TELnet6	234
Configuring the Internal Port	234
Configuring Multiple IPv6 Subnets	235
Configuring Neighbor Discovery	235
Configuring Static Routes	235
Configuring RIPNG Routing Policies	236
Configuring RIPNG Route Aggregation	237
Configuring Network Address Translation	244
Mapping the IPCP Address	249
Configuring Private Address Space	251
Configuring Load Sharing	254
Configuring Address Migration	255
Configuring Address Redirection	256
Configuring IP Multicast Routing for Local Area Networks and Point-to-Point Links	268

Verifying the IP Multicast Routing Configuration Overall Status	270
Viewing IP Multicasting Statistics	271
Controlling Local Group Membership Queries	274
Configuring IGMP Proxy	274
Configuring IGMP Proxy with Source Address Translation	274
Configuring a DVMRP Multicast Tunnel	277
Configuring DVMRP Scoping	278
Configuring DVMRP Multicasting over Frame Relay	280
Configuring a DVMRP Metric	281
Controlling the DVMRP Rate Limit for Multicast Traffic	281
Configuring DVMRP Route Aggregation	282
Controlling the Routing Table	283
Controlling the Forwarding Table	284
Configuring Interarea Multicasting	285
Configuring MOSPF Routing Policies	286
Configuring MOSPF Forwarding Policies	287
Configuring Multicast Border Router Functionality	288
Configuring PIM-SM	299
Setting Up RDP	304
Verifying the RDP Configuration	306
Troubleshooting the RDP Configuration	306
Configuring the Frame Relay Service	308
Configuring Congestion Control for NTTLMI Protocol Users	309
Configuring Congestion Control for Other LMI Protocol Users	310
Configuring Frame Relay SVC Connections	311
Verifying the Frame Relay Configuration	313
Configuring Transparent Bridging over Frame Relay	313
Configuring IP over Frame Relay	314
Adding an Aggregated Virtual Circuit	317
Configuring Realtime Priority Queueing	318
Configuring FRF.12 Fragmentation	319
Configuring a Primary PVC	323
Configuring a Backup PVC	323
Configuring a Backup Link	324
Establishing an SVC Connection	326
Setting Up PAP	345
Setting Up CHAP	346
Verifying Your PAP or MPPE Configuration	347
Setting Up EAP	347

Activating LAPB to Reduce Noisy Lines	348
Configuring Wide Area Networking with Async PPP	357
Setting Up a Path for Async PPP and AT Dial	357
Initializing the Modem	358
Adding a Number to the Phone Number List	359
Configuring Dial-Up Lines with A Modem Or TA	363
Configuring Leased Lines	365
Configuring ASPC Leased Lines	366
Configuring System Bandwidth Management Mode (DOD)	366
Configuring Bandwidth-on-Demand	367
Configuring the Dial List	367
Converting a Static Path to a Dynamic Path	370
Changing a Dynamic Path to a Static Path	370
Identifying Remote Sites with SCID and CLIP	370
Configuring the Path Preference List	371
Configuring Manual Bandwidth Management Mode	374
Verifying the Bandwidth Management Configuration	375
Troubleshooting the Bandwidth Management Configuration	376
Configuring Routing over DOD Links	377
Configuring Protocol Reservation for Bridged Traffic	387
Configuring Protocol Reservation for IP-Routed Packets	388
Configuring Tinygram Compression	395
Configuring History-Based or Per-Packet Compression	395
Enabling LAPB for a PPP Link	396
Verifying the Link-Level Compression Effectiveness	396
Configuring Data Prioritization	402
Configuring UDP/LAPB Packet Relay on MNR S6000 Routers	405
Configuring MNR S6000 Routers as CWR Peers	407
Configuring Basic Bridging	420
Verifying the Bridge Configuration	422
Getting Statistics	423
Troubleshooting the Bridge Configuration	424
Configuring Per-Port Transparent Bridging	425
Adding or Deleting Static Entries	426
Configuring Source Explicit Forwarding	426
Configuring Source Explicit Blocking	427
Configuring Destination Explicit Forwarding	428
Configuring Destination Explicit Blocking	429
Optimizing Bridge Performance	432

Configuring Filters	433
Configuring a Filter Between Station Groups	434
Creating a Mnemonic Filter by Using the PROToColRsrv TAG Action Option	445
Reserving Bandwidth For Packets of Different Lengths	446
Configuring UDP Broadcast Helper	448
Relaying BOOTP and DHCP Traffic	451
Verifying the UDP Broadcast Helper Configuration	453
Limiting the Number of Hops	453
Determining the Booting Order	454
Configuring Remote Access with Tunnel Terminator	467
Logging Messages	471
Specifying Log Content	472
Configuring PIM Protocol Authentication	474
Changing the Protocol Authentication Keyset (Rekeying)	475
Configuring PIM Authentication Rollback	476
Configuring X509 Certificates for the Router	477
Configuring X509 Certificates (Example)	478
Configuring a Dynamic-Key Security Policy for IPv4	479
Configuring a Dynamic Security Policy for IPv6	481
Configuring the GGSN and CBR	484
Setting Up the Dynamic Key Configuration with X509 Certificates	486
Setting Up a CWR Topology for Encryption	488
Configuring Routers in Fully Meshed Topology with a Dynamic Key	491
Creating an IPv6 IPSec Tunnel with an IKEv2 Dynamic Key	496
Creating a Hub and Spoke Topology	498
Configuring MD5 Authentication for OSPF	502
Configuring MD5 Authentication for BGP	503
Starting an SSH Session (EOS SSH Client and EOS SSH Server)	507
Starting and SSH Session (EOS SSH Client and Non-EOS SSH Server)	508
Starting an SSH Session (Non-EOS SSH Client and EOS SSH Server)	509
Stopping SSH Sessions	510
Configuring IP Security Parameters for End Systems	523
Configuring IP Security Options for Routers	524
Configuring Extended Security Option Labels	527
Verifying IP Security Options	528
Configuring a PP Tunnel Terminator	535
Configuring a Router-to-Router Tunnel Initiator/Terminator	537
Configuring a PPTP Virtual Leased Line	539
Configuring Tunnel Switching	542

Configuring NHRP	546
Verifying the Firewall Configuration	558
Checking the Overall Status	559
Blocking Unwanted Traffic	560
Enabling Security	566
Configuring RADIUS Authentication	574
Configuring LDAP for CCGW Clients	576
Configuring LDAP for Firewall Clients	578
Adding Attributes	580
Creating an Object Class	581
Configuring Basic GGSN Parameters	602
Configuring APNs on the GGSN	602
Configuring APNs to Use Static Subscriber IP Address Allocation	604
Configuring APNs to Use Dynamic Subscriber IP Address Allocation with Local Server	607
Configuring APNs to Use Dynamic Subscriber IP Address Allocation with a DHCP Server	610
Configuring Devices as HA GGSN Peers	614
Configuring the GGSN as a Backup VRRP Router	628
Configuring Border Routers to Support Conventional Data Redundancy (HA and DSR)	629
Configure the Border Router as a Backup VRRP Router	635
Configuring the GGSNs in the Sample Conventional Data Redundancy (HA and DSR) Topology	643
Configuring Border Routers in the Sample Conventional Data Redundancy (HA and DSR) Topology	650
Configuring GGSN Charging	659
Configuring GGSN Hardware Redundancy for Dimetra Systems	661
Configuring the SNMP Service for SNMPv1/v2c Management	663
Configuring Trap Profiles – Multiple Network Management Stations	667
Configuring Trap Profiles – Single Management Station	667
Opening the SNMP V3 User Manager Menu	669
Changing the MotoAdmin User Password	671
Changing the Privilege Level for the MotoMaster User	672
Creating MotoInformA (NoAuthNoPriv) Users Manually	676
Changing Credentials for Inform Users	677
Adding SNMPv3 USM Users	680
Deleting SNMPv3 USM Users	680
Changing Authentication Passphrases for SNMPv3 USM Users	681
Changing Encryption Passphrases for SNMPv3 USM Users	682
Viewing SNMPv3 USM Users	683
Configuring Remote Logging of Syslog Messages	694
Configuring SNMP for Notification Traps	698
Displaying the Software Version with the REMote Command	704

Preventing Remote Access	705
Restricting Remote Access	705
Restricting Telnet Access	706
Restricting SSH Access	706
Configuring VRRP – Gateway to a WAN	716
Configuring VRRP – One Master Router Connecting Two LANs	716
Configuring VRR – Load Sharing with Redundancy	717
Scheduling WAN Connections	722
Setting Up a Backup Port	723
Hanging Up a Port	723
Recovering from Port Loopback	723
Configuring Bidirectional Forwarding Detection (BFD) with Event-Based Macros	724
Configuring EOS Devices as a CCGW	735
Comparing the Inbound Call Unknown Start Count to Console Unknowns – Digital Channels	749
Comparing the Inbound Call Start Event Counts to Console Unknowns (MDC 1200 Channels)	750
Configuring EOS Devices for Overload Protection	752
Configuring IPLC Functionality	769
Adding and Configuring Channels for IPLC Standalone Topology – IP-Digital (GGM 8000 with IPLC Functionality at Core)	770
Adding and Configuring Channels for IPLC Standalone Topology – IP-Mixed	772
Adding and Configuring Channels for IPLC Paired Topology – IP-Mixed	775
Adding and Configuring Channels for IPLC Paired Topology – V24-Digital	779
Adding and Configuring Channels for IPLC Paired Topology – V24-Mixed	784
Creating a Configuration File with an IPLC Configuration Template	788
Downloading a Stored Configuration File to the GGM 8000	788
Configuring a Sample IPLC Standalone Topology – IP-Digital	791
Configuring a Sample IPLC Standalone Topology – IP-Mixed	792
Configuring a Sample IPLC Paired Topology – IP-Mixed (Core GGM 8000)	793
Configuring a Sample IPLC Paired Topology – IP-Mixed (Site GGM 8000)	797
Configuring a Sample IPLC Paired Topology – V24-Digital (Core GGM 8000)	800
Configuring a Sample IPLC Paired Topology – V24-Digital (Site GGM 8000)	801
Configuring a Sample IPLC Paired Topology – V24-Mixed (Core GGM 8000)	802
Configuring a Sample IPLC Paired Topology – V24-Mixed (Site GGM 8000)	803
Configuring a Sample IPLC Paired and Standalone Topology – Core GGM 8000 Supporting Multiple Channel Types	804
Configuring a Sample IPLC Paired Topology – V.24-Mixed (Site GGM 8000)	806
Configuring a Sample IPLC Paired Topology – IP-Mixed (Site GGM 8000)	807
Performing a Local Loopback Test	817
Performing a Remote Loopback Test	819

Performing a Loopback Diagnostic Test for Built-In ISDN Ports	820
Establishing Subnets: Example 1	827
Establishing Subnets: Example 2	829
Establishing Subnets: Example 3	831
Upgrading EOS Software	1035
Downgrading EOS Software	1038

About Enterprise OS Software Version 18.2 User Guide

This guide provides information about how to use Enterprise OS (EOS) software to operate and configure your device.



NOTE: The EOS software manuals are intended primarily as reference documents. You must verify command usage with your Motorola Solutions business/engineering representative before employing any of the commands described in this manual on a customer system or in documents for third party use.

If release notes were shipped with your product and the information in the release notes differs from the information in this guide, follow the instructions in the release notes.

This guide is intended for personnel who:

- Have experience in planning, maintaining, and troubleshooting local or wide area networks.
- Are familiar with network protocols, bridging and routing, and network management.

Before you use the information provided in this guide, you must first install the device according to the installation instructions in the appropriate hardware user guide, and install and configure Enterprise OS software on the device.

For the purposes of this manual, an EOS device is defined as a hardware platform running the EOS software. Supported EOS devices include:

- MNR S2500 routers (2021.1 through 7.18 system releases)
- MNR S2500 Conventional Channel Gateways
- MNR S6000 routers
- GGM 8000 gateways
- MCG 8000



NOTE: MCG 8000 is CCGW only platform. This means MCG 8000 does not support the routers functionality and GGSN functionality. MCG 8000 also does not support IPLC. Only commands used to configure CCGW are valid.

Each device listed above may not support all of the features described in this guide. For a list of the services supported by your hardware platform, see the appropriate hardware user guide. The following table lists the hardware user guides associated with each of the supported hardware platforms.

Related Information

Related Information	Purpose
<i>Standards and Guidelines for Communication Sites</i>	Provides standards and guidelines that should be followed when setting up a Motorola Solutions communications site. Also known as R56 manual.
<i>Enterprise OS Software Version 18.2 User Guide</i>	Provides information on how to use Enterprise OS (EOS) software to operate and configure system routers and gateways.

Related Information	Purpose
<i>MCG 8000 System Gateway Feature Guide</i>	Provides detailed instructions for the installation, configuration, troubleshooting, and recovery of the MCG 8000 that serves as a conventional channel gateway.
<i>GGM 8000 System Gateway Feature Guide</i>	Provides information relating to the installation, configuration, and management of the GGM 8000 Gateway used in various network locations.
<i>Motorola GGM 8000 Hardware User Guide</i>	Includes hardware installation, basic software configuration, and cabling instructions for the Motorola GGM 8000 gateway.
<i>Motorola Network Router (MNR) S2500 Hardware User Guide</i>	Includes hardware installation, basic software configuration, and cabling instructions for Motorola Network Router (MNR) S2500 routers.
<i>Motorola Network Router (MNR) S6000 Hardware User Guide</i>	Includes hardware installation, basic software configuration, and cabling instructions for Motorola Network Router (MNR) S6000 routers.
<i>S6000 and S2500 Routers Feature Guide</i>	Provides information relating to the installation, configuration, and management of the S6000 and S2500 routers used in various network locations.

To access the following manuals:

- *Motorola GGM 8000 Hardware User Guide*
- *Motorola Network Router (MNR) S2500 Hardware User Guide*
- *Motorola Network Router (MNR) S6000 Hardware User Guide*

you can extract the .tar file from the Motorola Network Router EOS or MNR GGSN CD-ROM onto a directory on your PC and navigate to: X:/<directory>/aa.b.c.dd/docs where X:/ is the drive letter for the directory, <directory> is the name of the directory into which you extracted the .tar file, and aa.b.c.dd is the specific build of EOS software provided on the CD-ROM. For example, if you extracted the .tar file into directory D:/EOS_SW, and the CD-ROM contains EOS version 16.6.0.26, you must navigate to: D:/EOS_SW/16.6.0.26/docs

These manuals are also available at <https://motonline.mot-solutions.com>. You can access them by selecting **Resource Center** → **Product Information** → **Manuals** → **Network Infrastructure** → **Routers and Gateways**.

The remaining manuals are available at <https://learning.motorolasolutions.com>.

Command Line Interface Symbol Conventions

Symbol	Description
angle brackets < >	Enclose a variable or value. You must specify the variable or value. For example, in the syntax: <code>DELeTe -IP ADDRess <IP address></code> you must supply a value for <IP address> when you enter the command. Do not type the angle brackets.
square brackets []	Enclose an optional value or a list of optional arguments. You can specify one or more values or arguments. For example, in the syntax:

Symbol	Description
	SET PRIvilege = [User NetMgr] you can specify either User or NetMgr when you enter the command. Do not type the square brackets.
parentheses ()	Enclose a list of values that can be assigned to a single parameter. You must enter at least one of the values. If you enter more than one value, you must type the parentheses and separate the values with commas. For example, in the syntax: SET CurrentPorts = ALL (<port>, <port>) you can specify two port numbers by enclosing them in parentheses and separating them with a comma when you enter the command.
vertical bar	Separates mutually exclusive values; you must enter one of the values. For example, in the syntax: SETDefault -PIM BSRPeriod = <secs> Default you can enter a value for <secs> or you can enter the word Default. Do not type the vertical bar.
hyphen -	Precedes service names. When you enter a service name as part of a command, you must precede the service name with a hyphen. For example, in the syntax: SHow [!<port> !*] -IP CONFIguration you must include the hyphen with the service name when you enter the command.
braces { }	Enclose a list of required values. You must enter one of the values. For example, in the syntax: SETDefault -PIM BSfragLimit = {(100-1480) Default} you must enter either a value (between 100 and 1480) or the word Default.
exclamation point !	Precedes a number that represents a port, virtual port, or path number. For example, in the syntax: SETDefault !<port> -IP RtphcHdrExtComp = Enable Disable you must include the exclamation point before the port number when you enter the command.
ellipsis ...	Indicate that you may specify one or more additional arguments on the same command line. For example, in the syntax: ADD! <filterid> -IP FIlters <condition> [,<condition>...] <condition> = <%offset>: [<operator>] <%pattern>

Symbol	Description
	you may specify multiple conditions when you enter the command. Do not type the ellipsis.

Text Conventions

Convention	Description
Syntax	Placeholders for values you must supply appear in angle brackets. Example: You can enable RIPIP by entering: <pre>SETDefault !<port> -RIPIP CONTrol = Listen</pre> In this example, you must supply a port number for <i><port></i>.
Commands	Enter the command exactly as shown in text and press RETURN or Esc. Example: To remove the IP address, you must enter: <pre>SETDefault !0 -IP NETaddr = 0.0.0.0</pre> This guide generally gives the full form of a command in uppercase and lowercase letters. However, you can abbreviate commands by entering only the uppercase letters and the appropriate value. Commands are not case-sensitive.
Screen displays	This typeface represents information as it appears on the screen.
User entries	This typeface represents text that you type or enter.

Chapter 1

Basic Ports and Paths Configuration

The terms *connector*, *path*, and *port* represent the fundamental interface syntax on the supported hardware platform.

Connector

Physical interface itself; for example, a FlexWAN connector, a BNC connector, or an RJ-45 connector. All connectors, by default, carry one path each. You can change this default on UltraWAN connectors so that the connector carries multiple paths, as is required in ISDN PRI operation.


Path

Connection on the physical interface that connects the device to a physical network medium such as Ethernet, a T1/E1 line, or serial line. Each path is associated with a connector. In an Integrated Services Digital Network (ISDN) environment, a path also represents the channel over which data is transmitted.

Port

Logical interface used by the software to represent a connection to a network.

Figure 1: Relationship Between Connectors, Paths, and Ports



The local and wide area interfaces available to you depend on your hardware platform and its configuration. For information on the types of interfaces your platform offers and how to configure paths and ports on a specific hardware platform, see the appropriate hardware user guide.

1.1

Path and Port Numbering

For software purposes, paths are numbered 1, 2, 3, and so on. The path number may be followed by a letter or a decimal and a channel number.

By default, each path is assigned to one port. For example, all network traffic received on physical path 1 is treated by the software as arriving on logical port 1, and all traffic that the software transmits through logical port 1 passes through physical path 1. This default configuration is called a *static port and path binding*. A *static path* is a path that is mapped to a port. All paths are static by default.

You can redefine the default mapping using software commands. For example, you can redirect network traffic that is routed through a particular logical port to a different physical path without manually switching cables on the connector.

You can also unbind paths from their ports and place them in a dial path pool to be shared by more than one port. The paths in the dial pool are called *dynamic paths*. A path in the dial pool can be dynamically bound to

a port running Point-to-Point Protocol (PPP) when the path is needed for data transfer events associated with dial-up. You can also bind a dynamic path to a port for dial backup purposes such as bandwidth-on-demand or disaster recovery. For more information about the use of the dial path pool, see [Configuring the Dial List on page 367](#).

Each logical port is usually associated with only one physical path. For Ethernet and Frame Relay, the path-to-port ratio is always one-to-one. But for paths connected to serial lines, multiple paths can be associated with and statically bound to a single port if the Point-to-Point Protocol (PPP) is running over the port.

1.1.1

Multiple Paths to a Wide Area Port Configuration


You can reconfigure the software so that multiple paths are mapped to one wide area port by entering the `ADD -PORT Paths` command.

If you assign multiple paths to a wide area port, the port must be running PPP. Configuring multiple paths to a wide area port is useful if you are using ISDN or switch-56 circuit services. You can map multiple paths to one port to take advantage of the disaster recovery and bandwidth-on-demand features and the dial path pools.

The following figure shows two sample topologies:

- In the topology shown on the left, one path is mapped to one port: path 2 is assigned by default to port 2.
- In the topology shown on the right, two paths are mapped to one port (for disaster recovery or bandwidth-on-demand). In this topology, the software was reconfigured so that paths 3 and 4 are mapped to port 3.

Figure 2: Possible Path-to-Port Assignments




1.1.2

S2500 Port and Path Numbering

The following figure shows the path/connector to port number mappings for MNR S2500s with an E&M module installed in the analog I/O slot. The MNR S2500 software reserves physical ports 4 through 7 for the analog conventional E&M ports.

Figure 3: S2500 Path-to-Port Mappings



Annotation	Description
1	LAN 1 maps path 1 to port 1
2	LAN/WAN 2 maps path 2 to port 2
3	LAN/WAN 3 maps path 3 to port 3
4	E&M analog paths 4-7 map to ports 4-7
5	Port 4 maps to channel 1
6	Port 5 maps to channel 2
7	Port 6 maps to channel 3
8	Port 7 maps to channel 4

The S2500 V.24 module can be installed in I/O slot A (port 2) or B (port 3). The MNR S2500 software reserves ports 2 or 3 A and B for digital conventional V.24 ports.

- If the V.24 module is installed in port 2, as shown in the following figure, the software uses port numbers 2A and 2B.
- If the V.24 module is installed in port 3, the software uses port numbers 3A and 3B.

 **NOTE:** You can install the V.24 module in either I/O slot A or I/O slot B, but the software supports only one V.24 module per MNR S2500.

Figure 4: S2500 Path-to-Port Mappings (V.24 Module in I/O Slot A)




Annotation	Description
1	LAN 1 maps path 1 to port 1
2	LAN/WAN 2 maps path 2 to ports 2A and 2B

1.1.3

S6000 Path-to-Port Mappings (FlexWAN Modules in Both Slots)

The following figure shows the path/connector to port number mappings for S6000s configured with FlexWAN modules in both I/O slots.

Figure 5: S6000 Path-to-Port Mappings (FlexWAN Modules in Both I/O Slots)



Annotation	Description
1	FlexWAN serial connectors
2	Maps path 4a to port 4a
3	Maps path 4b to port 4b
4	Maps path 4c to port 4c
5	Maps path 4d to port 4d
6	Maps path 5a to port 5a
7	Maps path 5b to port 5b
8	Maps path 5c to port 5c
9	Maps path 5d to port 5d
10	Ethernet RJ-45 connectors
11	LAN 1 maps path 1 to port 1
12	LAN 2 maps path 2 to port 2
13	LAN 3 maps path 3 to port 3


1.1.4

S6000 Path-to-Port Mappings (UltraWAN Modules in Both Slots)

The following figures show the path/connector to port mappings for S6000s configured with UltraWAN modules and UltraWAN II in both slots.


NOTE: The MNR S6000 supports two versions of the UltraWAN module. The functionality of the two module versions is the same but the UltraWAN II module (identified by a Roman numeral “II” on the front panel) requires EOS software version 15.4 or higher.

Figure 6: S6000 Path-to-Port Mappings (UltraWAN Modules in Both I/O Slots)



Annotation	Description
1	UltraWAN RJ-48C connectors
2	Maps path 4a to port 4a
3	Maps path 4b to port 4b
4	Maps path 4c to port 4c
5	Maps path 4d to port 4d
6	Maps path 5a to port 5a
7	Maps path 5b to port 5b
8	Maps path 5c to port 5c
9	Maps path 5d to port 5d
10	Ethernet RJ-45 connectors
11	LAN 1 maps path 1 to port 1
12	LAN 2 maps path 2 to port 2
13	LAN 3 maps path 3 to port 3

Figure 7: S6000 Path-to-Port Mappings (UltraWAN II Modules in Both I/O Slots)




Annotation	Description
1	UltraWAN RJ-48C connectors
2	Maps path 4a to port 4a
3	Maps path 4b to port 4b
4	Maps path 4c to port 4c
5	Maps path 4d to port 4d
6	Maps path 5a to port 5a
7	Maps path 5b to port 5b
8	Maps path 5c to port 5c
9	Maps path 5d to port 5d
10	Ethernet RJ-45 connectors
11	LAN 1 maps path 1 to port 1
12	LAN 2 maps path 2 to port 2
13	LAN 3 maps path 3 to port 3

1.1.5

S6000 Path-to-Port Mappings (T3/E3 Modules in Both Slots)

The following figure shows the path/connector to port number mappings for S6000s configured with T3/E3 modules in both I/O slots.

Figure 8: S6000 Path-to-Port Mappings (T3/E3 Modules in Both I/O Slots)




Annotation	Description
1	T3/E3 HSSI and BNC connectors
2	Maps path 4 to port 4
3	Maps path 5 to port 5
4	Ethernet RJ-45 connectors
5	LAN 1 maps path 1 to port 1
6	LAN 2 maps path 2 to port 2
7	LAN 3 maps path 3 to port 3

1.1.6

S6000 Path-to-Port Mappings (FlexWAN Module in Slot A and UltraWAN II Module in Slot B)

The following figure shows the path/connector to port number mappings for S6000s configured with a FlexWAN module in I/O slot A and an UltraWAN II module in I/O slot B.

Figure 9: S6000 Path-to-Port Mappings (FlexWAN Module in Slot A and UltraWAN II Module in Slot B)




Annotation	Description
1	FlexWAN serial connectors
2	Maps path 4a to port 4a
3	Maps path 4b to port 4b
4	Maps path 4c to port 4c
5	Maps path 4d to port 4d
6	UltraWAN RJ-48C connectors
7	Maps path 5a to port 5a
8	Maps path 5b to port 5b
9	Maps path 5c to port 5c
10	Maps path 5d to port 5d
11	Ethernet RJ-45 connectors
12	LAN 1 maps path 1 to port 1
13	LAN 2 maps path 2 to port 2
14	LAN 3 maps path 3 to port 3

1.1.7

S6000 Path-to-Port Mappings (UltraWAN II Module in Slot A and FlexWAN Module in Slot B)

The following figure shows the path/connector to port number mappings for S6000s configured with an UltraWAN II module in I/O slot A and a FlexWAN module in I/O slot B.

Figure 10: S6000 Path-to-Port Mappings (UltraWAN II Module in Slot A and FlexWAN Module in Slot B)




Annotation	Description
1	UltraWAN RJ-48C connectors
2	Maps path 4a to port 4a
3	Maps path 4b to port 4b
4	Maps path 4c to port 4c
5	Maps path 4d to port 4d
6	FlexWAN serial connectors
7	Maps path 5a to port 5a
8	Maps path 5b to port 5b
9	Maps path 5c to port 5c
10	Maps path 5d to port 5d
11	Ethernet RJ-45 connectors
12	LAN 1 maps path 1 to port 1
13	LAN 2 maps path 2 to port 2
14	LAN 3 maps path 3 to port 3

1.1.8

S6000 Path-to-Port Mappings (FlexWAN Module in Slot A and T3/E3 Module in Slot B)

The following figure shows the path/connector to port number mappings for S6000s configured with a FlexWAN module in I/O slot A and a T3/E3 module in I/O slot B.

Figure 11: S6000 Path-to-Port Mappings (FlexWAN Module in Slot A and T3/E3 Module in Slot B)




Annotation	Description
1	FlexWAN serial connectors
2	Maps path 4a to port 4a
3	Maps path 4b to port 4b
4	Maps path 4c to port 4c
5	Maps path 4d to port 4d
6	T3/E3 HSSI connector or BNC connector
7	Maps path 5 to port 5
8	Ethernet RJ-45 connectors
9	LAN 1 maps path 1 to port 1
10	LAN 2 maps path 2 to port 2
11	LAN 3 maps path 3 to port 3

1.1.9

S6000 Path-to-Port Mappings (T3/E3 Module in Slot A and FlexWAN Module in Slot B)

The following figure shows the path/connector to port number mappings for S6000s configured with a T3/E3 module in I/O slot A and a FlexWAN module in I/O slot B.

Figure 12: S6000 Path-to-Port Mappings (T3/E3 Module in Slot A and FlexWAN Module in Slot B)




Annotation	Description
1	T3/E3 HSSI connector or BNC connector
2	Maps path 4 to port 4
3	FlexWAN serial connectors
4	Maps path 5a to port 5a
5	Maps path 5b to port 5b
6	Maps path 5c to port 5c
7	Maps path 5d to port 5d
8	Ethernet RJ-45 connectors
9	LAN 1 maps path 1 to port 1
10	LAN 2 maps path 2 to port 2
11	LAN 3 maps path 3 to port 3

1.1.10

S6000 Path-to-Port Mappings (UltraWAN II Module in Slot A and T3/E3 Module in Slot B)

The following figure shows the path/connector to port number mappings for S6000s configured with an UltraWAN II module in I/O slot A and a T3/E3 module in I/O slot B.

Figure 13: S6000 Path-to-Port Mappings (UltraWAN II Module in Slot A and T3/E3 Module in Slot B)




Annotation	Description
1	UltraWAN RJ-48C connectors
2	Maps path 4a to port 4a
3	Maps path 4b to port 4b
4	Maps path 4c to port 4c
5	Maps path 4d to port 4d
6	HSSI and T3/E3 (BNC) connectors
7	Maps path 5 to port 5
8	Ethernet RJ-45 connectors
9	LAN 1 maps path 1 to port 1
10	LAN 2 maps path 2 to port 2
11	LAN 3 maps path 3 to port 3

1.1.11

S6000 Path-to-Port Mappings (T3/E3 Module in Slot A and UltraWAN II Module in Slot B)

The following figure shows the path/connector to port number mappings for S6000s configured with a T3/E3 module in I/O slot A and an UltraWAN II module in I/O slot B.

Figure 14: S6000 Path-to-Port Mappings (T3/E3 Module in Slot A and UltraWAN II Module in Slot B)



Annotation	Description
1	HSSI and T3/E3 (BNC) connectors
2	Maps path 4 to port 4
3	UltraWAN RJ-48C connectors
4	Maps path 5a to port 5a
5	Maps path 5b to port 5b
6	Maps path 5c to port 5c
7	Maps path 5d to port 5d
8	Ethernet RJ-45 connectors
9	LAN 1 maps path 1 to port 1
10	LAN 2 maps path 2 to port 2
11	LAN 3 maps path 3 to port 3


1.1.12

S6000 Path-to-Port Mappings (12-Port T1/E1 II Modules in Both Slots)

The following figure shows the path/connector to port number mappings for S6000s configured with 12-port T1/E1 II modules in both I/O slots.

 **NOTE:** The MNR S6000 supports two versions of the 12-port T1/E1 module. The functionality of the two module versions is the same but the 12-port T1/E1 II module (identified by a Roman numeral “II” on the front panel) requires EOS software version 15.4 or higher.

Figure 15: S6000 Path-to-Port Mappings (12-Port T1/E1 Modules in Both I/O Slots)




Annotation	Description
1	High-density T1/E1 Connector Maps path 4a to port 4a Maps path 4b to port 4b Maps path 4c to port 4c ... Maps path 4l to port 4l
2	High-density T1/E1 Connector Maps path 5a to port 5a Maps path 5b to port 5b Maps path 5c to port 5c ... Maps path 5l to port 5l
3	Ethernet RJ-45 connectors
4	LAN 1 maps path 1 to port 1
5	LAN 2 maps path 2 to port 2
6	LAN 3 maps path 3 to port 3

1.1.13

GGM 8000 Port and Path Numbering

Figure 16: GGM 8000 Path-to-Port Pairings (Base Unit)


This figure shows the path/connector to port number pairing for the GGM 8000 base unit.



Annotation	Description
1	Paths 5B and 5A are paired with Ports 5B and 5A
2	Paths 4, 3, 2, and 1 are paired with Ports 4, 3, 2, and 1

Figure 17: GGM 8000 Path-to-Port Pairings (Expansion Module with Analog/V.24 Interface Kit)

This figure shows the path/connector to port number pairings for a GGM 8000 configured with an expansion module and analog/V.24 interface kit. The GGM 8000 software reserves physical ports 8D through 8A for the E&M (analog) ports and ports 7B, 7A, 6B, and 6A for the V.24 ports.




Annotation	Description
1	Paths 8D, 8C, 8B, and 8A are paired with Ports 8D, 8C, 8B, and 8A path 8A is bound to channel 1 path 8B is bound to channel 2 path 8C is bound to channel 3

Annotation	Description
	path 8D is bound to channel 4
2	Paths 7B and 7A are paired with Ports 7B and 7A
3	Paths 6B and 6A are paired with ports 6B and 6A
4	Paths 5B and 5A are paired with Ports 5B and 5A
5	Paths 4, 3, 2, and 1 are paired with Ports 4, 3, 2, and 1

Figure 18: GGM 8000 Path-to-Port Pairing (FlexWAN Daughterboards in Expansion Module)


This figure shows the path/connector-to-port number pairing for a GGM 8000 configured with an expansion module with two FlexWAN daughterboards. The GGM 8000 software reserves physical ports 6 and 7 for the FlexWAN serial ports.



Annotation	Description
1	Path 7 is paired with Port 7
2	Path 6 is paired with Port 6
3	Paths 5B and 5A are paired with Ports 5B and 5A
4	Paths 4, 3, 2, and 1 are paired with Ports 4, 3, 2, and 1

Figure 19: GGM 8000 Path-to-Port Pairings (Low Density Enhanced Conventional Gateway Module)


This figure shows the path/connector to port number pairings for a GGM 8000 configured with a Low Density Enhanced Conventional Gateway module. The GGM 8000 software reserves physical ports 8D through 8A for the E&M (analog) ports and ports 7B, 7A, 6B, and 6A for the V.24 ports.



Annotation	Description
1	Path 8D is paired with Audio, E&M port 8D which may also be paired with Analog I/O port 9D
2	Path 8C is paired with Audio, E&M port 8C which may also be paired with Analog I/O port 9C
3	Path 8B is paired with Audio, E&M port 8B which may also be paired with Analog I/O port 9B
4	Path 8A is paired with Audio, E&M port 8A which may also be paired with Analog I/O port 9A
5	Paths 7B, 7A, 6B, and 6A are paired with Ports 7B, 7A, 6B, and 6A
6	Paths 5B and 5A are paired with Ports 5B and 5A
7	Paths 4, 3, 2, and 1 are paired with Ports 4, 3, 2, and 1

Figure 20: GGM 8000 Path-to-Port Pairings (High Density Enhanced Conventional Gateway Module)

This figure shows the path/connector to port number pairings for a GGM 8000 configured with a High Density Enhanced Conventional Gateway module. The GGM 8000 software reserves physical ports 8D through 8A and 12D through 12A for the E&M (analog) ports and ports 7B, 7A, 6B, 6A, 11B, 11A, 10B, and 10A for the V.24 ports.




Annotation	Description
1	Path 12D is paired with Audio, E&M port 12D which may also be paired with Analog I/O port 13D
2	Path 12C is paired with Audio, E&M port 12C which may also be paired with Analog I/O port 13C
3	Path 12B is paired with Audio, E&M port 12B which may also be paired with Analog I/O port 13B
4	Path 12A is paired with Audio, E&M port 12A which may also be paired with Analog I/O port 13A
5	Paths 11B, 11A, 10B, and 10A are paired with Ports 11B, 11A, 10B, and 10A

Annotation	Description
6	Path 8D is paired with Audio, E&M port 8D which may also be paired with Analog I/O port 9D
7	Path 8C is paired with Audio, E&M port 8C which may also be paired with Analog I/O port 9C
8	Path 8B is paired with Audio, E&M port 8B which may also be paired with Analog I/O port 9B
9	Path 8A is paired with Audio, E&M port 8A which may also be paired with Analog I/O port 9A
10	Paths 7B, 7A, 6B, and 6A are paired with Ports 7B, 7A, 6B, and 6A
11	Paths 5B and 5A are paired with Ports 5B and 5A
12	Paths 4, 3, 2, and 1 are paired with Ports 4, 3, 2, and 1

1.1.14

MCG 8000 Port and Path Numbering


Figure 21: MCG 8000 Path-to-Port Pairings (Base Unit)



Annotation	Description
1	10/100/1000 Ethernet Ports (RJ-45) and Status LEDs
2	Console Port (DB-9)
3	Reset Button
4	System LEDs

As illustrated in the following figure, the MCG 8000 software reserves physical ports 8D through 8A and 12D through 12A for the E&M (analog) ports and ports 7B, 7A, 6B, 6A, 11B, 11A, 10B, and 10A for the V.24 ports.

Figure 22: MCG 8000 Port and Path Numbering (High Density Enhanced Conventional Gateway Module)



Annotation	Description
1	Path 12D is paired with Audio, E&M port 12D which may also be paired with Analog I/O port 13D
2	Path 12C is paired with Audio, E&M port 12C which may also be paired with Analog I/O port 13C
3	Path 12B is paired with Audio, E&M port 12B which may also be paired with Analog I/O port 13B
4	Path 12A is paired with Audio, E&M port 12A which may also be paired with Analog I/O port 13A
5	Paths 11B, 11A, 10B, and 10A are paired with Ports 11B, 11A, 10B, and 10A
6	Path 8D is paired with Audio, E&M port 8D which may also be paired with Analog I/O port 9D
7	Path 8C is paired with Audio, E&M port 8C which may also be paired with Analog I/O port 9C
8	Path 8B is paired with Audio, E&M port 8B which may also be paired with Analog I/O port 9B
9	Path 8A is paired with Audio, E&M port 8A which may also be paired with Analog I/O port 9A
10	Paths 7B, 7A, 6B, and 6A are paired with Ports 7B, 7A, 6B, and 6A
11	Paths 3, 2, and 1 are paired with Ports 3, 2, and 1

1.2

Path, Port, and Connector Syntax

Path and port parameters have their own dedicated service type, but connector parameters do not have a dedicated service type:

- The P*A*th service is used to configure paths and connectors.
- The P*O*rt service is used to configure ports.

S2500 Path and Port Syntax

The following table lists the path and port numbering syntax rules used for default naming for the S2500.

Table 1: S2500 Path/Port Syntax Rules

Interface	Path/Port Syntax	Syntax Description
Ethernet (LAN) connectors (10BASE-T/100BASE-TX)	<n>	n = 1
Ethernet (LAN) connectors (10BASE-T)	<n>	n = 2 or 3
FlexWAN Serial connectors	<n>	n = 2 or 3
WAN/Telco (T1/E1 CSU/DSU) connectors	<n>	n = 2 or 3
V.24 connectors	<nA> and <nB>	n = 2 or 3 ¹
E&M Analog connectors	<n>	n = 4, 5, 6, 7 ²

¹ V.24 paths/ports 2A and 2B or 3A and 3B map to digital conventional channels 1 and 2 with the following mapping:

Port 2A -> Channel 1
Port 3A -> Channel 1
Port 2B -> Channel 2
Port 3B -> Channel 2

² E&M paths/ports 4-7 map to analog conventional channels 1-4 as follows:

Path/Port 4 maps to channel 1.
Path/Port 5 maps to channel 2.
Path/Port 6 maps to channel 3.
Path/Port 7 maps to channel 4.

S6000 Path and Port Syntax

The following table lists the path and port numbering syntax rules used for default naming for MNR S6000s.

Table 2: S6000 Path/Port Syntax Rules

Interface	Path/Port Syntax	Syntax Description
LAN connectors	<n>	n = 1, 2, or 3
FlexWAN Serial connectors	<n / m>	n = 4 or 5 m = a, b, c, or d
UltraWAN CSU/DSU connectors ¹	<nm.p>	n = 4 or 5

Interface	Path/Port Syntax	Syntax Description
		m = a, b, c, or d p = 1, 2, ...31
T3/E3 connectors	<n>	n = 4 or 5
¹ S6000 UltraWAN CSU/DSU interfaces support multi-path ISDN PRI dial as well as single-path leased unstructured operation.		

Table 3: S6000 UltraWAN Connector Syntax Rules

Interface	Connector Syntax	Syntax Description
UltraWAN CSU/DSU Connectors	4 <m> or 5 <m>	m = a, b, c, or d

GGM 8000 Path and Port Syntax

The following table lists the path and port numbering syntax rules used for default naming for the GGM 8000.



NOTE: Paths are numbered from right to left on the front panel of the GGM 8000.

Table 4: GGM 8000 Path/Port Syntax Rules

Interface	Front Panel Label	Path/Port Syntax	Syntax Description
Ethernet (LAN) connectors (10/100/1000) supported on base system	4, 3, 2, 1	<n>	n = 1, 2, 3, or 4
T1/E1 (CSU/DSU) connectors supported on base system	5B, 5A	<nm.p>	n = 5 m = a or b p = 1, 2, ... 31
E&M (Analog) connectors require optional E&M daughterboard and DSP SIMM (included in analog/V.24 interface kit) or Enhanced Conventional Gateway module	8D, 8C, 8B, 8A ¹ and 12D, 12C, 12B, and 12A ²	<nA>, <nB>, <nC>, or <nD>	n = 8 or 12 ²
V.24 connectors require optional V.24 daughterboard(s) (included in analog/V.24 interface kit) or Enhanced Conventional Gateway module	7B, 7A, 6B, 6A and 11B, 11A, 10B, 10A ³	<nA> or <nB>	n = 6, 7, 10, or 11 ³
FlexWAN serial connector requires optional FlexWAN daughterboard(s)	7B, 7A, 6B, 6A	<n>	n = 6 or 7

Interface	Front Panel Label	Path/Port Syntax	Syntax Description
¹ The Enhanced Conventional Gateway module supports pairs of RJ45 connectors. The Audio, E&M connector (ports 12D, 12C, 12B, 12A (supported on the High Density Enhanced Conventional Gateway only) and 8D, 8C, 8B, and 8A) can be used alone or paired with the corresponding Analog, I/O connector (ports 13D, 13C, 13B, 13A(supported on the High Density Enhanced Conventional Gateway only) and 9D, 9C, 9B, and 9A). ² E&M (analog) ports 12D, 12C, 12B, and 12A are supported only on the High Density Enhanced Conventional Gateway module. ³ V.24 ports 11B, 11A, 10B, and 10A are supported only on the High Density Enhanced Conventional Gateway module.			

MCG 8000 Path And Port Syntax

The following lists the path and port numbering syntax rules used for default naming for the MCG 8000.



NOTE: Paths are numbered from right to left on the front panel of the MCG 8000.

Table 5: MCG 8000 Path and Port Syntax

Interface	Front Panel Label	Path/Port Syntax	Syntax Description
Ethernet (LAN) connectors (10/100/1000)	3, 2, 1 ¹	<n>	n = 1, 2, 3
E&M (Analog) - supported only if equipped with ECCGW card	8D, 8C, 8B, 8A and 12D, 12C, 12B, and 12A ²	<n>A, <n>B, <n>C, or <n>D	n = 8 or 12 ²
V.24 connectors - supported only if equipped with ECCGW card	7B, 7A, 6B, 6A and 11B, 11A, 10B, 10A	<n>A or <n>B	n = 6, 7, 10, or 11

¹ Front panel has 4 LAN slots but only 3 are supported. MCG 8000 needs only slot 1 to communicate.
² The MCG 8000 supports pairs of RJ45 connectors. The Audio, E&M connector (ports 12D, 12C, 12B, 12A, 8D, 8C, 8B and 8A) can be used alone or paired with the corresponding Analog, I/O connector (ports 13D, 13C, 13B, 13A, 9D, 9C, 9B and 9A).

1.3

Configuring Local Area Interfaces

You can use this procedure to configure ports and paths on local area interfaces.

Procedure:

1. Optional: Assign a name to path 1 by entering:

```
SETDefault !<path> -PATH Name = "string"
```

Some restrictions apply to the name you assign by using the -PATH Name parameter. For more information, see "PATH Service Parameters" in the *Enterprise OS Software Reference Guide*.

Step example:

To name path 1, enter:

```
SETDefault !1 -PATH NAme = "FLOOR_1"
```

2. Disable all local area ports that you are not using by entering for each port:

```
SETDefault !<port> -PORT CONTrol = Disabled
```

Motorola Solutions recommends that you disable all ports you do not use. Disabling unused ports improves performance. All ports are enabled by default. If the port was disabled, you must re-enable the port to use it.

3. Repeat [step 1](#) to [step 2](#) for each local area port and local area path on your device.

By default, the Ethernet paths and ports are enabled.

1.4

Configuring Wide Area Interfaces

You can use this procedure to configure ports and paths on wide area interfaces.

Procedure:

1. Optional: Assign a name to the path by entering:

```
SETDefault !<path> -PATH NAme = "string"
```

Some restrictions apply to the name you assign by using the `-PATH NAme` parameter. For more information, see "PAth Service Parameters" in the *Enterprise OS Software Reference Guide*.

Step example:

To assign the path name SF-SJ, enter:

```
SETDefault !3 -PATH NAme = "SF-SJ"
```

2. If necessary, reconfigure the connector type for the path by entering:

```
SETDefault -PATH CONNector
```

The syntax varies depending on the platform you use.

On the platform with an RS-449 cable installed, the software cannot distinguish between the RS-449 cable and an V.35 cable. You must configure the `-PATH CONNector` parameter to RS-449, or the software assumes the cable is a V.35.

For MNR S2500s, Motorola Solutions recommends retaining the default setting of the `-PATH CONNector` parameter (Auto). When the `-PATH CONNector` parameter is set to Auto, the DTE connector type is automatically detected when the platform boots.

For more information on the `CONNector` parameter and the Autostartup feature, see "Port Service Parameters" in the Enterprise OS Software Version 18.2 Reference Guide and [Autostartup Configuration on page 98](#).

3. If necessary, reconfigure the transmit clock setting for the serial path by entering:

```
SETDefault !<pathlist> -PATH CLock = Internal | External | TestMode
```

The default setting is `External`. The EOS device usually derives its clock from an external modem, so in most cases you do not need to change the default setting. If you want the EOS device to derive the clock from the onboard oscillator, set the transmit clock to `TestMode`.

4. If necessary, reset the baud rate for the path by entering:

```
SETDefault !<path> -PATH BAud = <kbps>
```

For the default values and the range of baud rates available for each hardware platform, see “Path Service Parameters” in the Enterprise OS Software Version 18.2 Reference Guide.



NOTE: It is important to set the baud rate even if you use an external clock. The hardware platform uses the baud rate setting to allocate resources for a path, compute metrics, and select a forwarding path.

Step example:

To set the baud rate of path 3 at 256 kbps, enter:

```
SETDefault !3 -PATH BAud = 256
```

5. If the port is running PPP and you plan to use features that use multiple paths mapped to a port (see [Multiple Paths to a Wide Area Port Configuration on page 51](#)), assign a path or multiple paths to each port by entering:

```
ADD !<port> -PORT PATHs
```

The syntax varies depending on the platform you use.

For example, to assign paths 3 and 4 to port 3, you must enter:

```
ADD !3 -PORT PATHs 3,4
```

Assigning multiple paths to a port is supported only when PPP is the port owner. When you assign multiple paths to a port, a load-sharing algorithm is enabled. For more information, see [Load-Sharing Algorithm on page 418](#).

To receive incoming calls from a remote site, you also can assign a dial path pool to a port.

For example, the following command assigns all incoming calls from Boston using dial-up lines to port 2:

```
ADD !2 -PORT PATHs SCID "Boston"
```

If you use SCID to identify the incoming calls, you can only identify calls coming in from an EOS device at the remote site.

If you assign ISDN paths for dial-up lines, you also have the option of using Calling Line Identification Presentation (CLIP) to identify the incoming calls from a Motorola Solutions or other vendor hardware platform at the remote site to the port on the hardware platform at the central site. You set up CLIP with the `-PORT CLIPList` and `-PORT DialRcvrState` parameters. For details, see “Port Service Parameters” in the Enterprise OS Software Version 18.2 Reference Guide.

If you are configuring a virtual port for ISDN dial-up paths, modem dial-up paths, or for a leased line, you should configure the port for SCID and for CLIP to identify the port to the remote user.

- If you configure a port for both SCID and CLIP, the CLIP configuration overrides the SCID configuration for incoming ISDN dial-up path calls.
- If you configure the port for something other than ISDN dial-up paths, SCID is used and CLIP is ignored.

For more information about setting up dial-up lines, see [Configuring the Dial List on page 367](#).

6. If you disabled the path, changed the value of the `BAud`, `CLock`, or `CONNector` parameters, or assigned multiple paths to one port, re-enable the path by entering:

```
SETDefault !<path> -PATH CONTrol = Enabled
```

If multiple paths are assigned to a port, you should enable all paths assigned to the port so that the load-sharing algorithm takes effect.

7. Repeat [step 1](#) to [step 6](#) for each wide area path you configure.
8. If necessary, enable or disable the wide area port by entering:

```
SETDefault !<port> -PORT CONTrol = Enabled | Disabled
```

All ports are enabled by default. If the port was previously disabled, you must re-enable it.

9. Optional: Assign a name to the port by entering:

```
SETDefault !<port> NAme = "string"
```

Some restrictions apply to the name you assign with the `-PORT NAme` parameter. For more information, see "Port Service Parameters" in the Enterprise OS Software Version 18.2 Reference Guide.

Step example:

To assign wide area port 3 the name SanJose, enter:

```
SETDefault !3 -PORT NAme = "SanJose"
```

10. If necessary, change the default owner of the wide area port by entering:

```
SETDefault !<port> -PORT OWNeR = EThernet | PPP| FrameRelay |  
LoopBack | Auto
```

The default port owners for each wide area port type are as follows:

- MNR S2500s: Auto for DTE serial ports; PPP for ISDN ports
- MNR S6000s: Auto for serial ports; PPP for ISDN ports

For more information on the `-PORT OWNeR` parameter and the Autostartup feature, see "Port Service Parameters" in the Enterprise OS Software Version 18.2 Reference Guide and [Autostartup Configuration on page 98](#).

11. Repeat [step 8](#) through [step 10](#) for each wide area port you configure.

12. Disable each port that you are not using by entering:

```
SETDefault !<port> -PORT CONTroL = Disabled
```

Motorola Solutions recommends that you disable all ports you do not use. Disabling unused ports improves performance.

Chapter 2

Virtual Ports and Multiple Instance Lists Configuration

To configure multiple logical ports over one physical path, you create new logical interfaces called *virtual ports*. Virtual ports allow you to increase the port density on your device.

The following table lists the hardware platforms that support virtual ports and the maximum number of virtual ports that you can configure on each device. There is no per-path limit, except that the total number of virtual ports configured on all paths cannot exceed the maximum for the hardware platform.

Table 6: Number of Virtual Ports Supported per Hardware Platform


Hardware Platform	Number of Virtual Ports Supported
MNR S2500	64
MNR S6000	2048
GGM 8000	2048
MCG 8000	2048



NOTE: Virtual ports are used in MCG 8000 for IP channels, although not directly and cannot be used to create WAN tunnels.

Virtual ports are associated with a nonvirtual port called the *parent port*.

Figure 23: Parent Port and Virtual Port



A virtual port functions in the same way as a non-virtual port in that it is a logical interface that represents a connection to a network. The virtual port and its parent port share most of their properties. However, you can reference a virtual port and its parent port separately using port-oriented software features, such as route policy and packet filtering, and you can distinguish a virtual port and its parent port by distinct wide area addresses.

You can connect a virtual port to a network through a path providing a Frame Relay virtual circuit. You can also use virtual ports to make a connection by using PPP with dial-up features to achieve multideestination dialing (modem pooling).


The following table lists the topologies that require virtual ports and the node in the topology on which the virtual ports should be created.

Table 7: Topologies that Require Virtual Ports

Topology	Virtual Ports Required?	Node On Which Virtual Ports Should Be Created
Traditional routed environment: partially meshed or non-meshed Frame Relay topologies	Depends on bridging or routing protocol. See Virtual Ports Over Frame Relay on page 78 for more information.	"Hub" device
Multidestination dialing (modem pooling) over PPP	Yes, for dynamic dial-up lines.	Central node
Frame Relay topology with disaster recovery configured	Yes	Nodes on both ends of serial line running Frame Relay
Routing traffic for 802.1Q virtual LANs.	Yes, to recognize VLANs. Traditional 802.1D bridging forwards VLAN packets without recognizing them as such.	Node attached to LAN

In the following figure virtual ports are configured on the wide area ports of two devices on both sides of a Frame Relay network. Each device has a virtual port defined over the path directly connected to the Frame Relay network. On Router #1, virtual port V1 is defined over path 2. On Router #2, virtual port V2 is defined over path 3.

Figure 24: Virtual Ports on Both Sides of a Frame Relay Network



NOTE: For information on using IP/IP virtual ports in security policy configurations, see [IP Traffic Protection with Dynamic Security Policies on page 479](#). For information on using virtual ports with Layer 2 tunneling, see [L2Tunnel Connections Configuration on page 535](#).

2.1 Virtual Ports Over Frame Relay

Frame Relay is a peer-to-peer protocol that connects two nodes on the network. Internet Protocol-Open Shortest Path First (IP-OSPF) requires virtual ports because it does not provide a method for dealing with Frame Relay topologies in which devices are not directly connected to all other devices (full mesh).

Internet Protocol-Routing Information Protocol (IP-RIP) and IP-Integrated Intermediate System-to-Intermediate System (IIS-IS) can operate over partially meshed or non-meshed Frame Relay topologies without the use of virtual ports. The next-hop split horizon feature in IP-RIP allows communication between devices that are not directly connected to one another. To configure next-hop split horizon for this routing protocol, you must have a list of neighbors, which can be dynamically generated or manually configured in IP-RIP.

You do not need to further configure IP-Integrated IS-IS to run over partially meshed or non-meshed Frame Relay topologies; you only need to configure neighbors.

Although you are not required to define virtual ports on IP-RIP devices in partially meshed or non-meshed Frame Relay topologies, virtual ports do provide the following additional benefits:

- You can define a virtual port for each configured neighbor. This allows you to set up such features as filters and routing policies on a per-neighbor basis.
- Virtual ports provide greater control over your network.

The following table lists each bridging and routing protocol and the technique you must use to deal with the lack of connectivity in partially meshed and non-meshed Frame Relay topologies.

Table 8: Connectivity in Partially Meshed and Non-Meshed Topologies

Protocol	Technique
Bridging	Virtual port
IP-RIP	Next-hop split horizon
 NOTE: When configuring the IP-RIP protocol and another protocol that requires virtual ports over the same path, use virtual ports.	
IP-OSPF	Virtual port
IP-Integrated IS-IS	No special configuration required

2.2

Parent Ports for Frame Relay

When you configure a Frame Relay virtual port, it inherits the attributes of the path over which it is defined. It also inherits some of the attributes of its parent port. There are two kinds of inheritance: one is the inherited default for all VCs, and the other occurs when the port picks up the value of the parent port.

For example, if you create a Frame Relay virtual port associated with a wide area port, the virtual port inherits port attributes from the following sources:

- Default and configured values of POrt Service parameters specified for a wide area port, with the exception of the following parameters that are not related to Frame Relay virtual ports:

```

AutoDial
COMPRESSType
DialCONFig
DialCONTRol
DialDebounceTime
DialHistory
DialIdleTime
DialInitState
DialRcvrState
DialRetryCount
DialRetryTime
DialSamplPeriod
DialStatus
LinkCompStat
OWNer
PATHs
PathPreference

```

These parameters do apply to PPP virtual ports, such as SysCallerID virtual ports.

- Default and configured values of parameters from all other services specified for a wide area port.

To configure a virtual port, you must specify the virtual port and not the parent port. For example, if you use the `SETDefault !<port> -OSPF CONTROL = Enabled` syntax, you must specify the virtual port number instead of the parent port number for `<port>`.

2.3

Virtual Ports Over PPP

PPP virtual ports differ from Frame Relay virtual ports in the following ways:

- A PPP virtual port can potentially use any path in the dial pool.
Frame Relay virtual ports are always associated with a particular path.
- PPP virtual ports operate independently, without a parent port. No parent port exists for PPP virtual ports because the path was unbound from its port and placed into the dynamic dial path pool.
Frame Relay virtual ports inherit the attributes of the path over which they are defined. For more information, see [Virtual Ports for Ethernet WAN Links on page 80](#).
- PPP virtual ports can be used with dial-up related parameters.
Frame Relay virtual ports cannot be used with dial-up related parameters.

You can use virtual ports in a PPP environment to provide dial pooling at the central site router. With dial pooling, a set of dynamic paths is unbound from their default ports and waits in the dial pool for an incoming call. When a call is received, the dynamic path that answers is assigned to a virtual port, which is standing by with the appropriate configuration information for the calling network. Because not all sites using a dial pool call the central site at the same time, it is possible to share a small group of paths with a larger group of sites. Each site that can potentially call into the dial pool has its own virtual port defined, so there are usually more virtual ports configured for the dial pool than dynamic paths assigned to it.

2.4

Virtual Ports for Ethernet WAN Links

Ethernet WAN links are logical links that connect two EOS devices using the IPIP protocol on their Ethernet interfaces. These links can replace the Frame Relay links previously used to connect remote sites and zones in an ASTRO 25 system. Ethernet site links use bidirectional forwarding detection (BFD) and Ethernet interzone links use BGP keepalives instead of local management interface (LMI) to determine the status of the link.


When Ethernet WAN links are implemented, the EOS device does not perform traffic shaping or rate limiting, and it does not fragment routed packets; however, it does fragment packets sourced by the device itself. The EOS device marks the appropriate quality of service (QoS) bits and leaves it up to the ISP to ensure that latency and jitter limits are not exceeded. EOS devices determine the state of the Ethernet WAN link and report it via the SNMP interface following NMA fault object standards. In addition, the devices support a tool to measure link latency, jitter, and packet loss rate. For details, see [Delay and Jitter Performance Management Tool on page 1032](#).

2.4.1

Configuring Ethernet Site Links by Using IPIP with BFD

The IPIP with BFD design is used for Ethernet site links. In these configurations, you configure IPIP tunnels, and BFD runs encapsulated in the tunnel. When the BFD status changes, the EOS device sends an Ethernet WAN link object state change trap.

Figure 25: Ethernet WAN Links: Sample Single Site Link Configuration



Procedure:

1. Configure the IP addresses of WAN A, WAN B, and the Site WAN by entering:

```
SETDefault !<portnumber> -IP NETaddr = <IP_address> <subnet_mask>
```

These addresses are provided by the ISP and are routable by the ISP.

Step example:

To configure the IP addresses:

On the primary core router, enter:

```
SETDefault !3 -IP NETaddr = 25.1.1.1 255.255.255.0
```

On the secondary core router, enter:

```
SETDefault !3 -IP NETaddr = 25.1.1.2 255.255.255.0
```

On the site router or site gateway, enter:

```
SETDefault !2 -IP NETaddr = 26.1.1.1 255.255.255.0
```

2. Create virtual ports over which the tunnels connecting the primary core router and the site router or site gateway and the secondary core router and the site router or site gateway run by entering:

```
ADD !v<port> -PORT VirtualPort IPIP <local_ipaddr><remote_ipaddr>  
>[TTL <t1>]
```

Step example:

To create the virtual tunnel ports:

On the primary core router, enter:

```
ADD !v101 -PORT VirtualPort IPIP 25.1.1.1 26.1.1.1 TTL 10
```

On the secondary core router, enter:

```
ADD !v101 -PORT VirtualPort IPIP 25.1.1.2 26.1.1.1 TTL 10
```

On the site router or site gateway, enter the following commands:

```
ADD !v16 -PORT VirtualPort IPIP 26.1.1.1 25.1.1.1 TTL 10
```

```
ADD !v17 -PORT VirtualPort IPIP 26.1.1.1 25.1.1.2 TTL 10
```

3. Assign IP addresses to the tunnels that you created in [step 2](#) by entering:

```
SETDefault !v<port> -IP NETAddr = <IP_address> <subnet_mask>
```

Step example:

To assign IP addresses to the IPIP tunnels:

On the primary core router, enter:

```
SETDefault !v101 -IP NETAddr = 172.24.16.5 255.255.255.252
```

On the secondary core router, enter:

```
SETDefault !v101 -IP NETAddr = 172.24.20.5 255.255.255.252
```

On the site router or site gateway, enter the following commands:

```
SETDefault !v16 -IP NETAddr = 172.24.16.6 255.255.255.252
```

```
SETDefault !v17 -IP NETAddr = 172.24.20.6 255.255.255.252
```

4. Add static routes from the core routers to the site router or site gateway by entering:

```
ADD -IP ROUTe <IP_address> [<mask>] {<gateway> | !<port>} <metric> [Override]
[BFD[TTL <tTl>]]
```

The BFD option configures the static routes as BFD-controlled routes. BFD checks for reachability to the gateway address before importing the static route into the routing table.



NOTE: The optional TTL keyword and <tTl> value allow you to specify a single-hop or multi-hop BFD-controlled route. For details, see [Bidirectional Forwarding Detection \(BFD\) Configuration on page 138](#) and the description of the ROUTe parameter in "Configuring IP Service Parameters" in the Enterprise OS Software Version 18.2 Reference Guide.

Step example:

To configure the static routes:

On the primary core router, enter:

```
ADD -IP ROUTe 10.101.1.0 255.255.255.0 172.24.16.6 10 BFD
```

On the secondary core router, enter:

```
ADD -IP ROUTe 10.101.1.0 255.255.255.0 172.24.20.6 20 Override BFD
```

5. Add a security policy for the 10.101.0/24 subnet.

This ensures that the BFD packets to the WAN address are transmitted over an encrypted link. For details, see [Security Features Configuration on EOS Devices on page 473](#).

6. Add static routes from the site routers or site gateways back to the core by entering:

```
ADD -IP ROUTe <IP_address> [<mask>] {<gateway> | !<port>} <metric> [Override] [BFD
[TTL <tTl>]]
```

Step example:

To configure the static routes on the site router or site gateway, enter the following commands:

```
ADD -IP ROUTe 0.0.0.0 0.0.0.0 172.24.16.5 10 BFD
```

```
ADD -IP ROUTe 0.0.0.0 0.0.0.0 172.24.20.5 20 Override BFD
```

7. Optional: Configure a BFD delay timer for the static routes by entering:

```
SETDefault -IP BFDDelayTimer =
<UpStateTimer (0, 5-3600000 ms [Default = 0])
>
[<DownStateTimer 0, 5-3600000 ms [Default = 0]>]
```

Step example:

To configure a BFD delay timer with an UP state timeout value of 5 seconds and a DOWN state timeout value of 3 seconds, enter:

```
SETDefault -IP BFDDelayTimer = 5000 3000
```


BFDDelayTimer is a global value and applies to all static routes configured on the EOS device.

2.4.2

Configuring Ethernet Interzone Links by Using IPIP with BGP

The IPIP with BGP design is used for interzone links. In this configuration, you configure IPIP tunnels and BGP keepalive values. When the BGP peer status changes, the EOS device sends an Ethernet WAN link object state change trap.

Figure 26: Ethernet WAN Links: Sample Interzone Link Configuration



Procedure:

1. Configure the IP addresses of WAN A, WAN B, and WAN C by entering:

```
SETDefault !<portnumber> -IP NETaddr = <IP_address> <subnet_mask>
```

These addresses are provided by the ISP and are routable by the ISP.

Step example:

To configure the IP addresses:

On zone 1, exit router 1, enter:

```
SETDefault !3 -IP NETAddr = 198.16.1.1 255.255.255.0
```

On zone 2, exit router 1, enter:

```
SETDefault !3 -IP NETAddr = 198.16.1.2 255.255.255.0
```

On zone 3, exit router 1, enter:

```
SETDefault !3 -IP NETAddr = 198.16.1.3 255.255.255.0
```

2. Create virtual ports over which the tunnels connecting the zones run by entering:

```
ADD !v<port> -Port VirtualPort IPIP <local_ipaddr> <remote_ipaddr>
>[TTL <t1>]
```

Step example:

To create the virtual tunnel ports:

On zone 1, exit router 1, enter the following commands:

```
ADD !v2 -Port VirtualPort IPIP 198.16.1.1 198.16.1.2 TTL 10
```

```
ADD !v3 -Port VirtualPort IPIP 198.16.1.1 198.16.1.3 TTL 10
```

On zone 2, exit router 1, enter the following commands:

```
ADD !v1 -Port VirtualPort IPIP 198.16.1.2 198.16.1.1 TTL 10
```

```
ADD !v3 -Port VirtualPort IPIP 198.16.1.2 198.16.1.3 TTL 10
```

On zone 3, exit router 1, enter the following commands:

```
ADD !v1 -Port VirtualPort IPIP 198.16.1.3 198.16.1.1 TTL 10
```

```
ADD !v2 -Port VirtualPort IPIP 198.16.1.3 198.16.1.2 TTL 10
```

3. Assign IP addresses to the tunnels that you created in [step 2](#) by entering:

```
SETDefault !v<port> -IP NETAddr = <IP_address> <subnet_mask>
```

Step example:

To assign IP addresses to the IPIP tunnels:

On zone 1, exit router 1, enter the following commands:

```
SETDefault !v2 -IP NETAddr = 172.16.2.9 255.255.255.252
```

```
SETDefault !v3 -IP NETAddr = 172.16.2.13 255.255.255.252
```

On zone 2, exit router 1, enter the following commands:

```
SETDefault !v1 -IP NETAddr = 172.16.2.10 255.255.255.252
```

```
SETDefault !v3 -IP NETAddr = 172.16.4.13 255.255.255.252
```

On zone 3, exit router 1, enter the following commands:

```
SETDefault !v1 -IP NETAddr = 172.16.2.14 255.255.255.252
```

```
SETDefault !v2 -IP NETAddr = 172.16.4.14 255.255.255.252
```

4. Configure the BGP local autonomous system (AS) numbers by entering:

```
SETDefault -BGP LocalAS = <AS_number>
```

Step example:

To configure BGP local AS numbers:

On zone 1, exit router 1, enter:

```
SETDefault -BGP LocalAS = 1
```

On zone 2, exit router 1, enter:

```
SETDefault -BGP LocalAS = 2
```

On zone 3, exit router 1, enter:

```
SETDefault -BGP LocalAS = 3
```


By assigning different AS numbers to the exit routers in the three zones, you specify that the BGP peers are connected through an external BGP session.

5. Configure the external BGP peers by entering:

```
ADD -BGP PEER <IP_address> <AS_number>
```

Step example:

To configure external BGP peers:

On zone 1, exit router 1, enter the following commands:

```
SETDefault -BGP PEER 172.16.2.10 2
```

```
SETDefault -BGP PEER 172.16.2.14 3
```

On zone 2, exit router 1, enter the following commands:

```
SETDefault -BGP PEER 172.16.2.9 1
```

```
SETDefault -BGP PEER 172.16.4.14 3
```

On zone 3, exit router 1, enter the following commands:

```
SETDefault -BGP PEER 172.16.2.13 1
```

```
SETDefault -BGP PEER 172.16.4.13 2
```

6. Configure BGP control by entering the following commands:

```
SETDefault -BGP CONTROL = [All | None]
```

```
SETDefault [!<Peer ID>] -BGP PeerControl= [All | None]
```

Step example:

On zone 1, exit router 1, enter the following commands:

```
SETDefault -BGP CONTROL = All
```

```
SETDefault !172.16.2.10 -BGP PeerControl = All
```

```
SETDefault !172.16.2.14 -BGP PeerControl = All
```

On zone 2, exit router 1, enter the following commands:

```
SETDefault -BGP CONTROL = All
```

```
SETDefault !172.16.2.9 -BGP PeerControl = All
```

```
SETDefault !172.16.4.14 -BGP PeerControl = All
```

On zone 3, exit router 1, enter the following commands:

```
SETDefault -BGP CONTROL = All
```

```
SETDefault !172.16.2.13 -BGP PeerControl = All
```

```
SETDefault !172.16.4.13 -BGP PeerControl = All
```

7. Configure the BGP timers.

For details, see the descriptions of the `IntPolDefault`, `KeepAliveTime`, `HoldTime`, `IGPImportTime`, and `AdvertiseDelay` parameters in “BGP Service Parameters” in the Enterprise OS Software Version 18.2 Reference Guide.

Step example:

To configure the BGP timers, on all three exit routers, enter the following commands:

```
SETDefault -BGP IntPolDefault = (OSPFInt, STATIC)
```

```
SETDefault -BGP KeepAliveTime = 3
```

```
SETDefault -BGP HoldTime = 10
```

```
SETDefault -BGP IGPImportTime = 5
```

```
SETDefault -BGP AdvertiseDelay = 0
```

2.5

Virtual Ports for IPv6 Ethernet WAN Links

IPv6 Ethernet WAN links are logical links that connect two EOS devices using the IPIPv6 protocol on their Ethernet interfaces. These links can connect remote sites and zones in an ASTRO® 25 system. IPv6 site and interzone links use bidirectional forwarding detection (BFD) running over the inner IPv4 payload instead of local management interface (LMI) to determine the status of the link.


When IPv6 Ethernet WAN links are implemented, the EOS device does not perform traffic shaping or rate limiting, and it does not fragment routed packets; however, it does fragment packets sourced by the device itself. The EOS device marks the appropriate quality of service (QoS) bits and leaves it up to the ISP to ensure that latency and jitter limits are not exceeded.

2.5.1

Configuring Ethernet Site Links by Using IPIPv6 with BFD

The IPIPv6 with BFD design is used for Ethernet site links. In these configurations, you configure IPIPv6 tunnels, and BFD runs encapsulated in the tunnel. If encryption is required, you configure another tunnel and BFD runs inside that tunnel.

Figure 27: IPv6 Ethernet WAN Links: Sample Single Site Link Configuration



Procedure:

1. Enable IPv6 routing by entering:

```
SETDefault !<port> -IPv6 CONTrol = R0ute
```

Step example:

To enable IPv6 routing:

On the primary and secondary core routers, enter:

```
SETDefault !3 -IPv6 CONTrol = R0ute
```

On the site router or site gateway, enter:

```
SETDefault !2 -IPv6 CONTrol = R0ute
```

2. Configure the IP addresses of WAN A, WAN B, and the Site WAN by entering:

```
ADD !<portnumber> -IPv6 NETaddr <IPv6_address>
```

These addresses provide the outer header for the IPIPv6 tunnel. They are provided by the ISP and are routable by the ISP.

Step example:

To configure the IP addresses:

On the primary core router, enter:

```
ADD !3 -IPv6 NETaddr fd00::1/32
```

On the secondary core router, enter:

```
ADD !3 -IPv6 NETaddr fd00::2/32
```

On the site router or site gateway, enter:

```
ADD !2 -IPv6 NETaddr fd00::3/32
```

3. Perform one of the following actions:

- If encryption is not required, go to [step 7](#).
- If encryption is required, create virtual ports over which the IPv6 tunnels connecting the primary core router and the site router or site gateway and the secondary core router and the site router or site gateway run by entering:

```
ADD !v<port> -Port VirtualPort IPIPv6 <local_ipv6addr> <remote_ipv6addr> [TTL <t1>]
```

Step example:

To create the virtual tunnel ports:

On the primary core router, enter:

```
ADD !v300 -Port VirtualPort IPIPv6 fd00::1 fd00::3 TTL 10
```

On the secondary core router, enter:

```
ADD !v300 -Port VirtualPort IPIPv6 fd00::2 fd00::3 TTL 10
```

On the site router or site gateway, enter the following commands:

```
ADD !v300 -Port VirtualPort IPIPv6 fd00::3 fd00::1 TTL 10
```

```
ADD !v301 -Port VirtualPort IPIPv6 fd00::3 fd00::1 TTL 10
```

4. Assign IP addresses to the IPv6 tunnels that you created in [step 3](#) by entering:

```
SETDefault !v<port> -IP NETaddr = <IP_address> <subnet_mask>
```

The addresses configured here are significant only to the EOS device and the device's IPv4 encryption tunnel peer. They do not need to be advertised to or known by any other devices in the system.

Step example:

To assign IP addresses to the IPIPv6 tunnels:

On the primary core router, enter:

```
SETDefault !v300 -IP NETaddr = 5.5.5.1 255.255.255.252
```

On the secondary core router, enter:

```
SETDefault !v300 -IP NETaddr = 5.5.5.5 255.255.255.252
```

On the site router or site gateway, enter the following commands:

```
SETDefault !v300 -IP NETaddr = 5.5.5.2 255.255.255.252
```

```
SETDefault !v301 -IP NETaddr = 5.5.5.6 255.255.255.252
```

5. Create the IPv4 encryption tunnels by entering:

```
ADD !v<port> -Port VirtualPort IPIP <local_ipaddr> <remote_ipaddr>
> [TTL <t1>]
```

Step example:

To create the encryption tunnels:

On the primary core router, enter:

```
ADD !v101 -Port VirtualPort IPIP 5.5.5.1 5.5.5.2 TTL 10
```

On the secondary core router, enter:

```
ADD !v101 -Port VirtualPort IPIP 5.5.5.5 5.5.5.6 TTL 10
```

On the site router or site gateway, enter the following commands:

```
ADD !v16 -Port VirtualPort IPIP 5.5.5.2 5.5.5.1 TTL 10
```

```
ADD !v17 -Port VirtualPort IPIP 5.5.5.6 5.5.5.5 TTL 10
```

6. Skip to [step 8](#).

7. Create virtual ports over which the IPv6 tunnels connecting the primary core router and the site router or site gateway and the secondary core router and the site router or site gateway run by entering:

```
ADD !v<port> -Port VirtualPort IIPV6 <local_ip6addr> <remote_ip6addr> [TTL
<t1>]
```

Step example:

To create the virtual tunnel ports:

On the primary core router, enter:

```
ADD !v101 -Port VirtualPort IIPV6 fd00::1 fd00::3 TTL 10
```

On the secondary core router, enter:

```
ADD !v101 -Port VirtualPort IIPV6 fd00::2 fd00::3 TTL 10
```

On the site router or site gateway, enter the following commands:

```
ADD !v16 -Port VirtualPort IIPV6 fd00::3 fd00::1 TTL 10
```

```
ADD !v17 -Port VirtualPort IIPV6 fd00::3 fd00::2 TTL 10
```

8. Assign IP addresses to the tunnels you created in [step 5](#) or [step 7](#) by entering:

```
SETDefault !v<port> -IP NETAddr = <IP_address> <subnet_mask>
```

Step example:

To assign IP addresses to the tunnels:

On the primary core router, enter:

```
SETDefault !v101 -IP NETAddr = 172.24.16.5 255.255.255.252
```

On the secondary core router, enter:

```
SETDefault !v101 -IP NETAddr = 172.24.20.5 255.255.255.252
```

On the site router or site gateway, enter the following commands:

```
SETDefault !v16 -IP NETAddr = 172.24.16.6 255.255.255.252
```

```
SETDefault !v17 -IP NETAddr = 172.24.20.6 255.255.255.252
```

9. Add static routes from the core routers to the site router or site gateway by entering:

```
ADD -IP ROUTe <IP_address> [<mask>] {<gateway> | !<port>} <metric> [Override] [BFD
[TTL <t1>]]
```

The BFD option configures the static routes as BFD-controlled routes. BFD checks for reachability to the gateway address before importing the static route into the routing table.

The optional TTL keyword and <t1> value allow you to specify a single-hop or multi-hop BFD-controlled route. For details, see [Bidirectional Forwarding Detection \(BFD\) Configuration on page 138](#)

and the description of the `ROUte` parameter in “Configuring IP Service Parameters” of the Enterprise OS Software Version 18.2 Reference Guide.

Step example:

To configure the static routes:

On the primary core router, enter:

```
ADD -IP ROUte 10.101.1.0 255.255.255.0 172.24.16.6 10 BFD
```

On the secondary core router, enter:

```
ADD -IP ROUte 10.101.1.0 255.255.255.0 172.24.20.6 20 Override BFD
```

10. Add default static routes from the site routers or site gateways back to the core by entering:

```
ADD -IP ROUte <IP address> [<mask>] {<gateway> | !<port>} <metric> [Override] [BFD  
[TTL <tTL>]]
```

Step example: To configure the static routes on the site router or site gateway, enter the following commands:

```
ADD -IP ROUte 0.0.0.0 0.0.0.0 172.24.16.5 10 BFD
```

```
ADD -IP ROUte 0.0.0.0 0.0.0.0 172.24.20.5 20 Override BFD
```

11. If encryption is required, add dynamic security policies for the tunnels created in [step 5](#).


The `PreSharedKey <peer_ID>` is the tunnel peer IP address. For example, for the primary core router tunnel created in [step 5](#), the peer ID is 5.5.5.2. The dynamic security policies is bound to the virtual ports created in [step 5](#). For details, see [IP Traffic Protection with Dynamic Security Policies on page 479](#).

2.5.2

Configuring Ethernet Interzone Links by Using IPIIPv6 with BGP

The IPIIPv6 with BGP design is used for interzone links. In this configuration, you configure IPv6 tunnels and BGP keepalive values.

Figure 28: IPv6 Ethernet WAN Links: Sample Interzone Link Configuration



Procedure:

1. Enable IPv6 routing by entering:

```
SETDefault !<port> -IPv6 CONTrol = R0ute
```

Step example:

To enable IPv6 routing, on each of the routers, enter:

```
SETDefault !3 -IPv6 CONTrol = R0ute
```

2. Configure the IPv6 addresses of WAN A, WAN B, and WAN C by entering:

```
ADD !<portnumber> -IPv6 NETaddr <IPv6 address>
```

These addresses are provided by the ISP and are routable by the ISP. In this step, you are assigning IPv6 addresses to the physical ports.

Step example:

To configure the IP addresses:

On zone 1, exit router 1, enter:

```
ADD !3 -IPv6 NETAddr fd00::1/32
```

On zone 2, exit router 1, enter:

```
ADD !3 -IPv6 NETAddr fd00::2/32
```

On zone 3, exit router 1, enter:

```
ADD !3 -IPv6 NETAddr fd00::3/32
```

3. Create virtual ports over which the IPv6 tunnels connecting the zones run by entering:

```
ADD !v<port> -Port VirtualPort IIPv6 <local_ipv6addr> <remote_ipv6addr> [TTL <t1>]
```

Step example:

To create the virtual tunnel ports:

On zone 1, exit router 1, enter the following commands:

```
ADD !v2 -Port VirtualPort IIPv6 fd00::1/32 fd00::2/32 TTL 10
```

```
ADD !v3 -Port VirtualPort IIPv6 fd00::1/32 fd00::3/32 TTL 10
```

On zone 2, exit router 1, enter the following commands:

```
ADD !v1 -Port VirtualPort IIPv6 fd00::2/32 fd00::1/32 TTL 10
```

```
ADD !v3 -Port VirtualPort IIPv6 fd00::2/32 fd00::3/32 TTL 10
```

On zone 3, exit router 1, enter the following commands:

```
ADD !v1 -Port VirtualPort IIPv6 fd00::3/32 fd00::1/32 TTL 10
```

```
ADD !v2 -Port VirtualPort IIPv6 fd00::3/32 fd00::2/32 TTL 10
```

4. Assign IPv4 addresses to the tunnels you created in [step 2](#) by entering:

```
SETDefault !v<port> -IP NETAddr = <IP_address> <subnet_mask>
```

Step example:

To assign IP addresses to the IPIP tunnels:

On zone 1, exit router 1, enter the following commands:

```
SETDefault !v2 -IP NETAddr = 172.16.2.9 255.255.255.252
```

```
SETDefault !v3 -IP NETAddr = 172.16.2.13 255.255.255.252
```

On zone 2, exit router 1, enter the following commands:

```
SETDefault !v1 -IP NETAddr = 172.16.2.10 255.255.255.252
```

```
SETDefault !v3 -IP NETAddr = 172.16.4.13 255.255.255.252
```

On zone 3, exit router 1, enter the following commands:

```
SETDefault !v1 -IP NETAddr = 172.16.2.14 255.255.255.252
```

```
SETDefault !v2 -IP NETAddr = 172.16.4.14 255.255.255.252
```

5. Configure the BGP local autonomous system (AS) numbers by entering:

```
SETDefault -BGP LocalAS = <AS_number>
```

Step example:

To configure BGP local AS numbers:

On zone 1, exit router 1, enter:

```
SETDefault -BGP LocalAS = 1
```

On zone 2, exit router 1, enter:

```
SETDefault -BGP LocalAS = 2
```

On zone 3, exit router 1, enter:

```
SETDefault -BGP LocalAS = 3
```

By assigning different AS numbers to the exit routers in the three zones, you specify that the BGP peers are connected through an external BGP session.

6. Configure the external BGP peers by entering:

```
ADD -BGP PEER <IP_address> <AS_number>
```

Step example:

To configure external BGP peers:

On zone 1, exit router 1, enter the following commands:

```
SETDefault -BGP PEER 172.16.2.10 2
```

```
SETDefault -BGP PEER 172.16.2.14 3
```

On zone 2, exit router 1, enter the following commands:

```
SETDefault -BGP PEER 172.16.2.9 1
```

```
SETDefault -BGP PEER 172.16.4.14 3
```

On zone 3, exit router 1, enter the following commands:

```
SETDefault -BGP PEER 172.16.2.13 1
```

```
SETDefault -BGP PEER 172.16.4.13 2
```

7. Configure BGP control by entering the following commands:

```
SETDefault -BGP CONTrol = [All | None]
```

```
SETDefault [!<Peer_ID>] -BGP PeerControl= [All | None]
```

Step example:

To configure BGP control:

On zone 1, exit router 1, enter the following commands:

```
SETDefault -BGP CONTrol = All
```

```
SETDefault !172.16.2.10 -BGP PeerControl = All
```

```
SETDefault !172.16.2.14 -BGP PeerControl = All
```

On zone 2, exit router 1, enter the following commands:

```
SETDefault -BGP CONTrol = All
```

```
SETDefault !172.16.2.9 -BGP PeerControl = All
```

```
SETDefault !172.16.4.14 -BGP PeerControl = All
```

On zone 3, exit router 1, enter the following commands:

```
SETDefault -BGP CONTrol = All
```

```
SETDefault !172.16.2.13 -BGP PeerControl = All
```

```
SETDefault !172.16.4.13 -BGP PeerControl = All
```

8. Configure the BGP timers.

For details, see the descriptions of the IntPolDefault, KeepAliveTime, HoldTime, IGPimportTime, and AdvertiseDelay parameters in “BGP Service Parameters” in the Enterprise OS Software Version 18.2 Reference Guide.

Step example:

To configure the BGP timers, on all three exit routers, enter the following commands:

```
SETDefault -BGP IntPolDefault = (OSPFInt, STATIC)
```

```
SETDefault -BGP KeepAliveTime = 3
```

```
SETDefault -BGP HoldTime = 10
```

```
SETDefault -BGP IGPimportTime = 5
```



```
SETDefault -BGP AdvertiseDelay = 0
```

2.6

Configuring Frame Relay and PPP Virtual Ports

Prerequisites:

Ensure that the owner of the wide area parent port is set appropriately.

Procedure:

1. Create a virtual port for each remote network that is attached to a Frame Relay or ISDN cloud, or that is running the PPP Protocol by entering:

```
ADD !<vport> -PORT VirtualPort
{<path> (<FR_DLCI> | Vlan <vld>)}
|{IPIP [<local_ipaddr> !<port>](<remote_ipaddr> | P2MP) [TTL <t1>] [VRID
<vrid@port-no>, <vrid@port-no>]}
|IPIPV6 <local_ipv6addr> <remote_ipv6addr> [TTL <t1>]
|{SysCallerID "<callerID>"}
|{MFR <P1@D1> [, <P2@D2> .. <P8@D8>]}
| PPP | RAS | TunnelSwitch
```

You do not need to create virtual ports in numerical order. For instance, you can create virtual port v2 before v1.



IMPORTANT: You must use the !v<port> (as opposed to the !<port>) syntax when adding or deleting virtual ports.



NOTE: When you add a Frame Relay virtual port, you must precede the <FRDLCI> value with the following symbol: @

Step example:

If you have a remote network on interface 1 that uses Frame Relay data link connection identifier (DLCI) 35, to add virtual port V1, enter:

```
ADD !v1 -PORT VirtualPort 1@35
```

To create a PPP dial virtual port that uses the dynamic dial pool for its path resources in initiating and receiving calls from a remote device called NewYork, enter:

```
ADD !v3 -PORT VirtualPort SysCallerID "NewYork"
```

This command builds a mapping table entry between the virtual port and a remote device identifier (the site -SYS SysCallerID value of the remote device) and allows an incoming call to be mapped to a specific port. Additional configuration steps are required to use the dial pool. You can create multiple PPP virtual ports, but only one virtual port on a dynamic path can be active at a time.

To create a VLAN virtual port on port 1 for the vlan using 802.1Q identifier 42 (hex 3A), enter:

```
ADD !v9 -PORT VirtualPort 1 Vlan 42
```

2. If virtual port v3 is disabled, re-enable it by entering:

```
SETDefault !v3 -PORT CONTrol = Enabled
```

Virtual ports are enabled by default.

3. Optional: Assign a name to the virtual port by entering:

```
SETDefault !v3 -PORT Name = "<port_name>"
```

Some restrictions apply to the name you assign by using the -PORT Name parameter. For more information, see "Port Service Parameters" in the Enterprise OS Software Version 18.2 Reference Guide.

Step example:

To assign virtual port v3 the name First_St, enter:

```
SETDefault !v3 -PORT Name = "First_St"
```

4. Repeat [step 1](#) to [step 3](#) for each virtual port you want to configure.

The new virtual port settings take effect immediately. For information on Remote Authentication Service (RAS) virtual ports, see [Configuring Remote Access with Tunnel Terminator on page 467](#).

2.7

Configuring Virtual Ports for 802.1Q Virtual LANs

In a virtual LAN environment, stations on a single physical bridged LAN are divided into logical groups and communicate primarily with other stations in that group. Each such subdivision is a virtual LAN (VLAN).

VLANs:

- Create logical broadcast domains, thereby limiting broadcasting to certain ports on an Ethernet switch.
- Allow you to separate stations on the same switch into groups.
- Provide flexibility in security implementation since VLAN stations in one group do not see the traffic destined for other VLAN groups on the same switch.


The EOS software implements standard IEEE 802.1Q VLANs and uses virtual ports to attach to VLANs. When you create a VLAN virtual port, you specify the port that attaches to the physical LAN, and the 802.1Q VLAN identifier. Routing protocols can then use this virtual port to route traffic for the member stations of that VLAN. Each virtual port on a VLAN uses the same MAC addresses as the physical port.



NOTE: EOS devices do not support bridging between ports or virtual ports on the same physical LAN.

The EOS software also supports multicasting over VLANs as shown in the following example.

Figure 29: Multicasting Over VLANs



To set up this configuration, you must configure the VLAN router as described in the following procedure and configure the VLANs on the Ethernet switch according to the instructions of the switch manufacturer.

Procedure:

1. Configure a virtual port that accesses a VLAN by entering:

```
ADD !v<port> -Port VirtualPort <parent_port> vlan <vlid>
```

where <vlid> is the 802.1Q VLAN identifier.

Step example:

To configure virtual port V1 (with parent port 1) to be part of VLAN 2, enter:

```
ADD !V1 -Port VirtualPort 1 vlan 2
```

2. Repeat [step 1](#) to configure additional virtual ports.

Step example:

To configure a second virtual port (V2) to be part of VLAN 3, enter:

```
ADD !V2 -Port VirtualPort 1 vlan 3
```

3. Configure the IP address for each virtual port that you added by entering:

```
SETDefault !v<port> -IP NETAddr = <IP_address> <network_mask>
```

Step example: To configure IP addresses for virtual ports V1 and V2, enter the following commands:

```
SETDefault !V1 -IP NETAddr = 172.16.1.1 255.255.255.0
```

```
SETDefault !V2 -IP NETAddr = 172.16.2.5 255.255.255.0
```

4. Enable IP routing on the EOS device by entering:

```
SETDefault -IP CONTROL = route
```

5. Enable IP multicasting on the EOS device by entering:

```
SETDefault -MIP CONTROL = enable
```

6. Enable PIM sparse mode functionality on the virtual ports that you configured by entering:

```
SETDefault !<portlist> -PIM CONTROL = SparseMode
```

Step example:

To enable PIM sparse mode functionality on virtual ports V1 and V2, enter:

```
SETDefault !v1 !v2 -PIM CONTROL = sm
```

7. Configure the EOS device to advertise itself as a PIM Version2 candidate rendezvous point to the PIM-SM domain's bootstrap router by entering:

```
ADD -PIM CandidateRP <crp_ip_address> AddressList <addr_list_name>
```

where:

<crp_ip_address> is the primary IP network address of one of the device's interfaces

<addr_list_name> is the name of an address list previously created with the ADD -IP AddressList command

Step example:

To configure the VLAN router to advertise itself as a candidate rendezvous point, enter:

```
ADD -PIM CandidateRP 172.16.1.1 AddressList CRP1
```

8. Configure the EOS device to advertise itself as a candidate bootstrap router by entering:

```
ADD -PIM CandidateBSR <bsr_ip_address>
```

where <bsr_ip_address> is the primary IP network address of one of the device's interfaces.

Step example:

To configure the VLAN router to advertise itself as a candidate bootstrap router, enter:

```
ADD -PIM CandidateBSR 172.16.1.1
```

9. Enable PIM sparse mode candidate rendezvous point and bootstrap router functionality on the device by entering:

```
SETDefault -PIM SMControl = (crp, cbsr)
```

2.8

Multiple Instance Lists

The multiple instance list (MIL) feature allows you to apply a command to multiple instances at once, rather than applying the same command to each instance individually.

Instances

An instance can be a port, path, or slot number, a port or path name, or a group name. An instance is defined by ! followed by a path name, port name, or group name. A MIL can be either several instances (separated by a space or by a comma (,)), or a range separated by a hyphen (-).

The command syntax for commands using MIL is:

```
command {! <instance>|! <MIL>} [arguments...]
```

Table 9: MIL Syntaxes for Instances and Their Meanings

Syntax	Meaning
! <instance> - ! <instance>	Indicates a range of instances
! <instance> [,] ! <instance> ! <instance>	Indicates three instances
! <instance> [,] ! <instance> - ! <instance>	Indicates an instance and a range

Groups

A group is a user-defined list of instances using the MIL format.

You can define a group to cover instances that have common characteristics, observing the following rules:

- The group name cannot be *.
- The reserved group names include ALL, and SLOT1 through SLOT8. The reserved group name ALL contains only present instances.
- A group cannot contain any other group as a member.

The parameter name for the group in the System service is InstanceGRoup. The instance list is limited to 80 characters for each ADD -SYS InstanceGRoup command. However, you can append instances to an existing group for a total of up to 1024 characters. You can delete an entire group using the delete command, but you cannot delete a specific instance or instances from a group.

The MIL syntax for group only is:

```
! <Group_Name>
```

```
! <Group_Name> [ , ] ! <Group Name>
```

Table 10: MIL Syntax for Instances and Group and Their Meaning

Syntax	Meaning
! <Group_Name> [,] ! <instance> - ! <instance>	Indicates a group and a range
! <Group_Name> [,] ! <instance>	Indicates a group and an instance

After you defined a group, you can execute a command by using the group name. If a group already exists, the instances are added to that existing group. To add a group, you must use:

```
ADD [-SYS] InstanceGRoup <Group_Name> " <Instance_List>" [ANY|PAth|POrt|SLOt]
```

The group can be a port, path or slot.

Example:

```
ADD -SYS InstanceGRoup gp3 "!1, !3-!4"
```

```
SHow -SYS InstanceGRoup
```

```
=====RESERVED GROUPS=====
```

Group	Name	Type	List of instances
ALL		ANY	All existing ports/paths/slots in the system

```
=====USER DEFINED GROUPS=====
```

Group	Name	Type	List of instances
gp3		P0rt	!1, !3-!4

```
ADD -SYS InstanceGRoup gp3 "!5"
```

```
SHow -SYS InstanceGRoup
```

```
=====RESERVED GROUPS=====
```

Group	Name	Type	List of instances
ALL		ANY	All existing ports/paths/slots in the system

```
=====USER DEFINED GROUPS=====
```

Group	Name	Type	List of instances
gp3		P0rt	!1, !3-!4, !5

You can use the `SHow` command to view all the available groups. You can also use the `SHow` command with the `expand` option to view all existing groups, including reserved groups (if there is a range specified), except for the type `ANY`. The expanded list of instances does not display repetitions of the instances.

For a description of the `InstanceGRoup` parameter, see "SYS Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Chapter 3

Autostartup Configuration

The Autostartup feature enables Motorola Solutions Enterprise OS devices to start up and become operational with no software configuration on the device itself.

The Autostartup feature is a two-phase process. See [Autostartup Phase 1 on page 98](#) and [Autostartup Phase 2 on page 99](#).

For the Autostartup feature to work, you must set up a network management station at a central site location. On the central site network you must also set up a BootP and a Trivial File Transfer Protocol (TFTP) server.

You can set up your network to automatically start up the central site device. The central site device assists in automatically starting any peripheral nodes contained in the configuration on the network management station.

3.1

Autostartup Phase 1

Phase 1 of the Autostartup process initiates upon plugging in the EOS device. During Phase 1 an EOS device automatically detects the following wide area port and path attributes:

- WAN protocol (PPP or Frame Relay) that runs on a port
- Line type (leased or dial-up)

Phase 1 detects the path attributes instantly on initial system startup; however, if a connector is changed during normal operation, it can take several minutes for the autodetection software to sense the connector.

Automatic Attribute Detection for DTE Ports on Remote Routers

For certain EOS devices, the Autostartup phase 1 process also detects the DTE port that you cabled. This process works only on DTE ports and only when the `-PORT OWNER` parameter is set to `AUTO` (default). Autostartup can take several (three to five) minutes if you change the cable during normal system operation.

When establishing the physical and data links between themselves and the central node, the EOS devices attempt to detect the connector type, owner, and line type of the path associated with the cabled DTE port. The remote device detects these path characteristics by first attempting to try connector and line type combinations. The connector and line types are tried in the following order:

- RS-232 with leased line
- RS-232 with dial-up line
- RS-449/V.36 with leased line
- RS-449/V.36 with dial-up line
- V.35 with leased line
- V.35 with dial-up line

The EOS device scans each line quickly to determine the connector type most likely used. After it successfully detects the connector and line types, the device tries to detect the owner using a similar process. The scanning process continues periodically so that connector changes can be quickly determined. The possible owners are tried in the following order:

- PPP
- Frame Relay

Knowing the order in which the connector and line types and owners are tried can help you anticipate how long the establishment of the physical and data links takes. For example, the detection of a V.35 dial-up line running Frame Relay takes longer than the detection of an RS-232 dial-up line running Frame Relay because the V.35 dial-up line is tried later than an RS-232 dial-up line.

To determine the progress of physical and data link establishment, you can use [Determining the Progress of the Establishment of the Physical and Data Links on page 99](#).

After the wide area port and path attributes are detected, the EOS device establishes a physical link and then a data link between itself and the central node. Phase 1 is complete, and phase 2 of the Autostartup process begins. For more information, see [Autostartup Phase 2 on page 99](#).

3.1.1

Determining the Progress of the Establishment of the Physical and Data Links

Procedure:

1. Enter:

```
SHow -PORT DIAGnostics
```

A display shows you which connector type, port owner, and line type was tested and failed, and which is currently tested.

2. Check the LEDs associated with the DTE connectors on the EOS device to determine which connector is currently tried.

Because of a signal irregularity in the RS-449 connector, the Autostartup detection feature may detect the RS-449 connector as a V.35 connector. This condition eventually corrects itself.

3.2

Autostartup Phase 2

During phase 2 of the Autostartup process, the peripheral node obtains necessary configuration information from central site servers across the PPP or Frame Relay line, including:

- IP address
- Boot file location
- Configuration files

For phase 2 to work, you must configure the UDP Broadcast Helper feature on the central node and you must configure two servers on the central site network: a BootP server and a TFTP server; both server programs can reside on the same or different host machines. The BootP server “listens” for BootP requests and forwards an IP address and boot file location information to the requesting node. The TFTP server forwards configuration files to the remote node.

During phase 2:

- The remote node device’s BootP client sends a request out all available interfaces that do not have an IP address assigned to them. The “file” field of this request includes a model number that identifies the type of hardware from which the request is sent. This allows you to configure the BootP server to supply a generic or common filename for a group of clients. For example, you can configure the BootP server to provide one configuration file for all MNR S2500s, another configuration file for all MNR S6000s, and another configuration file for all GGM 8000s.
- The central node forwards the BootP request to the BootP server on the central site network using the UDP Broadcast Helper feature.

- The BootP server searches its database for an entry that matches the MAC address in the client's request and responds to the broadcast BootP request packet with the IP address of a TFTP server and a path to the appropriate configuration files.
 - The files can be in any directory in the TFTP server's file system, but `CONFFILE` and that files it lists must be in the same directory.
 - You can configure the server to supply a generic filename for a group of clients. For example, you might want your S2500s to obtain `fileA`, your S6000s to obtain `fileB`, and your GGM 8000s to obtain `fileC`.
- The remote device requests the file specified in the BootP reply from the TFTP server on the central site network, taking into account the hardware model information in the `file` field of the BootP reply. If the BootP client does not successfully download the requested file after up to 15 retries, it sends a new request (and up to 15 retries) to the same TFTP server.
- If the character string in the BootP reply ends in `CONFFILE`, the EOS TFTP client uses the same path as specified for `CONFFILE` to request all the files for which file names were listed in the `CONFFILE` just downloaded. In this case, the `CONFFILE` – assumed to be a plain ASCII text file containing a list of space-separated names of Compact Configuration Storage (CCS) files – is parsed to extract the names of binary CCS configuration files, which are then downloaded individually to the default file system in non-volatile memory.
- When the file transfer is complete, the remote node device automatically applies the newly-acquired configuration files. In the final steps of its self-initialization, EOS checks for the presence of a `boot.cfg` file in the default directory.
 - If EOS does not find a `boot.cfg` file, it reboots the device, providing an opportunity to read the CCS files that were previously downloaded.
 - If EOS finds a `boot.cfg` file, it does not reboot the device but rather assumes that the `boot.cfg` file contains user interface commands in ASCII text, which it executes one by one, causing configuration information to be saved in appropriate CCS files. Often the last command in `boot.cfg` is to rename or remove itself so it is not read as part of a subsequent bootstrap.

3.3

Preparing Autostartup

For the Autostartup feature to work, the central node requires certain software tools and some software configuration. Before beginning the Autostartup configuration procedure, you must complete the following actions on the central node if Autostartup is through a central node. If the autostarted device is connected to the central site network management station through a local area network, these actions are not necessary.

Prerequisites:

Configure a BootP server and a TFTP server on the central site network. For information on configuring the BootP and TFTP servers, see [Configuring the Central Node and the BootP/TFTP Servers on the Central Site Network on page 101](#).

Procedure:

1. Log on with Network Manager privilege.
2. Set up ports and paths. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
3. Note the port numbers and associated IP addresses that are used for the Autostartup phase 2 connection.
4. Configure a WAN port over which the peripheral nodes are going to communicate.

5. Examine your network and determine (or assign) the IP address and/or the MAC address for each device that requires Autostartup support.
6. Determine the server on which the BootP server is going to reside.
7. Determine the server on which the TFTP server is going to reside.
8. For Frame Relay configurations, determine the data link connection identifier (DLCI) for each remote device that requires Autostartup phase 2 support. The DLCI is assigned to the Frame Relay interface by the public data network (PDN) service vendor.

3.4

Configuring the Central Node and the BootP/TFTP Servers on the Central Site Network

You can use this procedure to configure the central node and the BootP/TFTP servers on the central site network. The Frame Relay connection between the central and peripheral nodes in the following figure has a DLCI of 30 assigned to it by the Frame Relay service vendor.

The following figures illustrate sample topologies in which Autostartup is configured for Frame Relay and PPP.

Figure 30: Autostartup Configuration for Frame Relay





Figure 31: Autostartup Configuration for PPP



The following items are the same in both sample topologies:

- The topologies included both a BootP server and a TFTP server. The BootP server has the IP address 129.213.201.25, and the TFTP server has the IP address 129.213.201.24. However, the BootP server and the TFTP server could reside on the same host.

- The central node is functioning as a router with an IP address of 129.213.201.21. The User Datagram Protocol (UDP) Broadcast Helper feature was enabled on the central node.

Procedure:

1. Configure the UDP Broadcast Helper on the central node by performing the following actions:

- a. Enable UDP Broadcast Helper by entering:

```
SETDefault -UDPHELP CONTro1 = Enable
```

- b. Add BootP server UDP port 67 to the active port list by entering:

```
ADD -UDPHELP ActivePorts {<UDP_port> | <name>}
```

BPSEVER is the name reserved for port 67.

Step example:

To add UDP port 67 to the active port list, enter:

```
ADD -UDPHELP ActivePorts 67
```

or

```
ADD -UDPHELP ActivePorts BPSEVER
```

- c. Add the IP address of the BootP server into the forward address list by entering:

```
ADD -UDPHELP ForwardAddress <UDP_port_or_name> <IP_address>
```

For more information on the UDP Broadcast Helper, see [User Datagram Protocol Broadcast Helper Configuration on page 447](#).

Step example:

To add the IP address of the BootP server (129.213.201.25) to the forward address list, enter:

```
ADD -UDPHELP ForwardAddress 67 129.213.201.25
```

2. Install a BootP server on your system.



NOTE: This procedure varies, depending on what BootP server program you are using. Refer to the documentation that accompanies your BootP server application for specific instructions.

You must specify the MAC address of the port on the peripheral node over which Autostartup is executed for each entry in the `bootptab` file. These entries contain the following information:

```
<server>:ip=<ip_address>:hd=<config_directory>:\:bf=<boot_file>:  
ht=<hardware_type>:ha=<mac_address>:\fs=<tftp_address>:
```

where:

<server> – the name of the remote device

<ip_address> – the IP address that is assigned to the system named <server>

<config_directory> – the pathname of the home directory on the TFTP server where the configuration files exist

<boot_file> – is the name of the boot file

<hardware_type> – the hardware type that specifies the type of MAC, Frame Relay, or PPP address of the peripheral node that is executing the startup

<mac_address> – the MAC address of the port over which Autostartup is executed

<tftp_address> – the IP address of the TFTP server

Step example:

```
remote: ip=129.213.201.22:hd=config-directory:\  
:bf=boot.68k:ht=ethernet:ha=080020011380:\  
fs=129.213.201.24
```

3. Set up the TFTP server by performing the following actions:

- a. Create the configuration file directory on the TFTP server.

If the file name does not end with `CONFFILE`, `.ppc`, `.68K`, or `.29K`, BootP first tries the path and file name you entered to retrieve configuration files from the TFTP server. If that fails, the BootP client tries IP-based addressing. Finally, if that fails, it tries MAC-based addressing. The BootP "Root Path" option may be used by the server.

When you use IP addressing to identify remote nodes, the configuration file directory could also be `/tftpbboot/autostartup/<ip-address-of-the-remote-node>`.

When you use MAC addressing, you must create a directory with the path name as `root_directory/CLIENTS/<MAC_address>`, where:

`<root_directory>` is the path of the grandparent directory of the host image files

`<MAC_address>` is the MAC address of the device



NOTE: For the TFTP server, the path name is case-sensitive. The directory `CLIENTS` must be uppercase. System-dependent path separators `/` or `\` are both acceptable characters.

- b. Create configuration files for each device (or each group of devices with the same configuration) for which you want to use the Autostartup feature.

Configuration files can be one of two types of files: a single ASCII text file or a set of encoded configuration files specific to the device(s) that is autostarted.

The most convenient and manageable method of generating a configuration file is to create an ASCII text configuration file using a text editor. An ASCII text configuration file contains ASCII text string of the commands required to configure the services that the device needs. Since this is an ASCII text file, you can modify easily as necessary. For information on the contents of the ASCII text file, see [ASCII Text Configuration File on page 107](#) or the `LoadConfigs` described "Commands" in the *Enterprise OS Software Version 18.2 Reference Guide*. The ASCII file must be named `boot.cfg` to be automatically read during bootup.

You can also use encoded configuration files called compact configuration store (CCS) files. Creating CCS files requires that you use an unconfigured device of the same type (and hardware configuration) as the device to be autostarted. You must enter the required configuration commands for the device to be autostarted on the unconfigured device as if it were the device to be autostarted. This process creates the necessary CCS files. As the CCS files are not text files, they are not directly editable.

- c. Copy the configuration files that you created in [step 3b](#) from the system you used to create them to the configuration directory you created in [step 3a](#).

- d. Perform one of the following actions:

- If you created an ASCII text file, go to [step 4](#).

- If you created CCS files, use a text editor to create an ASCII text file called `CONFFILE` in the configuration directory on the TFTP server.

You must create a `CONFFILE` for each device (or group of devices) that you want to autostart. `CONFFILE` must contain the names of the CCS files you created.

The configuration file may contain:

```
ip<sep>iprip<sep>rtmnet<sep>system<sep>
```

where:

`ip`, `iprip`, `rtmnet`, and `system` are the names of the CCS files contained in the same directory.

`<sep>` is the separator of each file. The `<sep>` can be a blank, a tab, a form feed, a carriage return, or a new-line character.

The configuration directory can be located anywhere in the TFTP server's file system, but `CONFFILE` and the file(s) it lists must be in the same directory.

4. When all the required configuration files and services are in place at the central site, set up the remote sites.

For information on installing and cabling the remote node, see the appropriate hardware user guide.

5. Plug in the remote node device.

The device starts up.

The initiation of the Autostartup phase 2 process depends on the following line types used for the physical link:

- If you use a leased line, you must plug in the appropriate cables between the EOS device and the modem to which the leased line is connected.
- If you use a dial-up line, you must configure the modem or channel service unit/digital service unit (CSU/DSU) to dial the central node so that a physical link can be established.

3.5

Autostarting the Central Site Node

Procedure:

1. Set up the central site node configuration file on the TFTP server and the BootP. See [Configuring the Central Node and the BootP/TFTP Servers on the Central Site Network on page 101](#).
2. Reset the central site node.

When the device cannot find configuration files to load, the system sends a BootP request trying each port until it receives a response from an external BootP server.

The BootP server provides the central site device with the address of the TFTP server that has its configuration file. The device then sends a TFTP request for a load of its configuration file. The node loads the configuration file. If the configuration files are CCS files, the node reboots in order for CCS configuration to take place. If the configuration file is an ASCII (`boot.cfg`) file, the node executes the EOS commands listed in the file and enters its normal operating state.

3.6

Sample Autostartup Configurations

In these sample configurations, the WAN link can be either PPP or Frame Relay. You can use any standard BootP and TFTP servers.




NOTE: The identifying tag in the `bootptab` file can only be the Ethernet address – the WAN port MAC address of the peripheral device over which the BootP request PDU travels to the central device.

3.6.1

Setting Up Configuration with Frame Relay WAN

The following figure demonstrates a typical configuration with an EOS device as the central node and a Frame Relay WAN.

Figure 32: Autostartup with Motorola Solutions Central Node and Frame Relay WAN



Procedure:

1. Plug in the peripheral node.
2. In the central node, perform the following actions:
 - a. Set an IP address for the Ethernet port, using the IP address 200.213.44.120.
 - b. Set an IP address for the WAN port, using the IP address 200.213.45.120.
 - c. Determine the DLCI over which the peripheral device needs to communicate, in this case 101.
 - d. Determine the port to which the DLCI is linked, in this case V1.
 - e. Enable IP routing.
 - f. Set up UDPHELP with 67 as the active port and the BootP server as the forward address, using 200.213.44.100 as the address.

3. At the BootP server, edit the /etc/BootPtab file and enter the following information:

```
Router1:ip=200.213.45.55:sm=255.255.255.0:hh=frame:hp=V1:\
bh=200.213.45.120:ci=101:fs=200.213.44.100:gw=200.213.45.120:
```

The BootP reply from the BootP server is sent to the gateway (200.213.45.120), so you must set 200.213.44.120 as the gateway for the BootP server for this address.

4. At the TFTP Server, create one of the following files:

- /my_root_dir/subdir1/subdir2/subdir3/boot.cfg
- /tftpboot/autostartup/200.213.45.55/CONFFILE
- /my_root_dir/subdir1/subdir2/CLIENTS/080000222512/CONFFILE

where: CONFFILE is a white space (space, return, tab)-separated list of CCS files and/or boot.cfg and "boot.cfg" is a list of EOS UI commands, one per line.

CONFFILE is a white space (space, return, tab)-separated list of CCS files and/or boot.cfg

boot.cfg is a list of EOS UI commands, one per line.

The last command should be a rename command so that the next time the device boots it does not execute the commands in boot.cfg again. A config.log file is created in the configuration directory to log the execution of the commands in the boot.cfg file. In the case that config.log already exists, the existing file is renamed as config.bak, and a new config.log file is created. For further details about working with ASCII configuration files, see [ASCII Text Configuration File on page 107](#).


capture.cfg does not capture the commands executed in the boot.cfg file (ASCII boot).

3.6.2

Setting Up Configuration with PP WAN

The following illustration, demonstrates a typical configuration with an EOS device as the central node and a PPP WAN.

Figure 33: Autostartup with Central Node BootP Server and PPP WAN



Procedure:

1. Plug in the peripheral node.
2. At the central site, perform the following actions:
 - a. Enable PPP on the WAN.
 - b. Set up UDPHELP: with 67 as the active port and the BootP server as the forwarding address by using 200.213.44.100 as the address.
3. At the BootP server, edit the `/etc/BootPtab` file and enter the following information:

```
Router2:ip=200.213.45.55:sm=255.255.255.0:hh=ppp:hp=3C:\
```

```
bh=200.213.45.120:fs=200.213.44.100:gw=200.213.45.120:
```

The BootP reply from the BootP server is sent to the gateway (200.213.45.120), so you must set 200.213.44.120 as the gateway for the BootP server for this address.

4. At the TFTP Server, create one of the following files:

- `/my_root_dir/subdir1/subdir2/subdir3/boot.cfg`
- `/tftpboot/autostartup/200.213.45.55/CONFFILE`
- `/my_root_dir/subdir1/subdir2/CLIENTS/080000222512/CONFFILE`

where: CONFFILE is a white space (space, return, tab)-separated list of CCS files and/or `boot.cfg` and "boot.cfg" is a list of EOS UI commands, one per line.

CONFFILE is a white space (space, return, tab)-separated list of CCS files and/or `boot.cfg`

`boot.cfg` is a list of EOS UI commands, one per line.

The last command should be a rename command so that the next time the device boots it does not execute the commands in `boot.cfg` again. A `config.log` file is created in the configuration directory to log the execution of the commands in the `boot.cfg` file. In the case that `config.log` already exists, the existing file is renamed as `config.bak`, and a new `config.log` file is created. For further details about working with ASCII configuration files, see [ASCII Text Configuration File on page 107](#).

`capture.cfg` does not capture the commands executed in the `boot.cfg` file (ASCII boot).

Chapter 4

ASCII Text Configuration File

When an EOS device is configured, encoded configuration files are created automatically and stored on the device's local storage. These encoded configuration files are automatically updated as additional configuration changes are made. If the device should reboot, these encoded configuration files are read to ensure that the device comes up with the same configuration as before the reboot.

When you upgrade the device to a new software version, you must upgrade these encoded configuration files for the new release. Because these configuration files are in an encoded format, however, you cannot easily edit them to make changes. In addition, if you want to configure another similar device, you must re-enter all of the configuration commands at the new device.

By using ASCII text files, rather than existing binary configuration files, whenever you load a software upgrade on a device, you avoid this inconvenience. An ASCII text configuration file contains the configuration commands for a specific device. You can create ASCII text configuration files using any simple text editor on a PC or UNIX workstation.

When GGSN is enabled on an EOS device, the device supports two configuration files:

- `boot.cfg` – routing, port, and path configurations
- `XGSN.cfg` – GGSN, virtual port, and APN configurations

The `XGSN.cfg` file is created by a network management application. It contains the GGSN configuration parameters and the APN configuration commands for the APNs configured on the device. These are the dynamic parts of the device configuration. For details, see [GGSN Configuration File Management on page 601](#).

4.1

ASCII Text Configuration File Download

To download the ASCII text configuration file to the device, you must use the `copy` (TFTP file transfer) or `GET` (FTP file transfer) command at the device.



NOTE: Some minimal configuration of the device is needed before the file transfer operation can occur.

The configuration commands for the various services are documented in the *Enterprise OS Software Version 18.2 Reference Guide*. In addition, see the appropriate hardware user guide for additional startup instructions.

4.2

Configuration File Execution with LoadConfigs

After you download the ASCII file to the device, you can execute the `LoadConfigs` command at the device to execute all of the configuration commands in the ASCII text configuration file. When you use `LoadConfigs` command to execute these configuration commands, the encoded configuration files are still created or updated, as necessary.

For further information about the `LoadConfigs` command, see “Commands” in the *Enterprise OS Software Version 18.2 Reference Guide*.

4.3

Additional Configuration File Uses

After you created an ASCII text configuration file, you can use it to configure additional devices by simply editing the file to contain configuration commands unique for each device.

ASCII text configuration files allow you to utilize the following extensions of the `LoadConfigs` command:

ASCII Boot

The ASCII Boot feature allows you to execute the equivalent of a `LoadConfigs` command automatically when the device is booted. In addition, when ASCII boot is executed, the device's previous configuration is deleted. The device's previous configuration is not deleted when `LoadConfigs` is executed. This feature is called ASCII Boot because you are booting your device with an ASCII file that is used to configure the device. For more information about ASCII Boot, see [Configuring a Device by Using ASCII Boot on page 109](#).

Autostartup

The Autostartup feature allows you to boot an unconfigured device and load it with its configuration without entering any configuration commands at the device. For more information about Autostartup, see [Autostartup Configuration on page 98](#).

ASCII Capture

The ASCII Capture feature allows you to capture all configuration changes after the device is booted, with the exception of the `UserManage` command, in an ASCII text file in the configuration directory. The configuration changes are captured in a UI command format that can be executed by the `LoadConfigs` command. Configuration commands are captured regardless of whether they were originally executed through menu interface, UI command, `REmote` command, scheduler, or SNMP SET requests. Macro definitions executed with the `DEfine` command and macro definitions deleted with the `UNDefine` command are captured in the `capture.cfg` file. However configuration commands executed by macros (using the `DO` command) are logged in a separate file (`macro.log`) and are not captured in the `capture.cfg` file. These `DO` commands are identified with a comment line in the `capture.cfg` file. For more information about ASCII Capture, see [ASCII Capture Feature on page 114](#).



NOTE: The configuration of passwords and other secure parameters is also captured, but not with the value you entered. This is because the capture file is a clear text file and for security reasons, the password/secure parameter values are not captured in clear text.

4.4

StaticRP.cfg File

In addition to the `boot.cfg` and `XGSN.cfg` files, EOS devices support a static RP configuration file (`StaticRP.cfg`). The `StaticRP.cfg` file is an ASCII text file that contains the EOS CLI commands required to configure the static RP set for a particular PIM domain. For details about configuring static RP sets for PIM devices, see [Static RP Sets on page 297](#). For details about the PIM service `StaticRP` parameter, used to add static RP sets, see "PIM Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

The `StaticRP.cfg` can be created with any simple text editor on a PC or UNIX workstation. You can include it in the PIM devices' `boot.cfg` file by using the `boot.cfg` file `include` command. For details, see [Include Command in boot.cfg Files on page 113](#).



NOTE: If you configure RP sets for a PIM domain using the `StaticRP.cfg` file, it is recommended that you disable BSR functionality (set the PIM service `BSRControl` parameter to `Disable`) on all the PIM devices so that the devices do not accept BSR packets.

4.5

ASCII Boot and Capture Features

The ASCII Boot feature is a special extension of the `LoadConfigs` command. It allows you to execute the equivalent of a `LoadConfigs` command automatically when you boot a device. The ASCII Boot feature allows to configure and maintain the complete configuration of a device in a single ASCII text configuration file.

When the device is booted and the `boot.cfg` file is detected in the configuration directory, all existing configuration files (except `IOXM`) are deleted from the configuration directory. The configuration commands in the `boot.cfg` file are then executed. If the `IgnoreErrors` option is enabled, the process does not stop if a configuration command fails to execute successfully, but rather continues with the next configuration command. Just as when you enter a configuration command on the command line, new encoded configuration files are created. As long as the `boot.cfg` file is detected in the configuration directory, this operation is repeated every time the device is rebooted.

By default, any additional configuration changes you make to the device after it is operational are not added in the ASCII Boot file (`boot.cfg`). This means that if the device is rebooted and the ASCII Boot feature is invoked again, the additional configuration changes are not effective. However, the ASCII Capture feature automatically captures the additional configuration changes in the ASCII Capture file (`capture.cfg`) in the configuration directory. You can set up the ASCII Capture feature to automatically append all additional configuration changes to the end of the `boot.cfg` file by configuring the `SYS services CAPTURE` parameter to the appropriate value.

4.6

Configuring a Device by Using ASCII Boot

Procedure:

1. By using any simple text editor on a PC or UNIX workstation, create an ASCII text file that contains all UI configuration commands needed to configure the device.



NOTE: Do not create the ASCII text file with a word processor that embeds text and format processing command codes in the file.

The ASCII text file must follow the same rules as any file executed by the `LoadConfigs` command. These rules are as follows:

- Use only one command per line. A line is zero or more ASCII characters terminated by a carriage return, line feed, carriage return followed by line feed, or line feed followed by carriage return.
- Use only the following commands: `SETDefault`, `ADD`, `DEL`, `SHow`, `SHowDefault`, `ReName`, and `SAVEbgp`.
- Include blank lines and comments (lines with `#` as the first character) if desired.
- If the command requires more than one line, use the backslash (`\`) to indicate that the command is to continue on the next line.

An example of an ASCII boot text file:

```
# SYS configuration #
SETD -SYS NMPrompt = "R-205"
SETD -SYS SysName = "ROUTER-205"
SETD -SYS WelcomeString = "Router 5 in bldg 200"
# IP configuration #
SETD !1 -IP NETaddr=101.101.101.101 255.255.255.0
SETD -IP CONTROL = (R0ute, SECurity)
SETD -IP SecCONTROL = (EXTended, LabelAdd)
# Firewall configuration #
SETD !1 -FireWall CONTROL=Filter
SETD !1 -FireWall DefAction=Log
```

```
ADD !1 -FireWall FTPIn Permit  
ADD !1 -FireWall FTPOut Permit
```

You can use the ASCII Boot process to configure additional devices by simply editing the ASCII text configuration file to generate configuration commands unique for each device.

2. Boot the device, if it is not already operational.

For information on how to boot the device, see the appropriate hardware user guide.

3. Configure the device for IP routing, if it is not configured.

For instruction on how to configure the device for IP routing, see [IP Routing Configuration on page 119](#).

4. To verify that you have connectivity between the network management station and the device, use the PING command.

You should try this at both the device and the network management station. For information about the PING command, see “Commands” in the *Enterprise OS Software Version 18.2 Reference Guide*.

5. Transfer the text file that you created in [step 1](#) to the device, and store it in the configuration directory with the `boot.cfg` file name, by performing the following actions:

- a. Set the default directory on the device for file transfer commands to the configuration directory, by entering:

```
ChangeDir
```

- b. To transfer the file by using TFTP, enter:

```
COPY <TFTP_server_IP_address>:<path>/<filename> boot.cfg
```

- c. To transfer the file using FTP, use:

```
GET <FTP_server_IP_address>:<path>/<filename> boot.cfg
```

Before you can use FTP to transfer the file, you must use the `SysconF` command to set up the FTP username and password for the device’s FTP client to use to connect to the FTP server. See “Commands” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: For the ASCII boot feature to work, the name of the ASCII text command file on the device must be `boot.cfg` and the file must be stored in the configuration file directory. File names on the EOS device are not case-sensitive.

6. Reboot the device.

When the device comes up, any pre-existing configuration files in the configuration directory are deleted. The following messages are displayed before and after the existing configuration files are deleted:

```
Deleting existing configuration files
```

```
Completed deleting configuration files
```

`LoadConfigs` then attempts to execute all of the UI commands in the `boot.cfg` file. The following message is displayed before the commands in the `boot.cfg` file are executed:

```
Executing configuration commands in BOOT.CFG
```

For larger ASCII boot files, the following message is displayed until all the commands in the `boot.cfg` file are executed:

```
System Initialized and Running
```

Do not enter any additional commands until all the commands in the `boot.cfg` file are executed, as indicated by the following message:

```
Completed executing configuration commands in BOOT.CFG
```

The commands executed and the output generated by executing the commands are displayed on the console and also written to the log file (`config.log`) in the configuration file directory.



NOTE: Unless the `InterAction` parameter was configured in the `boot.cfg` file to enable the display of `LoadConfigs` commands, the configuration commands in the `boot.cfg` file are not displayed as they are executed. For more information about the `InterAction` parameter. See “Commands” in the *Enterprise OS Software Version 18.2 Reference Guide*.

7. Check the log file to make sure all commands executed successfully by entering:

```
CAT primary/config.log
```

For details about displaying and interpreting the log file, see [Boot, Capture, and Log Files Contents Review on page 116](#).



NOTE: After booting with the `boot.cfg` file, any changes you make to the device via Telnet, console, or SNMP are not saved to the `boot.cfg` file. We recommend that you add all captured changes to the `boot.cfg` file at a PC or UNIX workstation and download the updated file to the EOS device in the `boot.cfg` file directly. You can also have the configuration commands captured by ASCII Capture and automatically appended to the `boot.cfg` file by configuring the `SYS services CAPTURE` parameter to the appropriate value. For details, see [ASCII Capture Feature on page 114](#).

If the log file indicates all of the commands executed successfully, the ASCII boot is complete.

8. If the log file indicates that any of the configuration commands failed, edit the ASCII text file (`boot.cfg`) at the network management station and repeat [step 2](#) to [step 7](#).

4.6.1

boot.cfg File Renaming

After all of the configuration commands executed and the software is operational, you may choose to rename the `boot.cfg` file. Created encoded configuration files can be used the next time the device is booted, instead of deleting all of these encoded files and executing all of the configuration commands. You can do this by making the last command in the `boot.cfg` file a `ReName` command, to rename the `boot.cfg` file to a different file name.

4.6.2

Macros in the boot.cfg File

With EOS version 11.7.4 and higher, macros are allowed in the ASCII `boot.cfg` configuration file. In previous releases, you could define macros only by typing the macro definition at the EOS prompt or by copying the encrypted `CCSMACRO` file to the device. The inclusion of macros in the `boot.cfg` file simplifies the management and serviceability of large networks of EOS devices, as configuration of a device can now be accomplished with a single editable file—the ASCII `boot.cfg` file.



IMPORTANT: You should rename your `CCSMACRO` file before loading Release 11.7.4 or higher. Beginning with release 11.7.4, EOS includes a new implementation that allows macros to be included in the `boot.cfg` file. If you load Release 11.7.4 or higher and use `boot.cfg`, your `CCSMACRO` file is deleted, and you need to recreate any existing macros.

Macros can only be defined and shown from the `boot.cfg` file; they cannot be executed by the `do <macro_name>` command; macro execution from inside `boot.cfg` would delay bootup time.

Defining Macros in the boot.cfg File

To define macros in the `boot.cfg` file, you can use the following syntax:

DEFINE *<macro_name>* = (*<text>*)

For example:

```
DEFINE wan_down = (  
  setd !1 -pim cont = nsm  
  setd !1 -vrrp cont = d 1  
  setd !1 -ip net = $1 255.255.255.0  
)
```

The DEFINE command signals the beginning of a macro definition, and the closing parenthesis terminates the macro definition.

- As soon as the closing parenthesis is encountered, the macro definition is placed in a macro cache, regardless of whether or not the definition is correct. If the macro definition is not correct, any errors are reported when the macro is executed via a do *<macro_name>* command.
- If the macro definition does not include a closing parenthesis within 1023 characters (excluding blank spaces), the following error message is reported:

macro definition too long

and the next line after 1023 characters is taken as the next boot.cfg command.

When you include a continued line (a single command that spans multiple lines) within a macro, you must use a backwards slash (\) preceded by a space to indicate that the command continues on the next line.

For example:

```
DEFINE continued_line = (  
  echo "this is \  
  a test macro"  
)
```

results in:

```
[1]EOS # show macro continued_line  
MACRO CONTINUED_LINE = (  
  echo "this is a test macro"  
)
```

Comments are allowed in EOS macro definitions. To identify a comment, you must use a hash mark (#) as the first character in the line.

For example:

```
define bringuphssi=(  
  #This macro will bring up the HSSI link  
  SETDefault !$1 -PAth CONTrol = Enable  
)
```

Storing Macros

When the boot.cfg file is executed, either during the bootup or through the LoadConfigs command from the UI, the macros definitions specified in the boot.cfg file are placed in a macro cache and can then be executed via a do *<macro_name>* command.

When macros are defined from the UI prompt, they are written both to the macro cache and to compact configuration storage (CCS) – the flash file system on EOS devices. When macros are defined from the boot.cfg file, on the other hand, they are not written to flash until the boot.cfg file is executed completely. When boot.cfg file is executed completely, the macros stored in the cache are written to a single CCSMACRO file on flash.

Whenever a device reboots, the device initialization service (IS) checks for the existence of the boot.cfg file in the configuration file directory. If the file exists, the device is configured from scratch with the configuration

commands contained in the `boot.cfg` file, and any existing encoded configuration files – including the `CCSMACRO` file – are deleted.

4.6.3

Include Command in `boot.cfg` Files

The `include` command in the `boot.cfg` file allows you to include device configuration files, such as the `StaticRP.cfg` file, in the `boot.cfg` file using the following syntax:

```
include <filename>.cfg
```

where:

<filename> is a filename of up to eight characters

<.cfg> is a file extension of up to three characters

Device configuration files that are included in the `boot.cfg` in this way are then stored permanently in the devices. When a device with a `boot.cfg` file boots up, EOS reads the file contents into memory, then parses the commands and executes each command. The `include` command directs the software to also read the new file (for example the `StaticRP.cfg` file), add its contents to memory, and parse and execute the commands contained in the file. During command execution the software displays a message when it is executing commands from the included file.

If a command in the included file fails, the EOS software displays the line number in the included file where the error occurred. In addition, the `config.log` file indicates the line(s) in the `boot.cfg` or included file that had errors.

The following is an example of the `boot.cfg` `include` command. In this example, the `boot.cfg` file includes a file called `test.cfg`, and both the `boot.cfg` and the `test.cfg` files contain errors.

`boot.cfg` file contents

```
setd !1 -ip net = 20.20.20.55 255.255.255.0
setd !2 -ip net = 66.66.66.66 255.255.255.0
include test.cfg
add -ip rou 10.79.134.0 255.255.255.0 20.20.20.34 5
setd -ip ser = fer-> Error!
```

`test.cfg` file contents

```
add -ip rou 10.79.131.8 255.255.255.255 20.20.20.96 5
add -ip nerty-> Error!
setd !1!2 -pim cont=sm
```

Corresponding partial output of the `config.log` file:

```
<snippet>
[3]BOOT.CFG: Including commands from test.cfg
[4]TEST.CFG: add -ip rou 10.79.131.8 255.255.255.255 20.20.20.96 5
[5]TEST.CFG: add -ip nerty
Parameter 'nerty' unknown
1502: The command did not execute successfully
[6]TEST.CFG: setd !1!2 -pim cont=sm
[7]BOOT.CFG: Completed executing commands from test.cfg
[8]BOOT.CFG: add -ip rou 10.79.134.0 255.255.255.0 20.20.20.34 5
[9]BOOT.CFG: setd -ip ser = fer
Parameter 'ser' unknown
1502: The command did not execute successfully
Warning - The following 3 configuration command(s) failed
5 8 12
Mapping from Log file --> Line Number
```

```
5 --> test.cfg:2
9 --> boot.cfg:5
```

4.6.4

Include Command Limitations

The current implementation does not support `include` commands in `LoadConfigs` commands that are run interactively (after the device booted) or in macros.

- If you run the `LoadConfigs` command with `include` commands interactively, the software returns a file not found error message.
- If you run a macro with an `include` command, the macro terminates after the `include` command is executed, the remaining commands in the macro are ignored, and the software returns the `boot.cfg` file execution errors.

4.7

ASCII Boot Feature Limitations

Since the ASCII Boot feature is an extension of the `LoadConfigs` command, it has the same command set limitations as the `LoadConfigs` command:

- Configuration commands that require user interaction are not supported. Because of that, the `AddUser`, `PassWord`, and `DiscoverRoute` commands are not supported by the ASCII Boot feature and `LoadConfigs`.
- The following menu-driven configuration commands are not supported in the `boot.cfg` file: `SysconF`, `SysInfo`, `SysPassWord`, and `UserManage`.
 - The `SysconF` configuration changes are not stored in configuration files and are thus not affected by the deletion of the existing configuration files. They always remain in effect regardless of whether you use the ASCII Boot feature.
 - There are `SYS` services parameters supported by `LoadConfigs` that you can configure to perform the equivalent function of the `SysPassWord` command.
 - The `UserManage` command is not supported at this time.

4.8

ASCII Capture Feature

The ASCII Capture feature captures all the successful configuration changes you make by using the `SETDefault`, `ADD`, and `DELeTe` commands on a device and saves them in an `ASCII.txt` file (`capture.cfg`) in the configuration directory. The `capture.cfg` file can then be executed by the `LoadConfigs` command.

Configuration commands are captured regardless of how the configuration change is originally made: `LoadConfigs` execution, `TELnet`, `MENu`, a Scheduler event, or an `SNMP SET` request.



NOTE: The standard encoded configuration files are also maintained.

Macro definitions executed with the `DEFine` command and macro definitions deleted with the `UNDefine` command are logged in the `capture.cfg` file. However, configuration commands executed by macros are not captured.

Configuration commands executed through a macro, by using the `D0` command, are logged in the `macro.log` file, and a comment is written to that file. The comment is displayed in the following format:

```
# MACRO command <expandedMacroCmd> executed dayOfWeek MOnth dd hh:mm:ss yyyy
```

where *<expandedMacroCmd>* is the macro command after the actual values are substituted for the variables.

For example:

```
# MACRO command do wan_down executed Tue Aug 21 10:25:35 2007
```

You can automatically append the captured commands to the ASCII Boot `boot.cfg` file by setting the `SYS services CAPTURE` parameter to the appropriate value:

```
SETDefault -SYS CAPTURE = (Enable, BootCfg)
```

To switch the capture file back to the default capture file, you must use the following command:

```
SETDefault -SYS CAPTURE = CaptureCfg
```

For security reasons, passwords or other secure parameters are not captured with their configured value. A no-op (non-operational) value is substituted in the captured command for the secure parameter. For more information about the password capture, see [Password Capturing Format on page 116](#).

By default, the ASCII Capture feature is enabled. You can disable the ASCII Capture feature by setting the `SYS services CAPTURE` parameter to the appropriate value:

```
SETDefault -SYS CAPTURE = Disable
```

When ASCII Capture is enabled, the captured configuration commands are first cached in RAM. The captured commands are written to the `capture.cfg` file in the configuration directory when the cache buffer is full, the ASCII capture feature is disabled, the capture file is switched, or a minute elapsed with no configuration changes. If the capture file does not exist, it is created. Commands are always captured with their full service and parameter names, even if you originally used `CurrentServices`, `Aliases`, or short forms of the service names and parameters.

Unlike the `boot.cfg` file, the `capture.cfg` and `macro.log` files are never overwritten, deleted, or renamed. However, the EOS stops writing to the `capture.cfg` and `macro.log` files when the file size reaches 100,000 bytes. Each time EOS is unable to write to the `capture.cfg` or `macro.log` file, a syslog message informing that the file size is too large is generated.

To enable EOS to begin writing to the `capture.cfg` or `macro.log` file again, you must perform one of the following actions:

- To delete the file, use the `RemoveFile` command.
- To rename the large file, use the `ReName` command.

4.9

ASCII Capture Feature Limitations

The ASCII Capture feature has the following limitations:

- Nonconfiguration commands including those supported by `LoadConfigs` (`SHOW`, `ReName`, `SAVEbgp`, `SHOWDefault`) are not captured.
- Configuration commands that require user interaction are not captured. Currently, these commands are: `AddUser`, `PassWord`, `UserManage`, and `DiscoverRoutes`.
- Macro definitions executed with the `DEFINE` command and macro definitions deleted with the `UNDEFINE` command are captured in the `capture.cfg` file. However, configuration commands executed by macros (by using the `DO` command) are logged in a separate, `macro.log`, file and are not captured in the `capture.cfg` file. These commands are identified with a comment line in the `macro.log` file. For more information about ASCII Capture, see [ASCII Capture Feature on page 114](#).
- For security reasons, passwords and other secure parameters are not captured with the entered value. Instead, the configuration commands are captured with 6 or more asterisks in place of the secure parameter. The value for the secure parameters is treated as a no-op (non-operational) value. Whenever a secure parameter is configured with the no-op value, the existing value of the secure parameter is not changed. For more information, see [Password Capturing Format on page 116](#).

- The ASCII Capture feature captures configuration commands executed through the Install utility, but it does not capture configuration commands executed during the initialization process.
- Commands executed by the ASCII Boot feature are never captured. However, the captured commands can be automatically appended to the ASCII Boot `boot.cfg` file.
- SNMP SET requests are saved in a special SNMP command that can only be executed by `LoadConfigs`. You should not attempt to create your own SNMP commands, because setting a MIB object to the wrong value can cause unpredictable behavior in your device. For more information on the SNMP command, see [SNMP Command on page 116](#).

4.10

Password Capturing Format

Passwords and other secure parameters are not captured in clear text. The ASCII Capture feature is designed to replace the password parameter value with a 6 asterisk format and a comment generated above the configuration command, notifying you that the subsequent password was not captured.

The example below, demonstrates how the password for the `AuthLocalUser` parameter in the PPP services is captured:

```
CAT b: /112/boot.cfg
# Command capture enabled Tue Nov 13 09:30:08 2001 via software initialization
# Product version = SW/S6000-PW , 15.5.0
setd -SYS NMPrompt = "rolls #"
setd !7 -PORT OWNeR = we
setd !7 -PATH Clock = e
setd !7 -PORT CONTroL = e
setd !v1 -PORT CONTroL = e
# Secure / Password value in the following command was not captured
setd !v1 -PPP AuthLocalUser = "*****", "*****"
```

4.11

SNMP Command

The SNMP command executes a configuration change that was originally executed through an SNMP SET request. When configuration commands are captured, SNMP SET requests are captured by converting the SNMP SET requests into an SNMP command and writing that command to the capture file.

The syntax of the SNMP command in the capture file is:

```
SNMP <MIB_object_id> <MIB_object_type> <MIB_object_value>
```



NOTE: To execute an SNMP command, you must put the command in an ASCII file and execute the ASCII file by using the `LoadConfigs` command.

4.12

Boot, Capture, and Log Files Contents Review

You can use the `CAT` command to examine the contents of the ASCII Boot `boot.cfg` file, the ASCII Capture `capture.cfg` file, or the log `config.log` file.

For more information on the `CAT` command, see "Commands" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Listing ASCII Boot File (boot.cfg)

To list the `boot.cfg` file, you must use:

```
CAT primary/boot.cfg
```


You can also specify the range of line numbers displayed in the `boot.cfg` file by the CAT command. For example, to display line numbers 100 and upward, you must use:

CAT start = 100 mode = ln a:/primary/boot.cfg

The resulting display is similar to the following:

```
[ 100]Setd -SYS CAPTure = (e,bc)
[ 101]Setd !5 -PATH Baud = 4096
[ 102]Add !v1 -PORT VirtualPort ppp
[ 103]Setd !v1 -PORT DialInitState = md
[ 104]Setd !v1 -PORT DialNoList "111" type = we pos = 1
[ 105]Setd !v1 -PORT CONTROL = e
```

Listing ASCII Capture File (capture.cfg)

To list the `capture.cfg` file, you must use:

CAT primary/capture.cfg

The resulting display is similar to the following:

```
# Command capture enabled Tue Aug 11 09:30:08 1998 via software initialization
# Product version = SW/NBDPE-DW , 11.2.0.25I
setd -SYS NMPrompt = "rolls #"
setd !7 -PORT OWner = we
setd !7 -PATH Clock = e
setd !7 -PORT CONTROL = e
```

Listing Log File (config.log)

When you run ASCII boot (LoadConfigs `boot.cfg`), the commands executed and the output generated by executing the commands are written to the log file (`config.log`) in the configuration file directory. To list the `config.log` file, enter:

CAT a:/primary/config.log

The following is an example of a log file for a device with both a `boot.cfg` and a `XGSN.cfg` (GGSN configuration) file:

```
[1]GGSN #cat a:/primary/config.log
[1]BOOT.CFG: setd -sys nmp = "GGSN #"
[2]BOOT.CFG: setd !1 -ip net = 30.1.1.1 255.255.255.0
[3]BOOT.CFG: setd !2 -ip net = 40.1.1.1 255.255.255.0
[4]BOOT.CFG: setd -ip cont = ro
[5]BOOT.CFG: add !v1 -po vp ipip 40.1.1.1 40.1.1.2
[6]BOOT.CFG: add !v2 -po vp ipip 40.1.1.1 40.1.1.3
[7]BOOT.CFG: add !v3 -po vp ipip 40.1.1.1 40.1.1.4
[8]BOOT.CFG: setd !v1 -ggsn apn = "moca.com"
[9]BOOT.CFG: setd !v2 -ggsn apn = "motorola.com"
[10]BOOT.CFG: setd !v3 -ggsn apn = "xzone.org"
[11]BOOT.CFG: setd !v1 -ip net = 60.1.1.1 255.255.255.0
[12]BOOT.CFG: setd !v2 -ip net = 61.1.1.1 255.255.255.0
[13]BOOT.CFG: setd !v3 -ip net = 62.1.1.1 255.255.255.0
[14]BOOT.CFG: add -ip ro 10.79.0.0 255.255.0.0 30.1.1.200 1
[15]BOOT.CFG: setd !15 -pa cont = dis
Invalid path identifier
1502: The command did not execute successfully
[16]BOOT.CFG: setd !14 -pathdjdd cont = en
Service name 'pathdjdd' unknown
1502: The command did not execute successfully
[17]BOOT.CFG:
Warning - The following 2 configuration command(s) failed
--<CR> to continue, Q to quit--
15 16
[1]XGSN.CFG: setd -ggsn cont =e
[2]XGSN.CFG: setd !moca.com -ggsn apc = en
[3]XGSN.CFG: setd !motorola.com -ggsn apc = en
[4]XGSN.CFG: setd !xzone.org -ggsn apc = en
[5]XGSN.CFG: add !moca.com -ggsn ap 50.1.1.1 - 50.1.255.255
```

```

[6]XGSN.CFG: add !motorola.com -ggsn ap 50.1.1.1 - 50.1.255.255
[7]XGSN.CFG: add !xzone.org -ggsn ap 50.1.1.1 - 50.1.255.255
[8]XGSN.CFG: setd !moca.com -ggsn ipa = rds
[9]XGSN.CFG: setd !moca.com -ggsn dhcpss = 10.1.1.101 50.1.1.101
[10]XGSN.CFG: setd !moca.com -ggsn AAAPrimAUTH = 10.1.1.100
[11]XGSN.CFG: setd !moca.org.com -ggsn ipa - rds
SETDefault !<AccessPointName> IPAssigner = [ LocalDhcpServer
| RADIUS | RemoteDhcpServer ]
1502: The command did not execute successfully
[12]XGSN.CFG: setd !moca.com -ggsn AAACONTrol = en
[13]XGSN.CFG: setd !moca.com -ggsn AAASecret = *****
[14]XGSN.CFG: setd !motorola.com -ggsn ipa = lds
[15]XGSN.CFG: setd !xzone.org -ggsn ipa = lds
Warning - The following 1 configuration command(s) failed
11
[3]GGSN #

```

The commands are written to the log file with a prefix that is the line number of the command in the `boot.cfg` file. If a configuration command fails, the log file includes the following diagnostic message after the failed command, as seen after line [8] in the example above:

The command did not execute successfully

- If the `LoadConfigs IgnoreErrors` option is enabled (the default setting), the log file includes a warning message at the end of the file. This message indicates the number of commands that failed and the line numbers of the first 32 commands that failed.



NOTE: If the device has specialized configuration files in addition to the `boot.cfg` file – a GGSN configuration file (`XGSN.cfg`), for example – the log file lists the boot errors for each file separately

- If the `LoadConfigs IgnoreErrors` option is disabled, the log file may include the following message:

`LoadConfigs terminated due to UI command failure`

This message indicates that `LoadConfigs` terminated prematurely without executing all of the commands.

If the ASCII Boot feature was executed, an additional status message is displayed in the SysconF Boot Statistics display. This message shows the number of commands in the `boot.cfg` file (and any other specialized configuration files) that failed to execute successfully.

Chapter 5

IP Routing Configuration

Depending on your network requirements, you can use the default values of the parameters in the various services, or you can further configure the device according to later sections in this chapter. This chapter describes how to configure, manage, and troubleshoot Internet Protocol (IP) routing on your system and describes how IP routing works.

To configure the device, you must set parameters in the following services:

- If your network uses the Routing Information Protocol (RIP), you must set parameters in the RIP Service.
- If your network uses the Open Shortest Path First (OSPF) routing protocol, you must set parameters in the OSPF Service.
- If your network uses OSI routing in an IP environment, you must set parameters in the Integrated IS-IS (ISIS) service.

The IP parameters enable the routing function and configure the networks connected to the device.

By default EOS drops all loosely sourced routed broadcast IP packets for security reasons. In certain network topologies the packets must be forwarded. In these cases, you can set the `FwdCONTRol` parameter to `FwdLooseSrcRouteBC` to ensure the packets are forwarded. In certain network topologies, IP source routing must be disabled for security reasons; however, certain elements in such networks require loosely sourced routed packets.

To implement this design, you can set the `FwdCONTRol` parameter to `StripLooseSrcRoute` so that the loose source routing IP option is stripped from the header before the packet is forwarded.

The `StripLooseSrcRoute` option takes precedence over the `FwdLooseSrcRouteBC` option. Loose source routing IP option is stripped from the header before the `FwdLooseSrcRouteBC` option is applied to the packet. In addition, the `StripLooseSrcRoute` option takes precedence over the `NoRelaySrcRoute` option (configured with the `CONTRol` parameter) on ports for which the `StripLooseSrcRoute` option is configured. For details, see the description of the `FwdCONTRol` parameter in the *Enterprise OS Software Version 18.2 Reference Guide*.

5.1

Configuring Basic IP Routing over LAN Points and PPP Links

Perform this procedure to configure basic IP routing over Local Area Networks (LAN) points and Point-to-Point Protocol (PPP) links.

A serial line running PPP can support IP routing without the assignment of IP subnets. This feature is called *unnumbered links*. An unnumbered PPP link is useful only between two routers; in other words, it cannot connect a router to a host.

You must configure a serial line running PPP as an unnumbered link using the `-IP NETAddr` parameter before the unnumbered link takes effect. When an update is sent over an unnumbered PPP link, the source IP address is borrowed from another interface. For this reason, a device configured as a router must have at least one IP address configured before you set up PPP unnumbered links.

When you run RIP over a PPP link, both ends of the link must be either unnumbered or numbered with the same IP subnet. Half-numbered links, or links with inconsistent IP subnets on both ends, are considered a configuration error.

When you run OSPF or IISIS over unnumbered PPP links, you are not limited in how you configure the PPP link. You can number either end of the link independently, or you can leave both ends unnumbered. If both ends are numbered, they need not be on the same IP subnet nor have the same subnet masks. You do not need to assign a network number to a Frame Relay cloud if you are using IISIS.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- You are familiar with the protocols supported by the device. This section describes the protocols only when the explanation is necessary for interpreting the parameters and screen displays used in the software.
- EOS devices fully support variable length IP subnetting. For information on IP addresses, subnets, subnet masks, and variable length subnet masks, see [Internet Addressing on page 822](#). EOS devices also supports multiple IP subnets. For more information on the multiple IP subnets feature, see [Configuring Multiple IP Networks/Subnets on page 126](#).

Obtain:

- IP address for each port that you want to configure (PPP links can be unnumbered).
- Subnet mask for the given network address, if it is different from the default or "natural" mask.

Procedure:

1. Assign an IP address for each LAN port that routes IP by entering:

```
SETDefault !<port> -IP NETAddr = <IP_address> [<subnet_mask> [Ones | Zeros [MTU]]]  
| UnNumbered | IPCAddress
```



CAUTION: An IP address assigned to port 0 is considered the IP address for all the interfaces. As a result, the device behaves as an IP host for Telnet access and network management and stops routing IP packets. Do not configure an IP address for port 0 if you want to route IP packets.

2. Assign an IP address or the value UnNumbered to each wide area port using PPP as the serial line protocol by entering:

```
SETDefault !<port> -IP NETAddr = <IP_address> [<subnet_mask> [Ones | Zeros [MTU]]]  
| UnNumbered | IPCAddress
```

PPP does not you to assign an IP address to each wide area port. Before configuring your device to route over PPP, determine if you want to assign an IP address to each wide area port. If you do not want to assign an IP address to a wide area port, you must set the value of the -IP NETAddr parameter to UnNumbered. An advantage of not assigning an IP address to each wide area port is that you conserve valuable network and subnet numbers.

The NETAddr parameter lets you enable Internet Protocol Control Protocol on the interface using the IPCAddress parameter. The IPCAddress parameter enables the device to automatically get an address from the internet service provider (ISP) and assign it to the local PPP WAN interface. After the IP address is obtained, address mapping can take place.

3. If you want to run OSPF as the routing protocol over dial-up circuits, configure a demand interface circuit by entering:

```
SETDefault !<port> -OSPF DemandInterface = Enable
```

With this setting, the device negotiates with the neighbor at the other end of the point-to-point link. If the neighbor agrees that the point-to-point link is a demand circuit, the device suppresses sending OSPF hello packets and routing refresh information, allowing the data link connection to be closed when not carrying application traffic. For the demand circuit to be cost-effective, make sure that

it is isolated from as many topology changes as possible because topology changes bring up the interface.

For more information on reducing the network costs when using demand interface circuits, see [Route Learning with OSPF on page 177](#).

4. Enable the dynamic routing protocols for IP routing using RIP, OSPF, or IISIS by performing one of the following actions:

- If you want to enable RIP operation on a specified port, set the `CONTROL` parameter in the RIP Service, by using its `TALK` and `Listen` values similarly to the following example:

```
SETDefault !<port> -RIP CONTROL = ([TALK | NoTALK],
[Listen | NoListen], [Poison | NoPoison], [TRigger | NoTRigger],
[NetAdvUnn | SubnetAdvUnn], [SubnetBcast | AllIsBcast],
[Aggregate | NoAggregate], [DeAggregate | NoDeAggregate],
[DynamicNbr | NoDynamicNbr], [FullMesh | NonMesh])
```

Setting the `CONTROL` parameter to the `TALK` and `Listen` values enables the device to send and receive routing information with other devices using RIP.

You can also configure RIP for networks with variable length subnet masks using an aggregate or deaggregate scheme or the range table mask scheme. For more information, see [Static Route Configuration on page 127](#).

If you set the value of the `-IP NETAddr` parameter to `UnNumbered` for a PPP serial link, make sure that you set the value of the `-RIP CONTROL` parameter to `NetAdvUnn` or `SubnetAdvUnn` depending on your network configuration.

- If you want to enable OSPF on a specified port, you must set the `CONTROL` parameter in the OSPF Service by entering:

```
SETDefault !<port> -OSPF CONTROL = Enable
```

After OSPF is enabled, the device exchanges routing information with other devices using OSPF.

- If you want to configure IISIS for Dual IP and OSI mode, you can use:

```
SETDefault -IISIS CONTROL = Enable
```

5. Enable IP routing by entering:

```
SETDefault -IP CONTROL = Route
```

For information on how to complete the configuration for PPP links, see [Wide Area Networking Configuration with PPP on page 341](#).

5.2

IP Routing for Wide Area Networks

IP routing over Frame Relay is supported over fully meshed, partially meshed, and non-meshed topologies.

- If you plan to use RIP over Frame Relay in a partially meshed or non-meshed topology, you must enable the next-hop split horizon feature by having a list of neighbors. You must set `-RIP CONTROL` to `NonMesh`. The list of neighbors can be dynamically generated by the system or manually configured.
- If you plan to use OSPF over Frame Relay in a partially meshed or non-meshed topology, you can create a virtual port for each remote network that is attached to a Frame Relay cloud. You must run the OSPF `NonMesh` mode (set `-OSPF CONTROL` to `NonMesh`) over the Frame Relay cloud.
- If you plan to use IISIS over Frame Relay in a partially meshed or non-meshed topology, no additional configuration is necessary. When you use IISIS, you do not need to assign a network number to the Frame Relay cloud.

For information on configuring Frame Relay, see [Wide Area Networking Configuration with Frame Relay on page 308](#). For information on configuring PPP, see [Wide Area Networking Configuration with PPP on page 341](#).

5.3

Verifying the IP Routing Configuration

To verify the configuration, examine network devices and send packets from one network to another by using the PING command.

Procedure:

1. Examine the network devices by performing the following actions:
 - a. Display information on the attached networks by entering:
`SHoW -IP NETaddr`
 - b. Determine which stations or networks are reachable from the device by entering:
`SHoW -IP AllRoutes`
 - c. Display information from the Address Translation Table by entering:
`SHoW -IP ADDRess`
2. Check with the PING command. See [Checking the Forwarded Packets with the PING Command on page 123](#).
3. Examine the statistics by entering:

`SHoW -SYS STATistics -IP`


You can collect statistics for a specific period by using the `SampleTime` and `STATistics` parameters. For more information, see "SYS Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

5.3.1

Checking the Forwarded Packets with the PING Command

You can use the following sample figure and procedure to check if packets are forwarded.

Figure 34: Wide Area Device Configuration



Procedure:

1. Determine whether all of the Router 1 network interfaces are up and running by entering the following commands from Router 6:

- a. To check port 1, enter:

```
PING 129.213.16.10
```

If you do not specify the amount of time in seconds that the device should attempt to ping, the device assumes 20 seconds.

For more information on the PING command, see "Commands" in the *Enterprise OS Software Version 18.2 Reference Guide*.

- b. To check port 3 of Router 1, enter:

```
PING 121.121.67.4
```

If a port is operational, a message similar to the following appears after each PING command:

```
pinging ... 121.121.67.4 is alive
```

If a port is not operational, a message similar to the following appears:

```
pinging ... 121.121.67.4 is not responding
```

If the message returned by the PING command indicates that the port is not responding, check the network connection to see if the cables are properly connected. If this is not the problem, contact your network supplier or Centralized Managed Support Operations (CMSO) for help.

2. After you determine that each port is operational, check that the device can forward packets from one network to another by performing the following actions:

- a. Enter the PING command on Router 6 to check if it can communicate with host 2:

```
PING 122.122.67.10
```

If host 2 is operational and the device functions properly, a message similar to the following appears:

```
pinging ... 122.122.67.10 is alive
```

- b. If you do not get this message, use the `TraceRoute` command on Router 6 to trace a path to your intended destination. Specify the IP address of the destination you want to trace.

For details about the `TraceRoute` command, see “Commands” in the *Enterprise OS Software Version 18.2 Reference Guide*.

- c. Verify the status of the router. See [Verifying the Overall Router Status on page 124](#).

You must verify that:

- Port 1 is properly configured on Router 1.
- Port 3 is properly configured on Router 1.
- Port 1 is properly configured on Router 2.
- Port 3 is properly configured on Router 2.
- The address used in the PING command is the correct address of host 2.
- Routing is enabled on Router 1, Router 2, and Router 6.
- The routing protocol is properly configured.

You may want to verify that routing protocols and static routes are configured properly by using the `TraceRoute` command.

For example, at Router 1, you can trace the route between Router 1 and Router 4, which verifies that Router 5 and Router 3 relayed the packets sent by Router 1.

For complete information on the `TraceRoute` command, see the *Enterprise OS Software Version 18.2 Reference Guide*.

If you cannot determine the cause of the problem, contact the Centralized Managed Support Operations (CMSO).

5.3.2

Verifying the Overall Router Status

Procedure:

1. Examine the path configurations by entering:

```
SHow -PATH CONFIguration
```

2. Examine the path configurations by entering:

```
SHow -PORT CONFIguration
```

3. Examine the IP configurations by entering:

SHoW -IP CONFIguration

4. Examine the RIPv2 configurations by entering:

SHoW -RIPv2 CONFIguration

5. Examine the OSPF configurations by entering:

SHoW -OSPF CONFIguration

6. Examine the ARP configurations by entering:

SHoW -ARP CONFIguration

7. Examine the BGP configurations by entering:

SHoW -BGP CONFIguration

8. Examine the ISIS configurations by entering:

SHoW -ISIS CONFIguration

5.4

Customizing Routing Configuration

After you set up and check routing configuration, the device is ready to perform packet routing. If desired, you can further customize your routing configuration.

Process:

1. Configure UDP Broadcast Helper. See [User Datagram Protocol Broadcast Helper Configuration on page 447](#).

UDP Broadcast Helper allows applications in the TCP/IP stack to forward broadcast packets through a gateway and to another network segment. For information on configuring UDP Broadcast Helper, see [User Datagram Protocol Broadcast Helper Configuration on page 447](#).

2. Configure multiple subnets. See [Configuring Multiple IP Networks/Subnets on page 126](#).
3. Configure logical networks over IP. See [Static Route Configuration on page 127](#).
4. Add a static IP address.

When you add a static IP address for a group port by using the `ADD -IP ADDRESS` parameter, and the group port has a member in one of the 6-port Ethernet cards, the IP traffic is not forwarded to those member ports. The traffic is forwarded when the MAC address is learned dynamically through the Address Resolution Protocol (ARP) on the member ports.

5. Configure static routes. See [Static Route Configuration on page 127](#).
6. Configure packet filtering. See [Configuring Packet Filtering on page 131](#).
7. Configure RTP header compression. See [RTP/UDP/IP Header Compression on page 137](#).
8. Configure bidirectional forwarding detection (BFD). See [Bidirectional Forwarding Detection \(BFD\) Configuration on page 138](#).
9. Configure routing policies. See [Configuring RIP Routing Policies on page 140](#), [Configuring OSPF Routing Policies on page 147](#), [Configuring ISIS Routing Policy on page 150](#), and [Interautonomous System Routing Configuration with BGP on page 152](#).
10. Configure the Border Gateway Protocol (BGP). See [Interautonomous System Routing Configuration with BGP on page 152](#).

5.4.1


Configuring Multiple IP Networks/Subnets

If your device supports multiple IP subnets, you can configure more than one IP network or subnet on any media. The procedure for configuring multiple IP subnets is the same for all interfaces. This procedure provides an example of configuring multiple IP subnets on an Ethernet network.

The following figure is an example of a topology in which two IP networks can be configured on an Ethernet. The first IP network is called 128.1.0.0; the second is called 128.2.0.0. Use the following example to configure these IP networks on Ethernet 1. Configure the two networks on port 1 of Router 1.

In this topology, the systems on Ethernet 1 are divided into two IP networks. Direct communication takes place among the hosts in network 128.1.0.0 and among the hosts in network 128.2.0.0. However, Router 1 must forward packets between a host on network 128.1.0.0 and a host on network 128.2.0.0.

Figure 35: Two IP Subnets Configured on the Same Ethernet



The ability to configure multiple IP subnets gives you the following advantages:

- You can maximize the use of your network media.
The structure of the IP address limits the number of systems that can be addressed on an IP network. Configuring multiple IP subnets on a single network media allows you to increase the number of systems that you can address on a single media.
- You can break down systems on the network media into subsets of virtual private networks (VPNs).
Direct communication occurs within these VPNs.

You must factor the advantages of being able to configure multiple IP subnets against the fact that traffic on a segment containing multiple IP subnets can increase significantly. Traffic from one network to another on the same segment must first go to the router and then back out on the same segment.

Procedure:

1. Configure the first network.

The first address that you configure is known as the primary address. This address is indicated by an asterisk when you enter the `SHoW -IP NETaddr` command.

Step example:

To set up network 128.1.0.0 with the IP address of 128.1.0.5, enter:

```
SETDefault !1 -IP NETaddr = 128.1.0.5
```

2. Configure any subsequent networks.

Step example:

To configure network 128.2.0.0 with the IP address of 128.2.0.5, enter:

```
ADD !1 -IP NETaddr 128.2.0.5
```

3. To delete an address, enter:

```
DELEte !<port> -IP NETaddr <IP address>
```

5.4.2

Static Route Configuration

To set a static route, you can use:

```
ADD -IP ROUTe <IP_address> [<mask>] {<gateway> | !<port>} <metric> [Override] [BFD [TTL <t1>]]
```

To delete a static route, you can use:

```
DELEte -IP ROUTe <IP_address> {<gateway> | !<port>}
```


Routing Between Gateways

You can add a static route for the Los Angeles network by using:

```
ADD -IP ROUTe 10.0.0.0 11.0.0.1 1
```

This example shows that network number 10.0.0.0 (the Los Angeles network) is reachable through gateway 11.0.0.1. The gateway address is the Internet address of port 3 on Router 2. Because a packet routed from Router 1 to the Los Angeles network has to go through one gateway, the metric is 1.

Figure 36: Routing Between Gateways



The gateway must be located on a network directly connected to the router on which you add the static route. In this example, Router 1 and Router 2 both have an interface to a common network.

If the outgoing interface is a PPP link (either numbered or unnumbered), you can add a static route using the outgoing port number instead of the next-hop gateway address.

For example, on Router 1, you can add a static route for the Los Angeles network by using:


```
ADD -IP ROUTe 10.0.0.0 !3
```

This command achieves the same results as the command in which you entered the gateway address 11.0.0.1 as explained in the previous example.

If the PPP link is unnumbered (no IP network address is configured), you **must** provide the outgoing interface port number because the next-hop gateway address is not available.

Subnet Masks

Figure 37: Adding a Route Statically in a Subnet Masked Environment



You can use a mask to add a route in Router 1 to a subnet connected to Router 4 by using:

```
ADD -IP ROUTe 130.10.112.0 255.255.255.0 130.10.109.5 3
```

This command adds the address 130.10.112.0 with subnet mask 255.255.255.0 to the routing table.

5.4.2.1

Static Route Override Feature

If a destination network is reachable with both a static route and a learned route, the EOS device uses the static route unless you specify the optional `Override` value in the `ADD ROUTe` command. In that case, if a learned route of higher precedence is available, it overrides the static route. For information on precedence,

see [Multipath Routing on page 168](#). You can add the `Override` value at the end of the command. For example, to add the same static route as described earlier with the `Override` value included, you can use:

```
ADD -IP ROUTe 130.10.112.0 255.255.255.0 130.10.109.5 3 Override
```

Static Route Override for Imported OSPF Routes

You can use the `Override` keyword at the end of the `ADD ROUTe` command to prioritize the routes learned from OSPF above the static override routes.

The scenario is different but when OSPF imports static routes into the OSPF domain, as configured with the following command:

```
ADD -OSPF StaticPolicy All <metric>
```

- If `<metric> = 0`, OSPF imports the route with a metric equal to the metric configured when the static route was added (using the `ADD ROUTe` command).
- If `<metric> > 0`, OSPF imports the route with the specified `<metric>`.

In addition, when OSPF imports static routes, EOS compares the metrics of the imported OSPF routes and the static override routes in the routing table and selects the route with the lowest metric.

For example, if three devices (Routers A, B, and C) advertise static route 1.1.1.0/24 with static override metrics 50, 80, and 100 respectively and `<metric> = 0` in the `ADD -OSPF StaticPolicy` command, the route tables in each device are as follows:

```
Router A:
NetworkGatewayMetricOwner
1.1.1.0/24Router A50Static Override
```

```
Router B:
NetworkGatewayMetricOwner
1.1.1.0/24Router A51*OSPF
Router B80 Static Override
```

```
Router C:
NetworkGatewayMetricOwner
1.1.1.0/24Router A51*OSPF
Router C100 Static Override
```

* The metric = the `<metric>` value plus the link cost (=1 for 100Mbps)

If Router A fails, each router recalculates, and Router B and Router C then choose Router B as the best path.




NOTE: If you add the `BFD` keyword at the end of the `ADD ROUTe` command, each route is added or deleted from the route table based on the BFD status.

Static Route Override for NULL Routes

The static override for NULL routes (static-sink override) feature allows you to configure a maximum metric up to which the OSPF route can be used to reach a destination. If the OSPF metric exceeds the specified metric, that route is considered “sunk” and is unreachable for unicast and PIM multicast traffic. This feature is useful if you want divert traffic from certain WAN links that may be in danger of oversubscription during certain failure conditions.

Figure 38: OSPF Meshed Network Example



The devices are part of the same OSPF area.

- Hub 1 and Hub 2 learned routes to the core, over the WAN1 link, and the metric for these routes is 200.
- Hub 1 and Hub 2 also learn of the route to the core through the other hub. This route traverses the WAN2 link to the site and back over the WAN2 link to the other hub, and the metric for this route is 400 (100 + 100 + 200). However this route is not preferred until a hub loses its connection to the core through the WAN1 link.

To prevent the hubs from using the WAN2 link to reach the core, you can set a static-sink override route with a metric of 250. This configuration allows the hub to use the WAN1 route with metric 200 but not the WAN2 route with metric 400, which is a “sunk” route since its metric is higher than that of the static-sink override route.



NOTE: “Sunk” routes are listed in the `SHoW -IP AllRoutes` display only if you use the long format:

`SHoW -IP AllRoutes LOng`

To configure a static-sink override route, you can use the following syntax:

```
ADD -IP ROUte <IP address> [<mask>] NULL <metric> Override
```

For static override NULL routes, the *<metric>* can be up to 65535.

For example:

```
ADD -IP ROUte 30.30.30.0 255.255.255.0 NULL 250 Override
```

Static-sink override routes are listed with source `Static-Sink-0` in the displays returned by the `SHoW -IP ROUte` and `SHoW -IP AllRoutes` commands.

You can use supernets in the `ADD -IP ROUte` command to sink multiple routes within a supernet. For example, the following command sinks all routes in the 10.1.0.0/16 supernet with metrics greater than 200:

```
ADD -IP ROUte 10.1.0.0/16 255.255.255.0 NULL 200 Override
```

EOS selects the best route to the destination from the routes in the routing table.

- If the best route is a dynamic (learned) route or a static override route (a static route added with the `Override` keyword), EOS compares the route against any matching static-sink override routes (which are stored in a separate table) and selects the route with the lowest metric.

- If the selected route is the learned or static override route from the routing table, that route is used, and both the selected route and the static-sink override route are listed in the `SHoW -IP ROuTe` and `SHoW -IP AllRoutes` displays.
- If the selected route is the static-sink override route, the route is considered “sunk” and the `SHoW -IP ROuTe` and `SHoW -IP AllRoutes` displays list an entry similar to:
no valid route found
- If the best route is a regular static route (a static route added without the `Override` keyword), EOS does not compare the route against the configured static-sink override routes.

For example, consider the following routing table entry:

```
10.1.253.81      255.255.255.255
***             100  0  0  Up  -- OSPFType1
172.19.32.5      100  0  0  Up  -- Static-0
***             305  0  0  Up  -- Static-Sink-0
```

In this example:

- The learned (OSPF) route is selected as the best route to the destination.
- Since the best route is a learned route, EOS compares the metric of this route against the metric of any matching static-sink override routes.
- Since the matching static-sink override route has a higher metric than the learned OSPF route (305 as opposed to 100), the learned OSPF route is used.

5.4.3

Configuring Packet Filtering

EOS devices supports packet filtering, which controls traffic on your IP network. For complete information on the parameters used in this procedure, see “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Procedure:

1. Set up a filter policy or policies by entering:

```
ADD -IP FilterAddrS <adr1> [<dir>] <adr2> [<action> [<protocol>
[<filterID>]]<action> = {PROToColRsrv=<tag>}|Discard | DODdiscard | Forward |
{QPriority = H | M | L}} <protocol> = FTP | IP | IPDATA | ICMP | SMTP | TCP |
TELNET | UDP
```

2. Create a filter or filters by entering:

```
ADD !<filterID> -IP Filters <condition> [, <condition...>] <condition> = <%offset>:
[<operator>]<%pattern>
```

3. Set the `FilterDefAction` parameter by entering:

```
SETDefault -IP FilterDefAction = [Forward | Discard]
```

4. Enable packet filtering by entering:

```
SETDefault -IP CONTROl = Filtering
```

5.4.3.1

Packet Filtering Components and Protocol Reservation

IP Packet Filtering Components

The IP packet filtering feature is composed of two components: setting up a filter policy and creating a filter. To configure IP packet filtering, you must set up a filter policy and, depending on your filtering needs, you may

or may not need to create a filter. If you want to filter packets based on general criteria such as protocol or IP address, you can configure this type of filtering by setting up a filter policy only. If you want to filter packets based on more specific criteria that requires the system to examine the bytes of a packet, you need to set up a filter policy and create a filter.

If you configure a filter policy only, use the protocol field of the `ADD -IP FilterAddrs` command to specify a protocol you want to filter. If you configure both a filter policy and a filter, use the protocol field to specify the starting point for the offset of a condition. For complete information on the `-IP FilterAddrs` parameter, see “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*

Protocol Reservation

One of the action options for the `-IP FilterAddrs` parameter is `PROTocolRsrv=<tag>`. You can use this option to set up protocol reservation. Protocol reservation assigns a percentage of bandwidth to designated packets that pass through a specified port and meet certain conditions. The conditions can be protocol type, packet length, packets destined for specified address, and so on.


You use different procedures to set up protocol reservation for different packet types. The IP filtering procedure, for example, is applied only to IP-routed packets. IP-routed packets are also filtered using the IP firewall feature. For detailed information about IP firewall, see [Internet Firewalls Configuration on page 551](#).

For a detailed description of the procedures used to configure protocol reservation for the different packet types, see [Protocol Reservation Configuration on page 383](#).

For examples of protocol reservation for designated IP-routed packets using the IP filtering procedure and the `PROTocolRsrv=<tag>` action option, see [IP Packet Filtering Examples on page 134](#).

The following flowchart describes how the IP packet filtering feature works. In the depicted scenario, IP filtering is enabled and at least one filter policy was configured.

Figure 39: IP Packet Filtering



Enabling the packet filtering feature can have a significant impact on device performance. The performance is affected because the verification and decision-making process that takes place after each packet is received requires significant amounts of processing power.

5.4.3.2

IP Packet Filtering Examples

The following examples show how to configure the IP packet filtering feature.

IP Packet Filtering Example 1

An EOS device with the IP address of 129.213.16.0 and the subnet mask of 255.255.252.0 connects a local company network to the Internet. You want the following routing operations:

- Allow outgoing Transmission Control Protocol (TCP) connections from hosts on the local network to any host on the Internet.
- Block incoming TCP connections except for electronic mail (Simple Mail Transfer Protocol (SMTP), destination port %19) from a host on the Internet to a mail server (129.213.16.9) on the local network.
- Allow Internet Control Message Protocol (ICMP) messages from the Internet for feedback.
- Block all other packets from passing through the device.

To establish a filter that allows hosts on the local network to make TCP connections to hosts on the Internet, you must use:

```
ADD -IP FilterAddrs 129.213.16.0/0.0.3.255 > ALL Forward TCP
```


To create a filter policy and filter (filter 1) that allows TCP connections from the Internet only to the local mail server provided that conditions specified in filter 1 are met, you must use:

```
ADD -IP FilterAddrs ALL > 129.213.16.9 Forward TCP 1
```

```
ADD !1 -IP Filters %2:%0019
```

The `ADD -IP FilterAddrs` command specified above tells the system to look for TCP packets. The system automatically adjusts for any IP options. After the system determines that it received a TCP packet, it looks at the packet's TCP header. The system looks at the 2-byte offset (in the destination port field) for the hexadecimal value %19 (SMTP port) as specified by the `ADD -IP Filter` command above. The following figure shows the TCP header of a packet that meets the criteria established by these commands.

Figure 40: TCP Header of a Compliant Packet



To establish a policy that allows ICMP messages from the Internet for feedback, you must use:

```
ADD -IP FilterAddrs all > 129.213.16.0/0.0.3.255 Forward ICMP
```

To discard all other packets that do not meet any of the criteria discussed previously, you must use:

```
SETDefault -IP FilterDefAction = Discard
```

IP Packet Filtering Example 2

You want to set up the following conditions for a host with the address of 129.213.128.1:

- Allow incoming Telnet connections for destination port %17.
- Do not allow incoming TCP connections to the host for well-known ports (ports below %400).
- Allow all other packets.

To set up a filter (filter 1) and corresponding policy that allows incoming Telnet connections for destination port hexadecimal %17 at a 2 byte offset, you must use:

```
ADD !1 -IP Filters %2:%0017
```

```
ADD -IP FilterAddrs ALL > 129.213.128.1 Forward TCP 1
```

To set up a filter (filter 2) and corresponding policy that does not allow incoming TCP connections for ports with a value of less than hexadecimal %400 at a 2 byte offset, you must use:

```
ADD !2 -IP Filters %2:<%400
```

```
ADD -IP FilterAddrs ALL > 129.213.128.1 Discard TCP 2
```

To forward all other packets that do not meet any of the criteria discussed previously, ensure that the `-IP FilterDefAction` parameter retains its default setting of `Forward`. If this parameter was set to `Discard`, you must use:

```
SETDefault -IP FilterDefAction = Forward
```

This example demonstrates that `Discard` and `Forward` filters can be combined. It also highlights the fact that the order in which filters are configured is important. As soon as the system finds the first match, it stops searching. In this example, it is important that the system examine and forward packets with destination port %17 (filter 1) before examining and discarding packets with a destination port below %400 (filter 2).

IP Packet Filtering Example 3

You want an EOS device to do the following routing operations:

- Discard all User Datagram Protocol (UDP) packets with a source port of %161 and a destination port of %162.
- Discard all UDP packets if the tenth byte of data has the value of %60.
- Forward all other packets.

To set up a filter (filter 1) and corresponding policy that discards all UDP packets with a source port of hexadecimal %161 at a 0-byte offset and a destination port of hexadecimal %162 at a 2-byte offset, you must use:

```
ADD !1 -IP Filters %0:%161, %2:%162
```

```
ADD -IP FilterAddrs ALL > ALL Discard UDP 1
```

When you create a filter by using the `ADD -IP Filters` command, separate conditions with a comma (,) as shown in this example, to indicate the creation of multiple conditions. If a filter has multiple conditions, all conditions must be satisfied for a match to take place.

To set up a filter (filter 2) and corresponding policy that discard all UDP packets if the tenth byte of data has the value of hexadecimal %60 at an offset of hexadecimal %a, you must enter:

```
ADD !2 -IP Filters %a:%60
```

```
ADD -IP FilterAddrs ALL > ALL Discard UDP 2
```

To forward all other packets that do not meet any of the criteria discussed previously, ensure that the `-IP FilterDefAction` parameter retains its default setting of `Forward`. If this parameter was set to `Discard`, you must use:

```
SETDefault -IP FilterDefAction = Forward
```

IP Packet Filtering Example 4

You want an EOS device to do the following routing operations:

- Forward all UDP packets from a host with the address 129.213.16.9 except for RIPIIP packets (destination port %208).
- Forward all other packets.

To set up a filter (filter 2) and corresponding policy that forwards all UDP packets from host 129.213.16.9 except for RIPIIP packets (destination port hexadecimal %208) with a 2-byte offset, you must use:

```
ADD !2 -IP Filters %2:%0208
ADD -IP FilterAddrs 129.213.16.9 > ALL Discard UDP 2
```

To forward all other packets that do not meet any of the criteria discussed previously, ensure that the `-IP FilterDefAction` parameter retains its default setting of `Forward`. If this parameter was set to `Discard`, you must use:

```
SETDefault -IP FilterDefAction = Forward
```

IP Packet Filtering Example 5

To assign a low priority to FTP packets going to and coming from host 129.0.0.2., you must use:

```
ADD -IP FilterAddrs ALL < 129.0.0.2 QPriority Low FTP
```

IP Packet Filtering Example 6

If the `DodDiscard` action in the `FilterAddrs` parameter is enabled, specified traffic is discarded if a dial-up path is down. If the dial-up path is up, the specified traffic is forwarded.

To mark ICMP traffic from host 10.0.0.1 to host 129.0.0.3 as `DodDiscard`, you must use:

```
ADD -IP FilterAddrs 10.0.0.1 > 129.0.0.3 DodDiscard ICMP
```

IP Packet Filtering Example 7

You can use [Setting Up the Filtering Options for IP Packet Filtering on page 136](#) to add and set up the following filtering for your device:

- Add an IP filter that assigns 20 percent of reserved bandwidth for all Telnet sessions, and 30 percent of reserved bandwidth for all FTP packets, sent out through port 2.
- Set the IP `FilterDefAction` parameter so that all packets that do not meet the filtering conditions are forwarded.

IP Packet Filtering Example 8

You can use [Setting Up the Filtering Options for IP Packet Filtering on page 136](#) to add and set up the following filtering for your device:

- Add an IP filter that assigns 10 percent of reserved bandwidth for all FTP packets sent out to the IP address 50.0.0.1 through port 3.
- Set the IP `FilterDefAction` parameter so that all packets that do not meet the filtering conditions are forwarded.

5.4.3.3

Setting Up the Filtering Options for IP Packet Filtering

Procedure:

EXAMPLE 7

1. Add an IP filter that assigns 20 percent of reserved bandwidth to a `PROTOCOLRsrv` tag of `Telnet-tag` for all Telnet packets sent out through port 2, and 30 percent of reserved bandwidth to a `PROTOCOLRsrv` tag of `FTP-tag` for all FTP packets sent out through port 2, by entering:

```
ADD -IP FilterAddrs all all PROTOCOLRsrv = Telnet-tag Telnet
ADD -IP FilterAddrs all all PROTOCOLRsrv = FTP-tag FTP
```

2. Add an IP filter default action that forwards any packets that do not satisfy the filter requirements by entering:

```
SETDefault -IP FilterDefAction = Forward
```

3. Enable the IP filtering feature by entering:

```
SETDefault -IP CONTROL = Filtering
```

4. Assign 20 percent of bandwidth to the PROTOcolRsrv name tag Telnet-tag and 30 percent of the bandwidth to the PROTOcolRsrv name tag FTP-tag for port 2 by entering:

```
ADD !2 -PORT PROTOcolRsrv Telnet-tag 20
```

```
ADD !2 -PORT PROTOcolRsrv FTP-tag 30
```

5. Set PROTOcolRsrv as the option for port 2 by entering:

```
SETDefault !2 -PORT QueueCONTROL = PROTOcolRsrv
```

After you made these entries, any packet sent out by the system through port 2 that has the name tag Telnet-tag is allocated 20 percent of the bandwidth, and all packets with the name tag FTP-tag are allocated 30 percent of the bandwidth.

EXAMPLE 8

6. Add an IP filter that assigns reserved bandwidth to a PROTOcolRsrv tag of FTP-tag for all FTP packets sent out to the IP address 50.0.0.1 by entering:

```
ADD -IP FilterAddrs all 50.0.0.1 PROTOcolRsrv = FTP-tag FTP
```

7. Add an IP filter default action that forwards any packets that do not fulfill the filter requirements by entering:

```
SETDefault -IP FilterDefAction = Forward
```

8. Enable the IP filtering by entering

```
SETDefault -IP CONTROL = Filtering
```

9. Assign 10 percent of bandwidth to the PROTOcolRsrv name tag "FTP-tag" for port 3 by entering:

```
ADD !3 -PORT PROTOcolRsrv FTP-tag 10
```

10. Set PROTOcolRsrv as the option for port 3 by entering:

```
SETDefault !3 -PORT QueueCONTROL = PROTOcolRsrv
```

After you made these entries, any packet forwarded by the system to port 3 that has the name tag FTP-tag is allocated 10 percent of the bandwidth.

5.4.4

RTP/UDP/IP Header Compression

The goal of RTP/UDP/IP header compression is to reduce header overhead, making RTP/UDP/IP a viable option for sending network audio/video applications, even over low-speed lines, despite the header size. RTP/UDP/IP header compression compresses headers on a link-by-link basis, in most cases reducing RTP/UDP/IP headers to two bytes when there is no UDP checksum and to four bytes with a UDP checksum.

The IP service supports the following parameters related to RTP/UDP/IP header compression:



NOTE: For this release, RTP/UDP/IP header compression applies to IPv4 headers on PPP and Frame Relay link layers. Compression of IPv6 headers is not implemented in this release, and multiaccess ports (such as Ethernet) are not supported.

- RTPHeaderComp – Enables or disables RTP/UDP/IP header compression on a port or specifies that RTP/UDP/IP header compression be enabled if the port receives compressed RTP packets.

- `RtphcHdrExtComp` – Enables or disables RTP/UDP/IP header extension compression on a port.
- `RtphcPortRange` – Adds or deletes a UDP port range that determines which packets are compressed when RTP/UDP/IP header compression is enabled on the port.

For further details about these parameters, including syntax, see “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

EOS is also compliant with FRF.iphc, a Frame Relay Forum (FRF) standard that proposes two peers negotiate various parameters before starting to compress IP packets. For details about Frame Relay IP Header Compression control (FRIPHC) negotiation, see [FRIPHC Negotiation on page 320](#). You can enable or disable FRIPHC negotiation using the FR service `FrFIPhc` parameter. For details, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

SSRC in RTP/UDP/IP Header Compression Context Establishment

RFC 2508 specifies that the following five fields be used for RTP/UDP/IP header compression context establishment:

- Source IP address
- Destination IP address
- Source UDP port
- Destination UDP port
- Synchronization source (SSRC)

EOS implementations prior to release 11.6.025 did not use SSRC for context establishment for the following reasons:

- During unit testing it was determined that rapidly-changing SSRC was causing context sessions to fill up very rapidly so that in a short time all possible (255) sessions were allocated to streams.
- In addition to RTP streams, CRTP compresses UDP streams. If a packet is a UDP rather than an RTP packet, the SSRC field falls in the data portion of the packet rather than the header. As a result, it is highly likely that the SSRC field is different in every UDP packet, giving each packet its own assigned context and causing CRTP to run out of sessions quickly.

Beginning with release 11.6.025, EOS includes SSRC as a criteria for context establishment, implementing a solution to the limitations described above as suggested in RFC 2508. A negative cache contains the contexts of streams that have rapidly-changing SSRC values. A stream is put in this cache if it fails to compress (because the SSRC field does not stay constant) for three attempts. This table does not rely on the value of SSRC for indexing.

5.4.5

Bidirectional Forwarding Detection (BFD) Configuration

Bidirectional forwarding detection (BFD) is a network protocol used to detect faults between two forwarding devices. As implemented in EOS, static routes can be added with BFD, using the IP service `ROUte` parameter `BFD option`.

You can create BFD sessions with session IDs using the IP service `BFDSEssion` parameter. Other services can then register with BFD (via their own UIs) using this session ID. For details, see the descriptions of the `ROUte` and `BFDSEssion` parameters in “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: If the BFD session starts from a VRRP-controlled address (VRRP is enabled on the interface), you must also set the IP service `BFDTRigger` parameter to configure BFD to register with VRRP on that interface to start and stop the BFD session. For details, see the description of the `BFDTRigger` parameter in “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

When a static route is added with BFD, the BFD protocol probes the gateway address. If there are no BFD Packets (with the UP status) received within a specified amount of time (determined by the configuration of the IP service `BFDTimers` parameter for both the local and remote peer), the status of the BFD session changes from UP to DOWN, the static route is deleted from the routing table, and OSPF (if configured) is triggered to withdraw the route. When a BFD Packet informing that the peer BFD session is UP is received, the static route is added back to the routing table, the status of the BFD session changes from DOWN to UP, and OSPF (if configured) is triggered to advertise the route.

The BFD delay timer (configured with the IP service `BFDDelayTimer` parameter) allows you to specify an UP (and optionally a DOWN) state timeout value for static routes configured with BFD. EOS waits the specified amount of time before changing the BFD session status for a static route from DOWN to UP and adding the route (UP state timeout value) or changing the BFD session status for the static route from UP to DOWN and deleting the route (DOWN state timeout value). During the timeout period, any bouncing of the BFD session status is not processed as a status change.

For example:

- If the BFD session status for a route is DOWN and the BFD delay timer UP state timeout value is set to 10 seconds, any change in BFD session status from DOWN to UP and back to DOWN that occurs within the 10-second timeout window is not processed as a status change and the route is not added. In order for the route to be added, the BFD session status must change from DOWN to UP and remain in the UP state until the configured UP state timeout value (10 seconds, in this example) expires.
- If the BFD session status for a route is UP and the BFD delay timer DOWN state timeout value is set to 10 seconds, any change in BFD session status from UP to DOWN and back to UP that occurs within the 10-second timeout window is not processed as a status change and the route is not deleted. In order for the route to be deleted, the BFD session status must change from UP to DOWN and remain in the DOWN state until the configured DOWN state timeout value (10 seconds, in this example) expires.

The BFD delay timer feature is disabled by default (both the UP and the DOWN state timeout values are set to 0). For details, see the description of the `BFDDelayTimer` parameter in “Configuring IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: When you configure BFD on a Frame Relay virtual port or on an IPIP virtual port tied to a Frame Relay virtual port, you must set the IPQOS service `CONTRol` parameter to `TosQueue` to ensure that BFD is transmitted on the correct Frame Relay priority queue. In the case of the IPIP virtual port, you must set IPQOS `CONTRol` to `TosQueue` for both the IPIP virtual port and the Frame Relay virtual port.

BFD is useful in redundant configurations where there is an encryption failure; for example, in a redundant configuration where there is an encryption failure between the primary core router and the primary site router or primary site gateway. When the BFD protocol is implemented, the encryption failure triggers the routing protocol to delete the static route between the primary core router and the primary site router or primary site gateway and to stop advertising this route. The static route on the secondary core router is then used by all routers in the network, and an encryption black hole is prevented.

Additionally, you can configure an event-based macro based on the change in BFD session status. For details, see [Configuring Bidirectional Forwarding Detection \(BFD\) with Event-Based Macros on page 724](#).

Multi-Hop BFD

Multi-hop BFD allows you to configure BFD sessions that detect failures in multi-hop paths. This implementation can be used on redundant links to preserve redundancy in dynamic system resilience (DSR) systems.

You can configure multi-hop BFD by specifying a TTL value from 2 to 254 when you add a static route using the IP service `ROUte` parameter:

```
ADD -IP ROUTe <IP address> [<mask>] {<gateway> | !<port>} <metric> [Override] [BFD [TTL <tTL>]]
```

For details, see the description of the `ROUte` parameter in “Configuring IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

You can also configure event-based macros based on multi-hop BFD session status. For details, see [Configuring Bidirectional Forwarding Detection \(BFD\) with Event-Based Macros on page 724](#).

To support conventional data redundancy (HA and DSR) functionality, EOS implements multi-hop BFD sessions between the VRRP addresses of the foreign agent and the home agent. For details, see [Conventional Data Redundancy \(HA and DSR\) Feature on page 592](#).

Support for IPDV as a "Passive" BFD Client

The EOS bi-directional forwarding detection (BFD) implementation supports registration of the IP Packet Delay Variation (IPDV) service as a "passive" BFD client (a client for which the peer IP address is unknown). A passive client can occur if an IPIP virtual port is created and not used (no other BFD client registers on that port).

The IPDV service creates and processes the Ethernet WAN link object and performs delay and jitter calculations for Ethernet WAN links. As a result of this implementation, the `SHoW -IP BFDState` display may now include a state of `Passive`. When a client such as IPDV, for which the peer IP address is unknown, registers with BFD, the `SHoW -IP BFDState` command returns an IP address of 0.0.0.0 and state of `Passive` (rather than `Down`) for the port on which the client registered. For details, see the description of the `BFDState` parameter in "IP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

5.5

Configuring RIP Routing Policies

The routing policies supported by RIP allow you to control the reporting of routing information on a per-port basis.

Table 11: RIP Routing Policies

Policy	Description	Parameter
Advertise	Controls which routes are reported regardless of the route source.	AdvertisePolicy
Static	Controls which static routes are reported in the IP routing environment.	StaticPolicy
Exterior	Controls which BGP routes are reported.	ExteriorPolicy
Interior	Controls which OSPF or IISIS routes are reported.	InteriorPolicy
Receive	Controls which RIP routes are received by a trusted neighbor.	ReceivePolicy

For more information on the parameters listed in this table, see "RIPIP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Prerequisites:

Ensure that:

- You familiarized yourself with the various policies that are available and determine which policies you want to configure.
- You determined the IP addresses of networks, if you plan to access or receive information on routes from specific networks as opposed to all or no networks.

Procedure:

1. Establish an advertise policy that controls the advertisement of routes through RIP regardless of the source from which the route is learned by entering:


```
ADD !<port> -RIIP AdvertisePolicy All | None | [~]<IP_address> [<metric> (0-15)]
```

Step example:

To configure a policy on port 1 that forwards information on all routes to network 10.0.0.0, enter:

```
ADD !1 -RIIP AdvertisePolicy 10.0.0.0
```

In this example, a metric associated with network 10.0.0.0 was not specified. If you decide not to specify a metric with the `AdvertisePolicy` parameter or to specify a metric of zero, a route is reported with a metric calculated from the routing table.

2. Establish a receive policy that accepts or refuses to accept information on routes learned by RIP from a trusted neighbor by entering:

```
ADD !<port> -RIIP ReceivePolicy All | None | [~]<IP_address> [<metric> (0-15)]
```

Step example:

To configure port 1 so that it accepts information on routes learned by RIP for network 10.0.0.0, enter:

```
ADD !1 -RIIP ReceivePolicy 10.0.0.0
```

In this example, a metric associated with network 10.0.0.0 was not specified. If you decide not to specify a metric with the `ReceivePolicy` parameter or specify a metric of zero, a route with the originally reported metric is stored in the routing table.

3. To control the reporting of routes learned from specific sources, establish the following policies:

- Exterior policy for routes learned from BGP
- Interior policy for routes learned from OSPF or IISIS
- Static policy for reporting static (user-) configured routes

You must use the `ExteriorPolicy`, `InteriorPolicy`, or `StaticPolicy` parameter to complete this step. The syntax for the `ExteriorPolicy`, `InteriorPolicy`, or `StaticPolicy` parameters is the same as the syntax for the `AdvertisePolicy` parameter. See [step 1](#).

The metric that you configure with the `ExteriorPolicy`, `InteriorPolicy`, or `StaticPolicy` parameter is static or unchanging. This is in contrast to the metric that is calculated from the routing table. You can additionally manipulate the formula that is used to calculate the metric. For these reasons, Motorola Solutions recommends using the metric that is calculated from the routing table.

Step example:

To configure a policy on port 1 that forwards routing information learned from BGP, OSPF or IISIS, and about static routes configured on network 10.0.0.0, enter:

```
ADD !1 -RIIP ExteriorPolicy 10.0.0.0
```

```
ADD !1 -RIIP InteriorPolicy 10.0.0.0
```

```
ADD !1 -RIIP StaticPolicy 10.0.0.0
```

In this example, a metric associated with network 10.0.0.0 was not specified in each of the commands. If you decide not to specify a metric with the `ExteriorPolicy`, `InteriorPolicy`, or `StaticPolicy` parameters or specify a metric of zero, a route is reported with a metric calculated from the routing table.

In RIPv1, the recipient of a RIP message has no way to distinguish whether a route is within the RIP domain or is imported from an exterior gateway protocol domain (for example, BGP) or another interior gateway protocol domain (for example, OSPF). The inclusion of route tag in RIPv2 allows this information to be propagated along with every route entry. The tag information is always preserved and accompanies the route when the route entry makes its way throughout the RIPv2 domain. This information can be useful for a device when it wants to export the route to another domain, by comparing the route tag (hence the type of route) against its export policy.

To add a route tag to use with this route entry, enter:

```
ADD !1 -RIIP ExteriorPolicy 10.0.0.0 tag 65
```

```
ADD !1 -RIPIP InteriorPolicy 10.0.0.0 tag 65
```

4. Perform one of the following actions:

- If you decide to have routes reported with a metric calculated from the routing table and manipulate the conversion formula, go to [step 5](#).
- If you do not want to manipulate the formula, end the procedure here. The RIP routing policies are configured.

5. If you configured the `ExteriorPolicy`, `InteriorPolicy`, or `StaticPolicy` parameter and want to manipulate the formula that is used to calculate the metric, enter:

```
ADD -RIPIP ImportMetric <from protocol> Multiply | Divide <operand>
```

Step example:

To manipulate the conversion formula used to report OSPF routes so that the metrics reported with these routes are imported into RIP without changes, enter:

```
ADD -RIPIP ImportMetric OSPF Divide 1
```

To manipulate the conversion formula used to report OSPF routes so that the metrics reported with these routes are divided by 16, enter:


```
ADD -RIPIP ImportMetric OSPF Divide 16
```

5.6

Migration to a RIPv2 Network

The following figure illustrates the starting point for a migration to a RIPv2 network. Initially all devices configured as routers are RIPv1 routers.

Figure 41: Migration to a RIPv2 Network



In the first upgrade, Router 2 and Router 3 are RIPv2 capable. Since Network B is now directly connected to the RIPv2 capable devices, you can activate RIPv2 capability on those devices at the interfaces directly connected to Network B by using:

```
SETD !2 -RIPIP V2CompatMode = RIPv2
```

Interface 1 of both devices remains at RIPv1 mode (default value).

In the next upgrade, Router 1 and Router 4 are upgraded to be RIPv2 capable. To activate RIPv2, you must use:

```
SETD !1 -RIPIP V2CompatMode = RIPv2
```

Network C still has an old device (Router 5). To maintain RIP updates among those devices, do not set the new devices to RIPv2 mode yet. To take advantage of RIPv2 while maintaining the update exchange with the old device, use the RIPv1Compatible mode for the new devices. For Router 3 and Router 4, you must use:

```
SETD !1 -RIPIP V2CompatMode = RIPv1Compatible
```

In the final upgrade, Router 5 is upgraded to RIPv2. Now you can apply RIPv2 on Router 5 by using:

```
SETD !1 -RIPIP V2CompatMode = RIPv2
```

For complete information on the commands and parameters discussed in this section, see "RIPIP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

5.7

RIP Configuration for Networks with Variable Length Subnet Masks

You can use variable length subnet masks in your network and use RIP as the routing protocol. By using variable length subnet masks, you can eliminate the need for additional IP addresses.

Because RIPv1 packets do not carry explicit subnet information, you can configure the device that receives or transmits a routing update to determine the appropriate subnet mask. You can implement the following schemes:

Aggregate/deaggregate scheme

The aggregate/deaggregate scheme primarily addresses the transmitter function and how the transmitting device translates routes from one mask length to another without confusing receiving devices.

Range table mask scheme

The range table mask scheme addresses the receiver operation and how the receiving device interprets an incoming route advertisement and assigns an appropriate subnet mask to it.

These two schemes are independent and can be used individually or together.

The routes can be RIP-learned routes, directly attached networks, or static or dynamic routes learned from other protocols (if the appropriate policy is enabled).

5.7.1

Aggregate/Deaggregate Scheme

To enable the aggregate/deaggregate scheme on a specified (outgoing) port, you can use the following commands:

```
SETDefault !<port> -RIPIP CONTrol = Aggregate
```

```
SETDefault !<port> -RIPIP CONTrol = DeAggregate
```




NOTE: Do not use the aggregate/deaggregate scheme with unnumbered PPP links. Use the SubnetAdvUnn | NetAdvUnn values with the -RIPIP CONTrol parameter or use the range table mask scheme.

When both values are selected, RIPIP performs route conversion by using the following algorithm:

- If it is a host route (with mask 255.255.255.255), you must propagate as is.
- If it is a network route using the natural mask (for example, 10.0.0.0 255.0.0.0), you must propagate as is.
- If it is a subnet route, you must do the following:
 1. If the outgoing interface is not subnetted, or if the outgoing interface belongs to a different IP network number, you must zero out the bits in the subnet portion and propagate the natural IP network number (aggregate to the natural mask).
For example, if the routing update is sending 10.1.0.0 to 11.1.0.0, then the subnet portion is zeroed out and the natural IP network number (10.0.0.0) is propagated.
 2. If the outgoing interface has the same IP natural network number as the propagated route and if the mask of the route is longer than the mask of the outgoing interface, you must adopt the shorter mask and zero out all the bits in the host field (aggregate to a shorter mask).
This step of the algorithm is used if `Aggregate` is selected on the outgoing port. For example, if the mask of the route is 255.255.255.0 and the mask of the outgoing interface is 255.255.0.0, the shorter mask is used and the host field bits are zeroed.
 3. If the outgoing interface has the same IP natural network number as the propagated route and if the mask of the route is shorter than the mask of the outgoing interface, you must adopt the longer mask and convert the route into a series of route advertisements that cover the full address space.
This step of the algorithm is used if `DeAggregate` is selected on the outgoing port. For example, if the route 10.1.0.0 255.255.0.0 is sent to neighbors with a longer mask (255.255.255.0), the route is expanded into a sequence of subnets 10.1.0.0 255.255.255.0 to 10.1.255.0 255.255.255.0. In this example, each route from the shorter side translates into 256 routes on the longer side.
 4. If 1, 2, and 3 are not true, you must propagate as is (the outgoing interface is subnetted from the same IP network as the receiving interface and has the same mask).

You can use the aggregate/deaggregate scheme in simple network topologies; for example, you may have a single router between the corporate backbone and stub networks (or remote offices) as shown in the following figure.

Figure 42: Route Aggregation/Deaggregation With RIPIP



In this configuration, the backbone has a shorter mask and no overlapping routes (10.2.0.0 255.255.0.0 and 10.2.2.0 255.255.255.0 are overlapping and cannot coexist). All subnets with the same aggregate

must be fully connected and contiguous; subnets with different aggregates can be located independently. For example, all 10.2.X.X subnets must be connected and contiguous, but 10.2.X.X and 10.3.X.X can be independent.

The aggregate/deaggregate scheme provides the following benefits:

- The aggregate scheme reduces the number of routes in the backbone routing table; a smaller routing table leads to smaller routing overhead and smoother network operation.
- All the work is performed by the border router that has interfaces to subnets with different masks. It is responsible for translating routing updates and using the correct subnet mask. No other routers need to be aware that variable length subnet masks are used.
- The scheme can be used easily and quickly in a simple network topology with minimal impact; it is compatible with older devices and major upgrades are not necessary.

Using the aggregate/deaggregate scheme has the following disadvantages:

- Deaggregation may not be a beneficial scheme in some topologies. We recommend that you do not use it for unnumbered PPP links or topologies more complex than that shown in Figure 5-9); use the default route advertisement as an alternative.
- If more than one device connects the longer subnets to the backbone, the RIP advertisements, after aggregation or deaggregation, may confuse each other. You may need to configure the `-RIPIP ReceivePolicy` parameter to filter out this type of information.

5.7.2

Range Table Mask Scheme

To configure the range table mask scheme, you can use:

```
ADD -RIPIP RcvSubnetMask <IP_address> <IP_address> <subnet_mask>
```

For example, you could specify that all subnets between 10.2.0.0 and 10.2.255.0 use subnet mask 255.255.255.0 by entering:

```
ADD -RIPIP RcvSubnetMask 10.2.0.0 10.2.255.0 255.255.255.0
```

You can configure any type of subnet mask to any network number. The subnet mask can be longer or shorter than its natural mask. For example, the 128.4.0.0 128.4.0.0 range with subnet mask of 255.252.0.0 means that network 128.4.0.0 is assigned 255.252.0.0 as the subnet mask.



NOTE: You cannot assign a subnet mask to the default route (0.0.0.0).

For each route received from a neighbor, RIPIP determines the appropriate mask using the following algorithm:

- If the route belongs to the same IP network number of the receiving interface, you must use the mask of the interface.
For example, if the device receives route 10.1.0.0 on interface 10.2.0.0 255.255.0.0, the device adopts the same subnet mask for the received route.
- If the route belongs to the same IP network number of any other interface, you must adopt the mask of that interface.
This action is useful when receiving routes over an unnumbered PPP link.
- You must use the natural mask based on the class (A, B, or C) of the received route.
After the mask was determined, the software checks to see if the route falls within any of the network ranges in the range table. If there is a match, the software overrides the mask with a user-configured subnet mask. If there are multiple matches in the table, the software picks the most specific match. The entries in the range table are organized from the more specific to the less specific.


For example, range 10.1.0.0–10.1.255.0 with a mask of 255.255.255.0 is more specific than 10.0.0.0–10.255.0.0 with a mask of 255.255.0.0. With these ranges in the table, network 10 is assigned subnet mask 255.255.0.0; networks that begin with 10.1.X.X are assigned a subnet mask of 255.255.255.0.

As the final step, the software compares the received route with the mask. If there are non-zero bits in the host field, the route is a host route, and the software converts the subnet mask to 255.255.255.255.

You can use the range table mask scheme for more complex topologies not covered by the aggregate/deaggregate scheme (overlapping routes exist, or routes are learned over unnumbered PPP links). With the range table mask scheme, there is no limit on the number of potential subnet masks.

The topology in the following figure has overlapping routes. In this situation, you need to configure the range table mask on Router A and Router B. You prevent Router A, for example, from adopting the shorter mask (255.0.0.0) of Network C when receiving route updates from Network B.

Figure 43: Range Table Mask



The range table mask scheme provides the following benefits:

- Classless addressing can be supported.
- No topology limitations exist.
- Overlapping routes are supported.

Using the range table mask scheme has the following disadvantages:

- Extensive configuration may be required on every device, leading to increased administrative overhead.
- When the scheme is first used, all EOS devices must be upgraded and synchronized at the same time. Because all device must be configured with identical information, the coexistence of non-Motorola Solutions routers may not be possible.

5.7.3

RIP V2 Compatibility

RIPv2-capable devices can inter-operate with RIPv1-capable devices. To facilitate this, as well as to offer a migration path for existing routers, RIPv2-capable devices have several modes of operation.

Four mode settings are available for transmit:

- **RIPv1 Only** – Only RIPv1 messages are sent.
- **RIPv1 Compatible** – RIPv2 messages are sent using directed broadcast if no AdvToNeighbor is set. Otherwise, the RIPv2 messages are sent to the neighbors individually using unicast. Note that these RIPv2 messages are picked up by RIPv1-capable devices.
- **RIPv2 Only** – RIPv2 messages are sent using multicast address 224.0.0.9 if no AdvToNeighbor is set. Since RIPv1-capable devices do not listen on that multicast address, they do not pick up the RIPv2 messages.
- **None** – No RIP messages are sent.

Four mode settings are available for receive:

- **RIPv1 Only** – Only RIPv1 messages are received.
- **RIPv1 Compatible** – Both RIPv1 and RIPv2 messages are received.
- **RIPv2 Only** – Only RIPv2 messages are received.
- **None** – No RIP messages are received.

Sets of modes for sending and receiving are available and are provided on a per interface basis. The None mode (where sending/receiving of RIP messages are stopped) is obtained by setting the corresponding `IPRIP Control` parameter to `NoTalk/NoListen`. You can set both the send and receive modes to the same value.

RIPv1 mode is the default mode of operation, which avoids problems with RIPv1 devices that may have trouble interpreting a RIPv2 message (although a properly implemented RIPv1 device would not have this problem). When the mode is set to RIPv1 Compatible, the device can exchange RIP information with both RIPv1 and RIPv2 devices. However, in the RIPv1 Compatible mode it is sending RIPv2 message to RIPv1 devices, so some issues involving subnet arise.

For more information about the `V2CompatMode` parameter, see “RIIP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

5.7.4

Authentication Block

Every RIPv2 message can contain an optional authentication block. The authentication block occupies the same size as a route entry (20 bytes) and is distinguished by a value of 0xFFFF in the Address Family Identifier field. If the authentication block is present, it must reside immediately following the RIP header, and there must only be one such authentication block in the entire RIPv2 message.

Currently, only the authentication type of 0x2 is specified, which denotes that the next 16 bytes contain the authentication string (password), expressed in ASCII characters. If the authentication string is less than 16 bytes, it is padded to the right with null. For more information about the `V2AUTHenticate` parameter, see “RIIP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

5.7.5

RIPv2 Message Authentication

RIPv2 introduces a function of authenticating a RIPv2 message. The function is optional and the recipient can ignore password authentication. The rules for accepting a RIPv2 message vary depending on the authentication setting of a device. For more information about the `V2PassWord` parameter, see “RIIP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

5.8

Configuring OSPF Routing Policies

The routing policies supported by OSPF allow you to control the reporting of routes learned from other sources.

Prerequisites:

Determine which policies you want to configure. The following table lists and briefly describes each policy and its associated parameter.

Table 12: OSPF Routing Policies

Policy	Description	Parameter
Exterior	Controls whether exterior routing protocol (BGP) learned routes are further advertised into the OSPF domain.	ExteriorPolicy
Interior	Controls whether interior routing protocol (RIP or IISIS) learned routes are further advertised into the OSPF domain.	InteriorPolicy
Static	Controls which static routes are further advertised into the OSPF domain.	StaticPolicy
Direct	Controls whether a locally attached network (with OSPF disabled on such interface) should be further advertised into the OSPF domain.	DirectPolicy

For more information on the parameters listed in this table, see the “OSPF Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

If you plan to access or receive information on specific routes as opposed to all or no routes, determine the IP addresses of these specific routes.

Procedure:


- To control the reporting of routes learned from other sources and advertised into the OSPF domain, see [Controlling the Reporting of Routes Learned from Other Sources and Advertised into the OSPF Domain on page 149](#).
- To enable OSPF route aggregation, see [Configuring OSPF Route Aggregation on page 150](#).

5.8.1

Controlling the Reporting of Routes Learned from Other Sources and Advertised into the OSPF Domain

Assuming your network topology is similar to that shown in the following figure, you can use the following procedure to control the reporting of routes learned from other sources and advertised into the OSPF domain.

Figure 44: OSPF Routing Policies



Procedure:

1. Enable the OSPF Protocol on the appropriate ports on the backbone devices.

Step example:

On Router 1, Router 2, and Router 3, enter the following commands:

```
SETDefault !1 -OSPF CONTrol = Enable
```

```
SETDefault !2 -OSPF CONTrol = Enable
```

2. Configure the backbone devices to learn routes from other interior routing protocols (such as RIPIP) within the same autonomous system.

The default setting of the `-OSPF InteriorPolicy` parameter is `None`. This means that if a device runs both the OSPF and RIPIP Protocols, the routes learned by one of these protocols are not reported to the other.

Step example:

To configure Router 1 to learn routes from RIPIP domain #1, on Router 1, enter:

```
ADD -OSPF InteriorPolicy All
```

To specify an IP address of the network in RIPIP domain #1, enter:

```
ADD -OSPF InteriorPolicy <IP_address>
```

3. Configure the backbone devices to learn routes from other exterior routing protocols, such as BGP, in another autonomous system.

Step example:

To configure Router 3 to learn routes from autonomous system 2, on Router 3, enter:

```
ADD -OSPF ExteriorPolicy All
```

5.8.2

Configuring OSPF Route Aggregation

In OSPF, you can import routes from an exterior routing protocol, such as BGP, an interior routing protocol, such as RIP, static routes, and directly connected networks. These imported routes become OSPF external routes. In some networks there are too many OSPF external routes.


OSPF route aggregation lets you define external route ranges to aggregate the area's routes, which reduces the routing information in the backbone and subsequently in all areas. If an external route is within the range and the range is configured as "advertise", then the range network is advertised. If several external routes are in the same external route range, then only one network is advertised. This reduces the number of routes in the backbone and regular areas.

To establish OSPF route aggregation, you must use:

```
ADD -OSPF ExtRouteraNges <IP address><mask> [Tag] [Advertise | DontAdvertise]
```

In the following example, RIP is enabled on Router A and port 1 on Router B. OSPF is enabled on port 2 of Router B and on Router C. You can use the following procedure to set up this configuration.

Figure 45: OSPF Route Aggregation



Procedure:

1. On Router B, enter:
2. On Router B and Router C, view the external LSAs by entering:

```
ADD -OSPF InteriorPolicy ALL 1
```

```
SHow -OSPF lsd external
```

Two external LSA are shown in the resulting display.

3. To establish OSPF route aggregation on Router B, enter:

```
ADD -OSPF ExtRouteraNges 10.0.0.0 255.0.0.0
```

When you enter the `SHow -OSPF LSD external` command on Router C and Router D, only one LSA appears on the resulting display.


5.9

Configuring IISIS Routing Policy

The routing policies supported by IISIS allow you to control the reporting of routes learned from other sources. Assuming that your network topology is similar to that shown in the following figure, you can use

this procedure to control the reporting of routes learned from other sources and advertised into the IISIS domain.

Figure 46: IISIS Routing Policies



Prerequisites:

Determine which policies you want to configure. The following table lists and briefly describes each policy and its associated parameter.

Table 13: IISIS Routing Policies

Policy	Description	Parameter
Exterior	Controls whether exterior routing protocol (BGP) learned routes are further advertised into the IISIS domain.	ExteriorPolicy
Interior	Controls whether interior routing protocol (RIP or OSPF) learned routes are advertised into the IISIS domain.	InteriorPolicy
Static	Controls which static routes are advertised into the IISIS domain.	StaticPolicy

For more information on the parameters listed in this table, see “IISIS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

If you plan to access or receive information on specific routes as opposed to all or no routes, determine the IP addresses of these specific routes.

Procedure:

1. Enable the IISIS Protocol on the backbone devices.

Step example:

On Router 1, Router 2, and Router 3, enter:

```
SETDefault -IISIS CONTrol = Enable
```

2. Configure the backbone devices to learn routes from other interior routing protocols (such as RIP or OSPF) within the same autonomous system.

The default setting of the `-IISIS InteriorPolicy` parameter is `None`. This means that if a device runs both the IISIS and RIP or OSPF Protocols, the routes learned by one of these protocols are not reported to the other.

Step example:

To configure Router 1 to learn routes from RIPv2 domain, on Router 1, enter:

```
ADD -IISIS InteriorPolicy All
```

You can specify an IP address of the network in RIPv2 domain by using:

```
ADD -IISIS InteriorPolicy <IP_address>
```

3. Configure the backbone devices to learn routes from other exterior routing protocols, such as BGP, in another autonomous system.

Step example:

To configure Router 3 to learn routes from autonomous system 2, enter the following command on Router 3:

```
ADD -IISIS ExteriorPolicy All
```

5.10

Interautonomous System Routing Configuration with BGP

To configure BGP as your interautonomous system routing protocol, you need to configure the following items:

- BGP peers
- Default route
- Route aggregation
- Route importing from an IGP to BGP domain (interior policy)
- Route importing from a BGP to IGP domain (exterior policy)
- Network number policies
- AS-path policies (permit, deny, and weight)


5.10.1

Configuring BGP Peers

For BGP to learn routes between autonomous systems (ASs) and determine the reachability of networks outside of its AS, you must configure BGP peers. For information on peers, see [External and Internal Peers on page 184](#).

You can use this procedure to set up the configuration shown in the following figure.

Figure 47: BGP Network Topology



Procedure:

1. Define the local AS number for the devices by entering:

```
SETDefault -BGP LocalAS = <AS Number>(1-65536)
```

This parameter defines the AS number used by this BGP speaker in the OPEN message and in all routing updates as the originating AS number. The local AS number also determines whether a peer is connected through an internal or external BGP session.

Step example:

Router A and Router B are internal peers and part of AS 1. On each device, enter:

```
SETDefault -BGP LocalAS = 1
```

2. Add a peer to each device by entering:

```
ADD -BGP PEER (<IP_address> | <IPv6_address>) [<NAME_ID>] <A>> [<"Sys-IP">]
```

AS numbers range from 1 to 65535.

The device must know the AS number for itself and the peer that is added before it can establish a BGP session with its peer.

Step example:

On Router B, specify Router A's IP address (10.0.0.1) and AS Number (1):

```
ADD -BGP PEER 10.0.0.1 1
```

On Router A, specify Router B's IP address (10.0.0.2) and AS Number (1):

```
ADD -BGP PEER 10.0.0.2 1
```

3. Enable BGP routing by entering on both devices:

```
SETDefault -BGP CONTROL = IPv4
```

4. For each peer that you added with the ADD -BGP PEER command, enable the peer by entering:

```
SETDefault [!<IP_address>] -BGP PeerControl = Enable
```

Step example:

On Router B, specify Router A's IP address for <IP address> by entering:

```
SETDefault !10.0.0.1 -BGP PeerControl = Enable
```

On Router A, specify Router B's IP address for <IP address> by entering:

```
SETDefault !10.0.0.2 -BGP PeerControl = Enable
```

Router A establishes a TCP connection with Router B (peer-to-peer communication). After the connection is established, both peers exchange BGP update packets indicating the networks each peer can reach.

5. Display routes learned through BGP by entering on any BGP device:

SHoW -BGP ROUte

For information on how to read the display, see “BGP Service Parameters” in the *Enterprise OS Software Reference Guide*.

6. Display peer information by entering on any BGP device:

SHoW -BGP PEER

The display shows the current mapping of peer ID to IP address to AS number and shows the current state of the peer (disabled, open, connecting). You can display peer-specific details by using the optional keyword `Long` with the `SHoW` command. If a peer was configured with the System IP address as its source address for BGP peering, the display that appears when you enter the `SHoW -BGP PEER Long` command lists the System IP address. You can display peer-specific BGP statistics by entering the optional keyword `Statistics` with the `SHoW` command.

5.10.2

Default Route Configuration

You can configure a default route in the BGP Routing Table to provide the IP address of a network that can be used as the default network to destinations that are not explicitly listed in the routing table.


Configuring a default route is helpful under the following circumstances:

- The routing policy of a peer does not permit the advertisement of a default route.
- When the local device is unable to maintain the complete BGP Routing Table due to memory limitations.

If a route for a particular destination address is not contained in the BGP Routing Table, BGP checks for a default route. To configure a default route, see the following figure and use the following command:

ADD -BGP DefaultNet <IP_address>

Figure 48: BGP Default Route



The configured IP address does not have to be a directly connected network. As long as the local device has a route to the IP address, it can forward all default route traffic to the IP address. The next-hop address in the BGP Routing Table is automatically calculated by the system software.


5.10.3

Configuring BGP Route Aggregation

BGP route aggregation uses the Classless InterDomain Routing (CIDR) address aggregation strategy to combine the characteristics of several different routes so that a single route can be advertised, as shown in

the following figure. By combining several networks into one supernet, the number of BGP messages sent to peers and the size of the routing table are reduced and unnecessary details about subnets are hidden from peers.

Figure 49: BGP Route Aggregation



Prerequisites:

Ensure that you configured BGP peers (making Router A and Router B peers, and Router A and Router C peers). See [Configuring BGP Peers on page 152](#).

Procedure:

1. Specify a list of networks that BGP advertises as a single supernet route by entering:

```
ADD -BGP AggregateRange <IP_address> <mask>
```



NOTE: Aggregation should never enclose Class D address space (224.0.0.0 through 239.255.255.255).

Step example:

On Router A, to combine the routes to sites 1, 2, and 3 into a range so that only a single route is advertised, enter:

```
ADD -BGP AggregateRange 192.1.0.0 255.255.0.0
```

2. Specify a list of routes that BGP explicitly advertises by entering:

```
ADD -BGP AggregateExcept <IP_address> <mask>
```

Step example:

If you do not want site 4 included in the aggregation range, enter:

```
ADD -BGP AggregateExcept 192.1.5.0 255.255.255.0
```

3. Enable route aggregation and the BGP routing protocol by entering:

```
SETDefault -BGP CONTrol = (IPv4, IPv4AGgregate)
```

Router A can advertise a single network (192.1.0.0/255.255.0.0) that summarizes each of the three connected sites and also explicitly advertises the exception route (192.1.5.0). Without the use of CIDR,

Router A advertises each route with a separate entry, and Router B's routing table grows in size. With route aggregation, Router B's routing table has an entry for 192.1.0.0 and 192.1.5.0.

4. If the route should be explicitly advertised to all internal peers, add a network filter (network address and mask) by entering:

```
ADD -BGP NetworkFilter <NetfilterID> <network_address> <mask>
```

Step example:

On Router A, to configure route 192.1.3.0 as network filter 1, enter:

```
ADD -BGP NetworkFilter 1 192.1.3.0 255.255.255.0
```

5. Apply a network policy for internal peers to advertise outgoing routes by entering:

```
ADD -BGP NetPolicyInt <NetfilterID> Explicit
```

You can also configure the `NetPolicyAll`, `NetPolicyExt`, and `NetPolicyPeer` parameters for explicit policies that are applied to outgoing constituent routes of aggregates. For more information, see [Route Aggregation on page 192](#), and "BGP Service Parameters" in the *Enterprise OS Software Reference Guide*.

Step example:

On Router A, to associate network filter 1 (192.1.3.0 255.255.255.0) to be explicitly advertised to Router C, enter:

```
ADD -BGP NetPolicyInt 1 Explicit
```

The Router C routing table has entries for 192.1.0.0, 192.1.5.0, and 192.1.3.0.

5.10.4

Importing Routes from an IGP Domain to a BGP Domain

To control how route reachability information is shared between devices in different domains, you can configure the BGP device to accept or reject Interior Gateway Protocol (IGP) routing information. You can configure an interior policy on your BGP device to control the import (also known as route leaking) and advertisement of routes from an IGP domain. IGP refers to protocols (such as RIP, OSPF, and ISIS) that operate within a domain.

You must control how you import routes from an IGP domain to BGP domain when your network is connected to the Internet. Motorola Solutions recommends that you statically map valid routes and have only these routes imported into the BGP domain. Otherwise, dynamically changing routes in the IGP domain are constantly imported into the BGP domain causing increased load on all core routers to process unnecessary route flaps (routes coming up and going down).

Procedure:

1. Define the network filters specifying the network address and mask by entering:

```
ADD -BGP NetworkFilter <NetfilterID> <network_address> <mask>
```

Step example:

To identify networks that have a value of 11.5.7 in the first three octets, enter:

```
ADD -BGP NetworkFilter 2 11.5.7.0 255.255.255.0
```

To identify networks that have a value of 193.4 in the first two octets, enter:

```
ADD -BGP NetworkFilter 5 193.4.0.0 255.255.0.0
```

To identify networks that have a value of 193.7.8 in the first three octets, enter:

```
ADD -BGP NetworkFilter 6 193.7.8.0 255.255.255.0
```

2. Advertise only the specified networks and block all others, or block only the specified networks and advertise all others by entering:


```
ADD -BGP InteriorPolicy <NetfilterID> <Permit / Deny>
```

By default, no IGP route (including static routes and directly connected routes) are imported into the BGP routing table. To set up whether IGP routes are imported into the BGP routing table, use the `SETDefault -BGP IntPolDefault` command. For more information about this parameter, see "BGP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: To avoid an invalid configuration and the interior policy from being ignored, do not configure the `InteriorPolicy` or `OspfExtPolicy` parameters with a mixture of permit and deny policies. You must specify the policy as either all permit or all deny policies.

Step example:

To import only networks that have a value of 11.5.7 in the first three octets (filter 2), all networks that have a value of 193.4 in the first two octet (filter 5), and all networks that have a value of 193.7.8 in the first three octets (filter 6), enter the following commands:

```
ADD -BGP InteriorPolicy 2 Permit
ADD -BGP InteriorPolicy 5 Permit
ADD -BGP InteriorPolicy 6 Permit
```

To block the import of only the specified networks and import all others (except OSPF external routes), enter the following commands:

```
ADD -BGP InteriorPolicy 2 Deny
ADD -BGP InteriorPolicy 5 Deny
ADD -BGP InteriorPolicy 6 Deny
```

To block the import of OSPF Type 1 external routes on the specified networks, enter the following commands:

```
ADD -BGP OspfExtPolicy 2 Deny ExType1
ADD -BGP OspfExtPolicy 5 Deny ExType1
ADD -BGP OspfExtPolicy 6 Deny ExType1
```

3. Restart all BGP sessions to recompute the route selection process by entering:

```
SETDefault -BGP CONTROL = IPv4
```

For more information, see [Policies on page 190](#).

5.10.5

Route Import from a BGP Domain to an IGP Domain

To control how route reachability information is shared between devices in different domains (BGP to IGP), you can configure the EOS device to accept or reject the injection of BGP routing information into the IGP routing domain. For stub and multi-homed ASs, you can configure a default route on your IGP device. For transit ASs, you can configure an exterior policy to control the import (also known as route leaking) and advertisement of routes from BGP, or just run an IGP between ASBRs. In most cases, importing routes is probably not required.

Stub Autonomous Systems

A stub AS has only one connection to another AS. In a stub AS using RIPv1, RIP cannot advertise both the network number and the mask. As a result, some BGP-derived routes may not be understood by RIP. You need to configure the `-RIPIP DefaultMetric` parameter on the border router to advertise a default route (to network 0.0.0.0) by using:

```
SETDefault !<port> -RIPIP DefaultMetric = <metric> (0-15)
```

In a stub AS using OSPF or IISIS, these protocols can advertise both the network number and the mask, so you can import BGP-derived routes and advertise them by OSPF or IISIS. However, the simplest solution is to configure the OSPF or IISIS `DefaultMetric` parameter on the Autonomous System Boundary Router (ASBR) to generate an external link state advertisement (LSA) for network 0.0.0.0 by using:

```
SETDefault -OSPF DefaultMetric = [Disable | <metric> (1-65535) [Type1 | Type2]]  
SETDefault -IISIS DefaultMetric = Disable | <metric> (1-63) [Internal | External]
```

Multi-Homed Autonomous Systems

A multi-homed AS has connections to more than one AS but does not carry transit traffic. All of the traffic in a multi-homed AS is considered local.

In a multi-homed AS executing RIP, some BGP-derived routes may not be understood by RIP; you should configure the `-RIPIP DefaultMetric` parameter on each border router to advertise the default route. The default route allows routers that are internal to the AS to select the least-cost default route to forward traffic destined for another AS.

In a multi-homed AS executing OSPF or IISIS, all BGP-derived routes can be advertised by OSPF, but it is not always necessary or desirable to import them into a multi-homed autonomous system. The easiest solution is to configure the OSPF or IISIS `DefaultMetric` parameter to `Type1` on each ASBR to generate an external LSA for network 0.0.0.0. When each internal router constructs its shortest path tree, the router selects the least-cost default route.

Transit Autonomous Systems

A transit AS has connections to more than one AS and carries both local traffic and transit traffic. Transit traffic is any traffic that does not originate or terminate within the local AS.

Every router within a transit AS must have explicit routing information for all networks that make up the internetwork. Each internal router must be able to forward a packet to any destination without relying on a default route. As a result, RIP cannot be used as the IGP for a transit AS. A transit AS requires a protocol that scales and supports the advertisement of BGP-derived routes. A transit AS must use BGP on all border routers, and OSPF, IISIS, or other protocols capable of conveying network masks. An IGP is the best choice, but may not scale well.

To allow the import of BGP routes into the IGP, you can use:

```
ADD -OSPF ExteriorPolicy All | None | [~]<IP_address> <metric> [Type1 | Type2]  
ADD -IISIS ExteriorPolicy All | None | [~]<IP_address> <metric> [Internal | External]
```

For more information, see [Policies on page 190](#).

5.10.6

Configuring Network Number Policies

You can control the receipt (import) or advertisement (export) of BGP routes based on the network address using permit or deny network policies. Permit or deny policies are applied to incoming or outgoing packets, or to both incoming and outgoing packets.

Procedure:

1. Identify the individual network number or block of network numbers to which the policy is applied by entering:

```
ADD -BGP NetworkFilter <Netfilter_ID> <network_address> <mask>
```

Step example:

To configure network filter 1 for network addresses starting with 192.2.1, enter:

```
ADD -BGP NetworkFilter 1 192.2.1.0 255.255.255.0
```

2. Define the policy: the type of peer (internal or external or both) the policy applies to, the type of policy (permit or deny), and if the policy applies to incoming packets, outgoing packets, or both by entering the following commands:

```
ADD -BGP NetPolicyAll <Netfilter_ID> {Permit | Deny [In | Out | Both]} | Explicit
ADD -BGP NetPolicyExt <Netfilter_ID> {Permit | Deny [In | Out | Both]} | Explicit
ADD -BGP NetPolicyInt <Netfilter_ID> {Permit | Deny [In | Out | Both]} | Explicit
ADD [!<IP_address>] -BGP NetPolicyPeer <Netfilter_ID> {Permit | Deny [In | Out | Both]} | Explicit
```

You can configure deny policies using the same syntaxes by specifying Deny instead of Permit.



NOTE: All policies in a specific direction (in/out) must be either all permit policies or all deny policies. A mix of permit and deny policies causes ambiguity resulting in ignoring the entire policy list.

Step example:

To permit the import of routes to 192.2.1.0 from all peers, enter:

```
ADD -BGP NetPolicyAll 1 Permit In
```

To permit the import of routes to 192.2.1.0 from external peers, enter:

```
ADD -BGP NetPolicyExt 1 Permit In
```

To permit the import of routes to 192.2.1.0 from internal peers, enter:

```
ADD -BGP NetPolicyInt 1 Permit In
```

To permit the import of routes to 192.2.1.0 to be accepted from peer 10.0.0.2, enter:

```
ADD -BGP !10.0.0.2 NetPolicyPeer 1 Permit In
```

If this is the only policy defined, all other routes from peer 10.0.0.2 are discarded.

3. Restart all BGP sessions to recompute the route selection process by entering:

```
SETDefault -BGP CONTROL = IPv4
```

For more information, see [Policies on page 190](#).

5.10.7

Configuring AS-Path Permit or Deny Policies

You can control the receipt (import) or advertisement (export) of BGP routes based on presence or absence of specific AS numbers in the AS-PATH attribute. Recall that the AS-PATH attribute is contained in each update message. For information about this attribute, see [Path Attributes on page 185](#).

Permit or deny policies are applied to incoming packets, outgoing packets, or to both inbound and outbound packets. These policies help filter incoming and outgoing routes.

Procedure:

1. Define the filter to which the policy applies by entering:

```
ADD -BGP AsFilter <AsfilterID> "<regular_expression>"
```

For examples of regular expressions, see [Regular Expression Examples on page 160](#).

2. Define the policy: which peer the policy applies to, the type of policy (permit or deny) and whether the policy applies to incoming packets, outgoing packets, or both by using the following syntaxes to assign policies to all peers, external peers, internal peers, or a specific peer:

```
ADD -BGP AsPolicyAll <AsfilterID> [[Permit | Deny [In | Out | Both]]  
[Weight <weight>]  
ADD -BGP AsPolicyExt <AsfilterID> [[Permit | Deny [In | Out | Both]]  
[Weight <weight>]  
ADD -BGP AsPolicyInt <AsfilterID> [[Permit | Deny [In | Out | Both]]  
[Weight <weight>]  
ADD [! <IP address>] -BGP AsPolicyPeer <AsfilterID> [[Permit | Deny  
[In | Out | Both]] [Weight <weight>]
```



NOTE: To maintain consistent routing information within an AS, do not apply permit or deny policies to internal BGP peers. The only type of policy that should be applied to internal peers is one that changes the preference by adding additional weight to selected paths.

3. Restart all BGP sessions to recompute the route selection process by entering:

```
SETDefault -BGP CONTrol = IPv4
```

For examples of deny and permit filters, see [Deny Filter Examples on page 160](#) and [Permit Filter Examples on page 161](#).

5.10.7.1

Regular Expression Examples

Blank spaces are represented in the examples as underscores (_). When two spaces are shown together, a space is inserted between the underscores, for example __. You must enter a blank space for each underscore shown in the examples.

Regular Expression Example 1

To create filter 1 that identifies an AS-PATH attribute containing AS 25, you must use:

```
ADD -BGP AsFilter 1 "_25_"
```

Regular Expression Example 2

To create filter 2 that identifies an AS-PATH attribute containing AS 35 and AS 50 (in this order), you must use:

```
ADD -BGP AsFilter 2 "_35_.*_50_"
```

Regular Expression Example 3

To create filter 3 that identifies an AS-PATH attribute containing AS 35 and AS 50 (in any order), you must use:

```
ADD -BGP AsFilter 3 "_35_.*_50_|_50_.*_35_"
```

The horizontal bar (|) indicates a logical OR operation.

Regular Expression Example 4

To create filter 4 that identifies an AS-PATH attribute containing the AS Sequence <AS5, AS46, AS32>, you must use:

```
ADD -BGP AsFilter 4 "<_5_ _46_ _32_>"
```

Regular Expression Example 5

To create filter 5 that identifies an AS-PATH attribute containing the AS Set [AS5, AS32, AS46], you must use:

```
ADD -BGP AsFilter 5 "[_5_ _32_ _46_]_"
```

5.10.7.2

Deny Filter Examples

Deny Filter Example 1

To block the import of routes containing AS 25 from all peers using filter 1, you must use:

```
ADD -BGP AsPolicyAll 1 Deny In
```

Deny Filter Example 2

To block the advertisement of routes containing AS 35 and AS 50 to external peers using filter 2, you must use:

```
ADD -BGP AsPolicyExt 2 Deny Out
```

Deny Filter Example 3

To block the import and advertisement of routes containing AS 35 and AS 50 to peer 10.0.0.2 using filter 2, you must use:

```
ADD !10.0.0.2 -BGP AsPolicyPeer 2 Deny Both
```

5.10.7.3

Permit Filter Examples

Permit Filter Example 1

To permit the import of routes containing AS 25 from all peers using filter 1, you must use:

```
ADD -BGP AsPolicyAll 1 Permit In
```

Permit Filter Example 2

To permit the advertisement of routes containing AS 35 and AS 50 to external peers using filter 2, you must use:

```
ADD -BGP AsPolicyExt 2 Permit Out
```

Permit Filter Example 3

To permit the import and advertisement of routes containing AS 25 to peer 10.0.0.2 using filter 1, you must use:

```
ADD !10.0.0.2 -BGP AsPolicyPeer 1 Permit Both
```

5.10.8

Configuring AS-Path Weight Policies

You can control the route selection process by assigning a specific weight to an AS, an AS path, or a subset of an AS path. The total weight for a given route is known as the *degree of preference* for the route and is calculated by summing all the individual AS-path weight expressions assigned to the route's AS-PATH attribute. If multiple routes exist for a destination, the route with the highest degree of preference is selected by the BGP route selection process. For more information, see [Path Selection on page 189](#).

Weight policies are only applied to incoming routing updates. These policies help control the route selection process based on AS numbers.

Procedure:

1. Define the filter to which the policy applies by entering:

```
ADD -BGP AsFilter <AsfilterID> "<regular expression>"
```

Step example:

The following commands create filters for AS 100, 200, 300, 400, and 500:

```
ADD -BGP AsFilter 1 "_100_"
```

```
ADD -BGP AsFilter 2 "_200_"
```

```
ADD -BGP AsFilter 3 "_300_"
```

```
ADD -BGP AsFilter 4 "_400_"
```

```
ADD -BGP AsFilter 5 "_500_"
```

Blank spaces are represented in the examples as underscores (_). When two spaces are shown together, a space is inserted between the underscores, for example __. You must enter a blank space for each underscore shown in the examples.

2. Define the policy (which peer the policy applies to) and the weight by entering the following commands:

```
ADD -BGP AsPolicyAll <AsfilterID> [Weight <weight>]
ADD -BGP AsPolicyExt <AsfilterID> [Weight <weight>]
ADD -BGP AsPolicyInt <AsfilterID> [Weight <weight>]
ADD [!<IP address>] -BGP AsPolicyPeer <AsfilterID> [Weight <weight>]
```

Additionally, the following syntaxes affect weight-based policies and the degree of preference in the route selection process:

```
SETDefault -BGP DefaultWeight = [IPv4 | Ipv6] <number>(-2000 to 2000)
SETDefault [!<IP address>] -BGP PeerWeight = [IPv4 | Ipv6] <number>(-2000 to 2000)
```

The `DefaultWeight` parameter configures a default weight that is added to each route when computing the degree of preference (LOCAL-PREF attribute) for the route. You can configure this parameter to give priority in the route selection process to routes received by one BGP speaker over routes received by other BGP speakers. The default value of this parameter is 0. For information about the LOCAL-PREF attribute, see [Path Selection on page 189](#).

The `PeerWeight` parameter configures a weight that is added to all routes received from the specified peer. The default value of this parameter is 0. For more information, see [Degree of Preference Calculations on page 162](#).

Both the `DefaultWeight` and `PeerWeight` parameters allow you to set weight by address family (IPv4 or IPv6). If you do not specify an address family, IPv4 is assumed.

3. Restart all BGP sessions to recompute the route selection process by entering:

```
SETDefault -BGP CONTro1 = IPv4
```

5.10.8.1

Weight Filter Examples

Weight Filter Example 1

To assign a weight of 10 to routes with AS 100 in the AS-PATH (filter 1) received from all peers, you must use:

```
ADD -BGP AsPolicyAll 1 Weight 10
```

You can assign a weight of 10 to routes with AS 100 in the AS-PATH (filter 1) received from external peers or internal peers using the `AsPolicyExt` and `AsPolicyInt` parameters, respectively.

Weight Filter Example 2

To assign a weight of 10 to routes with AS 100 in the AS-PATH (filter 1) received from peer 10.0.0.2, you must use:

```
ADD !10.0.0.2 -BGP AsPolicyPeer 1 Weight 10
```

5.10.8.2

Degree of Preference Calculations

The degree of preference, which must always be greater than or equal to 0, is calculated before the route selection process using the following formula:

Degree of preference = (Total ASPolicy Weight) + PeerWeight + DefaultWeight

In the following examples, assume that AS 100 has a weight of 10, AS 200 has a weight of 20, AS 300 has a weight of 30, and AS 500 has a weight of 50.

Degree of Preference Calculations Example 1

The local router receives a route with the following AS-PATH attribute: {_500_ _200_ _600_ _100_ _300_}. Assume that both the `PeerWeight` and `DefaultWeight` parameters have a value of 0. The degree of preference is equal to:

$$50 + 20 + 0 + 10 + 30 + 0 (\text{PeerWeight}) + 0 (\text{DefaultWeight}) = 110$$

If an AS in the AS-PATH attribute was not assigned a weight by using the `AsPolicyXXX` parameter, it is assumed to have a weight of zero (0).

Degree of Preference Calculations Example 2

Assume that peer 10.0.0.2 was assigned a `PeerWeight` of 100 and the local router's `DefaultWeight` value is 0. The local router receives a route from peer 10.0.0.2 with the following AS-PATH attribute: {_500_ _200_ _100_}. The degree of preference is equal to:

$$50 + 20 + 10 + 100 (\text{PeerWeight}) + 0 (\text{DefaultWeight}) = 180$$

Degree of Preference Calculations Example 3

Assume that peer 10.0.0.2 was assigned a `PeerWeight` of 100 and the local router was configured with a `DefaultWeight` value of -50. The local router receives a route from peer 10.0.0.2 with the following AS-PATH attribute: {_500_ _200_ _100_}. The degree of preference is equal to:

$$50 + 20 + 10 + 100 (\text{PeerWeight}) - 50 (\text{DefaultWeight}) = 130$$

For more information, see [Policies on page 190](#).

5.10.9

MED/LOCAL-PREF Value Configuration

EOS supports configurable multi-exit discriminator (MED) and local preference (LOCAL-PREF) path attributes. Both attributes provide a routing metric for BGP:

- The MED is the BGP metric for a route advertised by a device to a different autonomous system (AS). Lower MED values are preferred over higher values.
- The LOCAL-PREF is the BGP metric for a route advertised by a device to other devices in the same AS. Higher LOCAL-PREF values are preferred over lower values.

Both attributes are non-transitive (cannot be passed along beyond a single AS hop). When an EOS device receives a MED or LOCAL-PREF value from an external BGP peer, it passes the value along to its internal BGP peers but does not re-advertise the value to its external BGP peers.

When the device advertises a route to a BGP peer, EOS advertises the appropriate value (MED or LOCAL-PREF) depending on whether the route is external (MED) or internal (LOCAL-PREF). Routes without a MED/LOCAL-PREF attribute are considered to have the lowest possible MED value or the highest possible LOCAL-PREF value, meaning that they are preferred over other routes.

To add a MED/LOCAL-PREF value to a route, you must first assign the route a network filter ID using the BGP service `NetworkFilter` parameter. You can then:

- Add network policies to add MED and LOCAL-PREF values for this address by using the BGP service `NetPolicyAll`, `NetPolicyExt`, or `NetPolicyInt` parameter. For details, see [Configuring MED/LOCAL-PREF with a Network Policy on page 164](#).
OR
- Add OSPF export policies to add MED and LOCAL-PREF values to routes imported from OSPF external routes by using the BGP `OspfExtPolicy` parameter. For details, see [Configuring MED/LOCAL-PREF with an OSPF Export Policy on page 165](#).
OR

- Add interior policies to add MED and LOCAL-PREF values to routes imported from all IGP devices except OSPF external routes (for example, OSPF interior routes and static routes) by using the BGP service InteriorPolicy parameter. For details, see [Configuring MED/LOCAL-PREF with an Interior Policy on page 165](#).

EOS looks for all networks that match the given filter and rule. It then replaces the existing attribute with a new attribute that contains the specified MED and LOCAL-PREF values.

- The MED value is applied only to inter-AS advertisements and has no effect if it is configured for a route which is only advertised within the AS.
- The LOCAL-PREF value is applied only to intra-AS advertisements and has no effect if it is configured for a route which is advertised to a different AS.

To remove or change a MED or LOCAL-PREF value, you must delete the policy in which the value is configured and wait for the device to withdraw the route. You can then add a new policy without a MED or LOCAL-PREF value or with a different MED or LOCAL-PREF value.

5.10.9.1

Configuring MED/LOCAL-PREF with a Network Policy

Procedure:

1. Define a network filter specifying the network address and mask by entering:

```
ADD -BGP NetworkFilter <Netfilter_ID> <network_address> <mask>
```

Step example:

To add a network filter for network 4.4.4.0/24, enter:

```
ADD -BGP NetworkFilter 1 4.4.4.0 255.255.255.0
```

This command identifies networks that have a value of 4.4.4.0 in the first three octets with a network filter ID of 1.

2. Add network policies that apply MED and LOCAL-PREF values to all incoming and outgoing advertisements by performing one of the following actions:

- To define policies that are applicable to all peers, enter:

```
ADD -BGP NetPolicyAll [(IPv4 | IPv6)] <NetfilterID> {(Permit|Deny) [In|Out|Both]} | Explicit) [MED|LOCPREF <value>]
```

- To define policies that are applicable to all external peers, enter:

```
ADD -BGP NetPolicyExt [(IPv4 | IPv6)] <NetfilterID> {(Permit|Deny) [In|Out|Both]} | Explicit) [MED|LOCPREF <value>]
```

- To define policies that are applicable to all peers that are in the same AS, enter:

```
ADD -BGP NetPolicyInt [(IPv4 | IPv6)] <NetfilterID> {(Permit|Deny) [In|Out|Both]} | Explicit) [MED|LOCPREF <value>]
```

Any route that does not match the permit policy is not advertised.

Step example:

To define policies that apply the MED value 55 and the LOCAL-PREF value 100 to outgoing route advertisements for the route 4.4.4.0/24 (identified by network filter ID 1) for all peers, enter:

```
ADD -BGP NetPolicyAll IPv4 1 Permit Out MED 55
```

```
ADD -BGP NetPolicyAll IPv4 1 Permit Out LOCPREF 100
```

To define policies that apply the MED value 55 and the LOCAL-PREF value 100 to outgoing route advertisements for the route 4.4.4.0/24 (identified by network filter ID 1) for all external peers, enter:

```
ADD -BGP NetPolicyExt IPv4 1 Permit Out MED 55
```

```
ADD -BGP NetPolicyExt IPv4 1 Permit Out LOCPREF 100
```


To define policies that apply the MED value 55 and the LOCAL-PREF value 100 to outgoing route advertisements for the route 4.4.4.0/24 (identified by network filter ID 1) for all internal peers, enter:

```
ADD -BGP NetPolicyInt IPv4 1 Permit Out MED 55
ADD -BGP NetPolicyInt IPv4 1 Permit Out LOCPREF 100
```

3. Restart all BGP sessions to recompute the route selection process by entering:

```
SETDefault -BGP CONTrol = IPv4
```

5.10.9.2

Configuring MED/LOCAL-PREF with an OSPF Export Policy

Procedure:

1. Define a network filter specifying the network address and mask by entering:

```
ADD -BGP NetworkFilter <Netfilter_ID> <network_address> <mask>
```

Step example:

To add a network filter for network 4.4.4.0/24, enter:

```
ADD -BGP NetworkFilter 1 4.4.4.0 255.255.255.0
```

This command identifies networks that have a value of 4.4.4.0 in the first three octets with a network filter ID of 1.

2. Add OSPF export policies that apply MED and LOCAL-PREF values to routes imported from OSPF external routes by entering:

```
ADD -BGP OspfExtPolicy [IPv4|IPv6] <NetfilterID> (Permit|Deny) (ExType1|ExType2|All) [MED|LOCPREF <value>] [TAG <AS>]
```

Step example:

To define policies that apply the MED value 55 and the LOCAL-PREF value 100 to the route 4.4.4.0/24 (identified by network filter ID 1) for all OSPF Type1 routes, enter:

```
ADD -BGP OspfExtPolicy IPv4 1 Permit ExType1 MED 55
ADD -BGP OspfExtPolicy IPv4 1 Permit ExType1 LOCPREF 100
```

3. Restart all BGP sessions to recompute the route selection process by entering:

```
SETDefault -BGP CONTrol = IPv4
```

5.10.9.3

Configuring MED/LOCAL-PREF with an Interior Policy

Procedure:

1. Define a network filter specifying the network address and mask by entering:

```
ADD -BGP NetworkFilter <Netfilter_ID> <network_address> <mask>
```

Step example:

To add a network filter for network 4.4.4.0/24, enter:

```
ADD -BGP NetworkFilter 1 4.4.4.0 255.255.255.0
```

This command identifies networks that have a value of 4.4.4.0 in the first three octets with a network filter ID of 1.

2. Add interior policies that apply MED and LOCAL-PREF values to routes imported from all IGP routes except OSPF external routes by entering:

```
ADD -BGP InteriorPolicy [IPv4 | IPv6] <NetfilterID> (Permit|Deny) [MED|LOCPREF <value>]
```

Step example:

To define policies that apply the MED value 55 and the LOCAL-PREF value 100 to the route 4.4.4.0/24 (identified by network filter ID 1) for routes imported from all IGP devices except OSPF external routes, enter:

```
ADD -BGP InteriorPolicy IPv4 1 Permit MED 55
ADD -BGP InteriorPolicy IPv4 1 Permit LOCPREF 100
```

3. Restart all BGP sessions to recompute the route selection process by entering:

```
SETDefault -BGP CONTrol = IPv4
```

5.10.10

BGP Convergence Optimization

You can use the BGP `IGPImporttime` and `AdvertiseDelay` parameters to speed up the convergence, particularly when an existing device went down.

Time to get route changes from IGP

Using default values, when a route goes down, the Routing Table Manager informs BGP that a route changed, and BGP then stores that event for up to 30 seconds before reading the new routes.

To configure the time to get route changes from IGP to values less than 30 seconds, you must use:

```
SETDefault -BGP IGPimportTime = <5/10/15/20/25/30>
```

Time to send a route withdrawal/update

BGP gives priority to route withdrawals, and updates are sent only after all the pending withdrawals are completed. Using default values, BGP delays sending withdrawals/updates for 30 seconds. This delay is the delay *after* BGP is notified of the route change, so route withdrawals can conceivably be sent up to 60 seconds after the event. As the number of hops increases, the variability in these withdrawals may result in different parts of the network converging at different times.

To configure the time delay before BGP sends a route update/withdrawal to values other than 30 seconds, you must use:

```
SETDefault -BGP AdvertiseDelay = <0/5/10/15/20/25/30/35/40>
```

Time to withdraw routes from a failed peer


At the expiration of the BGP `HoldTime` parameter (which damps any loss of connectivity due to temporary microwave fades causing a loss of KeepAlives), all routes advertised by the peer must be withdrawn as soon as possible so peers to the device can use the alternate routes immediately. You can accomplish this by setting the `AdvertiseDelay` parameter to 0.

5.11

IP Network Topology

An IP network is configured on each interface where IP packets are received and sent. The interface can be either a local LAN interface or a serial line interface for a wide area network. The following figure shows a wide area device (Router 1) connecting two local Ethernet networks (Santa Clara buildings 1 and 2) to two wide area networks (Los Angeles and Santa Barbara). The Los Angeles network is connected by a point-to-point line, and the Santa Barbara network is connected by a Frame Relay link.

Figure 50: Wide Area Router Connecting Four IP Networks



Although the wide area ports that connect the Santa Clara network to the Los Angeles network are assigned IP addresses, PPP does not require that you assign an IP address to each wide area port. If you do not want to assign an IP address to a wide area port, you must set the `SETDefault -IP NETAddr` command to `UnNumbered`. For more information on this topic, see [Configuring Basic IP Routing over LAN Points and PPP Links on page 119](#).

CAUTION: Each IP address that you assign directly to a port must be unique; in other words, you cannot assign the same IP address to different ports.

A local network is referred to as an *attached network*. When two wide area routers are connected by one or more serial lines, their serial interfaces should be on the same network. For example, in the figure above, port 3 of Router 1 and port 3 of Router 2 are on the same network.


As an IP network grows, you can group contiguous routers into areas if you are using OSPF. You can interconnect two areas through a backbone area and group areas and backbone areas into autonomous systems. An autonomous system consists of routers and networks administered by a single authority. An autonomous system typically runs a single intra-autonomous system routing protocol, such as OSPF.

NOTE: RIP does not support areas.

Routing can take place between autonomous systems using an interautonomous system protocol, such as BGP.

The following figure shows two areas within an autonomous system connected by a backbone area. It also shows an autonomous system connected to two other autonomous systems.

Figure 51: Typical IP Network Running OSPF



A router must check its routing table to determine where to route a packet. If the destination is on an attached network, the router can send it directly to the network. But if the destination is farther away in the internetwork, the router must route the packet to another router (called a gateway) that is closer to the destination. The route to a network beyond a gateway can be statically configured or dynamically learned through RIP, OSPF, IISIS, and BGP. When two routers are located on the same network (each of them has at least one interface to the network), they are considered neighbors or neighboring gateways.

5.12
Multipath Routing

EOS devices support multipath routing, which means that up to four routes for each destination address can be stored in the routing table. The following are the advantages of multipath routing:

- The device can still route a packet using an alternative route if the primary route fails. It is more responsive to network topology changes than with the existence of only one route to a destination.
- The device can distribute the load among the available equal-cost best paths.

Routes learned by different routing protocols are assigned a different precedence in the routing table. Some routing protocols such as OSPF and IISIS have multiple route classes. The classes are also assigned a different precedence.

When multiple routes for a destination exist, the device uses the route with the highest precedence. The following table lists the types of routes used by decreasing precedence (precedence level 1 is the highest).

Table 14: Route Precedence

Precedence Level	All Protocols without BGP	All Protocols with BGP
1	Static (added without Override)	Static (added without Override)
2	OSPF Intra Area	OSPF Intra Area

Precedence Level	All Protocols without BGP	All Protocols with BGP
3	OSPF Inter Area	OSPF Inter Area
4	IISIS Intra Area	IISIS Intra Area
5	IISIS Inter Area	IISIS Inter Area
6	RIP	RIP
7	OSPF Type 1 External	OSPF Type 1 External
8	OSPF Type 2 External	OSPF Type 2 External
9	IISIS External	IISIS External
10	ICMP redirect (not applicable for a router; only for host mode)	ICMP redirect (not applicable for a router; only for host mode)
11	Static (with override)	Static (with override)
12a		BGP
12b		Route configured with DefaultNets
13		Non-BGP default route

When the IP Protocol routes a packet and BGP is enabled, the software looks up the route in the following order:

1. The software looks for a route learned by any protocol except BGP in the All Protocols Routing Table. The software searches in the order specified in the **All Protocols without BGP** column in the table above. BGP has its own routing table separate from all the other protocols. The software does not consider the default route yet.
2. If a route to the destination is not found, the software looks for the route in the BGP Routing Table (precedence level 12a in the table above).
3. If a route to the destination is not found in the BGP Routing Table, the software looks for any configured default networks configured with the `ADD -BGP DefaultNet <IPaddress>` syntax.
4. If no default networks are configured, the software again searches the All Protocols Routing Table, looking for default routes.

If BGP is disabled, the software follows steps 1 and 4.

OSPF Type 1 and 2 external metrics allow you to define how you want Autonomous System Boundary Routers (ASBRs) to report metrics.

- A Type 1 external metric is the sum of the metric learned within the autonomous system by OSPF plus the metric learned outside of an autonomous system by BGP.
- A Type 2 external metric is the metric learned outside of an autonomous system by BGP only.

In an OSPF environment, you can set your ASBR to report Type 1 or 2 external metrics by using the `InteriorPolicy`, `ExteriorPolicy`, `StaticPolicy`, and `DefaultMetric` parameters. For complete information on these parameters, see “OSPF Service Parameters” chapter in the *Enterprise OS Software Version 18.2 Reference Guide*.

If the route with the highest precedence fails, the route with the next highest precedence is used. A route in the routing table is deleted in the following situations:

- OSPF and IISIS both compute routes based on link state information from all routers. The entire routing table is recomputed each time the topology changes.
- A dynamic route learned through RIP is deleted when an EOS device times out and goes through the HOLD-DOWN and GARBAGE COLLECTION states. A device times out when it fails to hear from a neighbor

for a period that is six times the value of the `UpdateTime` parameter. For example, if the value of the `UpdateTime` parameter is 45 seconds, the device times out if it does not hear from its neighbor in 270 seconds.

- A `DELeTe ROUte` command removes a static route.
- A lowest precedence route is deleted when four routes of higher precedence are available. This situation occurs when a fifth route is learned and has a higher precedence than the lowest precedence route.

You can remove dynamic routes learned by RIP by using the `FLush -IP AllRoutes` command. However, this command does not flush routes learned by OSPF, IISIS, or BGP from the routing table.

Route Selection and Load Splitting

If two or more routes with the same route source precedence are available to reach a destination, the EOS device always selects the route with the lowest metric (measured in hops for RIP and in administrative cost for OSPF or IISIS). If there is more than one route learned by the same routing protocol with the same equal-cost, low metric, you can split the load between these routes on a round-robin basis. The `-IP CONTROl` parameter (`SplitLoad` | `NoSplitLoad`) determines whether load splitting is performed.

Because load splitting balances the load among different routes, it is recommended if two or more routes are available to reach a destination and the routes have similar metrics. However, if the routes connecting various networks have different metrics (there is only one route with the fewest hops or lowest cost to a destination), load splitting is not necessary.

Route Selection Examples

The following table is an example of a routing table and shows how a route is chosen.

Routes are selected on the basis of precedence, lowest metric, or in cases where multiple routes have the same precedence and metric, through load splitting or the first route discovered. The examples in the following table demonstrate these criteria and can be applied to all types of routes.

Table 15: Routing Table Containing Multiple Paths

Network	Gateway	Metric ¹	Route Source
10.0.0.0	129.213.1.1	100	OSPF – Intra
	129.213.1.2	1	RIP
20.0.0.0	129.213.16.1	1	RIP
	129.214.1.1	2	RIP
30.0.0.0	129.213.16.1	100	OSPF – Intra
	129.213.16.2	100	OSPF – Intra
	129.213.16.3	100	OSPF – Intra

¹ RIP uses hop count as its metric. The OSPF metric is computed from total administrative cost between EOS device and destination.

Route Selection Example 1

For network 10.0.0.0, there are two routes available, but these routes are not comparable. The first route is learned by OSPF, and the second route is learned by RIP. Because routes learned by OSPF take precedence over routes learned by RIP, gateway 129.213.1.1 is selected.

Route Selection Example 2

For network 20.0.0.0, there are two routes available through RIP. Because gateway 129.213.16.1 requires one hop and gateway 129.214.1.1 requires two, the device always selects gateway 129.213.16.1 because it requires the fewest hops or the lowest metric to reach its destination.

Route Selection Example 3

The routing table entry for network 30.0.0.0 has three available routes to reach it. All are dynamic routes learned through OSPF, and all require an administrative cost of 100. The device chooses the route as described in the following table:

Table 16: Route Selection With and Without Load Splitting

Route Selection	Description
With load splitting	The route is chosen on a round-robin basis. The gateways are used in the following order: 1. 129.213.16.1 2. 129.213.16.2 3. 129.213.16.3 If one of these routes becomes invalid, it is no longer considered in the selection procedure.
Without load splitting	The route recorded earliest is always used. In this case, the gateway 129.213.16.1 is used.

5.13

Default Routes

When a device needs to route a packet destined for an address for which there are no entries in the routing table, it uses the default route if one exists. The network 0.0.0.0 represents the default route. The device supports up to four default routes; when more than one default route is available, the same selection rules apply. If load splitting is enabled, the load is distributed among equal-cost best paths. For additional information, see [Multipath Routing on page 168](#).



NOTE: An advantage of a device using a default route is that network overhead in an autonomous system can be reduced. The reduction in overhead occurs because the device does not need to advertise all external routes.

Default Routes Example

Router A receives a RIPNG update packet from Router B, which has an entry indicating that network 0.0.0.0 is reachable with metric 3. Router B is a default gateway for Router A; if Router A needs to route a packet with a destination not found in its routing table, it sends the packet to Router B.

You can configure the interior routing protocols for IP (RIP, OSPF, and IISIS) to advertise a default route by assigning a non-zero value to the `DefaultMetric` parameter of the routing protocol's service. You do not need to configure the `DefaultMetric` parameter on every EOS device throughout the domain. The default route learned on one interface is propagated to neighbors on the other interfaces (unless inhibited by the `NetworkPolicy` parameter).

Each interior routing protocol propagates the advertisement of the default route as the normal operation. For the RIP Protocol, you can suppress propagation of the default route using the `AdvertisePolicy` parameter. Since OSPF and IISIS are both link state routing protocols, they cannot suppress any routing information within the bounds of their routing system. However, they can control information that they import from other routing systems.

If more than one interior routing protocol is in operation on a network, the routes from one system can be introduced into the other system using the `InteriorPolicy` parameter of the protocol that is importing the routes.

Suppose that RIP and IISIS are both in operation, and that RIP needs to import routes from the IISIS system. Use the `InteriorPolicy` parameter in the RIP Service to achieve this routing. For more information on this parameter, see the “RIP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

When an EOS device is operating both an interior protocol (RIP, OSPF, or IISIS) and an exterior protocol (BGP), then the `ExteriorPolicy` parameter in the service of the interior routing protocol is used to control the import of routes from the exterior routing protocol into the interior routing protocol. The `InteriorPolicy` parameter in the service of the exterior routing protocol is used to control the import of routes from the interior routing protocol into the exterior routing protocol.

You can configure the default route in one of two ways:

- On the exit router of a domain, you must configure a static override default route with a metric of 1 that points to the first hop outside the domain. Then use the `StaticPolicy` parameter of RIPNG to import this route into the routing protocol and advertise it into the domain.
- At the top level, you must set the `DefaultMetric` parameter in RIPNG to instruct the device to originate a default route.



NOTE: The second method may also be useful at a device that interconnects a RIP domain with an OSPF domain, instead of importing all routing information from each domain into the other.

If the device is instructed to originate the default route (the `DefaultMetric` parameter is set to a nonzero value), it does not accept another device’s advertisement of the default route. Consider two devices in parallel, both originating a default route and each accepting the other’s advertisement of the default route. Any packet received by one device is forwarded to the other device, and back again, until the time-to-live timer is exhausted and the packet is dropped.


5.14

Route Learning Within an Autonomous System

EOS devices fully support RIP according to RFC 1058. They also support OSPF Version 2 according to RFC 1583.

If you plan to use both RIP and OSPF when expanding your IP network, Motorola Solutions recommends that autonomous systems using OSPF make up the core of the network and that autonomous systems using RIP surround those using OSPF.

Figure 52: Recommended Autonomous System Topology



If a device runs both OSPF and RIP Protocols, the routes learned by one of these protocols are not reported by the other according to the default settings of the `InteriorPolicy` parameter in both the RIP and OSPF Services. If you want cross-reporting between these protocols, set the `InteriorPolicy` parameter in the RIP, IISIS, and OSPF Services accordingly.

If you use OSPF in a topology with end stations, you must configure a default gateway on the end stations. Many end stations learn RIP routes dynamically, but they usually do not learn OSPF routes dynamically.

5.14.1

Route Learning with RIP

Normally, every 30 seconds or every time it learns a route change for a network, the EOS device uses broadcast packets to report to its neighboring gateways the following types of information:

- The networks it can reach
- The metric associated with each network it can reach

By default, the information in update packets pertains only to learned routes. Static route information is not reported.

You can configure some routing parameters (see [Route Learning with OSPF on page 177](#)) to determine how the device sends out the updates and what is included in them. For example, you can configure the parameters for the following purposes:

- To change the frequency of the broadcast traffic (`UpdateTime` parameter)
- To prevent the device from sending or receiving update and request packets (`CONTRol` parameter)
- To control the set of neighboring devices from which the EOS device receives updates and to which it sends them (`AdvToNeighbor` and `RcvFromNeighbor` parameters)
- To prevent the EOS device from sending out a trigger update response upon a route change for a network (`CONTRol` parameter)
- To enable the EOS device to report static routes (`StaticPolicy` parameter)
- To enable the EOS device to report routes learned with other interior routing protocols, such as IISIS and OSPF (`InteriorPolicy` parameter)
- To cause some routes not to be reported or to be reported with the infinity metric, that is, by using poison reverse (`CONTRol` parameter)

5.14.2

Network Reachability

The following types of networks are considered **reachable** when an EOS device broadcasts its RIP update packets:

- All directly connected networks, unless the network is shared by its neighbor and itself
- All static routes (as controlled by the `StaticPolicy` parameter)
- All dynamic routes learned through RIP and either OSPF or IISIS in the routing table (as controlled by the `InteriorPolicy` parameter)
- All dynamic routes learned through BGP (as controlled by the `ExteriorPolicy` parameter)

5.14.3

Solution to Slow Convergence Problem with Split Horizon


Ideally, all routers learn of new routes and discard obsolete routes immediately. That is, the contents of their respective routing tables converge rapidly so that all routing tables always contain correct information. An undesirable side effect of RIP is the possibility that the time is prolonged during which the unreachable network is considered reachable. One solution to this problem of slow convergence is called *split horizon*.

In a WAN environment, the Motorola Solutions implementation of next-hop split horizon (`-RIP IP CONTro1 = NonMesh`) eliminates the need for a fully meshed network when using RIP. In next-hop split horizon, the router learning of a network records the IP address of the neighbor from which the network was learned instead of recording the port through which the network was learned. When the router advertises its own reachable networks, it advertises to all neighbors except the one from which it learned of the network that is advertised.

Non-meshed WAN Networks

The following figure shows a non-meshed Frame Relay network using RIP on which Router R is the root router and Routers A, B, and C are remote devices that are configured as neighbors on Router R. Router R sends RIP updates individually to its neighbors, remote Routers A, B, and C. When sending RIP packets, Router R advertises to neighbors all networks it knows about (in this example, networks 20.0.0.0, 30.0.0.0, 40.0.0.0, and 50.0.0.0) if next-hop split horizon is not used. Network 10.0.0.0, common to all routers in the diagram, is automatically excluded from RIP updates between these routers.

Figure 53: Route Advertisement Over Non-meshed Frame Relay Network



By applying next-hop split horizon, Router R does not advertise network 30.0.0.0 to Router A, because it learned of 30.0.0.0 from Router A. Router R also does not advertise network 40.0.0.0 to Router B, nor does it advertise 50.0.0.0 to Router C, because it learned of those networks from those routers.




NOTE: You must enable the next-hop split horizon feature by setting `-RIP IP CONTro1` to `NonMesh`.

Meshed WAN Networks

In the following figure, the WAN network is meshed because all routers are directly connected to one another. You must set `-RIP IP CONTro1` to `FullMesh` so that next-hop split horizon is disabled. With Frame Relay networks, RIP neighbors can be dynamically learned.

Figure 54: Route Advertisement Over Meshed Frame Relay Network



LAN NETWORKS

On a LAN network, it is not necessary to configure neighbors. If you do not configure neighbors, RIP broadcasts the updates over the LAN. If you configure neighbors, RIP unicasts the updates.

5.14.4

Solution to Slow Convergence Problem with Poison Reverse

You can configure poison reverse or no poison reverse using the `Poison` or `NoPoison` value for the `-RIP` `CONTROL` parameter.

If poison reverse is enabled, the EOS device advertises all routes to all neighbors, but when advertising a route to a neighbor that advertised the same route, the device sets the metric to infinity (0xFFFF) to prevent the recipient from adding the route to its routing table. Poison reverse speeds convergence but adds to network overhead.

If poison reverse is disabled, the EOS device omits routes learned from one neighbor from RIP updates sent to that neighbor. No poison reverse has the advantage of minimizing network overhead in large network configurations at the expense of slower convergence.

5.14.5

RIP Configuration Parameters

The following table shows how to change the way the EOS device broadcasts or processes RIP update packets.

 **NOTE:** The table includes only the parameters that were not discussed in previous sections. For complete information on the parameters listed in this table, see “RIP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Table 17: Configuring the EOS Device for RIP Updates Using RIP

Parameter	Result
UpdateTime ¹	Changes the frequency of the update packets.
CONTROL parameter option:	
TRigger NoTRigger	Determines whether a route change for a network triggers an update packet from the device.

Parameter	Result
AdvToNeighbor	Determines to which gateways on the directly connected networks the device sends the update packets.
RcvFromNeighbor	Determines to which gateways on the directly connected networks the device should listen for routing information.


¹ The parameter applies to the entire device. All devices exchanging RIP information should have the same value for this parameter. Otherwise, routing loops or loss of connectivity in the network may occur.

5.14.6

Different States of RIP-Learned Routes

To avoid routing loops, new information about a route is ignored for a designated period before it is used. The following figure summarizes how a route learned through RIP changes states.

Figure 55: Different States of a RIP Route



- **GARBAGE COLLECTION state**

When the timer for a route that is in the **HOLD-DOWN** state expires, that route changes to **GARBAGE COLLECTION** state. This happens when no update packets are received to indicate that the route is still reachable. In this state, if a neighboring gateway reports the route with a non-infinity metric within 120 seconds, the route can go back to the **UP** state. If no updates are received within 120 seconds (garbage-collection timer), the route is deleted from the routing table. It is possible to go into **GARBAGE COLLECTION** state if no updates are received within 180 seconds.

- **UP state**

A route is considered **UP** if it is reachable with a non-infinity metric (15 or fewer hops). Whether it is reachable is determined by the last update received from the neighboring gateways. It remains **UP** for 180 seconds (the route timer). The timer is reset each time a new update for the route is received.

- **HOLD-DOWN state**

A route in **UP** state changes to **HOLD-DOWN** state if an update received from the original gateway indicates that the route is associated with an infinity metric (16 hops). In this state, all update information received from other gateways for that route is ignored.

However, if an update is received from the original gateway within 60 seconds (the hold-down timer), and it associates a non-infinity metric with the route, the route goes back to **UP** state.

If the hold-down timer expires, the route goes from **HOLD-DOWN** state to **GARBAGE COLLECTION** state for 120 seconds.

When you display the routing table with the `SHoW -IP AllRoutes` command, the state of each route is displayed under the **STATUS** heading.

5.14.7

Route Learning with OSPF

Normally, every 30 minutes or every time the EOS device learns a route change for a network, it uses multicast packets to report to its neighbors the following types of information:

- The networks and the directly connected devices
- The metric associated with each directly connected device and network

In an unchanging topology, OSPF only sends updates every 30 minutes while RIP sends updates every 30 seconds. OSPF provides a significant savings in network overhead when compared to RIP.

5.14.7.1


Different Functions of OSPF Devices

In an autonomous system running OSPF, devices can be assigned several different functions. An OSPF device can be assigned to route within an area (intra-area), between areas (interarea), or between autonomous systems (interautonomous system).

By default, the device performs intra-area routing. A device that routes between areas is an Area Border Router (ABR). Routing between autonomous systems is performed by a device that acts as an Autonomous System Boundary Router (ASBR). In addition to its routing function, a device can function as the designated router (DR) or backup designated router (BDR) on a multiaccess network. (A multiaccess network is any network other than a point-to-point link, such as Frame Relay, or a LAN.)

The following figure is an example of an autonomous system running OSPF, with devices configured as described in the preceding paragraphs. A stub network is a network that only has one OSPF device; a multiaccess network has more than one OSPF device.

Figure 56: Autonomous System With Multifunctional OSPF Devices



5.14.7.2

Area Border Router

An Area Border Router (ABR) is a device that has interfaces in more than one area.

For example, in [Figure 56: Autonomous System With Multifunctional OSPF Devices on page 178](#), ABR2 interfaces Network T3, which is part of Area 1. It also interfaces Network T4, which is part of the backbone area. (A device automatically acts as an ABR when different area numbers are assigned to different ports.)

Each ABR maintains a distinct database for each area to which it belongs. In the example shown in the figure above, ABR2 maintains databases for Area 1 and the backbone area.


All ABRs should have at least one interface connected to the backbone area. However, if there are no interfaces of an ABR connected to the backbone area, you can configure a virtual link to provide complete connectivity.

A virtual link is established between two ABRs. One of the ABRs must be directly connected to the backbone area, which provides a link for the other ABR. Also, both ABRs must be part of at least one common nonbackbone area for the virtual link to be established.

For example, in the following figure, ABR1 is an ABR for Areas 1 and 2. However, it is isolated from the backbone area. Because ABR1 and ABR2 are both connected to Area 1 and ABR2 is connected to the backbone area, a virtual link can be established between ABR1 and ABR2. The `VirtualLink` parameter

allows you to establish a virtual link between two ABRs. For more information on this parameter, see “OSPF Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Figure 57: Backbone Area With One Isolated ABR



5.14.7.3

Autonomous System Boundary Router

An autonomous system boundary router (ASBR) is a device that interfaces one or more devices in other autonomous systems.

For example, in [Figure 57: Backbone Area With One Isolated ABR on page 179](#), ASBR1 interfaces networks in Autonomous Systems 1 and 2. (A device automatically acts as an ASBR when different routing protocols [RIP, OSPF, IISIS, or BGP] are enabled on different ports.)

An ASBR can also function as an ABR if it is connected to more than one area in addition to the connection to another autonomous system.

Typically, an ASBR runs an interautonomous system routing protocol, such as BGP, on the interface that connects the other autonomous systems. On interfaces within the autonomous system that it is part of, the ASBR runs an intra-autonomous routing protocol, such as RIP, OSPF, or both.

In addition to its routing function, a device that is elected as the designated router (DR) on the multiaccess network performs the function of flooding for the network. A device that is elected as the backup designated router (BDR) on the multiaccess network should be adjacent to the same devices to which the DR is adjacent. The BDR has a subset of the DR’s responsibilities. The BDR takes over the DR’s role in the event of its failure. For details on the functions performed by the DR and the BDR, see [Routes Learning and Network Reachability on page 179](#).

The value of the `RouterPriority` parameter determines which devices on a multiaccess network function as the DR and the BDR. For more information on the `RouterPriority` parameter, see “OSPF Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

5.14.7.4


Routes Learning and Network Reachability

In an autonomous system running OSPF, intra-area routes, interarea routes, and interautonomous system routes are learned as described in this section. Each EOS device periodically exchanges a hello packet with its neighbor. The hello packet includes a list of all devices from which the originating device recently received a hello packet. The exchange of hello packets establishes a bidirectional relationship between neighbors. The `HelloTime` parameter allows you to set the frequency at which hello messages are sent. If you modify the setting of the `HelloTime` parameter, you must check the setting of the `RouterDeadTime` parameter. If the `HelloTime` parameter is set to a value larger than that of the `RouterDeadTime` parameter, the devices do not become fully adjacent. For more information on the `HelloTime` and `RouterDeadTime` parameters, see the “OSPF Service Parameters” chapter in the *Enterprise OS Software Reference Guide*.

After bidirectional relationships are established between neighboring devices, each pair of neighboring devices must decide if they should form an adjacency (a formalized bidirectional relationship). Neighboring devices connected by a point-to-point network always form adjacencies. (A point-to-point network is a network in which two devices are connected through a single network connection.) However, on a multiaccess network, adjacencies are formed only between DRs and BDRs and each of their neighbors.

For example, in the following figure, DR1 is the DR and BDR1 is the BDR for Network T1. Adjacencies are formed between DR1 and RT1, between BDR1 and RT1, between DR1 and RT3, between BDR1 and RT3, between RT1 and RT2, and between DR1 and BDR1. However, RT1 and RT3 did not fully establish an adjacency with each other. They are in a state known as *a two-way state*.

Figure 58: Forming Adjacencies on a Multiaccess Network



If adjacencies were formed between each device and its neighbor on a multiaccess network, it can be shown mathematically that the amount of traffic on the network would be significantly heavier. By minimizing these adjacencies, the DRs and BDRs reduce this traffic to manageable proportions.

You can use the `SHoW -oSPF NeighborStatus` command to display an OSPF neighbor status table, which shows the status of direct connect neighbor adjacencies for your device. For more information on the `SHoW NeighborStatus` command and an explanation of OSPF neighbor status table entries, see “OSPF Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

After an adjacency is formed between a pair of devices, each device on a point-to-point network and the DRs and BDRs on a multiaccess network send out a link state advertisement to its neighbor every 30 minutes or whenever a change in topology occurs. External link state advertisements are flooded throughout the device’s autonomous system. Other link state advertisements are flooded within a single area. For details on link state advertisements, see [Link State Advertisements on page 181](#).

Each device in an area maintains an identical database of the area’s topology. The database contains both the topology of the device’s area and routes to networks outside of the device’s area. This database is used to build a shortest path tree. The device doing the computation uses itself as the root of the tree and builds each node of the tree based on the metric advertised in the link state advertisement.

The EOS device always selects the path with the lowest metric. For details on metrics, see [Metrics on page 182](#).

If there is more than one equal cost path, the EOS device can use multipath routing and load splitting. For more information on these features, see [Multipath Routing on page 168](#). The device stores information on all reachable networks in its routing table.

The following types of networks are considered reachable:

- All directly connected networks, unless the network is shared by its neighbor and itself
- All static routes (if configured)
- All dynamic routes learned through RIP, OSPF, and IISIS in the routing table
- All dynamic routes learned through BGP (if configured)

5.14.7.5

Network Costs Reduction with Demand Interface Circuits

In a remote office internetworking environment, many remote offices are connected to a central site through *demand circuits*, such as ISDN circuits or dial-up lines with modems, or Frame Relay SVC neighbors. The cost of these demand circuits depends on the connection time or line usage.

OSPF periodically sends hello packets to refresh routing information, requiring the circuit to be constantly open, which results in unwanted usage charges. To reduce the cost of running OSPF over demand circuits and increase bandwidth, OSPF was modified to operate more efficiently over demand circuits. When no network topology changes occur, OSPF sends no routing information traffic at all, allowing the data link connection to be closed when not required for application data traffic. As soon as data is sent, a data link connection is attempted. If the connection is successful, the data is sent and the circuit stays open. After a period of inactivity, the circuit is closed again to conserve cost and resources.

Using the `-OSPF DemandInterface` parameter, you can configure an interface to be a demand interface. The neighboring device must agree that the point-to-point link is a demand circuit by setting the DC bit defined in the OSPF Options field of device LSAs, OSPF hello packets, and database description packets as follows:

- In a device's self-originating LSAs, the DC bit is set if and only if the device can properly process LSAs having the `DoNotAge` bit set.
If the `DoNotAge` bit is set, only truly changed LSAs are flooded over demand circuits. If a newly received LSA is only a periodic refresh, it is not flooded on attached demand circuits.
LSAs are not aged while they are held in the link state database, meaning they do not have to be refreshed, further reducing the routing traffic and the amount of time the circuit must remain up.
- For hello and database description packets, the DC bit is set in outgoing packets if and only if the device wants to treat the attached network as a demand circuit and tries to negotiate with the neighboring device for the suppression of hellos on point-to-point demand circuits.
Over point-to-point demand circuits, both end points must agree to suppress sending hello packets by setting the DC bit in OSPF hellos and database description packets. Receiving a packet with this setting indicates agreement, and OSPF hello packets are sent only until initial link state database synchronization is achieved with the neighbor. After the state of the neighbor connection reaches "full," hellos are suppressed and the data link connection to the neighbor is assumed to be available.
For OSPF broadcast and nonbroadcast multiaccess (NBMA) networks that are configured as demand circuits, the exchange of hello packets remains periodic for the proper operation of the DR election algorithm.

5.14.7.6

Link State Advertisements

A link state advertisement identifies the state of a device's interfaces and adjacencies. The types of link state advertisements sent out by a device depend on the function that the device was configured to perform. Multiple link state advertisements can be contained in a link state update packet. There are four types of link state advertisements:

Router link state advertisements

Each OSPF device sends out a router link state advertisement. This advertisement describes its links to stub and multiaccess networks as well as links to other devices for a given area. (A stub network is a network that only has one OSPF device; a multiaccess network has more than one OSPF device.) The advertisement is flooded throughout the area the originating device belongs to.

Network link state advertisements

DRs on multiaccess networks send out network link state advertisements. These advertisements describe the devices on the network that are fully adjacent with the DRs. These advertisements are flooded throughout a single area.

Summary link state advertisements

ABRs send out summary link state advertisements. These advertisements summarize all the interarea routes for all the areas to which the device is attached. (Each available interarea route is summarized in a separate summary link state advertisement.) These advertisements are flooded throughout the area that the ABR interfaces.

External link state advertisements

ASBRs send out external link state advertisements. These advertisements contain information on destinations outside of the autonomous system in which the device resides, including static and dynamic routes learned by RIP, BGP, and IISIS. (Each destination outside the autonomous system is described in a separate external link state advertisement.) These external routes are described using Type 1 or 2 external metrics. These advertisements are flooded throughout the autonomous system in which the device resides. For more information on Type 1 and 2 external metrics, see [Metrics on page 182](#).

To view the short version of the link state database, you must use:

```
SHOW -OSPF LinkStateData
```

5.14.7.7

Metrics

OSPF uses an administrative cost as its metric. The `SETDefault -OSPF Cost` command allows you to set the cost for a specific path. The default value of the `Cost` parameter is 108 /bandwidth of the medium that interfaces a port. For 10 Mb Ethernet, the default value is 10; for T1 lines, the default value is 65. For example, to set the cost on port 3, you must use:

```
SETDefault !3 -OSPF Cost = 58
```

In this example, the T1 serial line that interfaces port 3 was assigned the cost of 58. For more information on the `Cost` parameter, see “OSPF Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

The device running OSPF selects the route with the lowest total administrative cost to reach its destination. For example, imagine that OSPF learns about two intra-area routes to reach a particular destination. Route 1 has the administrative cost of 200, while Route 2 has the administrative cost of 300. The OSPF device selects Route 1 because it has the lowest administrative cost.

5.14.7.8

OSPF Configuration Customization

The following table summarizes the OSPF parameters that allow you to customize the configuration of your OSPF device. For complete information on these parameters, see “OSPF Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Table 18: OSPF Configuration Parameters

Parameter	Operation
AreaID	Determines the area to which a specified port on a device belongs.
Cost	Determines metrics (cost and type of service) associated with a specified port.
DEBUG	Determines the level of OSPF tracing that is performed at the local console port.
DefaultMetric	Enables or disables advertisement of the default route and determines the metric for the default route.

Parameter	Operation
Delay	Determines the delay in seconds for the specified port. The delay time is added to all link state advertisements before it is sent on an interface.
DirectPolicy	Determines whether a locally attached network on a non-OSPF port should be advertised into the OSPF domain.
ExteriorPolicy	Determines which networks learned through BGP are reported in external link state advertisements and what metric and metric type to use.
HelloTime	Changes the frequency at which hello packets are exchanged between neighbors on a network
InteriorPolicy	Determines which networks learned through RIP and IISIS are reported in external link state advertisements and what metric and metric type to use.
Neighbor	Determines to which device on the directly connected network a device sends packets. Also determines how packet is addressed.
PassWord	Determines the password for a specified port to authenticate packets.
ReceivePolicy	Determines which networks from external link state advertisements are stored in routing tables and what metric to use.
RouterDeadTime	Changes the frequency for determining when a neighbor device is down.
ROUTerPriority	Determines the priority for a device on a specified port. The device with the highest priority becomes the DR for a multiaccess network.
StaticPolicy	Determines which static routes are advertised for a particular network and what metric is used.
StubDefaultMetric	Specifies whether or not the device should generate the default route and metric into the stub areas.
VirtualLink	Determines whether the specified port of a device acts as a virtual link between an area and a backbone area.

5.15

Autonomous System Routing with BGP

Border Gateway Protocol (BGP) is an interautonomous system routing protocol that is used to exchange routing information between different autonomous systems (ASs). An EOS device can use BGP to determine the reachability of networks outside of its AS.

BGP provides the following advantages:

- Consumes less bandwidth
BGP uses incremental updates to reduce the amount of routing information. When a BGP session is first established, the peers exchange the entire contents of their routing tables. After this initial data exchange, BGP peers only exchange changes to their routing tables, effectively reducing the size of their routing tables and consuming less bandwidth.
- Allows the detection of routing loops

BGP minimizes the occurrence of routing loops. In addition to network reachability information, BGP requires that update messages contain a list of the ASs the routing information traversed. Routing loops are eliminated because a device never selects a path that contains its own AS.

- Selects routes based on performance and policy constraints.
 BGP allows a default weight to be added to all internal and external routes before computing the degree of preference for a route. If there are multiple routes to the destination networks, the route with the highest weight is chosen, allowing some routes to have higher priority than others.

BGP allows user-configured AS-path policies and network number policies to be implemented. These policies determine whether a BGP speaker accepts and distributes routing information based on an AS-PATH attribute or IP network number.

5.15.1
External and Internal Peers


Two devices that exchange routing information using BGP are peers. The two following kinds of peers exist:

Internal peers
Two devices residing in the same AS are internal peers. Internal peers do not need to be attached to the same network.

External peers
Two devices residing in adjacent ASs are external peers. External peers must be attached to the same network. BGP uses the common network to exchange messages between external peers.

The following figure illustrates a network topology in which Router A and Router B are configured as internal peers. Router B and Router C are configured as external peers.

Figure 59: Internal and External Peers in an Autonomous System



Annotation	Description
1	Router A and Router B are internal peers
2	Router A and Router D are external peers

Each peer establishes a BGP connection with the other peer. After the connection is established, the peers exchange update packets that indicate the networks each peer can reach. A peer also may report the networks that other gateways in other ASs can reach.

For example, Router B reports to Router C all the networks that are reachable within Autonomous System 1, and Router C does the same for Autonomous System 2. In addition to reporting the networks reachable within Autonomous System 1, Router B also reports to Router C all the networks that are reachable through Router A.

5.15.2

Peer-to-Peer Communication

After BGP peers are configured, three peer-to-peer communication states can be established between two peers:

- Connection establishment state
- Confirm state
- Established state

The EOS device enters the connection establishment state immediately after you configure BGP and set up peers. In this state, Router B tries to establish a TCP connection with a configured gateway (Router C).

After a TCP connection is established, the devices exchange open messages in which the following information is exchanged:

- The version of BGP that a device wants to “speak”
- The hold time (maximum time for which a connection is kept open without receiving any keepalive or update packets)
- The AS number
- The Router ID

If the open messages are satisfactory, each peer enters the confirm state in which they exchange keepalive packets. When keepalive packets are received, the peers reach the established state, in which they exchange routing information in update messages.

The connection between peers is assumed to be a reliable TCP connection. After the peers exchanged routing tables, the only packets regularly exchanged (every 30 seconds) are keepalive packets. Routing updates occur only when new routes are reachable or previously advertised routes became unreachable, which greatly reduces the amount of routing traffic and the time required to exchange and process information.

Update messages contain all reachable network addresses and the corresponding distances associated with each gateway, as well as the complete AS path for each network. Update messages may also contain explicit unreachable routes.


5.15.3

Path Attributes

For each route, the BGP uses a set of path attributes to describe the route. These attributes help to eliminate looping of routing information, assist with policy-based routing decisions, indicate the original source of the path information as well as the IP address of the next-hop device to the destination, provide routing metrics, simplify the route selection process, and perform route aggregation.

Path attributes are classified as either well-known or optional, as shown in the following figure.

Figure 60: Path Attributes




- A well-known attribute must be recognized by all BGP implementations.
The well-known attributes are further divided into mandatory and discretionary. A well-known, mandatory attribute (ORIGIN, AS-PATH, NEXT-HOP) must be included in every route description. A well-known, discretionary attribute (LOCAL-PREF, ATOMIC-AGGREGATE) may or may not be included in a route description depending on whether the attribute is implemented by the BGP speaker. A BGP speaker that receives a well-known path attribute is required to forward the attribute to its peer in update messages.
- An optional attribute may or may not be recognized by a BGP implementation.
The optional attributes are divided into transitive and non-transitive. An optional transitive attribute (AGGREGATION) may be passed along unchanged by a BGP device that did not implement the attribute. An optional non-transitive attribute (MULTI-EXIT-DISC) may not be passed along by a BGP device that did not implement the attribute.

You can display detailed information about path attributes associated with AS paths by entering the `SHOW -BGP ASPath Debug` command.

AS-PATH

The BGP Protocol uses the AS-PATH attribute to eliminate the occurrence of routing loops. As reachability information for a network traverses the internetwork, BGP creates a list of the ASs through which the routing information passed. Each BGP speaker adds its own AS to the list before advertising network reachability to a peer, as shown in following figure. The list of the ASs along the path to a destination network is called the AS-PATH attribute.

Figure 61: AS-PATH Attribute Example



The AS-PATH attribute is composed of a sequence of AS path segments. Each AS path segment may be either an AS SEQUENCE or an AS SET:

- An **AS_SEQUENCE** is an ordered set of ASs that the route in the update message traversed.
- An **AS_SET** is an unordered set of ASs that the route in the update message traversed. AS SETs are used by the route aggregation algorithm to reduce the size of the AS path information. An AS_SET lists each AS number only once, regardless of how many times it appeared in the multiple AS paths that were aggregated. An AS_SET indicates that the destinations can be reached through paths that traverse at least some of the listed autonomous systems. AS_SETs provide enough information to eliminate the looping of routing information.

The **AS-PATH** attribute helps suppress routing loops. A device never accepts a route with its own AS in the AS-PATH list. The **AS-PATH** attribute can also be used to make policy-based routing decisions. For more information, see [Policies on page 190](#).

ORIGIN


The **ORIGIN** path attribute defines the original source of the path information. The **ORIGIN** path attribute may contain the values **IGP** or **Incomplete**. **IGP** indicates that the destination network was learned by the original BGP speaker from the Interior Gateway Protocol (IGP) running in the original AS; this routing information is considered to be trustworthy.

Incomplete indicates that the routing information was obtained from some means other than an IGP. For example, the route is learned by using a static configuration.

NEXT-HOP

The **NEXT-HOP** path attribute defines the IP address of the border router that should be used as the next-hop to the destination networks. A BGP speaker can use its own IP address or the IP address of another device attached to the same subnet. The **NEXT-HOP** attribute allows a BGP speaker to advertise routes through another border router attached to the same subnet. For example, as shown in the following figure, Router B and Router C are both border routers for AS 200. However, only Router C is a BGP speaker and has a session with Router A. Router C advertises the route to network 130.7.0.0 with Router B as the next-hop router.

Figure 62: NEXT-HOP Router




MULTI-EXIT-DISC

The **MULTI-EXIT-DISC (MED)** attribute is the BGP metric for a route advertised by a device to a different autonomous system (AS). Lower MED values are preferred over higher values. The MED value is non-transitive (cannot be passed along beyond a single AS hop). It is assigned administratively to outbound routes (via the device command line interface) to inform external BGP peers about the degree of path preference for routes entering the local AS via the advertising device.

If the border routers of an AS receive different metrics for the same network, they compare the different metrics. The result of the comparison along with other factors determines the best route. If the MED attribute is received over an external link, it may be propagated over internal links to other BGP speakers. However, it is never propagated to other BGP speakers in neighboring ASs.

In the following figure, Router A and Router B advertise a route to network 30.0.0.0, which is completely contained within AS 300. Router A advertises the route with a MED value of 10. Router B advertises the route with a MED value of 5. The border routers in AS 200 make a comparison between the two metrics to select a better entry point in AS 300. However, the BGP speakers in AS 200 never propagate this metric to AS 100.

Figure 63: BGP Routing Metric



LOCAL-PREF

The `LOCAL-PREF` attribute simplifies the route selection process. It advertises a degree of preference for each external route to BGP peers in the same AS so that a route with a higher degree of preference is selected over a route with a lower degree of preference.

For example, assume Routers A, B, and C are internal BGP peers in AS 100. Each device receives an advertisement containing prefix 1.1.1.1/32 from other external BGP peers, and the path through Router A is the preferred route to exit AS 100 and reach prefix 1.1.1.1/32.

- Router A is configured to set the `LOCAL-PREF` to a high value for prefix 1.1.1.1/32 before that prefix is advertised to Router A's internal BGP peers.
- Routers B and C select the path to 1.1.1.1/32 advertised by Router A rather than the path learned from their external BGP peers due to the high `LOCAL-PREF` value associated with the route advertised by Router A.

`LOCAL-PREF` is set through a policy on a BGP device for a given set of prefixes and is included as an attribute for those prefixes in advertisements to the device's internal BGP peer. This attribute is not included in advertisements sent to the device's external BGP peers.

ATOMIC-AGGREGATE

The `ATOMIC-AGGREGATE` attribute is attached to a less specific route before propagating it to other BGP speakers to ensure that the aggregate is not deaggregated by other BGP speakers.

If a BGP speaker is presented with a set of overlapping routes from one of its peers, the more specific route takes precedence. If the BGP speaker selects the less specific route, the device attaches the `ATOMIC-AGGREGATE` attribute.

AGGREGATION

The `AGGREGATION` attribute allows a BGP speaker performing route aggregation to advertise the AS that performed the aggregation. Aggregation is the process of combining several different routes so that a single route with a shorter mask can be advertised.

5.15.4

Path Selection

One of the most important tasks of BGP is to select the best path to a destination network based on the AS topology. In traditional routing protocols, each path has only a single metric to represent its cost. To evaluate two paths, the EOS device compares the two metrics and selects the path with the lowest cost metric.

In interdomain routing, no universally agreed-upon metric among ASs can be used to evaluate different paths to a network. Therefore, each AS may implement its own set of criteria for path selection.

Path selection for the Motorola Solutions BGP-4 implementation is based on the following criteria in order of priority:

User-defined policies

Policies that are configured to control the distribution of routing information affect the paths that are available and the path selection process.

User-defined policies

Policies that are configured to control the distribution of routing information affect the paths that are available and the path selection process.

AS weight factor

When multiple paths exist to a given network, you can assign weights for AS paths or subsets of AS paths. The weight for a path is calculated by summing all the individual AS-path weight expressions that are applicable for the path. The path having the highest weight is selected.

LOCAL-PREF

The `LOCAL-PREF` attribute is the BGP metric for an external route advertised by a BGP device to its internal peers. Higher `LOCAL-PREF` values are preferred over lower values. The `LOCAL-PREF` attribute is non-transitive (cannot be passed along beyond a single AS hop) and is applied only to intra-AS advertisements. `LOCAL-PREF` has no effect if it is configured for a route which is advertised to a different AS.

AS count

If competing paths have the same weight, the path with a lowest total AS count is preferred over paths that have a larger AS count. BGP considers the path with a lower number of AS hops to be shorter.

Link type

If competing paths have the same AS count, BGP prefers paths that arrive over external links to those that arrive over internal links.

Path origin

If competing paths arrive over the same type of link, paths that are originated by an IGP Protocol has precedence.

MULTI-EXIT-DISC

If competing paths have the same `ORIGIN` attribute, BGP selects the path with the lowest `MULTI-EXIT-DISC` attribute value.

BGP-ID

If all of the above result in a tie, the selection process gives preference to the path from the peer with the lower IP address (BGP-ID). The IP addresses are compared as unsigned 32-bit integers.

5.15.5


Policies

Policies determine whether AS-level routing information is accepted and distributed by a BGP speaker. Policies control routing information in the following ways:

- Routing information can be received by a BGP speaker but not added to the routing table.
- Routing information can be received and added to the routing table but only advertised to some of the device's BGP peers.

BGP stores its routes in a routing information database (RIB), which is conceptually divided into three distinct parts, as shown in the following figure.

Figure 64: BGP Routing Selection and Policies



The Adjacent RIBs-In contains unprocessed routing information that was received by the local BGP speaker from its peers. The information is learned from inbound update messages and represents routes that are available for input to the route-selection process of the local BGP speaker.

The Local RIB contains routes selected for use by the route-selection process of the local BGP speaker.

The Adjacent RIBs-Out organizes the routes that the local BGP speaker selected for advertisement to its neighboring BGP speakers using outbound update messages.

Incoming policies are applied as part of the path selection process to manage the flow of information from the Adjacent RIBs-In to the Local-RIB. Outgoing policies are applied to manage the flow of information from the Local-RIB to the Adjacent RIBs-Out. The redistribution of routes is performed only on routes that are placed in the Adjacent RIBs-Out.

The three important effects of this sequence of operations are as follows:

- Incoming filters are implicitly applied before re-advertisement takes place.
- Only those routes used by the local BGP speaker are considered for re-advertisement.
- Only one outgoing route is advertised for each network even if many incoming paths for that route are learned from a variety of peers.

Interior Policies

BGP obtains network reachability information from the following sources:

- Internal or external BGP speakers

- IGP speaker (RIP, OSPF, IISIS) residing in the same device
- Static route configurations
- Directly attached networks

To control the flow of routing information into BGP, you can configure an interior policy using the `InteriorPolicy` parameter. This parameter controls the blocking of IGP routes (RIP, OSPF, IISIS), static routes, and directly connected networks into BGP for advertisement to BGP speakers residing in adjacent autonomous systems.

By default BGP does not import any of the IGP, directly connected, or static routes.

Exterior Policies

Each IGP Protocol controls the import of BGP routes through the configuration of its `ExteriorPolicy` parameter (in the RIP, OSPF, and IISIS Services).

BGP aggregated routes can only be leaked into IGP domains if the IGP routing protocol supports a mask along with each network route. OSPF and IISIS support this feature, but RIP does not provide this information in update packets. You cannot export aggregated routes into domains that run RIP.

Network Number-Based Policies

Network number policies provide filtering of incoming and outgoing BGP advertisements based on IP network numbers. You can configure the following types of network filters:

- Do not accept routes for network x.x.x.x from BGP peer A.
- Do not advertise a route for network z.z.z.z to BGP speaker B.

The network number specified in these examples can be a single network number or a range of network numbers specified by a CIDR address prefix.

You configure network number policies using the `NetworkFilter`, `NetPolicyAll`, `NetPolicyExt`, `NetPolicyInt`, and `NetPolicyPeer` parameters. The `NetPolicy` parameters allow you to configure the policy on all peers, external peers, internal peers, or on a specific peer.

AS-PATH-Based Policies

The AS-path policy provides filtering based on information contained in the AS-PATH attribute in each update message. Typical policies contain a combination of the following elements:

- Source AS
- Destination AS
- AS presence (within the AS-PATH attribute)
- Advertise or receive

Using these elements, you can configure the following policies:

- Distribute routes from AS 2 only to ASs 3, 6, and 7.
- Accept only those routes from AS 4 that have AS 7 contained in the path.
- Do not accept routes requiring a path through AS 3 from AS 10.
- Distribute routes containing AS 5 only to ASs 3 and 4.

To maintain consistent routing information within an AS, do not apply accept or deny policies to internal BGP peers.

You configure AS-path policies using the `AsFilter`, `AsPolicyAll`, `AsPolicyExt`, `AsPolicyInt`, and `AsPolicyPeer` parameters.

For more information about the AS-PATH attribute, see [Path Attributes on page 185](#).

5.15.6

Route Aggregation

BGP route aggregation uses the Classless InterDomain Routing (CIDR) route aggregation strategy to combine several different routes so that a single route with a shorter mask can be advertised. By combining several networks into one supernet, the number of BGP messages sent to peers and the size of the routing table is reduced. Unnecessary details about subnets are hidden from peers. CIDR is a method of using IP addresses without regard to traditional address classes that helps reduce routing table growth by summarizing several networks or subnets with a single routing update.



NOTE: Supernetting can only be understood and supported by protocols that carry mask information along with routes, such as OSPF and IISIS. The RIP Protocol always interprets routes as Class A, B, or C and cannot fully interpret routes with supernet masks. Do not configure BGP route aggregation in this situation.

BGP devices learn all the subnet routes through an intradomain routing protocol, such as OSPF, or static configuration. The BGP device may advertise to its BGP peers a single aggregate route that describes all the destinations connected to it. When a BGP device performs route aggregation, it needs to know the range of block of IP addresses to be aggregated or not aggregated.

The BGP device should aggregate as many routes as possible except those that cannot be treated as part of a single unit due to multi-homing, policies, or other constraints.

Aggregation should never encompass Class D address space (224.0.0.0 through 239.255.255.255).

5.16

Integrated IS-IS (IISIS)

IISIS is a protocol that provides integrated OSI-type routing for IP and OSI environments; it is the IP extension added to the original OSI IS-IS Protocol. IISIS routing simplifies network topology, reduces network management complexity, and reduces routing traffic overhead. In IP environments, IISIS is an alternative to other IP routing protocols, such as RIP and OSPF.

The original IS-IS routing protocol was developed by ISO to provide network layer connectivity in OSI environments. IS-IS is designed to work with CLNP and ES-IS. IS-IS is an international standard.

You can use IISIS for IP environments only, and you can use IISIS in OSI mode for routing in pure OSI environments. You can also use IISIS in dual IP and OSI mode for routing in environments where both types of networks are used. In dual mode, for example, one device can serve IP and OSI subnets simultaneously and IISIS routes traffic between the two subnets.

5.17

Address Resolution

To resolve Internet addresses with associated Ethernet addresses when routing, EOS devices use the Address Resolution Protocol (ARP) as described in RFC 826.

Configure the `-ARP CONTrol` parameter to decide whether the device supports proxy ARP requests on the specified interface. A proxy request is a request for a target Internet address that is not on the subnet where the request originated. If the device generates proxy replies, it replies with its own Ethernet address, provided that it has a route in the routing table for the target subnet.

ARP determines the destination's Ethernet header format to be used by sending out ARP requests that include the format. Specify the format by configuring the `RequestFormat` parameter in the ARP Service. The system replying to the request then uses the Ethernet header format it supports, and the device records the IP address, the Ethernet address, and the Ethernet header format of the replying system. The information is valid for the time specified by the `HoldTime` parameter in the ARP Service. Configure the time if you want the

device to hold the information for more than or less than 24 hours, which is the default value of the `HoldTime` parameter.

Inverse ARP

Inverse ARP is an adaptation of ARP that resolves DLCIs on Frame Relay networks to IP addresses, as described in RFC 1293.

5.18

Other Global Routing Configurations

After you determine how the EOS device should route packets on each of its interfaces, you can influence the global routing operation in several areas, as follows:

- Treatment of Internet Control Message Protocol (ICMP) request packets and generation of packets
Configure the `-IP ICMPReply` parameter to control whether the device responds to Address Mask Request and Information Request packets. The device supports all the ICMP messages described in RFC 1009. Configure the `-IP ICMPGenerate` parameter to control whether the device originates ICMP ReDirect, Destination Unreachable, TimeExceed, and TimeStamp reply packets.
- How long the IP layer waits for all IP fragments
Configure the `-IP ReassemblyTime` parameter to control the length of time the IP layer waits for all IP fragments of an IP datagram to be received. This parameter applies only to packets specifically destined for the local device.
- The level of security
Use the parameters beginning with "Sec" in the IP Service to configure the system for IP security options processing.
- The value of time-to-live (TTL)
Configure the `-IP DefaultTTL` parameter to specify the value the device puts in the TTL field of an IP packet when it generates the packet.
- Prioritization of packets within the IP Protocol
Some actions of the `-IP FilterAddrs` parameter enable your device to do special processing of IP packets over WAN links, improving the IP WAN traffic management. When you use the `FilterAddrs` parameter with the `-IP Filters` parameter, you can specify the packets to which the special processing applies.

For more information, see the "IP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Chapter 6

Quality of Service (IPQoS)

Policy-based QoS management enables you to control bandwidth allocation and service levels on IP traffic flows, thereby ensuring that mission-critical and real-time application traffic is allocated adequate network resources to traverse the network regardless of the competing demands for bandwidth by other applications. Flows can also be policed to ensure correct marking of the IP/TOS-byte in the IP header. You can define an IP traffic flow as either an aggregated flow or a specific application flow between two end systems using the IP service `ClassifierList` parameter. For details, see the description of the `ClassifierList` parameter in the “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*

You configure QoS control via the IPQoS service as port-specific policies which are associated with the flows defined by the `ClassifierList` parameter. QoS policies can be applied to the inbound traffic at the ingress port and/or the outbound traffic at the egress port. Policies are stored in the precedence order that you define when you add them using the IPQoS service `POLICY` parameter. The policy action associated with the first matching policy found for the packet is applied.

You can define IPQoS based on the following criteria:

- Priority queueing
- Bandwidth control
- Class-based queueing (CBQ)
- Protocol reservation
- IEEE 802.1P prioritization
- TOS-to-IPIP mapping

You can also configure a default IPQoS policy to be used as a “catch all” policy if traffic does not match the other policies you defined. For details, see [Defining a Default IPQoS Policy on page 210](#).

6.1

IPQoS Priority Queueing

EOS supports IPQoS priority queueing for outgoing packets on:

- IPIP and IPsec interfaces
- Physical and VLAN interfaces

Each outgoing packet is put into an appropriate queue (urgent, high, medium, normal, or low) based on its IP TOS value. The packets are transmitted in priority order, providing absolute preferential treatment to high priority traffic and ensuring that mission-critical traffic is transmitted before lower-priority traffic. When congestion occurs on a queue, all the lower priority queues are flushed.



NOTE: IPQoS priority queueing is supported only on virtual ports for outgoing traffic on Ethernet interfaces.

When one or more VLAN interfaces are configured on a physical port, all those interfaces share the same instance of Priority Queues with the common Rate Limit configuration, configured primarily for the physical interface.

All queues have a limit, which is fixed at the following number of packets:

- Queue 0 (Urgent): 300 packets
- Queue 1 (High): 300 packets
- Queue 2 (Medium): 300 packets

- Queue 3 (Normal): 300 packets
- Queue 4 (Low): 300 packets

Optionally, you can define a queue depth caution threshold using the `CautionTrap PERcent` action attribute in the IPQoS `POLicy` parameter. The queue depth is a percentage value derived from the queue limit for each queue. When the number of packets in the queue exceeds the specified percentage of the queue limit, EOS sends a caution trap (`mnrWANIfQueueDepthCaution`), and this trap sets a time stamp. If the queue depth exceeds the specified percentage of the queue limit again, EOS checks to see if the interval specified by the `POrt` service `TrapTimeInterval` parameter (300 seconds by default) expired.

- If `TrapTimeInterval` did not expire, no additional trap is sent.
- If `TrapTimeInterval` expired, a second trap is sent and its time stamp becomes the new starting point for `TrapTimeInterval`.

For example if you set the `CautionTrap PERcent` to 50 for queue 1, the queue depth caution threshold is 150 packets (50% of 300) for that queue. When the 151st packet is received, EOS drops the packet, sends an `mnrWANIfQueueDepthCaution` trap and sets a time stamp. If the queue exceeds 150 packets again, EOS drops any packets above the threshold, checks to see if the `TrapTimeInterval` expired since the first trap and either ignores the second trap or sends it and makes its time stamp the new starting point for the `TrapTimeInterval`, as described above.

If you set the `CautionTrap PERcent` to 0, or do not specify a `CautionTrap PERcent` value, generation of the `mnrWANIfQueueDepthCaution` trap is disabled.

The device continues receiving packets as long as the queue limit is not exceeded. If the queue limit is exceeded, EOS drops subsequent packets and sends an `mnrWANIfQueueDepthExceeded` trap.

For example, the 301st and subsequent packets on queue 1 are dropped and EOS sends an `mnrWANIfQueueDepthExceeded` trap. If queue 1 exceeds 300 packets again, EOS checks to see if the interval specified by the `POrt` service `TrapTimeInterval` parameter (300 seconds by default) expired.

- If the `TrapTimeInterval` did not expire, the 301st and subsequent packets are dropped, but no trap is sent.
- If the `TrapTimeInterval` expired, the 301st and subsequent packets are dropped, a second trap is sent, and its time stamp becomes the new starting point for the `TrapTimeInterval`.

If all the queues for all priority-queue-enabled ports reach their limits, EOS sends an `mnrWANIfTotalQueueDepthExceeded` trap.

For details about the `CautionTrap PERcent` value see the description of the `POLicy` parameter in the “IPQoS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. For details about the `TrapTimeInterval` parameter, see the description of that parameter in the “POrt Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. For details about the `mnrWANIfQueueDepthCaution`, `mnrWANIfQueueDepthExceeded`, and `mnrWANIfTotalQueueDepthExceeded` traps, see [Supported MIBs and Traps on page 845](#).

6.1.1

Configuring IPQoS Based on Priority Queueing for IPIP and IPsec Interfaces



NOTE: When IPQoS priority queueing is implemented, address resolution protocol (ARP) requests and replies and internet key exchange (IKE) packets are not metered, so it is advisable to reserve some bandwidth for this traffic.

Procedure:

1. Configure IP routing, ports, and IP addresses as necessary for your system.
2. Configure classifier lists corresponding to IP precedence bits by entering:

```
ADD -IP ClassifierList <name> <priority> TOS <value>/<mask>
```

Step example:

```
ADD -IP ClassifierList cl_pq1 1 tos e0/70
ADD -IP ClassifierList cl_pq2 2 tos c0/70
ADD -IP ClassifierList cl_pq3 3 tos a0/70
ADD -IP ClassifierList cl_pq4 4 tos 80/70
ADD -IP ClassifierList cl_pq5 5 tos 60/70
ADD -IP ClassifierList cl_pq6 6 tos 40/70
ADD -IP ClassifierList cl_pq7 7 tos 20/70
ADD -IP ClassifierList cl_pq8 8 tos 00/70
```

3. Specify priority queueing as the queueing policy on the desired IPIP virtual port(s) by entering:

```
SETDefault !<port> -Port QueueCONTRol = PriorityQueues
```

Step example:

```
SETDefault !v2 -Port QueueCONTRol = PriorityQueues
```

4. Specify that the priority queueing are applied only on outgoing traffic by entering:

```
SETDefault !<port> -IPQoS CONTRol = OutFilter
```

Step example:

To enable the outbound IPQoS policy filter on virtual port 2, enter:

```
SETDefault !v2 -IPQoS CONTRol = OutFilter
```

5. Specify that EOS ignores the don't fragment (DF) bit in the IP header of incoming packets and fragments and forwards packets that exceed the configured MTU size by entering:

```
SETDefault !<port> -Port FwdCONTRol = IgnoreDFBit
```

Step example:

```
SETDefault !v2 -Port FwdCONTRol = IgnoreDFBit
```

6. Specify that fragmented packets received on the IPIP tunnel are reassembled before forwarding to their destination by entering:

```
SETDefault !<port> -Port FwdCONTRol = ReassembleFragPkts
```

Step example:

```
SETDefault !v2 -Port FwdCONTRol = ReassembleFragPkts
```

7. Define IPQoS policies per queue for the classifier lists you defined in step 2 by entering:

```
ADD !<port> -IPQoS POLicy <name> <priority> <classifier_list_name>
<direction> (PriorityQue Urgent | High | Medium | Normal | Low
[CautionTrap <[PERcent] (0-100)>])
```

If desired, you can use the optional CautionTrap PERcent value to specify the queue depth above which a caution trap is generated. If you set the CautionTrap PERcent value to 0 or do not specify a CautionTrap PERcent value, generation of caution traps is disabled.

Step example:

To define IPQoS policies that assign priority queues to outbound IP traffic matching classifier lists cl_pq1 through cl_pq8 on virtual port 2, enter:

```
ADD !v2 -IPQoS POLicy qos1 1 cl_pq1 OUTbound PriorityQue Urgent
ADD !v2 -IPQoS POLicy qos2 2 cl_pq2 OUTbound PriorityQue High
ADD !v2 -IPQoS POLicy qos3 3 cl_pq3 OUTbound PriorityQue High
ADD !v2 -IPQoS POLicy qos4 4 cl_pq4 OUTbound PriorityQue Medium
```



```
ADD !v2 -IPQoS POLicy qos5 5 cl_pq5 OUTbound PriorityQue Medium
ADD !v2 -IPQoS POLicy qos6 6 cl_pq6 OUTbound PriorityQue Normal
ADD !v2 -IPQoS POLicy qos7 7 cl_pq7 OUTbound PriorityQue Normal
ADD !v2 -IPQoS POLicy qos8 8 cl_pq8 OUTbound PriorityQue Low
```

8. Rate limit the outbound virtual port traffic by entering:

```
ADD !<port> -POrt RateLimit <name> <CIR> <PIR> <CBS> [TransmitExcess] [Default]
```

For IPQoS priority queueing, only the <CIR> and <CBS> values are used.

The <PIR> (peak rate) should always be set to 0.

In general, the burst duration (T) used for bandwidth calculation and round robin processing for IPQoS priority queues is calculated as $T = CBS / CIR$. However, if you want to use a default burst duration of 7 milliseconds, set <CBS> to 0, and EOS does not perform the CBS/CIR calculation.

The minimum value for <CBS> is the IPIP tunnel MTU configured for the virtual port with the IP service NETaddr parameter. If the <CBS> is less than the MTU, no packets are transmitted.

Step example:

To define rate limit name "PQ" that rate limits outbound IP traffic on virtual port 2 to a CIR (average rate) and CBS (burst size) of 256 kbps and 2048 bytes respectively, and to specify that this is the default rate limit for the virtual port, enter:

```
ADD !v2 -POrt RateLimit PQ 256 0 2048 Default
```

9. If you want to bundle a virtual port with up to 10 other virtual ports to share a rate limit (CIR) and priority queues, use:

```
ADD !<port> -POrt PortBundle <!port1, !port2 ... !port10>
```

Step example:

To bundle virtual port 2 with virtual ports 3, 4, and 5, enter:

```
ADD !v2 -POrt PortBundle !v3,!v4,!v5
```

6.1.2

Configuring IPQoS Based on Priority Queueing for Physical and VLAN Interfaces

This example shows how to configure IPQoS Priority Queueing on both the physical interface !4 and the related VLAN interface !v40

Procedure:

1. Configure a virtual port that accesses a VLAN by entering:

```
ADD !v40 -POrt VirtualPort 4 Vlan 401
```

2. Configure the IP address for the virtual port by entering:

```
SETDefault !v40 -IP NETaddr = 50.50.50.1 255.255.255.0
```

3. Configure classifier lists corresponding to IP precedence bits by entering:

```
ADD -IP ClassifierList c1 1 TOS e0/e0
ADD -IP ClassifierList c2 2 TOS c0/e0
ADD -IP ClassifierList c3 3 TOS a0/e0
ADD -IP ClassifierList c4 4 TOS 80/e0
ADD -IP ClassifierList c5 5 TOS 60/e0
ADD -IP ClassifierList c6 6 TOS 40/e0
```

```
ADD -IP ClassifierList c7 7 TOS 20/e0
```

```
ADD -IP ClassifierList c8 8 TOS 00/e0
```

4. Define IPQoS policies the physical interface !4 by entering:

```
ADD !4 -IPQOS POLicy qos1 1 c1 OUTbound PriorityQue Urgent CautionTrap PErcent 28
```

```
ADD !4 -IPQOS POLicy qos2 2 c2 OUTbound PriorityQue High
```

```
ADD !4 -IPQOS POLicy qos3 3 c3 OUTbound PriorityQue High
```

```
ADD !4 -IPQOS POLicy qos4 4 c4 OUTbound PriorityQue Medium
```

```
ADD !4 -IPQOS POLicy qos5 5 c5 OUTbound PriorityQue Medium
```

```
ADD !4 -IPQOS POLicy qos6 6 c6 OUTbound PriorityQue Normal
```

```
ADD !4 -IPQOS POLicy qos7 7 c7 OUTbound PriorityQue Normal
```

```
ADD !4 -IPQOS POLicy qos8 8 c8 OUTbound PriorityQue Low
```

5. Define IPQoS policies the VLAN interface !v40 by entering:

```
ADD !v40 -IPQOS POLicy qos1 1 c1 OUTbound PriorityQue Urgent CautionTrap PErcent 28
```

```
ADD !v40 -IPQOS POLicy qos2 2 c2 OUTbound PriorityQue High
```

```
ADD !v40 -IPQOS POLicy qos3 3 c3 OUTbound PriorityQue High
```

```
ADD !v40 -IPQOS POLicy qos4 4 c4 OUTbound PriorityQue Medium
```

```
ADD !v40 -IPQOS POLicy qos5 5 c5 OUTbound PriorityQue Medium
```

```
ADD !v40 -IPQOS POLicy qos6 6 c6 OUTbound PriorityQue Normal
```

```
ADD !v40 -IPQOS POLicy qos7 7 c7 OUTbound PriorityQue Normal
```

```
ADD !v40 -IPQOS POLicy qos8 8 c8 OUTbound PriorityQue Low
```

6. Specify that the priority queuing for !4 is applied only on outgoing traffic by entering:

```
SETDefault !4 -IPQOS CONTrol = OutFilter
```

7. Specify priority queueing as the queueing policy on the !v4 interface by entering:

```
SETDefault !4 -port qc = PriorityQueues
```

8. Specify that the priority queuing for !v40 is applied only on outgoing traffic by entering:

```
SETDefault !v40 -IPQOS CONTrol = OutFilter
```

9. Rate limit the outbound !v4 interface traffic by entering:

```
ADD !4 -PORT RateLimit pq 100000 0 0 DEFAULT
```

6.1.3

Interpreting Buffer Overrun Errors When IPQoS Priority Queuing Is Enabled

Buffer overrun discards that appear in the port statistics may be caused by network traffic that is dropped because the IPQoS priority queuing rate limit is exceeded.

Procedure:

1. To view the IPQoS priority queue statistics, enter:

```
SHow -Port QueueStatistics
```

2. Add up the Discard and Drop counts for all the priority queues for the port showing buffer overruns and compare this to the buffer overrun error count noted in the

SHow -Port QueueStatistics display in [step 1](#).

If the total of the Discard and Drop counts for the port is about the same as the buffer overrun error count, the buffer overrun error count is reflecting IPQoS priority queuing rate limit discards. Eliminating the rate limit discards (reducing network traffic) also eliminates the buffer overrun errors.

6.1.4

IPQoS Port Bundling

The IPQoS port bundling feature allows you to bundle a virtual port with up to 10 other virtual ports. Bundled ports share a rate limit (CIR) and priority queues. The virtual port on which the bundle is created is called the primary port. For all packets sent out the bundled ports, priority queueing and bandwidth throttling is applied at the primary port, and all bundled ports share the primary port priority queues and bandwidth.

To bundle ports, you can use the POrt service PortBundle parameter:

```
ADD !<port> -Port PortBundle <!port1, !port2 ... !port10>
```

For details, see the description of the PortBundle parameter in the “POrt Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: You must define the virtual ports and define IPQoS policies for them separately before you bundle the ports.

6.2

Configuring Rate Limiting with IPQoS Policy

If rate limiting is specified in a QoS policy, the associated traffic flow is metered and policed to ensure bandwidth consumption does not exceed the defined rate limits, and the conforming and excess traffic is handled per policy. Rate limiting can be applied to traffic transmitted or received on an interface. You can also define actions, such as forward/discard/remark TOS-byte, to handle traffic that conforms to or exceeds the rate limit. When multiple flows are aggregated into a service class, rate limiting protects conforming flows from aggressive flows consuming network resources and potentially leading to a denial of service.

Procedure:

1. Configure IP routing, ports, and IP addresses as necessary for your system.
2. Configure classifier lists by entering:

```
ADD -IP ClassifierList <name>(1-15 char) <priority> (1-9999)[Exclude]
[From (<ipaddr> [/<prefix>(0-32)]) | <addr_list_name>]
[To (<ipaddr> [/<prefix>(0-32)]) | <addr_list_name>]
[TOS <phb> | (<value>(hex0-FF)) [/ <mask>(hex 0-FF)]]]
[(ServiceList <service_list_name>) |
(Protocol <protocol>(0-255))|
(TCP ([Src <compare> <> | < | => <port> | <port1-port2>]
Dst <compare> <> | < | => <port> | <port1-port2>))] |
(UDP ([Src <compare> <> | < | => <port> | <port1-port2>]
Dst <compare> <> | < | => <port> | <port1-port2>))] |
ARCHIE | BGP | ICMP | DNS | FINGER | FTP | GOPHER | GRE | HTTP | NFS | NNTP | NTP |
OSPF | POP | RIP | RPLAYER | RSVP | SECHTTP | SMTP | SNMP | SOCKS | SYSLOG | TELNET
| WAIS | WHOIS
```

Step example:

To create classifier lists that matches packets for the subnets 151.0.0.0 and 152.0.0.0, enter:

```
ADD -IP ClassifierList c1 20 to 151.0.0.0/16
```

```
ADD -IP ClassifierList c2 30 to 152.0.0.0/16
```

To create classifier lists that matches all FTP and HTTP packets, enter:

```
ADD -IP ClassifierList cl_ftp 40 FTP
```

```
ADD -IP ClassifierList cl_http 50 HTTP
```

3. Define IPQoS policies to apply rate limit actions to the classifier lists you defined in step 2 on a specified port by entering:

```
ADD !<port> -IPQOS POLIcy <name> <priority> <classifier_list_name> <direction>
RateLimit <average_rate> <peak_rate> <burst> [ConformAction Transmit | Drop |
SetTOS <value>[/<mask>]] [ExcessAction Transmit | Drop | SetTOS <value>[/<mask>]]
```

If you use the optional ConformAction and ExcessAction attributes, you can specify that one of the following actions be performed on traffic that meets (ConformAction) or exceeds (ExcessAction) the configured RateLimit:

- Transmit: Transmit the packet.
- Drop: Drop the packet.
- SetTOS <value> [/<mask>]: Remark the TOS byte to the specified value and transmit the packet. If you specify the optional mask, the 0-bits in the mask specify bits in the TOS byte that must not be changed and the 1-bits specify those that must take on the corresponding bits in the specified TOS value.



NOTE:

Since neither a ConformAction or an ExceedAction is specified in these policies, the default actions are used: Transmit packets that conform to the specified rate limit and Drop packets that exceed the specified rate limit.

When configuring rate limiting for priority queuing based on IPQoS, only the Transmit and Drop actions are supported.

Step example:

To limit all packets received on port 3 with destination subnet 151.0.0.0 and 152.0.0.0 to an average of 128 kbps, at a peak rate of 256 kbps, and a maximum burst of 3000 bytes, enter:

```
ADD !3 -IPQOS POLIcy ratelimit1 20 c1 INbound RateLimit 128 256 3000
```

```
ADD !3 -IPQOS POLIcy ratelimit2 30 c2 INbound RateLimit 128 256 3000
```

To limit all FTP and HTTP packets transmitted on port 3 to an average of 128 kbps, at a peak rate of 256 kbps, and a maximum burst of 3000 bytes, enter:

```
ADD !3 -IPQOS POLIcy ftp1 40 cl_ftp OUTbound RateLimit 128 256 3000
```

```
ADD !3 -IPQOS POLIcy http1 50 cl_http OUTbound RateLimit 128 256 3000
```

4. Specify the per-port QoS filter by entering:

```
SETDefault !<port> -IPQOS CONTrol = ([InFilter | NoInFilter] ,[OutFilter |
NoOutFilter])
```

Step example:

To enable the inbound quality of service filter for the policies that limit packets received on port 3, enter:

```
SETDefault !3 -IPQOS CONTrol = InFilter
```

To enable the outbound quality of service filter for the policies that limit packets transmitted on port 3, enter:

```
SETDefault !3 -IPQOS CONTrol = OutFilter
```

6.2.1

Configuring Rate Limiting – Example 1

A device with a 100 Mbps port 1 interface is sending too much traffic to ports 3 and 4, with subnets 151.0.0.0 and 152.0.0.0 respectively.

Procedure:

1. Create a classifier list which matches packets for the subnets 151.0.0.0 (port 3) and 152.0.0.0 (port 4):

```
ADD -IP ClassifierList CL1 20 to 151.0.0.0/16
ADD -IP ClassifierList CL2 30 to 152.0.0.0/16
```
2. Define the policy on port 1, applying the rate limit action associated with the classifier list you defined in [step 1](#).

```
ADD !1 -IPQos POLicy POL1 20 CL1 INbound RateLimit 128 256 3000
ADD !1 -IPQos POLicy POL2 30 CL2 INbound RateLimit 128 256 3000
```

All packets with destination subnet 151.0.0.0 and 152.0.0.0 will be limited to an average of 128 kbps, at a peak rate of 256 kbps, and maximum burst of 3000 bytes.
3. Enable the inbound Quality of Service Policy filter with the following command:

```
SETD !1 -IPQos CONTrol=InFilter.
```

6.2.2

Configuring Rate Limiting – Example 2

This example shows rate limiting outbound traffic based on a protocol. In this example, the amount of traffic transmitted on port 3 is limited. The type of traffic which is limited, is defined by the type of protocols. FTP traffic is limited to a maximum rate of 256 kbps, and all HTTP traffic is limited to a maximum rate of 128 Kbps. All FTP and HTTP packets that exceed the defined rates are discarded.

Procedure:

1. Create a classifier list which matches all FTP and HTTP packets.

```
ADD -IP ClassifierList CL_ftp 20 FTP
ADD -IP ClassifierList CL_http 30 HTTP
```
2. Define the OUTbound policy on port 3, applying the rate limit action associated with the classifier list you defined in [step 1](#).

```
ADD !3 -IPQos POLicy POL1 20 CL_ftp OUTbound RateLimit 256 256 4000
ADD !3 -IPQos POLicy POL2 30 CL_http OUTbound RateLimit 128 128 4000
```

All FTP and HTTP outbound packets on port 3 are limited to 256 and 128 kbps respectively. All packets that exceeds the specified rate are discarded.
3. Enable the outbound Quality of Service Policy filter with the following command:

```
SETD !3 -IPQos CONTrol=OutFilter
```

6.3

Class-Based Queueing

EOS supports class-based CBQ on Frame Relay parent ports, Frame Relay virtual ports, and PPP parent ports. CBQ is a link-sharing packet scheduler that combines priority queuing with weighted round-robin scheduling.

When you configure CBQ, you partition link bandwidth by hierarchically structured classes. Each class is regulated to its allocated bandwidth, and you can optionally configure a child class to borrow bandwidth from its parent class with excess bandwidth. Packet scheduling is priority based, with traffic from higher priority classes transmitted ahead of traffic from lower priority classes. For classes with the same priority, traffic is weighted round-robin between classes.



NOTE: If CBQ is configured on both a Frame Relay parent port and a Frame Relay virtual port, then only the classes defined for the virtual port are used.

The EOS implementation of CBQ also supports optional random early detection (RED) and explicit congestion notification (ECN).

Random early detection (RED)

Manages CBQ class queue size by dropping packets by probability based on the estimated average queue size. The drop probability increases linearly as the average queue size grows. Packet drops begin when the average queue size exceeds the system-defined minimum threshold and all arriving packets are dropped at the maximum drop probability when the average queue size reaches the system-defined maximum threshold. All incoming packets are dropped when the absolute class queue limit is reached. The minimum and maximum thresholds are based on the class bandwidth and the packet drop precedence value as marked by the CBQ rate limit (if rate limiting is enabled for the virtual port traffic). Packets marked green (which are at or below the average rate) have a higher threshold than yellow packets (which exceed the average rate but are at or below the peak rate), which in turn have a higher threshold than the RED packets (which exceed the peak rate). You enable RED for a class-based queue by including the RED value with the PORT service `ClassBasedQueue` parameter. For details, see the description of the `ClassBasedQueue` parameter in the “PORT Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: RED is most effective on TCP flows, which decrease transmission rate when packet loss is detected.

Explicit congestion notification (ECN)

Can be used as an active queue management mechanism to prevent links from becoming congested. When this feature is enabled, some classes of traffic are marked with IP ECN-capable transport (ECT) bits to indicate to the devices that these packets are eligible for ECN settings. If a device with ECN for CBQ enabled discovers that a link is congested, or about to become congested, it marks a certain percentage of these packets as congestion-experienced packets. The receiving device can then take measures to reduce traffic on the link.

ECN marking is based on the current bandwidth utilization. At low bandwidth (less than the configured low water mark), no congestion marking occurs. When a certain bandwidth is reached (the low water mark), every third IP packet is ECN marked. When a higher bandwidth is reached (the high water mark), all packets are ECN marked. You configure the high and low water marks with the PORT service `ClassBasedQueue` parameter ECN option. For details, see the description of the `ClassBasedQueue` parameter in the “PORT Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. Each class in class-based queueing has its own ECN setting. The ECN bit settings are checked and set on an egress port depending on the average bandwidth utilization for that class.

By allowing devices to use the ECN bit setting in a packet header as an indication of congestion, instead of relying solely on packet drops, the ECN for CBQ feature can potentially reduce or eliminate packet loss on loss-sensitive data streams.



NOTE:

Motorola Solutions recommends that you do not use CRTP on ports where ECN for CBQ is enabled and vice versa. Due to the way ECN is implemented in EOS, ECN settings nullify the advantages of CRTP. For details about CRTP, see [CRTP Twice Algorithm on page 331](#).

In addition to CBQ classes and the optional RED and ECN features, you can configure CBQ rate limiting on PPTP/L2TP/IPIP and Frame Relay virtual ports. You do this by using the PORT service `RateLimit` parameter to define a token bucket rate limiter for either the port traffic or for one or more of the CBQ classes configured on the associated egress physical port. For details, see the description of the `RateLimit` parameter in the


“Port Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. Packets in excess of the token bucket peak rate can either be discarded or transmitted downstream to the physical port with a high drop precedence. If RED is enabled on the target CBQ class, the excess packets with the higher drop precedence values are preferentially discarded to free up queue slots for packets that conform to the specified token bucket rate.

6.3.1 CBQ Classes

You can define CBQ classes by using the POrt service ClassBasedQue parameter:

```
ADD !<port> -POrt ClassBasedQue <class_name> <parent_name> <bandwidth> [RED] [Borrow]
[AvgPacketSize <bytes>] [Priority <value>] [DEfault] [ECN <LowWaterMark> <HighWaterMark>]
```

Figure 65: Sample CBQ Class Hierarchy



To define CBQ classes on port 3 that correspond to the hierarchy illustrated above, you must enter:

```
ADD !3 -POrt ClassBasedQue ROOT NULL 1024
ADD !3 -POrt ClassBasedQue Voice ROOT 128 Priority 6
ADD !3 -POrt ClassBasedQue RSVP ROOT 128 Priority 4
ADD !3 -POrt ClassBasedQue Premium ROOT 384 Priority 5
ADD !3 -POrt ClassBasedQue DEFAULT ROOT 384 RED DEfault
ADD !3 -POrt ClassBasedQue BusApp Premium 256 RED Priority 6
ADD !3 -POrt ClassBasedQue Video Premium 128 Priority 4
```

6.3.2 Configuring IPQoS Based on CBQ

You can define IPQoS policies that assign CBQ classes to IP packets at the outbound PPP, Frame Relay, or tunnel port.

 **NOTE:** If you do not specify a CBQ class, the Differentiated Services Code Point (DSCP) value in the TOS byte determine the CBQ class assigned to the outbound packet based on the configured TOS-to-CBQ class mapping. For details, see [Configuring TOS-to-IPIP Mapping on page 210](#). An unrecognized DSCP value maps to the default CBQ class.

Procedure:

1. Configure IP routing, ports, and IP addresses as necessary for your system.

2. Configure a classifier list by entering:

```
ADD -IP ClassifierList <name>(1-15 char) <priority> (1-9999)[Exclude]
[From (<ipaddr> [/<prefix>(0-32)]) | <addr_list_name>]
[To (<ipaddr> [/<prefix>(0-32)]) | <addr_list_name>]
[TOS <phb> | (<value>(hex0-FF)) [/ <mask>(hex 0-FF)]]]
[(ServiceList <service_list_name>) |
(Protocol <protocol>(0-255)) |
(TCP ([Src <compare> <> | < | => <port> | <port1-port2>]
Dst <compare> <> | < | => <port> | <port1-port2>))] |
(UDP ([Src <compare> <> | < | => <port> | <port1-port2>]
Dst <compare> <> | < | => <port> | <port1-port2>))] |
ARCHIE | BGP | ICMP | DNS | FINGER | FTP | GOPHER | GRE | HTTP | NFS | NNTP | NTP |
OSPF | POP | RIP | RPLAYER | RSVP | SECHTTP | SMTP | SNMP | SOCKS | SYSLOG | TELNET
| WAIS | WHOIS
```

Step example:

To create a classifier list that matches all packets for the default gateway, enter:

```
ADD -IP ClassifierList CL_any 10 to 0.0.0.0/0
```

3. Define CBQ classes for the parent port by entering:

```
ADD !<port> -Port ClassBasedQueue <class_name> <parent_name> <bandwidth> [RED]
[Borrow] [AvgPacketSize <bytes>] [Priority <value>] [Default] [ECN <LowWaterMark>
<HighWaterMark>]
```

Step example:

To create a CBQ class named "CBQ1" under the "ROOT" class on port 3, enter:

```
ADD !3 -Port ClassBasedQueue ROOT NULL 1000
ADD !3 -Port ClassBasedQueue CBQ1 ROOT 1000 Default
```

As configured above, CBQ class CBQ1 is allocated 1 Mbps of bandwidth and is designated default class for the egress physical port.

4. Define IPQoS policies at the egress ports that assign CBQ classes to the packets that match the specified classifier list by entering:

```
ADD !<port> -IPQoS POLicy <name> <priority> <classifier_list_name> <direction>
(ClassBasedQueue <class_name>)
```



NOTE: If the <class_name> does not exist, the default CBQ class on the port is assigned to the port instead.

Step example:

If you have IPIP tunnels (virtual ports) bound to parent port 3, to assign outgoing traffic that matches classifier list CL_any to CBQ class CBQ1, at the virtual port ports, enter:

```
ADD !v1-!v3 -IPQoS POLicy CBQPol_1 10 CL_any OutBound ClassBasedQueue CBQ1
```

5. Enable the outbound QoS filter on the virtual ports by entering:

```
SETDefault !<port> -IPQoS CONTrol = OutFilter
```

Step example:

To enable the outbound QoS filter on virtual ports 1-3, enter:

```
SETDefault !v1-!v3 -IPQoS CONTrol = OutFilter
```

6. Configure the queuing policy for the parent port to be CBQ by entering:


```
SETDefault !<port> -Port QueueControl = ClassBasedQue
```

Step example:

To configure the queuing policy for port 3 to be CBQ, enter:

```
SETDefault !3 -Port QueueControl = ClassBasedQue
```

7. Rate limit the outbound virtual port traffic by entering:

```
ADD !<port> -PORT RateLimit <name> <CIR> <PIR> <CBS> [TransmitExcess] [Default]
```

Step example:

To define a CIR (average rate), PIR (peak rate), and CBS (burst size) of 64 kbps, 128 kbps, and 2000 bytes respectively on the virtual ports associated with CBQ class CBQ1, and to specify that CBQ1 is the default class defined for the egress physical port, enter:

```
ADD !v1-!v3 -Port RateLimit CBQ1 64 128 2000 Default
```

8. Enable the CBQ rate limit defined for the CBQ classes on the specified port by entering:

```
SETDefault !<port> -Port QueueControl = CBQRateLimit
```

Step example:

To enable the CBQ rate limit on virtual ports 1-3, enter:

```
SETDefault !v1-!v3 -Port QueueControl = CBQRateLimit
```

6.3.3

Configuring IP/TOS to CBQ Class Mapping

Procedure:

1. Configure IP routing, virtual ports, and IP addresses as necessary for your system.
2. Enable IP/TOS-to-CBQ class mapping on the desired egress port by entering:

```
SETDefault !<port> -IPQOS Control = TosMap
```

Step example:

To enable IP/TOS-to-CBQ class mapping on port 3, enter:

```
SETDefault !3 -IPQOS Control = TosMap
```

3. Define one or more TOS-byte values that can be mapped into a CBQ class that could realize the behaviors specified by the TOS value by entering:

```
ADD !<port> -IPQOS TosMap <class_name> Default | ([Mask <mask>] <tos>... <tos_n>)
```

where:

<class_name> specifies a CBQ class name (defined with the Port service ClassBasedQue parameter)

<tos>...<tos_n> specifies the TOS value(s) in the range 0-255

Default specifies the default CBQ class name for all TOS values with no defined class mapping

If you specify the optional <mask>, 0s in the 8-bit mask (in hexadecimal) specify bits in the TOS byte that are to be ignored in the matching operation. The mask is ANDed with the TOS byte in the IP header of the outbound packet and the result is compared against <tos>. The default mask is 0xFc, which corresponds to the Differentiated Services Code Point (DSCP) in the upper 6 bits of the IP/TOS byte.

Step example:

To map TOS values 10, 20, and 30 to CBQ class CBQ_1 for outgoing IP traffic on port 3, enter:

```
ADD !3 -IPQOS TosMap CBQ_1 10 20 30
```

6.3.4

Class-Based Queueing Statistics View

You can use the `QueueStatistics` parameter to display CBQ statistics:

```
SHoW [!<port>] -Port QueueStatistics
```



NOTE: If RED is enabled for a CBQ class, the drop statistics are also included.

The `SHoW -Port QueueStatistics` command displays the following statistics for each defined class, as well as the “ROOT” class and the “DEFAULT” class (if a “DEFAULT” class was designated):

- `XmitPackets` – The number of packets transmitted.
- `XmitBytes` – The number of bytes transmitted.
- `DropPkts` – The number of packets dropped (tail drop or by RED).
- `Borrows` – The number of (bandwidth) borrow operations executed.
- `Delays` – The number of timer delays initiated due to over-limit conditions.

In addition, the display lists the following ECN statistics for each defined class:

- `LWM` – The ECN low water mark, in kbps, specified for this class with the `ClassBasedQue` parameter `ECN` option. This is the bandwidth at or above which 33% of the IP packets (every third packet) is marked for ECN. If ECN control is disabled for this class, the `LWM` is listed as 0.
- `OverLWM` – The number of packets received when the bandwidth utilization for the class was over the low water mark.
- `HWM` – The ECN high water mark, in kbps, specified for this class with the `ClassBasedQue` parameter `ECN` option. This is the bandwidth at or above which all IP packets are marked for ECN. If ECN control is disabled for this class, the `HWM` is listed as 0.
- `OverHWM` – The number of packets received when the bandwidth utilization for the class was over the high water mark.
- `ECT Pkts` – The number of ECT packets seen by this device in this class on this egress port. The ECT packets can be ECT(0) or ECT(1), as defined in RFC 3168.
- `ECN Marked` – The number of packets which are marked with ECN bit (11) by this device. This indicates that the link is experiencing congestion. The actual percentage of packets marked with ECN indicates how much congestion.

6.4

Configuring IPQoS Based on Protocol Reservation

Protocol reservation enables you to reserve bandwidth for lower bandwidth usage, interactive, response-time-sensitive, or transaction-oriented network applications, which are normally crowded out by heavy bandwidth usage applications such as file transfer or mail. IPQoS based on protocol reservation allocates a specified percentage of bandwidth for designated packet types (identified by their protocol reservation tags) transmitting through a designated WAN port and meeting specified conditions.

You set the tag to identify the packet types and the percentage of reserved bandwidth using the `Port` service `PRoTocolRsrv` parameter. The procedures to configure protocol reservation can vary with different packet types. For a complete description of the procedures to configure a WAN port for protocol reservation for

all packet types, see the “Configuring Protocol Reservation” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: For Frame Relay parent ports with the QueueControl parameter set to ProtocolRsrv, the parent port reservation database configured by the ProtocolRsrv parameter is inherited by a Frame Relay virtual circuit if the QueueControl parameter is set to Default for that virtual circuit port.

Procedure:

1. Configure IP routing, ports, and IP addresses as necessary for your system.
2. Configure protocol reservation.

You can use different procedures to configure protocol reservation depending on the protocol type. For detailed instructions, see [Protocol Reservation Configuration on page 383](#).

Step example:

To add IP filters that designate a tag name of TELNETTAG to identify Telnet packets and a tag name of FTPTAG to identify FTP packets, enter:

```
ADD -IP FilterAddrs ALL ALL PROTOcolRsrv = TELNETTAG Telnet
ADD -IP FilterAddrs ALL ALL PROTOcolRsrv = FTPTAG FTP
```

To add an IP filter default action that forwards any packets that do not satisfy the filter requirements, enter:

```
SETDefault -IP FilterDefAction = Forward
```

To enable the IP filtering feature, enter:

```
SETDefault -IP CONTROL = Filtering
```

3. Assign protocol reservation tags by entering:

```
ADD !<port> -Port PROTOcolRsrv <name_tag> <percentage share>
```

where:

<name_tag> specifies the name used by the system to identify the packets to which to allocate the configured bandwidth amount

<percentage share> specifies the bandwidth amount that is to be allocated in reserve to the packets with the name tag that matches the specified <name_tag>

Step example:

To assign 20 percent of bandwidth to the PROTOcolRsrv name tag TELNETTAG and 30 percent of the bandwidth to the PROTOcolRsrv name tag FTPTAG, enter:

```
ADD !4 -PORT PROTOcolRsrv TELNETTAG 20
ADD !4 -PORT PROTOcolRsrv FTPTAG 30
```

4. Configure classifier lists that match specified packet types by entering:

```
ADD -IP ClassifierList <name> <priority> (ARCHIE | BGP | ICMP | DNS | FINGER | FTP
| GOPHER | GRE | HTTP | NFS | NNTP | NTP | OSPF | POP | RIP | RPLAYER | RSVP |
SECHTTP | SMTP | SNMP | SOCKS | SYSLOG | TELNET | WAIS | WHOIS)
```

Step example:

To define classifier lists named <c1_telnet> and <c1_ftp> that match Telnet and FTP packets, enter:

```
ADD -IP ClassifierList c1_telnet 20 TELNET
ADD -IP ClassifierList c1_ftp 30 FTP
```

5. Specify that queueing for the specified WAN port is set for protocol reservation by entering:

```
SETDefault !<port> -Port QueueControl = PROTOcolRsrv
```

Step example:

```
SETDefault !4 -Port QueueControl = ProtocolRsrv
```

6. Specify that the priority queueing is applied only on outgoing traffic at the specified WAN port by entering:

```
SETDefault !<port> -IPQoS Control = OutFilter
```

Step example:

```
SETDefault !4 -IPQoS Control = OutFilter
```

7. Define IPQoS policies to assign specified protocol reservation tags to outgoing packets at the specified port that meet the criteria defined in the specified classifier list(s) by entering:

```
ADD !<port> -IPQoS Policy <name> <priority> <classifier_list_name> <direction>  
ProtocolRsrv <name_tag>
```

Step example:

To define an IPQoS policy named `prot_res1` that assigns the `TELNETTAG` protocol reservation tag to outgoing packets at WAN port 4 that meet the criteria specified in classifier list `cl_telnet` and an IPQoS policy named `prot_res2` that assigns the `FTP` protocol reservation tag to outgoing packets at WAN port 4 that meet the criteria specified in classifier list `cl_ftp`, enter:

```
ADD !4 -IPQoS Policy prot_res1 20 cl_telnet Outbound ProtocolRsrv TELNETTAG  
ADD !4 -IPQoS Policy prot_res2 30 cl_ftp Outbound ProtocolRsrv FTP
```

6.5

IEEE 802.1P Prioritization

You can configure IEEE 802.1P prioritization for either incoming or outgoing IP traffic on a specified port connected to a VLAN network.

- If you enable 802.1P-to-TOS mapping on the ingress VLAN port, the 802.1P priority in 802.1Q tag of the incoming IP packet determines the IP/TOS value based on a default or user-configured mapping.
- If you enable TOS-to-802.1P mapping on the egress VLAN port, the IP/TOS value determines the 802.1P priority of the outgoing packet based on a default or user-configured mapping. You can also use a QoS policy to tag outgoing IP traffic with a specific 802.1P priority.

6.5.1

Configuring 802.1P-to-TOS Mapping on an Ingress VLAN Port

Procedure:

1. Configure IP routing, virtual ports, and IP addresses as necessary for your system.
2. Enable 802.1P-to-TOS mapping on the ingress VLAN port by entering:

```
SETDefault !<port> -IPQoS Control = InVlanMap
```

Step example:

To enable 802.1P-to-TOS mapping on VLAN port `v101`, enter:

```
SETDefault !v101 -IPQoS Control = InVlanMap
```

3. If needed, specify the 802.1P-to-TOS mapping by entering:

```
SETDefault !<port> -IPQoS InVlanMap = (<tos_0>...<tos_n>)
```

where `<tos_0>...<tos_n>` specifies the TOS-byte value (0-0xff) that corresponds to the 802.1P priority value `n` (0-7). You must define eight instances of TOS-byte value in hexadecimal in ascending order of the 802.1P priority values.

Step example:

To map 802.1P priority values 0, 1, 2, 3, 4, 5, 6, and 7 to TOS-byte values 00, 20, 40, 40, 40, a0, a0, and e0 for incoming IP traffic on VLAN port 101, enter

```
SETDefault !v101 -IPQOS InVlanMap = 00 20 40 40 40 a0 a0 e0
```

If you do not define a mapping with the `InVlanMap` parameter, the default mapping is used: the 3-bit 802.1P value is mapped directly into the 3-bit IP precedence field in the TOS byte, with the remaining TOS bits set to zero.

6.5.2

Configuring TOS-to-802.1P Mapping on an Egress VLAN Port

Procedure:

1. Configure IP routing, virtual ports, and IP addresses as necessary for your system.
2. Enable TOS-to-802.1P mapping on the egress VLAN port by entering:

```
SETDefault !<port> -IPQOS CONTROL = OutVlanMap
```

Step example:

To enable TOS-to-802.1P mapping on VLAN port v102, enter:

```
SETDefault !v102 -IPQOS CONTROL = OutVlanMap
```

3. Specify the TOS-to-802.1P mapping by performing one of the following actions:

- To specify the mapping by using the `OutVlanMap` parameter, enter:

```
ADD !<port> -IPQOS OutVlanMap <vlan_priority> Default | ([Mask <mask>] <tos> ..  
<tos_n> )
```

where:

`<vlan_priority>` specifies the 802.1P priority value (0-7) for one or more of the specified hexadecimal TOS-byte values (`<tos>...<tos_n>`)

`Default` specifies the default 802.1P priority for all TOS-byte values without a specified mapping

If you specify the optional `<mask>`, 0s in the 8-bit mask (in hexadecimal) specify bits in the TOS byte that are to be ignored in the matching operation. The mask is ANDed with the TOS byte in the IP header of the outbound packet and the result is compared against `<tos>`. The default mask is 0xFc, which corresponds to the Differentiated Services Code Point (DSCP) in the upper 6 bits of the IP/TOS byte.

Example:

To map TOS-byte values 20, 28, and 2a to 802.1P priority value 1, and to specify 5 as the the default 802.1P priority for all TOS-byte values without a specified mapping for outgoing IP traffic on VLAN port 102, enter:

```
SETDefault !v102 -IPQOS OutVlanMap = 1 20 29 2a
```

```
SETDefault !v102 -IPQOS OutVlanMap = 5 Default
```

If you do not specify a `Default` priority, the following mapping is used for TOS-byte values without a specified mapping: the IP precedence (bits 0-2) value in the TOS byte is used as the 802.1P priority value.

- To specify the mapping by using the `POLicy` parameter, enter:

```
ADD !<port> -IPQOS POLicy <name> <priority> <classifier_list_name> OUTbound  
VlanPriority <802.1P (0-7)>]
```

Example:

To create a policy named `VLAN_pol` that assigns 802.1P priority value 6 to outbound IP traffic on VLAN port 102 that matches the criteria specified in the classifier list named `<CL_2>`, enter:

```
ADD !v102 -IPQOS POLICY VLAN_pol 1 CL_2 OUTbound VlanPriority 6
```

6.6

Configuring TOS-to-IPIP Mapping

TOS-to-IPIP mapping allows you to map the well-known TOS precedence value (0-7) from the inner IP header of the IPIP packet to a specified hexadecimal TOS/Diffserv byte value in the outer IP header for outbound IP traffic.

Procedure:

1. Configure IP routing, virtual ports, and IP addresses as necessary for your system.
2. Enable TOS-to-IPIP mapping on the desired egress port by entering:

```
SETDefault !<port> -IPQOS CONTROL = OutIPIPtoTosMap
```

Step example:

To enable TOS-to-IPIP mapping on virtual port 101, enter:

```
SETDefault !v101 -IPQOS CONTROL = TosMap
```

3. Map the well-known TOS precedence value (0-7) from the inner IP header of the IPIP packet to a specified hexadecimal TOS/Diffserv byte value in the outer IP header for outbound IP traffic by entering:

```
SETDefault !<port> -IPQOS OutIPIPtoTosMap = Default | (<IpIpMap_0> ... <IpIpMap_7>)
```

where:

`<IpIpMap_n>` specifies the Diffserv byte value (in hexadecimal format) for TOS value `n`

`Default` specifies the default DiffServ/TOS-byte values for all TOS values without a specified mapping

Step example:

To map a TOS precedence value of 3 to a TOS/DiffServ byte value of 0x60 on virtual port 101, enter:

```
SETDefault !v101 -IPQOS OutIPIPtoTosMap = 00 20 40 60 80 A0 C0 E0
```

Each value after the “=” represents the precedence in order from 0 through 7, so the fourth value (60) maps to TOS precedence value 3.

6.7

Defining a Default IPQoS Policy

You can create a default IPQoS policy to be used as a “catch all” policy if traffic does not match the other policies you defined. To do this, you must use the `Exclude` value of the IP service `ClassifierList` parameter to create a profile to match the “left over” traffic.

For example, assume that you defined policies for traffic destined to IP addresses 10. 125. 126. 147 and 10. 127. 126. 117. Your “left over” traffic is all traffic *other* than traffic destined to these IP addresses. You can use the `Exclude` value to define a “default” policy for this “left over” traffic.



IMPORTANT: When you use the `Exclude` value of the `ClassifierList` parameter to define a default policy that applies to all traffic *except* the traffic that matches the policies you already defined, you must include all IP addresses in the “default” policy, as well as the excluded addresses.

Procedure:

1. Create an address list that can be used to reference the IP addresses you want to exclude by entering:

```
ADD -IP AddressList <name>(1-15 characters) <IP_address>
```

Step example:

To create an address list called “default_al” that references IP addresses 10. 125. 126. 147 and 10. 127. 126. 117, enter:

```
ADD -IP AddressList default_al 10.125.126.147
ADD -IP AddressList default_al 10.127.126.117
```

2. Create a classifier list to exclude IP traffic that is destined for the address list you created in [step 1](#) by entering:

```
ADD -IP ClassifierList <name>(1-15 char) <priority> (1-9999) [Exclude] To
<ipaddr_list_name>
```

Step example:

To create a classifier list called “default_cl” to exclude packets destined for the address list you created in [step 1](#) (default_al), enter:

```
ADD -IP ClassifierList default_cl 10 Exclude To default_al
```

3. Add all IP addresses to the classifier list with a lower priority (lower priority = higher number) than the excluded addresses by entering:

```
ADD -IP ClassifierList <name>(1-15 char) <priority> (1-9999) To 0.0.0.0/0
```

Step example:

For example, if you added the excluded addresses with priority 10, you could add all IP addresses to the default policy with priority 20 by entering:

```
ADD -IP ClassifierList default_cl 20 To 0.0.0.0/0
```

4. Create a default IPQoS policy for the classifier list you created in [step 3](#) by entering:

```
ADD !<port> -IPQOS POLicy <policy_name> <priority> <cl_name> <direction>
[ RateLimit <avg rate> <peak rate> <burst>]
[ ExcessAction Transmit | Drop | SetTOS <tos>[/<mask>] ]
[ ConformAction Transmit | Drop | SetTOS <tos>[/<mask>]]
[ (ClassBasedQueue <class_name>) | ProtocolRsrv <name_tag>) |
(PriorityQueue Urgent | High | Medium | Low )]
[ TOSQueue <priority_queue(0-7)> ]
[ VlanPriority <802.1P>]
[ NextHop (<ipaddr> [<ipaddr> [<ipaddr>]]) | (<port>
[!<port> [!<port>]])]
```

Step example:

To define the “default” policy on port 3, applying a rate limit action to the classifier list you defined in [step 3](#), enter:

```
ADD !3 -IPQOS POLicy default_pol 10 default_cl INbound RateLimit 128 256 3000
```

All packets in classifier list default_cl other than those you specifically excluded (defined as packets with destination IP address 10. 125. 126. 147 or 10. 127. 126. 117) are limited to an average of 128 kbps, at a peak rate of 256 kbps, and maximum burst of 3000 bytes.

Chapter 7

System IP (SYSIP) Configuration

The System IP is a software-only interface that emulates an IP interface. It is always active (up) and allows IP operations on the device, as long as at least one interface on the device is active.



NOTE: Only one System IP is allowed per device, and only IP-based applications are supported by this feature.

A System IP is similar to a regular IP interface except that it is not attached to any physical port, and it is always up as long as at least one interface on the device is active. By configuring a System IP in your EOS software, you do not need to know the specific address of the active interface to be able to Telnet to the device; you need only remember the System IP.

After you create a System IP, you can perform most IP-based operations using the IP address you configured for the System IP. The following applications benefit from the use of a System IP:

- IP applications, such as Telnet, SNMP, and FTP
- Tunnel endpoints, such as GRE
- OSPF router ID

7.1

System IP Configuration

To configure the System IP, you can use:

```
SETDefault -IP SystemIP = <IP_address>
```

For example, to create a System IP with address 1.1.1.1, you can use:

```
SETDefault -IP SystemIP = 1.1.1.1
```

For the System IP to be valid (“up”), at least one interface on the device must be active. The chosen IP address can be:

- One of the IP addresses of an active IP interface.
- Part of an active IP interface’s subnet.
- Different from all active IP interfaces.



IMPORTANT: If you change the system IP address on an EOS device, the SNMPv3 credentials must be re-established with that device. Therefore if SNMPv3 users were configured on the device prior to the system IP address change, you must issue the `ResetV3` command to reset the SNMPv3 data, then reconfigure the SNMPv3 users with the appropriate privilege levels. For details about how to configure SNMPv3 users, see [Predefined SNMPv3 USM Users Management on EOS Devices on page 670](#) and [General SNMPv3 USM Users Management on EOS Devices on page 679](#).

When you create the System IP address, an entry is added to the route table. The route table must contain this entry, so that a ping command can be issued from the device to the System IP. The following table lists the information contained in the entry added to the route table.

Table 19: Information Added to the Route Table by System IP

Value	Description
Gateway	Any active (up) IP interface
Metric	1 (the “owner” is System IP)

Value	Description
NB2	Added to the ARP table to allow the device to respond to ARP requests on the System IP. The MAC address in the ARP response is the active IP interface that has a network number that matches the network number of the System IP. This is employed only if the System IP has the same network number as one of the active interfaces.

The following are examples of RIP and ARP table entries based on System IP address usage.

When the System IP is set to 4.4.4.4 and there is no routing entry for network 4.0.0.0, the following entry is automatically added to the routing table:

```
4.4.4.4 255.255.255.255 3.3.3.52 1 Up 0 SystemIP
```

When the System IP is set to 3.3.3.3 and there is an interface which already has 3.3.3.52 configured, the following new entry in the ARP table is made:

```
3.3.3.52 2F Local %08000205734A Ethernet Static
```

```
3.3.3.3 2F Local %08000205734A Ethernet SystemIP
```

To delete the System IP, you must use:

```
SETDefault -IP SystemIP = 0.0.0.0
```

When you issue this command, the System IP's route table entry and ARP table entry are removed. The routing advertisements that System IP had sent out are flushed automatically by the neighboring devices.



NOTE: You can also use the `SETDefault -IP SystemIP` command to change the IP address of an existing System IP.

7.2

Routing Issues

The EOS routing engine takes care of outgoing packets that have the System IP as their source IP address. For incoming packets that have the System IP as their destination IP address, the System IP's route needs to be advertised to neighboring networks.

To advertise to neighbors, routing packets must be sent out periodically. There are several methods to accomplish this:

- If the System IP address is the same as the active IP interface or has the same subnet as the active IP interface:

The necessary action depends on whether the port is up or down. The following table summarizes the necessary actions for all routing protocols supported.

Table 20: System IP with Same Subnet as an IP Address of an Active IP Interface

Value	Port "UP" (Active)	Port "DOWN" (Inactive)
OSPF	If OSPF is enabled: do nothing. If OSPF is disabled: ○ If DIR NET, do nothing. ○ If not DIR NET, send Host Route advertisement to other interfaces.	Send Host Route advertisement to other interfaces.

Value	Port "UP" (Active)	Port "DOWN" (Inactive)
RIP/RIPv2	RIP Enabled: do nothing. RIP Disabled: Send subnet route advertisement to other interfaces.	Send Host Route advertisement to other interfaces.

- If the System IP address is different than all active IP interfaces:
Send out a route stub (host) advertisement for all IP interfaces with the gateway IP address of the IP interface's address. The following table summarizes the necessary actions for all routing protocols supported.

Table 21: System IP Different from All Active IP Interface IP Addresses

Value	Description
OSPF	Send Host route advertisement to all interfaces.
RIP/RIPv2	Send Host route advertisement to all interfaces.

7.3

ARP Packets

ARP packets for the System IP MAC address are generated only if the System IP address has the same subnet as one of the IP interfaces. In this case, the MAC address of that interface is used.

7.4

NAT/Firewall Issues

Currently, all NAT/Firewall policies are on a per-port basis, which creates a lack of NAT/Firewall issues on the System IP. All NAT/Firewall policies are applied at the port level.

7.5

VRRP Issues

The System IP address of a device should not be used as a VRRP Virtual Router IP (VIP). Doing so may cause two devices to respond to the same IP address. When sending a request (such as SNMP) to the VIP, the response can be returned from either of the devices.

Chapter 8

Resource Reservation Protocol Configuration

Resource Reservation Protocol (RSVP) provides network consistency for realtime traffic by permitting applications to request Quality of Service assurances from the network. Without this network consistency, realtime traffic can experience information loss, jitter, loss of synchronization, and not enough bandwidth.

Data applications need a relatively small amount of bandwidth, but multimedia applications demand high bandwidth. RSVP allows multicast applications such as video conferencing, multimedia, and virtual private network (VPN) network management to reserve resources (bandwidth) for consistent data delivery.

RSVP features three participants: the sender(s), the network, and the receiver(s). Each sender application periodically sends an RSVP Path message to a receiver for each data flow it originates. Provided in the Path message are the characteristics of the data traffic the application expects to generate: the data rate (bandwidth), the queue size, and the maximum packet size (MTU).

The Path message travels from a sender to receiver(s) along the same route(s) used by data packets. A device in the network that does not implement RSVP routes the Path message through as if it were a data packet. An RSVP-capable device processes the information in the Path message and uses it later in the reservation request message(s) sent back in the reverse direction to the sender.

The receiver application initiates reservation requests based on information it receives from the Path message. Each device that receives a reservation request message (Resv) reserves the requested bandwidth, if there is sufficient bandwidth, and sends the Resv message to its previous hop which is the next device in the route toward the sender. If there is no sufficient bandwidth, the Resv message goes no further and an error message is sent back to the receiver application. Another requirement for a successful reservation request is that, on a per flow basis, bandwidth greater than the user-configurable `MaxFlowRate` parameter cannot be requested.

8.1

Configuring the RSVP

Perform this procedure to configure a sample Resource Reservation Protocol (RSVP) configuration.

For information on the commands and parameters used to configure RSVP, see "RSVP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Procedure:

1. Configure and enable IP routing on LAN port 1 and Frame Relay virtual port V1 by entering the following commands:


```
ADD !1 -IP NETaddr = <ipaddr1>
ADD !V1 -IP NETaddr = <ipaddr2>
SETDefault -IP CONTrol = ROUTe
```
2. Configure and enable IP routing protocol OSPF for unicast routing.
3. Configure and enable IP multicast routing protocol DVMRP or MOSPF by entering the following commands:


```
SETDefault -MIP CONTrol = Enable
ADD !V1 -DVMRP Neighbor @dlci
```

```
SETDefault !1 -DVMPR CONTROL = Enable
```

```
SETDefault !V1 -DVMPR CONTROL = Enable
```

4. Configure the committed information rate of the Frame Relay virtual port by entering:

```
SETDefault !V1 -FR CIR = <vcid>
```

5. Configure protocol reservation on the Frame Relay port by entering the following commands:

```
ADD !V1 -PORT PROTOcolRsrv RSVP 60 # reserve 60% port bandwidth for RSVP #
```

```
SETDefault !V1 -PORT QueueCONTROL = PROTOcolRsrv
```

This sample configuration requests 60% of the available bandwidth for RSVP.

6. Configure and enable RSVP by entering the following commands:

```
SETDefault -RSVP CONTROL = ENable
```

```
SETDefault !V1 -RSVP MaxFlowRate = 100 # limit per flow bandwidth to 100 bytes/sec  
#
```

8.2

RSVP Proxy Sender and Receiver


RSVP allows host applications to make bandwidth reservations in devices along the data path from senders to receivers. However, for “dumb” devices, such as IP telephone handsets, which are not connected to a PC or other host devices, it is not possible to reserve bandwidth by using RSVP. The RSVP Proxy sender and receiver feature solves this problem by emulating RSVP senders or receivers on behalf of these devices.

8.2.1

Proxy Sender: Unicast Destination and One Sender Port

The following figure illustrates a sample configuration featuring an RSVP Proxy sender with a unicast destination and one sender port.

Figure 66: RSVP Proxy Sender with a Unicast Destination and One Sender Port



To configure the example shown above, you must use:

```
ADD -RSVP ProxySENDER EX1_PS SESSion 192.0.0.1/2500 UDP SENDER 10.0.0.1/3000 RATE 2000  
1000 TimeOut 40
```

The ADD command emulates a proxy sender for the non-RSVP host whose IP address is 10.0.0.1.

- Data transmitted by the non-RSVP host on UDP port number 3000 and destined for unicast address 192.0.0.1 port 2500 initiate an RSVP session.
- The RATE parameter specifies the characteristics of the data traffic expected to be generated by the sender host: 000 bytes per second with a burst size of 1000 bytes.
- The TimeOut value (40 seconds) specifies that the RSVP session is torn down if the sender stops transmitting data for that length of time. The TimeOut parameter is optional; if you do not specify a value,


the `TimeOut` parameter is set to 300 seconds by default. To disable the timer, set the `TimeOut` parameter to 0.

8.2.2

Proxy Receiver: Unicast Destination and One Sender Port

The following figure illustrates a sample configuration featuring an RSVP Proxy receiver with a unicast destination and one sender port.

Figure 67: RSVP Proxy Receiver with a Unicast Destination and One Sender



To configure the example shown above, you must use:

```
ADD -RSVP ProxyREceiver EX1_PR SESSION 192.0.0.1/2500 UDP SENDER 10.0.0.1/3000 RATE 2000  
1000 STYLE FixedFilter
```

The `ADD` command emulates a RSVP receiver for non-RSVP host, 192.0.0.1.

- A RSVP path message received for session 192.0.0.1 port 2500 with UDP protocol ID causes an RSVP reservation request message to be sent by the device on behalf of the non-RSVP host.
- The `RATE` parameter specifies that the request is for a bandwidth of 2000 bytes per second and burst size of 1000 bytes for data transmitted by sender 10.0.0.1 on port 3000.
- The requested bandwidth is for the exclusive use of data sent by 10.0.0.1 on port 3000, as specified by the `FixedFilter` reservation style.

8.2.3

Proxy Sender: Multicast Destination with a Range of Sender Ports

To configure RSVP Proxy sender for a multicast destination with a range of sender ports, you must use:

```
ADD -RSVP ProxySENDER EX1_PSM SESSION 239.0.0.2/2500 TCP SENDER 10.0.0.1/3000-3017 RATE  
2000 1000 TimeOut 40
```

The `ADD` command is the same as the `ADD` command for a unicast session, described in [Proxy Sender: Unicast Destination and One Sender Port on page 216](#), except for the multicast destination address in the `SESSION` parameter, the TCP protocol ID and the range of sender ports.

Although not shown in the `ADD` command above, the `SESSION` parameter may also contain a port range, as in the sender port range, which initiates multiple RSVP sessions, one for each port number within the range.

8.2.4

Proxy Receiver: Multicast Destination with a Range of Sender Ports

To configure proxy receiver for a multicast destination with a range of sender ports, you must use:

```
ADD -RSVP ProxyREceiver EX1_PRM SESSion 239.0.0.1/2500 TCP SENDER 10.0.0.1/3000-3017 RATE
2000 1000 STYLE SharedExplicit
```


The values for this command are the same as those for the unicast proxy receiver ADD command, described in [Proxy Receiver: Unicast Destination and One Sender Port on page 217](#) except for the multicast destination address, the TCP protocol ID, the sender port range and the SharedExplicit reservation style. A reserved bandwidth of 2000 bytes per second and 1000 byte burst size is requested and shared by data transmitted by the sender, 10.0.0.1, on ports 3000 to 3017 inclusive.

8.2.5

Configuring RSVP with L2TP Tunnel

The following figure illustrates a topology in which there is an L2TP VLL between Router A and Router B. This configuration reserves a large percentage of the tunnel's bandwidth capacity for L2TP traffic, with the rest of the tunnel available for Internet access.

Figure 68: RSVP for L2TP Tunnel



In this configuration, both routers are RSVP-aware, but the host PCs are not. L2TP uses UDP port 1701 to send packets. Also, the outer IP addresses of the WAN link (192.0.0.x) are used as sender and destination addresses.

Assuming the tunnel bandwidth is 64K bit/s, you must reserve 80% for RSVP/L2TP, and the average traffic on the tunnel is 10K bit/s full duplex, perform this procedure to create the topology illustrated above.

Procedure:

1. Set up L2T VLL. See [L2Tunnel Connections Configuration on page 535](#).
2. On Router A, enter the following commands:

```
SETDefault !2 -P0rt qcont = protr
ADD !2 -P0rt protr RSVP 80
```

The WAN port is set up to use protocol reservation for queuing and to reserve 80% of the port's bandwidth for RSVP.

3. Add a proxy sender to Router A by entering:

```
ADD -RSVP psend L2TPA sess 192.0.0.2/1701 UDP sender 192.0.0.1/1701 rate 1000 6400
```

The bandwidth is reserved in bytes.

Data transmitted by the PC host to destination address 192.0.0.2 port 1701 causes a RSVP PATH message to be sent to Router B.

4. Add a proxy receiver on Router B by entering:

```
ADD -RSVP prec L2TPB sess 192.0.0.2/1701 UDP sender 192.0.0.1/1701 rate 1000 6400
style FixedFilter
```

A RSVP PATH message received by Router B for the session from sender 192.0.0.1 on port 1701 causes an RSVP RESV message to be generated, which requests the amount of bandwidth as specified by the `RATE` value in the `ADD` command.

Chapter 9

Dynamic Host Configuration Protocol Configuration

Dynamic Host Configuration Protocol (DHCP) dynamically assigns an IP address to the requesting client for a specified duration of time. When the client no longer needs the IP address, or when the assigned duration expires, the address is returned to the pool. DHCP eliminates the need to manually assign an address to each client and provides efficient use of a finite number of IP addresses, especially for remote access clients.

When your TCP/IP environment already exists, you can use the existing IP address scenario and set up a DHCP address pool using the network mask that matches the primary network address. You can define a DHCP profile to supply configuration information as well as addressing to the DHCP client. Profile information can include gateway, DNS servers, WINS name server, Domain Name, lease times, renew timer, and rebind timer information.



NOTE: The `AddressPool` and `DefAddrPool` parameters apply to LANs only. The `RasAddrPool` parameter is used for RAS only.

DHCP configuration consists of three major steps:

- Enabling the DHCP service
- Configuring the address pools
- Selecting the configuration options

9.1

Configuring DHCP on EOS Devices

Procedure:

1. Enable Dynamic Host Configuration Protocol (DHCP) on the device by entering:

```
SETDefault !1 -DHCP CONTROL=Enable
```

For further details about enabling the DHCP service, see [Address Pools Configuration on page 221](#).

2. Set up the DHCP IP address pool by entering the following commands:

```
SETDefault !1 -DHCP CONTROL= AddressPool
```

```
ADD !1 -DHCP AddressPool 129.0.0.1 - 129.0.255.254 !P1
```

This IP address range is now assigned to the DHCP address pool and is available for all client requests.

For further details about setting up address pools, see [Address Pools Configuration on page 221](#).

3. Optional: Map a static address which associates a fixed IP address with the server MAC address.
4. Create a RAS address pool for RAS clients by entering:

```
ADD !1 -DHCP RasAddressPool 129.000.00.000 - 129.000.00.000 !P4
```

This setup allows remote access clients to acquire an IP address when they dial in for the duration of the session. DHCP manages a separate RAS IP address pool, associated with a configuration set, for the RAS clients. The allocation of a valid address and configuration is managed locally in the device.

Step example:

```
ADD !1 -DHCP StaticAddress %080002188E57 129.000.00.00 !p1
```


9.2

DHCP Service Activation

You can use the `CONTRol` parameter to enable the DHCP Service.

`CONTRol` is a per-port parameter that contains three options:

- Enabled/Disabled
- AddressPool/DefAddressPool
- ICMPCheck/NoICMPCheck

To make DHCP operational on a port, you must set the `CONTRol` parameter to `Enabled` and set the other two `CONTRol` parameter options as required for your application. For each port, DHCP maintains an active address pool – either `AddressPool` or `DefAddressPool` – from which to allocate an available address for its external client when requested. You can choose either `AddressPool` or `DefAddressPool` as the active address pool on a port by setting the `CONTRol` parameter to the appropriate option (`AddressPool` or `DefAddressPool`). The `IcmpCheck/NoIcmpCheck` option determines whether DHCP executes ICMP ECHO checking (ping that address) before it assigns an address to its client. This option is equivalent to a ping command.

For example, when you enter:

```
SETDefault !2 -DHCP CONTRol = (Enabled,AddressPool,IcmpCheck)
```

DHCP operation is enabled on port 2.

`AddressPool` is the active address pool and DHCP retrieves an available address and assigns it to the client when requested.

DHCP executes ICMP ECHO checking to make sure no other network device is using this address before it assigns the address to the client.

9.3

Address Pools Configuration

For external DHCP clients, DHCP uses the `AddressPool` and `DefAddressPool` address pools. For internal RAS clients, DHCP uses another address pool called the `RasAddressPool`.

When you set the `CONTRol` parameter to `AddressPool`, the `AddressPool` is considered the active address pool. When you set the `CONTRol` parameter to `DefAddressPool`, the `DefAddressPool` is considered the active address pool. At any given time only one address pool (either `AddressPool` or `DefAddressPool`) can be configured as an active address pool on a particular port.

In addition to address pools, DHCP also supplies a set of configurations, called the DHCP profile, to its clients. For details about DHCP profiles, see [DHCP Profiles on page 222](#).

For complete information on the commands and parameters, see “DHCP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

AddressPool

Before DHCP assigns an IP address to its external client, it checks the validity of the particular network. Currently, DHCP only operates on the network number that matches the primary `-IP NETaddr` and the network mask.

For example, when you enter:

```
ADD !2 -DHCP AddressPool 129.213.201.152 - 129.213.201.173 !P1
```

The address block of (129.213.201.152 - 129.213.201.173) is added with the associated profile `!P1` into the `AddressPool` of port 2.

When you enter:

```
ADD !2 -DHCP AddressPool 129.213.201.185 - 129.213.201.193 !P2
```

The address block of (129.213.201.185 - 129.213.201.193) is added with the associated profile !P2 into the AddressPool of port 2.

Each configured address block is associated with a DHCP profile which contains a set of DHCP options to offer the client when a member IP address is selected. If no *<profile_id>* is configured with an address block, the P1 profile is selected as the default profile.

You can configure multiple address blocks in the address pool for a particular port. When you configure multiple address blocks, it is important that the associated profile of each address block use the same profile set, so as to keep a consistent DHCP configuration allocation when addresses are selected from different address blocks.

DefAddressPool

DefAddressPool is generated by the system on a LAN interface. You can view it using the DefAddressPool parameter, but you cannot modify it. The DefAddressPool is always associated with the profile P1.

System-Generated DefAddressPool Display for All LAN Interfaces

To display the system-generated DefAddressPool on all LAN interfaces, you must enter:

```
SHoW !* -DHCP DefAddressPool
```

A display similar to the following appears:

```
-----Default Address Pool-----
Port    Default Address Pool    Profile
0       10.0.0.50 - 10.0.255.254   P1
1       10.1.0.50 - 10.1.255.254   P1
2       10.2.0.50 - 10.2.255.254   P1
3       10.3.0.50 - 10.3.255.254   P1
3B      10.11.0.50 - 10.11.255.254 P1
```

System-Generated DefAddressPool Display for DHCP-Enabled LAN Interfaces

To display the system-generated DefAddressPool only on those LAN interfaces for which DHCP is enabled on the interface (the CONTROl parameter is set to Enabled), you must enter:

```
SHoW -DHCP DefAddressPool
```

A display similar to the following appears:

```
-----Default Address Pool-----
Port    Default Address Pool    Profile
2       10.2.0.50 - 10.2.255.254   P1
```

When DefAddressPool is selected as the active address pool, the network number of DefAddressPool matches the network number (-IP NETAddr) configured on that port. In the previous example, DefAddressPool on port !2 was automatically generated as (10.2.0.50 - 10.2.255.254) because the IP network number was previously configured to be 10.2.0.0/16 on that port with:

```
SETDefault !2 -IP NETAddr = 10.2.0.1 255.255.0.0
```

9.4

DHCP Profiles

A DHCP profile contains a set of network configurations or attributes that DHCP offers (with IP address) to its clients. Supported DHCP configurations include:

- ProfDefGateWay (DHCP option 3)
- ProfDNS (DHCP option 6)
- ProfDomainName (DHCP option 15)
- ProfLEASE (DHCP option 51)

- ProfNetBiosNs (DHCP option 44)
- ProfReBindTimer (DHCP option 59)
- ProfRenewTimer (DHCP option 58)
- ProfSubnetMask (DHCP option 1).

One default profile (P1) is provided by the system, but you can modify its contents.

For complete information on the commands and parameters, see “DHCP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

9.5

Default DHCP Profile (Profile P1) Contents Display

To display the contents of the P1 default profile shipped with Enterprise OS software, you must use:

```
SHow -DHCP PProfile !P1
```

The following information is displayed:

```
Profile: 1
ProfDefGateWay      = (AUTO)
ProfDNS              = (AUTO) 129.213.128.98 128.9.0.107 198.41.0.4 192.33.4.12
ProfDomainName       = (NONE)
ProfLEASE            = 28800 (secs)
ProfNetBiosNs        = (NONE)
ProfReBindTimer      = 21600 (secs)
ProfRenewTimer       = 7200 (secs)
ProfSubnetMask       = (AUTO)
```

The contents of profile P1 is the default profile. This profile is also used as a template when you add a new profile. For details about each field (option) in the profile, see “DHCP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Some profile fields support AUTO. AUTO in each option field means automatically derived by the system. Some profile fields support NONE. NONE in each option field means not available, and that associated option is not offered to the client when addresses are assigned to the client.

9.6

Changing DHCP Profile Configuration Values

Procedure:

- To change the ProfDefGateWay to 129.213.128.122 in an existing DHCP profile, enter:
SETDefault !P1 -DHCP ProfDefGateWay = 129.213.128.122
- To eliminate ProfDefGateWay as an option in DHCP profile P1 when DHCP assigns an address, enter:
SETDefault !P1 -DHCP ProfDefGateWay = NONE
- To automatically derive the ProfDefGateWay in DHCP profile P1 when profile P1 is used when assigning an address, enter:
SETDefault !P1 -DHCP ProfDefGateWay = AUTO
For more information about ProfDefGateWay, see “DHCP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.
- To specify 129.213.128.16 and 129.213.128.2 as the ProfDNS in DHCP profile P1, enter:
SETDefault !P1 -DHCP ProfDNS = 129.213.128.16 129.213.128.2
- To eliminate ProfDNS in DHCP profile P1 when profile P1 is used to assign an address, enter:

```
SETDefault !P1 -DHCP ProfDNS = NONE
```

- To automatically derive the ProfDNS in DHCP profile P1, enter:

```
SETDefault !P1 -DHCP ProfDNS = AUTO
```

The four values are set as AUTO: 129.213.128.98, 128.9.0.107, 198.41.0.4, and 192.33.4.12 . When profile P1 is used to assign an address, the system offers ProfDefDNS with these values.

- To specify moca.mot.com as the ProfDomainName in DHCP profile P1, enter:

```
SETDefault !P1 -DHCP ProfDomainName = "moca.mot.com"
```

- To eliminate ProfDomainName as an option in DHCP profile P1 when profile P1 is used, enter:

```
SETDefault !P1 -DHCP ProfDomainName = NONE
```

- To eliminate ProfDomainName in DHCP profile P1, enter:

```
SETDefault !P1 -DHCP ProfDomainName = ""
```

The empty string "" is equivalent to NONE. ProfDomainName is not configured and not offered to its client.

- To specify 129.213.128.36 and 129.213.128.22 as the ProfNetBiosNs in DHCP profile P1, enter:

```
SETDefault !P1 -DHCP ProfNetBiosNs = 129.213.128.36 129.213.128.22
```

- To eliminate ProfNetBiosNs as an option in DHCP profile P1, enter:

```
SETDefault !P1 -DHCP ProfNetBiosNs = NONE
```

- To specify the ProfLEASE value in DHCP profile P1 as 7200 (seconds), enter:

```
SETDefault !P1 -DHCP ProfLEASE = 7200
```

- To specify the ProfReNewTimer value in DHCP profile P1 as 3600 (seconds), enter:

```
SETDefault !P1 -DHCP ProfReNewTimer = 3600
```

The ProfReNewTimer indicates the renew time that DHCP client the uses to start its renewal cycle.

- To specify the ProfReBindTimer value in DHCP profile P1 as 5400 (seconds), enter:

```
SETDefault !P1 -DHCP ProfReBindTimer = 5400
```

The ProfReBindTimer specifies the renew time that DHCP client uses to start its rebind cycle. The ProfReNewTimer should be less than ProfReBindTimer and less than ProfLEASE.

For example, you should configure values with the following relationship:

```
ProfReNewTimer = 1/2 * ProfLEASE
```

```
ProfReBindTimer = 3/4 * ProfLEASE
```

- To specify 255.255.255.0 as the ProfSubnetMask in DHCP profile P1, enter:

```
SETDefault !P1 -DHCP ProfSubnetMask = 255.255.255.0
```

- To eliminate ProfSubnetMask in DHCP profile P1, enter:

```
SETDefault !P1 -DHCP ProfSubnetMask = NONE
```

- To automatically derive the ProfSubnetMask in DHCP profile P1 so that when profile P1 is used when assigning an address, the system offers ProfDefGateWay with a value derived by the system, enter:

```
SETDefault !P1 -DHCP ProfSubnetMask = AUTO
```

9.7

Custom DHCP Profiles Configuration

You can create up to 32 custom DHCP profiles and change the contents of each profile. The contents of profile P1 are provided by the Enterprise OS software, but you can modify configuration values this profile. Profile P1 is used as the default profile throughout the DHCP service.

When you add a new profile, the EOS device copies the contents of profile P1 into the new profile as template contents. You can then change each option of the custom profile after it is created.

- You can add a new profile (P2 up to P32) to be associated with the address assignments. For example, you can use:

```
ADD -DHCP PProfile P7
```

- You can delete a profile (P2 up to P32) except default profile P1. For example, you can use:

```
DELEte -DHCP PProfile P7
```

- You can display how many DHCP profiles are currently configured. For example, you can use:

```
SHow -DHCP PProfile
```

The following display appears:

```
[5]DPE # SHow -DHCP PProfile
Profile P1
Profile P7
```

- You can display the contents of a specific DHCP profiles. For example, you can use:

```
SHow -DHCP PProfile P7
```

The following display appears:

```
[6]DPE # SHow -DHCP PProfile P7
Profile: 7
ProfDefGateWay    = (AUTO)
ProfDNS           = (AUTO) 129.213.128.98 128.9.0.107 198.41.0.4 192.33.4.12
ProfDomainName    = (NONE)
ProfLEASE         = 28800 (secs)
ProfNetBiosNs     = (NONE)
ProfReBindTimer   = 21600 (secs)
ProfReNewTimer    = 7200 (secs)
ProfSubnetMask    = (AUTO)
```

- You can map ProfClassIdent with a custom profile. For example, you can use:

```
ADD !2 -DHCP ProfClassIdent EngHost P4
```

This profile overrides the associated profile when the address is selected and the client identifies itself with a Class Identifier option. The following information appears:

```
[49]DPE # SHow !* -DHCP ProfClassIdent
-----Class Identifier Profile Mapping Table-----
Port  Class Identifier  Profile
2     DavidsHost       P4
5     DCH               P3

[50]DPE # SHow -DHCP ProfClassIdent
-----Class Identifier Profile Mapping Table-----
Port  Class Identifier  Profile
2     DavidsHost -    P4
```

- You can assign a static address. For example, you can use:

```
ADD !2 -DHCP StaticAddress %0020AF735FB1 129.213.201.167
```

StaticAddress allows you to maintain a special static mapping with the client's MAC address. By default, each dynamic address is selected from the active address except when a static address mapping is found. The following information is displayed:

```

SHow !* -DHCP StaticAddress
-----Static Address Mapping Table-----
Port      Host Address      IP address      Profile
2         %0020AF735FB1      129.213.201.167 P1
5         %0020AF4165EE      13.13.13.13     P1

```

With these configurations, a client with MAC address of %0020AF735FB1 requesting received from port ! 2 always gets IP address of 129.213.201.167 allocated. A client with MAC of %0020AF4165EE requesting received from port ! 5 always gets the 13.13.13.13 IP address allocated.

- You can delete a static address. For example, you can use:

```
DELeTe !5 -DHCP StaticAddress %0020AF4165EE 13.13.13.13
```

This command deletes this static mapping from the table.

- You can maintain an internal address pool to be used by the Remote Access Server (RAS) client: For example, you can use:

```
ADD !1 -DHCP RasAddressPool 20.20.20.25 - 20.20.20.55 P1
```

When the RAS client makes an address request to DHCP service, it is this address pool that DHCP uses to allocate the RAS address. This command adds the address block (20.20.20.25 - 20.20.20.55) associated with profile P1 into the RasAddressPool of port !1.

- You can delete the address block. For example, you can use:

```
DELeTe !1 -DHCP RasAddressPool 20.20.20.25 - 20.20.20.55
```

This command deletes the address block (20.20.20.25 - 20.20.20.55) from !2 RAS.

- You can display the log. For example, you can use:

```
SHow -DHCP Log
```

The following information is displayed:

```
Log = (NoSyslog,NoConsole)
```

- You can turn on the DHCP syslog feature. For example, you can use:

```
SETDefault -DHCP Log = Syslog
```

When Log is configured with Syslog, key operational information is sent to the system log server. The sys log server must be configured through the AUDIT LOG service, before DHCP can send system log messages.

- You can turn on the DHCP 'Console' feature. For example, you can use:

```
SETDefault -DHCP Log = Console
```

For details, see [Console Logging Feature on page 226](#).

9.8

Console Logging Feature

When the Log parameter is set to Console, key operational information is sent to the console. Console logging lists more detailed information than syslog messages. You can use the Console logging feature for diagnostics and debugging.

- You can display the current DHCP allocation status. For example, you can use:

```
SHow DHCP STATUS
```

The following information is displayed:

Assigned/ Port	Remaining IP Address	Client Address	Prof	Renewed	Lease	State
2	129.213.201.152	%0020AF735FB1	P1	02/09/98 14:50:08	23585	ASSIGNED
2	129.213.201.153	%0020AF4165EE	P1	02/09/98 15:42:16	26715	ASSIGNED

- You can display the current DHCP configuration. For example, to view the DHCP configuration for port 2, you must use:

SHow !2 -DHCP CONFIguration

A display similar to the following appears:

```
>>>>>>>> DHCP (global) CONFIguration <<<<<<<<<
Log = (NoSyslog,Console)
>>>>>>>> DHCP (port 2) CONFIguration <<<<<<<<<
Port 2 CONTrol = (Enabled,AddressPool,IcmpCheck)
-----Address Pool table-----
Port          Address Pool          Profile
2            129.213.201.152 - 129.213.201.153    P1
-----Default Address Pool-----
Port          Default Address Pool      Profile
2            10.2.0.50 - 10.2.255.254          P1
-----Class Identifier Profile Mapping Table-----
Port  Class Identifier  Profile
2     DavidsHost        P4
No Ras Address Pool configured for port 2
-----Static Address Mapping Table-----
Port  Host Address      IP address      Profile
2     %0020AF735FB8    129.213.201.166    P1
>>>>>>>> DHCP PProfiles <<<<<<<<<<<<<<<<<<<<<<
Profile P1
Profile P7
>>>>>>>> DHCP PProfiles Contents <<<<<<<<<<<<<<<<
Profile: 1
ProfDefGateWay    = (AUTO)
ProfDNS           = (AUTO) 129.213.128.98 128.9.0.107 198.41.0.4 192.33.4.12
ProfDomainName    = (NONE)
ProfLEASE         = 28800 (secs)
ProfNetBiosNs     = (NONE)
ProfReBindTimer   = 21600 (secs)
ProfReNewTimer    = 7200 (secs)
ProfSubnetMask    = (AUTO)
Profile: 7
ProfDefGateWay    = (AUTO)
ProfDNS           = (AUTO) 129.213.128.98 128.9.0.107 198.41.0.4 192.33.4.12
ProfDomainName    = (NONE)
ProfLEASE         = 28800 (secs)
ProfNetBiosNs     = (NONE)
ProfReBindTimer   = 21600 (secs)
ProfReNewTimer    = 7200 (secs)
ProfSubnetMask    = (AUTO)
[58]DPE #
```

Chapter 10

IPv6 Routing Configuration

To configure the device as an IPv6 router, you must set parameters in the following services:

- If your network uses the Routing Information Protocol Next Generation (RIPNG) as the interior gateway protocol for IPv6 intra-autonomous system routing, you must set parameters in the RIPNG Service.
- If your network uses BGP as the exterior gateway protocol for IPv6 inter-autonomous system routing, you must set parameters in the BGP service.

The IPv6 parameters enable the IPv6 routing function and configure the networks connected to the device.

10.1

Configuring IPv6 for Local Area Networks

Perform this procedure to configure basic IPv6 routing over LAN (Ethernet) ports and Point-to-Point Protocol (PPP) links.

By default, a link-local address is assigned to an IPv6 routing interface. For a transit device with no attached hosts, you may omit the interface-specific subnet prefix assignment. In this case, the device can communicate with other nodes on the same link using the interface's link-local address as its source address.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up.
- You are familiar with the protocols supported by the device.

Obtain an IPv6 address for each port you want to configure.

Procedure:

1. Assign an IPv6 address for each port that performs IPv6 routing by entering:

```
ADD !<port> -IPv6 NETaddr <IPv6_address>
```

The IPv6 address format and its textual representation must be consistent with the Internet draft on IPv6 Addressing Architecture, version 2.

The local interface identifier is generated using the EUI-64 format. For backward compatibility, static routes and dynamically learned routes that conform to RFC 1884 are also supported.

2. Enable the dynamic routing protocols (RIPNG) for IPv6 routing by entering:

```
SETDefault !<port> -RIPNG CONTROL = ([Talk | NoTalk], [Listen | NoListen], [Poison | NoPoison], [Trigger | NoTrigger])
```

Setting the CONTROL parameter to the TALK and Listen values enables the device to send and receive routing information with other devices using RIPNG.

3. Enable IPv6 routing on a specified port by entering:

```
SETDefault !<port> -IPv6 CONTROL = R0ute
```


10.2

IPv6 Configuration for Wide Area Network Connectivity

EOS devices support the following types of IPv6 tunnels:

- IPv4 over IPv6 tunnels
- IPv6 over IPv6 tunnels



NOTE: IPv6 routing over Frame Relay is currently not supported.

10.2.1

Configuring IPv4 over IPv6 Tunnels

An IPv4 over IPv6 tunnel is a point-to-point link between IPv6 nodes that allows payloads of IPv4 data packets to be encapsulated and transported over an IPv6 backbone. The IPv6 protocol functions as a link layer protocol on this link, and the two IPv6 nodes play specific roles. One node encapsulates the original packets received from the LAN or from itself and forwards the resulting tunnel packets through the tunnel. The other node decapsulates the received tunnel packets and forwards the resulting original packets towards their destinations, which may be the node itself.

Procedure:

1. Enable IPv6 routing on the physical port by entering:

```
SETDefault !<port> -IPv6 CONTrol = R0ute
```

Step example:

```
SETDefault !3 -IPv6 CONTrol = R0ute
```

2. Assign an IPv6 address to the physical port by entering:

```
ADD !<port> -IPv6 NETaddr <IPv6_address>
```

Step example:

```
ADD !3 -IPv6 NETaddr fd00::215:9aff:fe85:5e67
```

3. Create a virtual port over which the IPv6 tunnel runs by entering:

```
ADD !v<port> -Port VirtualPort IPIPV6 <local_ipv6addr> <remote_ipv6addr> [TTL <t1>]
```

Step example:

```
ADD !v3 -Port VirtualPort IPIPV6 fd00::215:9aff:fe85:5e67 fd00::a00:3eff:fedc:5273  
TTL 10
```

4. Assign an IPv4 address to the virtual port over which the IPv6 tunnel runs by entering:

```
SETDefault !v<port> -IP NETaddr = <IP_address> <subnet_mask>
```

Step example:

```
SETDefault !v3 -IP NETaddr = 3.3.3.133 255.255.255.0
```

10.2.2

Configuring IPv6 over IPv4 Tunnels with an Encrypted Payload

If you want to encrypt the IPv4 payload, you must configure an additional IPv4 tunnel, which runs inside the IPv6 tunnel.

Procedure:

1. Enable IPv6 routing on the physical port by entering:

```
SETDefault !<port> -IPv6 CONTrol = R0ute
```

Step example:

```
SETDefault !3 -IPv6 CONTrol = R0ute
```

2. Assign an IPv6 address to the physical port by entering:

```
ADD !<port> -IPv6 NETaddr <IPv6_address>
```

Step example:

```
ADD !3 -IPv6 NETaddr fd00::5:6:7:8
```

3. Assign the virtual port to the IPv6 tunnel by entering:

```
ADD !v<port> -P0rt VirtualPort IPIPv6 <local_ipv6addr> <remote_ipv6addr> [TTL <t1l>]
```

Step example:

```
ADD !v2 -P0rt VirtualPort IPIPv6 fd00::5:6:7:8 fd00::1:2:3:4 TTL 10
```

4. Assign an IPv4 address to the virtual port over which the IPv6 tunnel runs by entering:

```
SETDefault !v<port> -IP NETaddr = <IP_address> <subnet_mask>
```

Step example:

```
SETDefault !v2 -IP NETaddr = 4.4.4.132 255.255.255.0
```

5. Enable security policy checking on the virtual port over which the IPv6 tunnel runs by entering:

```
SETDefault !v<port> -CRYPTO CONTrol = Enable
```

Step example:

```
SETDefault !v3 -CRYPTO CONTrol = Enable
```

6. Define a dynamic security policy by entering:

```
ADD !v<port> DynamicPOLicy <policy_name> <priority> <mode> <selctrlist_name>  
<xfrm1ist_name> WinSize <replay_win_size>
```

Step example:

```
ADD !v2 -CRYPTO DynamicPOLicy p3 3 Tunnel s1 t2 WinSize 256
```

7. Define a pre-shared key entering:

```
ADD -CRYPTO PreSharedKey [<version>] <peer_ID> <pre-shared_key>
```



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

Step example:

```
ADD !v3 -CRYPTO PreSharedKey IKEv2 fd00::1:2:3:4 "12345"
```

10.2.3

Configuring IPv6 over IPv6 Tunnels

An IPv6 over IPv6 tunnel is a point-to-point link between IPv6 nodes that allows payloads of IPv6 data packets to be encapsulated and transported over an IPv6 backbone. The IPv6 protocol functions as a link layer protocol on this link, and the two IPv6 nodes play specific roles. One node encapsulates the original packets received from the LAN or from itself and forwards the resulting tunnel packets through the tunnel. The other node decapsulates the received tunnel packets and forwards the resulting original packets towards their destinations, which may be the node itself.

Procedure:

1. Enable IPv6 routing on the physical port by entering:

```
SETDefault !<port> -IPv6 CONTrol = ROute
```

Step example:

```
SETDefault !3 -IPv6 CONTrol = ROute
```

2. Assign an IPv6 address to the physical port by entering:

```
ADD !<port> -IPv6 NETaddr <IPv6_address>
```

Step example:

```
ADD !3 -IPv6 NETaddr fa00:0:0:2000:0022/64
```

3. Create a virtual port over which the IPv6 tunnel runs by entering:

```
ADD !v<port> -Port VirtualPort IPIPv6 <local_ipv6addr> <remote_ipv6addr>
```

Step example:

```
ADD !v3 -Port VirtualPort IPIPv6 fd00:0:0:2000:0022 fd00:0:0:2000:0021
```

4. Assign an IPv6 address to the virtual port over which the IPv6 tunnel runs by entering:

```
SETDefault !v<port> -IPv6 NETaddr = <IPv6_address>
```

Step example:

```
SETDefault !v3 -IPv6 NETaddr = fc00:0:0:2000:0022/64
```

10.2.4

Configuring IPv6 over IPv6 Tunnels with an Encrypted Payload

Procedure:

1. Enable IPv6 routing on the physical port by entering:

```
SETDefault !<port> -IPv6 CONTrol = ROute
```

Step example:

```
SETDefault !3 -IPv6 CONTrol = ROute
```

2. Assign an IPv6 address to the physical port by entering:

```
ADD !<port> -IPv6 NETaddr <IPv6_address>
```

Step example:

```
ADD !3 -IPv6 NETaddr fa00:0:0:2000:0022/64
```

3. Create a virtual port over which the IPv6 tunnel runs by entering:

```
ADD !v<port> -Port VirtualPort IPIPv6 <local_ipv6addr> <remote_ipv6addr>
```

Step example:

```
ADD !v3 -Port VirtualPort IPIPv6 fd00:0:0:2000:0022 fd00:0:0:2000:0021
```

4. Assign an IPv6 address to the virtual port over which the IPv6 tunnel runs by entering:

```
SETDefault !v<port> -IPv6 NETAddr = <IPv6_address>
```

Step example:

```
SETDefault !v3 -IPv6 NETAddr = fc00:0:0:2000:0022/64
```

5. Enable security policy checking on the virtual port over which the IPv6 tunnel runs by entering:

```
SETDefault !v<port> -CRYPTO CONTrol = Enable
```

Step example:

```
SETDefault !v3 -CRYPTO CONTrol = Enable
```

6. Define a dynamic security policy by entering:

```
ADD !v<port> DynamicPOLicy <policy_name> <priority> [<version>] <mode>  
<selctrlist_name> <xfrmlist_name> WinSize <replay_win_size>
```



NOTE: You must set the <version> value to IKEv2 for IPv6 over IPv6 tunnels.

Step example:

```
ADD !v3 -CRYPTO DynamicPOLicy p3 3 IKEv2 Tunnel s1 t2 WinSize 256
```

7. Define a pre-shared key by entering:

```
ADD -CRYPTO PreSharedKey [<version>] <peer_ID> <pre-shared_key>
```



NOTE: You must set the <version> value to IKEv2 for IPv6 over IPv6 tunnels.

If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the pre-shared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

Step example:

```
ADD !v3 -CRYPTO PreSharedKey IKEv2 fd00:0:0:2000:0021 "12345"
```

10.3

Verifying the Configuration

To verify the configuration, you must examine network devices and send packets from one network to another to determine if they are properly forwarded.

Process:

1. Examine network devices. See [Examining Network Devices on page 233](#).
2. Retrieve IPv6 statistics. See [Retrieving IPv6 Statistics on page 233](#).
3. Check the overall status of an EOS device configured as an IPv6 router. [Checking the Overall Status on page 233](#).
4. Check if the device can properly forward packets. See [Checking with PING6, TraceRoute6, and TELnet6 on page 234](#).

10.3.1

Examining Network Devices

Perform this procedure to examine the status of an EOS device configured as an IPv6 router.

Procedure:

1. Display information on the attached networks by entering:
`SHoW -IPv6 NETAddr`
2. Determine which stations or networks are reachable from the device by entering:
`SHoW -IPv6 AllRoutes`
3. Display information from the Address Translation Table by entering:
`SHoW -IPv6 ADDRess`
4. Display information on the configured tunnels by entering:
`SHoW -IPv6 TUnneL`
5. Examine the neighbor discovery status by entering:
`SHoW -IPv6 NbrDiscovery`

10.3.2

Retrieving IPv6 Statistics

You can collect statistics for a specific period using the `SampleTime` and `STATistics` parameters. For more information, see “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

For more information about the supported statistics, see [Statistics Displays on page 967](#).

Procedure:

1. Retrieve the IPv6 statistics by entering:
`SHoW -SYS STATistics -IPv6`
2. Retrieve the RIPNG statistics by entering:
`SHoW -SYS statistics -RIPNG`
3. Retrieve the BGP statistics by entering:
`SHoW -sys statistics -BGP`

10.3.3

Checking the Overall Status

Perform this procedure to check the overall status of an EOS device configured as an IPv6 router.

Procedure:

1. Examine the path configurations by entering:
`SHoW -PATH CONFIguration`
2. Examine the port configurations by entering:
`SHoW -PORT CONFIguration`
3. Examine the IPv6 configurations by entering:
`SHoW -IPv6 CONFIguration`

4. Examine the RIPNG configurations by entering:

`SHoW -RIPNG CONFIguration`

5. Examine the BGP configuration by entering:

`SHoW -BGP configuration`

10.4

Checking with PING6, TraceRoute6, and TELnet6

After the configuration of a device for IPv6 is completed, perform this procedure to determine if the device can properly forward packets.

Procedure:

1. To verify proper routing to each of the other routers or hosts, use the `PING6` command.
2. If the `PING6` command fails, use the `TraceRoute6` command to determine where the route fails.
3. To access the last device that responds to the `TraceRoute6` command, use the `TELnet6` command.

You can examine the configuration, IPv6 routing table, and status of the device, and use the basic tools (`PING6`, `TraceRoute6`, and `TELnet6`) to troubleshoot the problem.

10.5

Customizing the IPv6 Routing Configuration

After you set up and check the basic IPv6 routing configuration, an EOS device configured as an IPv6 router can perform packet routing. You can further customize the IPv6 routing configuration by performing the following procedures.

Process:

- Configure the internal port. See [Configuring the Internal Port on page 234](#).
- Configure multiple subnets. See [Configuring Multiple IPv6 Subnets on page 235](#).
- Configure neighbor discovery. See [Configuring Neighbor Discovery on page 235](#).
- Configure static routes. See [Configuring Static Routes on page 235](#).
- Configure the RIPNG routing policies. See [Configuring RIPNG Routing Policies on page 236](#).
- Configure RIPNG route aggregation. See [Configuring RIPNG Route Aggregation on page 237](#).

10.5.1

Configuring the Internal Port

The IPv6 link-local subnet exists by default once IPv6 routing is enabled on an interface. Therefore, you are not required to configure a IPv6 subnet on an interface to perform IPv6 routing. Since a local address is required to support end-to-end communication with remote nodes for applications such as `PING6` and `TELnet6`, you must assign a node address to the device by configuring an internal port.

The internal port stays up if IPv6 is active on one or more interfaces. The address assigned to the internal port is used as a local address for end system packets originated from the EOS device. The internal port address is advertised by the routing protocol as a reachable subnet prefix.

Procedure:

1. Configure a single subnet prefix on the internal port by entering:

```
ADD !0 -IPV6 NETAddr 3ffe:1:1:1/64 EUI-64
```

The MAC address of the CPU board is used to generate the EUI-64 interface identifier and append to the configured prefix to form the global IPv6 address.

2. Activate the internal port by entering:

```
SETD !0 -IPV6 CONTrol=R0ute
```

10.5.2

Configuring Multiple IPv6 Subnets

EOS devices support multiple IPv6 subnets. You can configure more than one IPv6 network or subnet on a port. The procedure for configuring multiple IPv6 subnets on all interfaces is the same. Perform this procedure to configure multiple IPv6 subnets on an Ethernet network.

Procedure:

1. To configure two IPv6 networks on an Ethernet network, enter:

```
ADD !<port> -IPV6 NETAddr <IPv6_address>
```

2. To delete a subnet, enter:

```
DELeTe !<port> -IPV6 NETAddr <IPv6_address>
```

10.5.3

Configuring Neighbor Discovery

EOS devices support IPv6 Neighbor Discovery only on LAN ports according to RFC 1970.

Procedure:

To configure neighbor discovery, enter:

```
SETD !<port> NbrDiscovery = ([DupAddrDetect | NoDupAddrDetect],  
[SendRouterAdv | NoSendRouterAdv]  
[StatefulAddrConf | NoStatefulAddrConf],  
[OtherStatefulConf | NoOtherStatefulConf],  
[AdvHopLimit | NoAdvHopLimit], [AdvMTU | NoAdvMTU]  
[AdvReachableTime | NoAdvReachableTime],  
[AdvRetryTime | NoAdvRetryTime])
```

Link layer address resolution is enabled by default.

You can enable duplicate address detection with the `NbrDiscovery` parameter. If the interface MAC address already exists when the link comes up, routing is disabled on the interface and a system message is logged.

You can also set up autoconfiguration of adjacent hosts with the `NbrDiscovery` parameter. In conjunction with the `NetAddress` parameter options, the neighbor hosts acquire the configured subnet prefix and link-specific network parameters advertised by the EOS device.

10.5.4

Configuring Static Routes

A static route is a user-defined route by which a network can be reached. You can configure as many static routes as desired.

Procedure:

1. To add a static route, enter:

```
ADD !<port> -IPv6 Route <IPv6_address> [<gateway>] <metric> [Override]
```

Step example:

The subnet 3ffe:1:1:3/64 is reachable on port 2 through the gateway fe80:200:81ff: fed5:0892. The metric is 3.

To configure this example, enter:

```
ADD !2 -IPv6 route 3ffe:1:1:3/64 fe80:200:81ff:fed5:0892 3
```



NOTE: The gateway must be located on a network directly connected to the device on which you add the static route.

If the outgoing interface is a point-to-point configured IPv4 tunnel, you can add a static route without the next-hop gateway address:

```
ADD !t1 -IPv6 Route 3ffe:1:1:5/64 2
```

2. To delete a static route, enter:

```
DELeTe -IPv6 ROute <IPv6_address> [<gateway>]
```

10.5.5

Configuring RIPNG Routing Policies

The routing policies supported by RIPNG allow you to control the reporting of routing information on a per-port basis.

Prerequisites:

Ensure that:

- You familiarized yourself with the various policies that are available, and determine which policies you want to configure. The following table lists and briefly describes each policy and its associated parameter.
For more information on the parameters listed in the table, see “RIPNG Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.
- You determined the IPv6 addresses of these specific networks if you plan to access or receive information on routes from specific networks as opposed to all or no networks.

Table 22: RIPNG Routing Policies

Policy	Description	Parameter
Advertise	Controls which routes are reported regardless of the route source.	AdvertisePolicy
Static	Controls which static routes are reported in the IPv6 routing environment.	StaticPolicy
Exterior	Controls which BGP routes are reported	ExteriorPolicy
Receive	Controls which RIPNG routes are accepted by a trusted neighbor.	ReceivePolicy

Procedure:

1. Establish an advertise policy that controls the advertisement of routes through RIPNG regardless of the source from which the route is learned, by entering:

```
ADD !<port> -RIPNG AdvertisePolicy All | None | [~]<IPv6_address> [<metric> (0-15)]
```

Step example:

To configure a policy on port 1 that forwards information on all routes to network 3FFE:1:1:1/64, enter:

```
ADD !1 -RIPNG AdvertisePolicy 3FFE:1:1:1/64
```

In this situation, a metric associated with network 3FFE:1:1:1/64 is not specified. If you decide not to specify a metric with the `AdvertisePolicy` parameter or to specify a metric of zero, a route is reported with a metric calculated from the routing table.

2. Establish a receive policy that accepts or rejects routes learned by RIPNG from trusted neighbors by entering:

```
ADD !<port> -RIPNG ReceivePolicy All | None | [~] <IPv6_address> [<metric> (0-15)]
```

Step example:

To configure port 1 so that it accepts information on routes learned by RIPNG for network 3FFE:1:1:1/64, enter:

```
ADD !1 -RIPNG ReceivePolicy 3FFE:1:1:1/64
```

In this situation, a metric associated with network 3FFE:1:1:1/64 is not specified. If you decide not to specify a metric with the `ReceivePolicy` parameter or specify a metric of zero, a route with the originally reported metric is stored in the routing table.

3. Establish a static policy for reporting static routes.

Step example:

To configure a policy on port 1 that forwards routing information about static routes configured on network 3FFE:1:1:1/64, enter:

```
ADD !1 -RIPNG StaticPolicy 3FFE:1:1:1/64
```

In this situation, a metric associated with network 3FFE:1:1:1/64 was not specified. If you decide not to specify a metric with the `StaticPolicy` parameter or specify a metric of zero, a route is reported with a metric calculated from the routing table.

4. Establish an exterior policy for reporting routes learned by BGP.

Step example:

To configure a policy on port 1 that advertises all the BGP routes except 3ffd/64, enter:

```
ADD !1-RIPNG ExteriorPolicy 3ffd/64
```

5. Specify the list of neighbor devices to exchange routing information through RIPNG using the `AdvToNeighbor` parameter.

Step example:

```
ADD !V1 -RIPNG AdvToNeighbor fe80::1
```

6. Perform [step 1](#) to [step 5](#) on all peer neighbor devices.

10.5.6

Configuring RIPNG Route Aggregation

Route aggregation allows multiple specific routes to be advertised as a single prefix route. By combining several network/host routes into one subnet route, the update message is smaller and the size of the routing table is reduced.

Procedure:

To configure RIPNG route aggregation, specify a prefix route that RIPNG advertises as a simple supernet route on a port by entering:

```
ADD !<port> -RIPNG AggregateRoute <IPv6_prefix> [metric]
```

When RIPNG route advertisement is enabled with the `SETD !<port> -RIPNG CONTROL=TA` command, all configured aggregate routes are advertised while all the associated sub-routes advertisements are suppressed.

10.6

IPv6 Network Topology

An IPv6 network is configured on each interface where IPv6 packets are received and sent. The interface can be either a local Ethernet interface or a PPP link.



CAUTION: Each IPv6 subnet that you assign directly to a port must be unique; in other words, you cannot assign the same IPv6 subnet to different ports.

A router must check its routing table to determine where to route a packet. If the destination is on an attached network, the router can send it directly to the network. But if the destination is farther away in the internetwork, the router must route the packet to another router (called a *gateway*) that is closer to the destination. The route to a gateway can be statically configured or dynamically learned through RIPNG. When two routers are located on the same network (that is, each of them has at least one interface to the network), they are considered *neighbors* or *neighboring gateways*.



NOTE: Only RIPNG is supported for interautonomous system routing.

10.7

Multipath Routing

EOS devices supports multipath routing, which means that up to four routes for each destination address can be stored in the routing table. The advantages of multipath routing are as follows:

- The device can still route a packet using an alternative route if the primary one fails; it is more responsive to network topology changes than if only one route to a destination exists.
- The device can distribute the load among the available equal-cost best paths.

When multiple routes for a destination exist, the device uses the route with the highest precedence. The following table lists the types of routes used (by decreasing precedence).

Table 23: Route Precedence

Precedence Level	All Protocols without BGP	With BGP
1	Static (added without Override)	Static (added without override)
2	RIPNG	RIPNG
3	ICMP redirect (host only)	ICMP redirect (host only)
4	Static (with override)	Static (with override)
5		BGP
6	Default route	BGP default route

The routing table displays routes with a high precedence first.

If the route with the highest precedence fails, the route with the next highest precedence is used. A route in the routing table is deleted in these situations:

- A dynamic route learned through RIPNG is deleted when an EOS device times out and goes through the HOLD-DOWN and GARBAGE COLLECTION states. A device times out when it fails to hear from a neighbor for a period that is six times the value of the `UpdateTime` parameter. For example, if the value of the

`UpdateTime` parameter is 45 seconds, the device times out if it does not hear from its neighbor for 270 seconds.

- A static route is deleted via the `DELeTe ROUte` command.
- A lowest precedence route is deleted when four routes of higher precedence are available. This situation occurs when a fifth route is learned and has a higher precedence than the lowest precedence route.

You can remove dynamic routes learned by RIPNG using the `FLush -IPv6 AllRoutes` command.

Route Selection and Load Splitting

If two or more routes with the same route source precedence are available to reach a destination, the EOS device always selects the route with the lowest metric (measured in hops for RIPNG). If there is more than one route learned by the same routing protocol with the same equal-cost, low metric, you can split the load between these routes on a round-robin basis. The `SplitLoad` parameter (`Enable | Disable`) determines whether or not load splitting is performed.

Because load splitting balances the load among different routes, Motorola Solutions recommends it if two or more routes are available to reach a destination and the routes have similar metrics. However, if the routes connecting various networks have different metrics (that is, there is only one route with the fewest hops or lowest cost to a destination), load splitting is not necessary.

10.8

Default Routes

Router A receives a RIPNG update packet from Router B, which has an entry indicating that network `::/0` is reachable with metric 3.

EOS devices support up to four default routes; when more than one default route is available, the same selection rules apply. If load splitting is enabled, the load is distributed among equal-cost best paths.

An advantage of a device using a default route is that network overhead in an autonomous system can be reduced. The reduction in overhead occurs because the device does not need to advertise all external routes.

Default Routes Example

Router A receives a RIPNG update packet from Router B, which has an entry indicating that network `::/0` is reachable with metric 3. Router A considers Router B its default gateway. That is, if Router A needs to route a packet whose destination is not found in its routing table, it sends the packet to Router B.

To configure RIPNG to advertise a default route, assign a non-zero value to the `DefaultMetric` parameter. You do not need to configure the `DefaultMetric` parameter on every EOS device throughout the domain. The default route learned on one interface is propagated to neighbors on the other interfaces.

RIPNG propagates the advertisement of the default route as the normal operation. For the RIPNG Protocol, you can suppress propagation of the default route using the `AdvertisePolicy` parameter.

You can configure the default route in one of two ways:

- On the exit router of a domain, configure a static override default route with a metric of 1 that points to the first hop outside the domain. Then use the `StaticPolicy` parameter of RIPNG to import this route into the routing protocol and advertise it into the domain.
- At the top level, set the `DefaultMetric` parameter in RIPNG to instruct the device to originate a default route.

10.8.1

Learning Routes with RIPNG

EOS devices support RIPNG according to RFC 2080. Normally, every 30 seconds or every time it learns a route change for a network, the device uses multicast packets to report to its neighboring gateways the following types of information:

- The networks it can reach
- The metric associated with each network it can reach

By default, the information in update packets pertains only to learned routes. Static route information is not reported.

You can configure some parameters to determine how the device sends out the updates and what is included in them. For example, you can configure the parameters for the following purposes:

- To change the frequency of the multicast traffic (`UpdateTime` parameter)
- To prevent the device from sending or receiving update and request packets (`CONTRol` parameter)
- To control the set of neighboring devices from which the EOS device receives updates and to which it sends them (`AdvToNeighbor` and `RcvFromNeighbor` parameters)
- To prevent the device from sending out a trigger update response upon a route change for a network (`CONTRol` parameter)
- To enable the device to report static routes (`StaticPolicy` parameter)
- To cause some routes not to be reported or to be reported with the infinity metric (using the poison reverse `CONTRol` parameter)

10.8.2

Network Reachability

The following types of networks are considered **reachable** when a device multicasts its RIPNG update packets:

- All directly connected networks
- All static routes (as controlled by the `StaticPolicy` parameter)
- All dynamic routes learned through RIPNG

10.8.3

Solution to the Slow Convergence Problem with Split Horizon

When the device advertises its own reachable networks, it advertises to all neighbors except the one from which it learned of the advertised network.

On a LAN network, it is not necessary to configure neighbors. If you do not configure neighbors, RIPNG multicasts the updates over the LAN. If you configure neighbors, RIPNG unicasts the updates to each of the neighbors on the `AdvToNeighbor` list.

10.8.4

Solution to the Slow Convergence Problem with Poison Reverse

You configure poison reverse using the `Poison` or `NoPoison` value for the `-RIPNG CONTrol` parameter.

If poison reverse is enabled, the device advertises all routes to all neighbors, but when advertising a route to a neighbor that advertised the same route, the device sets the metric to infinity (0xFFFF) to prevent the recipient from adding the route to its routing table. Poison reverse speeds convergence but adds to network overhead.


If poison reverse is disabled, the device omits routes learned from one neighbor from RIPNG updates sent to that neighbor. Disabling poison reverse has the advantage of minimizing network overhead in large network configurations at the expense of slower convergence.

10.8.5

Different States of RIPNG-Learned Routes

To avoid routing loops, new information about a route is ignored for a designated period before it is used. The following figure summarizes how a route learned through RIPNG changes states.

Figure 69: Different States of a RIPNG Route



- **GARBAGE COLLECTION state**

When the timer for a route that was in the **HOLD-DOWN** state expires, that route changes to **GARBAGE COLLECTION** state. This happens when no update packets are received to indicate that the route is still reachable. In this state, if a neighboring gateway reports the route with a noninfinity metric within 120 seconds, the route can go back to the **UP** state. If no updates are received within 120 seconds (garbage-collection timer), the route is deleted from the routing table. It is possible to go into **GARBAGE COLLECTION** state if no updates are received within 180 seconds.

- **UP state**

A route is considered UP if it is reachable with a noninfinity metric (15 or fewer hops). Whether it is reachable is determined by the last update received from the neighboring gateways. It remains UP for 180 seconds (the route timer). The timer is reset each time a new update for the route is received.

- **HOLD-DOWN state**

A route in UP state changes to HOLD-DOWN state if an update received from the original gateway indicates that the route is associated with an infinity metric (16 hops). In this state, all update information received from other gateways for that route is ignored.

However, if an update is received from the original gateway within 60 seconds (the hold-down timer), and it associates a noninfinity metric with the route, the route goes back to UP state.

If the hold-down timer expires, the route goes from HOLD-DOWN state to GARBAGE COLLECTION state for 120 seconds.

When you display the routing table with the `SHoW -IPV6 AllRoutes` command, the state of each route is displayed under the `STATUS` heading.

Chapter 11

Network Address Translation (NAT) Configuration

The following terms apply to Network Address Translation (NAT):

Inside network

The network that includes addresses you want to map.

Outside network

The network, such as the Internet, to which you want to connect the inside network.

Inside addresses

The untranslated addresses of the inside network.

Outside addresses

The addresses that are mapped to the inside addresses.

You can use NAT for the following purposes:

- Private address space – You want to connect to the Internet, but your network does not use globally-routable IP addresses. If your network uses private addresses, you can use NAT to translate them to access the outside network.
- Load sharing – You want to do load sharing of incoming TCP traffic. For example, traffic destined for a web server identified by one IP address can be redirected to multiple servers with duplicate websites.
- Address migration – You must change your inside addresses. You may change Internet service providers (ISPs), for example, and have to change your numbering scheme. Instead of changing IP addresses on every device in your network, you can use NAT to translate the current addresses into new addresses.
- Address redirection – You want to redirect traffic from one host to another. For example, you want to stop Telnet traffic to Server A but want to allow traffic to a new Server B. Translate the address of Server A to that of Server B to transparently redirect users to the new server.

11.1

NAT Guidelines

The following guidelines must be followed before configuring Network Address Translation (NAT):

- Traffic that has embedded IP addresses, such as DNS requests and responses, is not translated by NAT.
- You must use private IP address space on your inside network recommended by the Internet Assigned Numbers Authority (IANA). The following address ranges are designated as private networks that should not be advertised:

10.0.0.0 – 10.255.255.255 (10.0.0.0/8)

172.16.0.0 – 172.31.255.255 (172.16.0.0/12)

192.168.0.0 – 192.168.255.255 (192.168.0.0/16)

If you use a different address range that can be validly assigned to someone else's network, you are not able to communicate with that network. For example, if you FTP from the inside host 1.1.1.1 to a valid outside network host with the IP address 1.1.1.7, the device does not forward the FTP request outside the network.

- NAT may not be practical if large numbers of hosts in the inside network communicate with the outside network, because NAT is slower than untranslated addresses. Most traffic should originate from or go to

hosts within the domain. Because most hosts never communicate with an outside network, only a subset of inside addresses need to be translated into outside addresses.

11.2

Configuring Network Address Translation

Perform this procedure to configure Network Address Translation (NAT).

Prerequisites:

Ensure that:

- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#).
- You are logged on to the system with Network Manager privileges.
- You assigned the appropriate IP addresses to the relevant ports on the device. See [IP Routing Configuration on page 119](#).

Procedure:

1. Enable NAT on a port connected to an outside network, by entering:

```
SETDefault !<port> -NAT CONTrol = Enable
```

2. To define the address mapping, enter:

```
ADD !<port> -NAT AddressMap <LHS Address(es)> <RHS Address(es)> [InBound | OutBound  
| BiDirectional | LoadShare] [Log [0..7]] (<LHS and RHS Address(es)> = <IPaddr/  
mask> | <IPaddr> | <IPaddr>-<IPaddr> [, <IPaddr/mask> | <IPaddr> | <IPaddr>-  
<IPaddr>..])
```

where LHS Address(es) and RHS Address(es) is a single address or a range of addresses.



NOTE: If you specify BiDirectional, you can only use a one-to-one map.

Address maps take effect immediately. When you remove a map using the DElete command, all session table entries based on the map are deleted.

If you have static and dynamic maps, the EOS device checks the static map first and then the dynamic map.

For examples of address mapping, see [Address Mapping Examples on page 245](#).

For more information about static and dynamic maps, see [Address Mapping on page 248](#).

3. To map an address and TCP/UDP port to another address and TCP/UDP port, enter:

```
ADD !<port> -NAT TcpUdpPortMap <IPaddr>, <TCP/UDP_port#> <IPaddr> [, <TCP/  
UDP_port#>] [InBound | Outbound | BiDirectional] [Log[0..7]]
```

If you do not specify the second TCP/UDP port number, the software uses the same port number specified on the first address.



NOTE: The software cannot differentiate between TCP and UDP ports.

If you defined address maps and TCP/UDP maps, the EOS device checks the TCP/UDP maps first and then the address maps.

Step example:

To allow inbound telnet traffic to host 10.0.0.1 on port 1, enter:

```
ADD !1 -NAT TcpUdpPortMap 144.195.48.20,23 10.0.0.1 Inbound
```

4. To log messages to the AuditLog Service, enable the service by entering:

```
SETDefault -AuditLog CONTrol = (COnfig, MESSAGES, SEcurity)
```


The AuditLog Service controls the delivery of messages to the syslog server(s). Only the start-of-connection packets are logged, to avoid flooding logging messages. No more than 10 log messages per second are generated. All messages over 10 are suppressed and the next log message (generated after the one second window expires) contains a counter of how many previous messages were suppressed.

5. To specify whether messages are logged to the AuditLog Service, the local console, or both, enter:

```
SETDefault -NAT Log = [Syslog | NoSyslog] [Console | NoConsole] [SessionFail | NoSessionFail] [SessionSuccess | NoSessionSuccess] [LogDetail | NoLogDetail]
```

6. To display the NAT session information, enter:

```
SHoW [!<port> | !*] -NAT SESSions [TCP | UDP | FTP | Others]
```

You can display the address translations, usage, and idle time statistics for active NAT sessions. Each translation is identified by a session ID. When no option is specified, all active NAT sessions are displayed. The FTP session display can have two TCP connections active at the same time. One is the control channel, for passing commands and responses between client and server. The other is the data channel for the actual data transfer.

7. To determine the maximum time-out period allowed for NAT sessions to remain idle before the session is terminated, enter:

```
SETDefault !<port> -NAT SessionTimeout [TCP | Others] <minutes> (0-99999) [ : <seconds> (0-59)]
```

8. If the address translation fails at the start of a session, to specify how to handle the failed packet, enter:

```
SETDefault !<port> -NAT XlateFailAction = PassThrough | Drop | GenerateICMP
```

Translation failure can occur in the following cases:

- The device ran out of mappable IP addresses, TCP/UDP port resources, or other internal resources such as memory.
- The session direction was not permitted, even though there was a match for the address in the map.

9. To set up a dynamic address map for IP addresses assigned per session, enter:

```
SetD !<portlist> -NAT IPCPAddressMap = < Enable > [ LHS Address(es)] [Log {0...7}]
```

When this address map is enabled, the port control must be toggled. The default address (0.0.0.0.0./0) is used if not specified otherwise.

11.3

Address Mapping Examples

One-to-One Address Mapping

The following command establishes a one-to-one map for BiDirectional connections on port 1. The LHS address can be inside or outside.

```
ADD !1 -NAT AddressMap 144.195.48.20 144.195.40.17 BiDirectional
```

Many-to-One Address Mapping

The following command establishes a many-to-one map for OutBound connections on port 1. The LHS addresses are inside, and the RHS address is outside.

```
ADD !1 -NAT AddressMap 192.168.0.0/16 144.195.18.4 OutBound
```

Many-to-Many Address Mapping

The following command establishes many-to-many map for OutBound connections on port 1. The command maps four inside addresses to 16 outside addresses.

```
ADD !1 -NAT AddressMap 192.168.0.1-192.168.0.3, 192.168.10.17 144.195.40.0/28 OutBound
```

One-to-Many Address Mapping

The following command establishes a one-to-many map for LoadShare connections on port 1. The LHS address is outside, and the RHS addresses are inside.

```
ADD !1 -NAT AddressMap 144.195.18.4 192.168.0.0/16 LoadShare
```

Entering Address Maps Directly

You can also enter address maps directly into the natmap file. The natmap file in your configuration file directory is in ASCII text format that can be edited with any text editor.

If you choose to manually edit the natmap file, you must enter the NATReStart command or reboot the device to load the maps listed in the file.

The NATReStart command examines the natmap file for syntax errors. If a syntax error is encountered in the natmap file, processing stops, an error message is displayed, and no further maps are initialized.

11.4

Basic NAT Operation

A device that uses Network Address Translation (NAT) has at least one port connected to the inside network and one port connected to the outside network. Enable NAT on the port that is directly connected to the outside network. If you have multiple connections to outside networks, enable NAT on each outside port. You must have separate address maps for each NAT port. The inside network is always the network that contains the address you need to map.

When configuring address maps by using the `ADD !<port> -NAT AddressMap` command, you specify the source address (left-hand side (LHS) address) and the translated address (right-hand side (RHS) address). You also specify the direction of the translation: outbound, inbound, bidirectional, or load sharing.


The device identifies the first packet and the direction of a TCP session and creates a NAT session if a map exists for that direction. However, the device cannot identify the first packet of a UDP session, so every UDP packet is considered the first packet of a session, which results in a new NAT session for every UDP packet.

11.5


Traffic Directions

Depending on which direction you specify, the LHS addresses can be either the inside or the outside addresses.

When you specify outbound, the inside addresses are the LHS addresses, which are translated into the outside RHS addresses. Outbound traffic is defined as traffic that leaves the device through the NAT port, as illustrated in the following figure.


Figure 70: Outbound Traffic

When you specify inbound or load sharing, the LHS addresses are the outside addresses, which are translated into the inside RHS addresses. Inbound or load sharing traffic is defined as traffic that enters the device through the NAT port, as illustrated in the following figure.

Figure 71: Inbound or Load Sharing Traffic

When you specify bidirectional, you can use either the LHS or RHS address for the inside address, as illustrated in the following figure.

Figure 72: Bidirectional Traffic



11.6 Address Mapping

Static mapping

Establishes a one-to-one map between two addresses. Static mapping is useful when a host on the inside network must be reachable by a specific IP address, for example, a web server. Use static mapping with outbound, inbound, or bidirectional connections. Bidirectional translation can use only static mapping.

Dynamic mapping

Establishes a many-to-one, a many-to-many, or a one-to-many map. Many-to-one map, used by outbound or inbound connections, translates multiple addresses into a single address. In an outbound connection, multiple inside hosts can connect using the same outside address because the device appends a TCP or UDP port number to the outside address for each connection, as shown in the following table.

Table 24: Dynamic Many-to-One Mapping

Inside Address	Outside Address and Port
10.0.0.5	144.195.23.10:1031
10.0.0.18	144.195.23.10:1032
10.0.0.25	144.195.23.10:1033

In an inbound connection, multiple outside addresses are translated into a single inside address.

A many-to-many map, used by outbound or load sharing connections, translates multiple addresses by assigning addresses from a pool. Each subsequent connection is assigned the next available address from the pool. If the device cannot assign an address because it ran out of addresses in the pool, it drops the packet. Make sure you have enough addresses in the address range if you want to use a many-to-many map.

A one-to-many map, used by load sharing or outbound connections, translates a single address to one of many addresses. In a load sharing connection, each subsequent connection to the outside address is translated into the next available address from the inside address pool. In an outbound connection, each subsequent connection from a single inside host is assigned a new IP address, making it appear as though a single user is multiple users.

NAT Proxy ARP

NAT Proxy ARP is a mechanism that allows you to map addresses for which no actual device exists. NAT address mapping involves mapping addresses in one IP address range to addresses in another IP address range. The translated address could be either one of the addresses assigned to the device interface or it could be a different address. If the translated IP address is different from the directly connected interface address, it is possible that the NAT session setup could fail because no device exists to respond to the initial ARP request.

To allow such configurations, the NAT proxy ARP process checks each NAT mapping entry and adds a static entry into the ARP table. Before the entry is added the following conditions are checked:

- Whether the added IP address does not belong to the interface's own address.
- Whether the added IP address is a directly connected address for that port.

The device then responds to all the ARP requests for which it finds a match in its proxy table with the target IP address in the ARP request. After the ARP request is resolved, the normal NAT session proceeds.

To view the content of the proxy ARP table, you must use:

```
SHow -IP ProxyAddress
```

Mask

When you use the `ADD ! <port> -NAT AddressMap` command, you can specify an address block using a mask: `<IPaddr/mask>`. `<mask>` is a number in the range of 0-32, which indicates the number of bits in the IP address that remain unchanged for the IP addresses in that block. The remaining bits in the IP address should all be 0s. The address block includes all addresses except for the first address and the last (x.x.x.255) address.

Table 25: Mask Examples

Address Block	Description
144.195.0.0/16	All the addresses in the range from 144.195.0.1 to 144.195.255.254.
144.195.1.2/32	The host itself 144.195.1.2.
0.0.0.0/0	All the IP addresses in your network.
224.0.0.0/4	All the class D multicast addresses, from 224.0.0.1 through 239.255.255.254.

11.7


Mapping the IPCP Address

IPCP Address Mapping is useful if you have Internet connectivity through an Internet Service Provider (ISP), and the ISP provides IP addresses dynamically (that is the ISP provides different IP addresses every time you dial in).

If you have only one PC, there is no problem. If you have more than one device or a private network as shown in the following figure, then you must change your NAT mapping every time you dial into the ISP.

You can avoid having to change your NAT mapping every time by using the IPCP Address Mapping feature.

Figure 73: IPCP Address Mapping Network Diagram



Procedure:

1. Set the IP address to IPCPAddress on port 2 by entering:

```
SETDefault !2 -IP NETAddr = IPCPAddress
```

The device gets the IP address from the ISP through IPCP negotiation and assigns it to port 2.

2. Enable NAT on port 2 by entering:

```
SETDefault !2 -NAT CONTrol = Enable
```

3. Set port 2 as the default gateway by entering:

```
ADD -IP ROute 0.0.0.0 !2 0
```

The device forwards all IP traffic for unknown networks to the Internet.

4. Enable NAT IPCPAddress Mapping on port 2by entering:

```
SETDefault !2 -NAT IPCPAddressMap = Enable
```

Result:

The LHS is set to the default value of 0.0.0.0/0, and the log level is set to Log6. The direction of the mapping is always set to OutBound. The device creates a NAT Address Mapping on port 2 with the IP address received from the ISP as the RHS address.

After dialing into the ISP you receive the IP address of 123.4.5.6. The IP address for port 2 is assigned as 123.4.5.6 and a NAT mapping is created on port 2. The NAT mapping is presented as demonstrated on the following example:

```
LHS Address      : 0.0.0.0/0
RHS Address      : 123.4.5.6
Direction       : OutBound
Log              : Log6
```

All the devices on the private network are able to access the Internet. After the connection to the ISP is disconnected, the existing NAT sessions remain. These sessions are removed only when the sessions time out or are deleted. When the connection is restored, the existing NAT sessions are either deleted or not deleted depending on the IP address obtained from the ISP. If you get the same address as before, in this case 123.4.5.6, then the NAT sessions and the NAT Address Mapping are not removed. Otherwise, all the sessions and the mapping are removed and a new NAT Address Mapping is added.

11.8

TCP/UDP Port Mapping

Mapping an address and TCP/UDP port to another address and TCP/UDP port allows more control of the type of traffic NAT translates.

Most TCP and UDP servers use the same port number range, 0-1023, to listen for incoming connections. In addition, most servers use a fixed, well-known port number for listening to a particular service.

The following table lists the major services and their port numbers. For a detailed list of reserved services and port numbers, see RFC 1700.

Table 26: TCP/UDP Port Numbers and Services

Service	TCP/UDP Port Numbers
DNS	53
finger	79
FTP	20, 21
Gopher	70
HTTP	80
NNTP	119
NTP	123
POP	109, 110
RIP	520
SMTP	25
SNMP	161, 162
syslog	514
talk	517, 518
Telnet	23
TFTP	69
UUCP	9540
WAIS	210
whois	43



NOTE: TCP/UDP clients use the port number range 1024-65535.

11.9

Configuring Private Address Space

NAT Example 1


The Human Resources (HR) network is an isolated network using the address block 10.0.0.0/8, which is a recommended private address block not advertised outside the domain.

You can use the information from the following figure to:

- Enable any host in the HR network to connect to the corporate network.

- Allow access to an HR server from the corporate network.

Figure 74: Private Network Connecting to the Corporate Network




NAT Example 2

The Human Resources (HR) network is an isolated network using the address block 10.0.0.0/8, which is a recommended private address block that is not advertised outside the domain.

You can use the information from the following figure to:

- Enable any host in the HR network to connect to the corporate network.
- Allow access to an HR server from the corporate network.
- Limit access to the corporate network to only one server.

**Procedure:**

1. If you want to configure Private Address Space for NAT Example 1, perform the following actions:
 - a. Enable NAT on port 1 of Device A by entering:

```
SETDefault !1 -NAT CONTROL = Enable
```

The NAT port must be connected to the outside network.
 - b. To access the corporate network from the HR network, map the HR address block 10.0.0.0/8 to a single outside address, 144.195.23.10, and specify outbound by entering:

```
ADD !1 -NAT AddressMap 10.0.0.0/8 144.195.23.10 OutBound
```

Any address from the HR network is translated into the outside address 144.195.23.10.
 - c. To make the HR server address 10.0.0.33 available to the corporate network, map it to the outside address 144.195.23.11, and specify bidirectional by entering:

```
ADD !1 -NAT 10.0.0.33 144.195.23.11 BiDirectional
```

The bidirectional option allows the server to access the corporate network as well.
2. If you want to configure Private Address Space for NAT Example 2, perform the following actions:
 - a. Enable NAT on port 1 of Device A by entering:

```
SETDefault !1 -NAT CONTROL = Enable
```

This NAT port is connected to the outside network with regard to inside Network A.

- b. To access the corporate network from the HR network, map the HR address block 10.0.0.0/8 to a single outside address, 144.195.23.10, and specify outbound by entering:

```
ADD !1 -NAT AddressMap 10.0.0.0/8 144.195.23.10 OutBound
```

Any address from the HR network is translated into the outside address 144.195.23.10.

- c. To make the HR server address 10.0.0.33 available to the corporate network, map it to the outside address 144.195.23.11, and specify BiDirectional by entering:

```
ADD !1 -NAT 10.0.0.33 144.195.23.11 BiDirectional
```

The BiDirectional option allows the server to access the corporate network as well.

- d. Enable NAT on port 2 of Device A by entering:

```
SETDefault !2 -NAT CONTROL = Enable
```

This NAT port is connected to the outside network with regard to inside Network B.

- e. To limit access from the HR network to only server 144.195.0.10, map all destination addresses in the corporate network to the server address, and specify InBound by entering:

```
ADD !2 -NAT 144.195.0.0/16 144.195.0.10 InBound
```


11.10

Configuring Load Sharing

The Motorola Solutions web server is replicated on five different servers, as illustrated in the following figure. The URL www.mot.com is accessed thousands of times a day. The domain name server (DNS) advertises the address 192.156.136.22 for www.mot.com, even though there is not an actual server at that address.

You can use this procedure to forward traffic directed to the virtual server 192.156.136.22 evenly to five web servers.

Figure 75: Load Sharing



Load sharing only translates inbound connections. No translation is required for outbound sessions originating from the servers.

The EOS device does not detect if a load sharing host is down. However, a log message indicating the possibility that a host may be down is logged. If a host is down, and you want to remove the host from the list, you can delete the original load sharing map (this causes all active sessions to be flushed) and create a new load sharing map without the host.

Procedure:

1. Enable NAT on port 1 of the device by entering:

```
SETDefault !1 -NAT CONTROL = Enable
```

The NAT port must be connected to the outside network.

2. To distribute traffic evenly to the five web servers, map the virtual server address 192.156.136.22 to all five server addresses, and specify load sharing by entering:

```
ADD !1 -NAT AddressMap 192.156.136.22 192.156.136.1, 192.156.136.5, 192.156.136.8,  
192.156.136.33, 192.156.136.61 LoadShare
```


11.11

Configuring Address Migration

Your company changed ISPs and the new ISP must reassign your IP addresses to work with their network. Reconfiguring every host with a new address requires extensive effort, time, and interruption to users. Moreover, if you change service providers frequently, the process have to be repeated every time.

You can use this procedure to translate each old IP address into a new address.

Figure 76: Address Migration



If you do not need bidirectional access, you can map a range of addresses to a new range using the OutBound option.

Procedure:

1. Enable NAT on port 1 of Device A by entering:

```
SETDefault !1 -NAT CONTROL = Enable
```

The NAT port must be connected to the outside network.

2. Map each old address to a new address separately, and specify bidirectional by entering the following commands:

```
ADD !1 -NAT AddressMap 115.2.3.1 140.45.67.4 BiDirectional
```

```
ADD !1 -NAT AddressMap 115.2.3.2 140.45.67.5 BiDirectional
```

```
ADD !1 -NAT AddressMap 115.2.3.3 140.45.67.6 BiDirectional
```


11.12

Configuring Address Redirection

The new, fast FTP server butterfly.isp.com was bought to replace the old FTP server, tarantula.isp.com.

You can use this procedure to transparently redirect traffic from tarantula to butterfly.

Figure 77: Address Redirection



Procedure:

1. Enable NAT on port 1 of Device A by entering:

```
SETDefault !1 -NAT CONTROL = Enable
```

The NAT port must be connected to the outside network.

2. Map tarantula FTP server (144.195.48.20) to butterfly FTP server (144.195.40.17), and specify inbound by entering:

```
ADD !1 -NAT AddressMap 144.195.48.20 144.195.40.17 InBound
```

11.13

NAT Redundancy

When Network Address Translation (NAT) registers as a client of the replication protocol, TCP connections are replicated between the border routers (home agents) in a conventional data redundancy (HA and DSR) configuration. Replication of TCP connections enables the border routers to share the NAT sessions table and avoid dropping packets in TCP traffic streams since both border routers may be active at a given time. For more information on NAT redundancy, see [Replication Protocol on page 595](#).

11.14

Limitation on One-to-One Mapping

If you want to map any TCP/UDP services whose port number is less than 1024, you must use the TcpUdpPort Mapping. This is because there is a limitation on the one-to-one InBound address mapping for Network Address Translation (NAT). If the LHS address is the same as the IP address of the interface, any connection to TCP/UDP services whose port number is less than 1024 it is not translated. NAT does not translate these packets; rather it passes them through to the EOS device.

For example:

An InBound NAT Address Mapping for port 1:

```
LHS Address      : 123.4.5.6
RHS Address      : 10.0.0.1
Direction        : InBound
IP Address on port 1: 123.4.5.6
```

If the following IP packet is received on port 1, NAT does not translate this packet.

```
IP source address : 123.4.5.7
IP destination address : 123.4.5.6
TCP source port#   : 2033
TCP destination port# : 23
```

This packet is received by the EOS device, and as a result the client Telnets into the device and not to 10.0.0.1. To solve this, you must define the following TcpUdpPortMap:

```
LHS Address      : 123.4.5.6,23
RHS Address      : 10.0.0.1,23
Direction        : InBound
```

Chapter 12

IP Multicast Routing Configuration

IP Multicast is a network service in which a single IP packet can be received by multiple hosts. Using this service, an application can send one copy of each packet and address it to a group of computers that want to receive it. IP Multicast is a more efficient way to use network resources, especially for bandwidth-intensive services such as audio and video applications.

For example, using regular IP (unicast) to send an audio and video conference lecture from one computer to multiple hosts would result in a set of individual packet streams, one each from the originating computer to the receiving computers. This large amount of network traffic could overload network segments and packet switches.

If IP multicast packets are used instead, then the transmitting computer generates only a single stream of multicast packets that is sent to a multicast session address. The multicast routing on campus directs this packet stream to only those subnets on campus where some computer(s) indicated an interest in receiving the multicast session.

All receiving computers program their network interface to listen for packets sent to the specific multicast session address that is associated with the conference lecture. Even if there are many listeners on a single Ethernet LAN, only one multicast packet stream is transmitted on the LAN for that session address. All receiving computers listening to that multicast session receive the same packet stream, which can result in a large reduction in the amount of network traffic that must be sent to the receiving LAN.

Membership in a multicast group is dynamic. A host may join or leave a group at any time. A host may be a member of an arbitrary number of multicast groups; group members can span multiple subnets. A host may send datagrams to a multicast group without being a member.

Each multicast group has a unique multicast (Class D) address. Some multicast addresses are assigned by the Internet Addressing and Naming Authority (IANA) and correspond to groups that always exist even if they have no current members. Such addresses are said to be well-known. Typically, packets transmitted to these addresses use a TTL of 1. Other multicast addresses are available for temporary use. They correspond to transient multicast groups that are created when needed and discarded when the membership reaches zero. For more information, see [Multicast Addresses on page 260](#).

Special gateways forward multicast packets, but hosts do not need to explicitly know about these devices. It is the responsibility of the multicast routing device to receive the multicast packet from the host and correctly forward it to those members of the group.

Hosts and routers must run the IGMP Protocol for multicast connectivity. In addition, the router must run one or more of the following multicasting routing protocols:

- DVMRP
- MOSPF Version 2
- Core-Based Trees (CBT)
- Protocol Independent Multicast (PIM) (Sparse and/or Dense Mode)

The EOS software includes the DVMRP, MOSPF, and PIM-SM routing protocols, which are user configurable, and the IGMP Protocol, which requires no configuration.

For more information about DVMRP and MOSPF, see [Distance Vector Multicast Routing Protocol on page 262](#) and [Multicast Open Shortest Path First Protocol on page 263](#).

For more information about PIM-SM see [Protocol Independent Multicast-Sparse Mode \(PIM\) Configuration on page 290](#).

For more information about IGMP, see [Internet Group Management Protocol on page 261](#).

Depending on your network requirements, you can use the default values of the parameters in the various services, or you can further configure the device according to later sections in this chapter.

To configure a device as an IP multicast router, you must set parameters in the MIP Service. You must also set parameters in the DVMRP Service if your network uses the Distance Vector Multicast Routing Protocol (DVMRP), or in the MOSPF Service if your network uses the Multicast Open Shortest Path First (MOSPF) routing protocol.

12.1

Multicast Routing Terms

The following terms are used in this chapter to explain multicast routing:

Child link

The set of interfaces on which to forward multicast packets.

Core-Based Trees (CBT)

With this protocol, all the group members share a multicast delivery tree. CBT builds a multicast tree based on a core router instead of the source. It builds one multicast tree per group instead of one per (source, group) pair.

Designated router

The device that is responsible for sending IGMP Host Membership Query datagrams. For DVMRP devices, the designated router is the device with the lowest IP address on the subnet. For MOSPF devices, the designated router is the device with the highest priority. For PIM-SM, the designated router is the device with the highest DR priority or, in the case of a tie, the device with the highest IP address.

Distance Vector Multicast Routing Protocol (DVMRP)

This distance-vector protocol builds a shortest-path source-based delivery tree between each source (sender) and multicast group (receivers). It builds one multicast delivery tree per (source, group) pair.

Dominant router

One of several devices on a link that is elected for a particular source. It is responsible for forwarding multicast datagrams on a subnet for a source.

Downstream router

The device to which a multicast packet is forwarded.

Forwarding table

A table containing group information (source and group pairs). This group information is applied to the routing table's shortest-path tree and helps the device forward multicast datagrams to each member of the group.

Leaf link

A child link that no device uses to reach a source.

Multicast

A technique that allows copies of a single packet to be passed to a selected subset of all possible destinations.

Multicast Open Shortest Path First (MOSPF)

This protocol is an extension to OSPF that builds a shortest-path source-based delivery tree on demand.

Parent link

The expected interface that receives packets from a source; also the interface that leads to the previous-hop router back to the source.

Protocol Independent Multicast (PIM)

This protocol is not an extension of any unicast routing protocol; it relies on the unicast routing protocols and is suited for large heterogeneous internetworks. It contains two modes: Dense and Sparse. Dense mode uses a similar algorithm to the one used by DVMRP.

Reverse Path Broadcasting (RPB)

A refinement of the RPF algorithm that eliminates duplicate broadcast packets.

Reverse Path Forwarding (RPF)

An algorithm that forwards multicast datagrams by computing the shortest (reverse) path tree from the source network to all possible recipients. The device forwards the packet further only if it considers the sending device as the next-hop address of the source multicast.

Reverse Path Multicasting (RPM)

A refinement of the TRPB algorithm that provides on-demand pruning of the shortest-path multicast tree.

Routing table

A table containing a list of routes that are learned from other multicast devices' route report messages.

Truncated Reverse Path Broadcasting (TRPB)

A refinement of the RPB algorithm that only forwards packets to where they are wanted by pruning the shortest-path tree. This algorithm prunes branches of nonmember leaf networks.

Upstream router

The device from which a multicast packet is received.

12.2

MBONE Connectivity with Multicasting

The MBONE is a virtual network running on top of the Internet that is composed of a cooperative set of workstations and devices with multicast capability. The MBONE is in existence since 1992, primarily as a research and collaboration tool using multimedia applications. It was greatly expanded from the original Internet Engineering Task Force (IETF) video and audio multicasts, and now includes 24-hour world news audio sessions and NASA space missions, which use real-time audio and video transmissions.

The Motorola Solutions implementation of IP multicasting allows you to access the MBONE across the Internet using tunneling and provides the following additional capabilities:

- Obtain audio and video transmissions using your existing infrastructure, over Ethernet and on any media over which Motorola Solutions supports IP routing.
- Enable the development of entirely new classes of IP-based applications.
- Ease the migration of existing LAN-based multicast applications and distributed systems to an IP-based environment.
- Conserve bandwidth by reducing traffic and protect the host from receiving unwanted datagrams (only members of the group receive the multicast packet).
- Extend the benefits of multicast delivery beyond the confines of a single subnetwork as more multicast-capable IP devices are used.
- Experience complete compatibility with the UNIX program, mrouterd 3.5 and above (less compatibility with previous releases), the UNIX program implementing DVMRP that runs on most systems on the MBONE.

12.3

Multicast Addresses

IP multicasting uses the destination address of the datagram to specify multicast delivery using Class D addresses in the range of 224.0.0.0 through 239.255.255.255.

The following Class D addresses are reserved:

- 224.0.0.0 – This address cannot be assigned to any group.

- 224.0.0.1 – This address is permanently assigned to the “all hosts” group, which includes all hosts and gateway participating in IP multicasting on a local network. No IP multicast address exists that refers to all hosts in the Internet.
- 224.0.0.2 – This address is assigned to all routers on a local network.
- 224.0.0.4 – This address is assigned to DVMRP routers on a local network.
- 224.0.0.5 – This address is assigned to all OSPF routers on a local network.
- 224.0.0.6 – This address is assigned to all OSPF designated routers and backup designated routers on a local network.
- 224.0.0.13 – This address is assigned to all PIM-SM routers on a local network.
- 224.0.0.0 to 224.0.0.255 – These addresses are reserved for multicast applications that do not multicast more than one hop. Multicast packets addressed to these addresses are not forwarded outside the local network.
- 239.0.0.0 to 239.255.255.255 – These addresses are reserved for scoping purposes (a device is configured as a boundary router and multicast traffic does not cross the boundary) and for private multicast groups (traffic is not routed across the Internet).

IP multicast addresses can only be used as the destination address; they can never appear in the source address field of a datagram, nor can they appear in a record route option. For more information about IP addressing, see [Internet Addressing on page 822](#).

12.4

Internet Group Management Protocol

To participate in IP multicasting, multicast hosts and routers must have the IGMP operating. This protocol is the group membership protocol used by hosts to inform routers of the existence of members on their directly connected networks, and allows them to send and receive multicast datagrams.

Multicast routers learn about group membership when a host joining a new group sends an IGMP message to the group address declaring its membership. If the DVMRP Protocol is running, the local multicast router receives the group membership message and sends a DVMRP Graft message to its upstream router if it ever sent a DVMRP Prune message. If the MOSPF Protocol is running, the local multicast router receives the group membership message, establishes routes, and propagates the group membership information to other multicast routers throughout the internetwork.

Because membership is dynamic, local multicast routers periodically query hosts on the local network with Host Membership Query messages to determine which hosts remain members of which groups. These messages are periodically sent by the designated router (the one with the lowest IP address in DVMRP or the one with the highest router priority in MOSPF) to refresh their knowledge of membership present on a particular subnet.

Hosts respond with Host Membership Report messages. If no host reports membership in a group after a query, the multicast router assumes that no host on the network remains in that group. If the DVMRP is running, the router sends a Prune message to its upstream router for the next data packet destined to this group and assumes that no other downstream routers are interested in this group. If the MOSPF Protocol is running, the EOS device stops advertising group membership to other multicast routers. Hosts can also send Host Leaves Group messages whenever they want to leave a multicast group.

The information learned by the IGMP is stored in a local group membership database and is used by both the DVMRP and MOSPF Protocols.

12.5

Distance Vector Multicast Routing Protocol

To propagate routing information among multicast routers, you can configure a multicast routing protocol such as DVMRP on an EOS device. Multicast routers use DVMRP to pass source subnet information among themselves, using the information to establish routes to deliver a copy of the multicast datagram to every subnet containing a member of the multicast group.

Like RIP, DVMRP passes information about known subnets and the cost to route between gateways. For each possible multicast group, the EOS device imposes a routing tree on top of the graph of the physical interconnections. When the device receives a datagram destined for an IP multicast address, it sends a copy of the datagram over the network links that correspond to branches in the routing tree.

The Motorola Solutions implementation of DVMRP applies the Reverse Path Multicasting (RPM) algorithm that allows for the shortest-path multicast tree to be pruned on demand. Pruning preserves bandwidth by removing multicast routers from the tree when no members for that group are on any directly-connected subnets and no downstream routers are interested in that group (multicast packets do not need to be received and are discarded by this router because no group members are attached). DVMRP uses a number of messages to discover neighboring routers, including:

Probe

Discovers neighbors that support multicast routing.

Route Report

Contains route information.

Prune

Destined to the parent router to detach it from the delivery tree if no members for that group are on any directly connected subnets.

Graft

Sent to an upstream router when a new member joins the group after a Prune message had previously been sent.

Graft Acknowledge

Sent to the downstream router to acknowledge the previous Graft message.

Routing Table

Each DVMRP multicast router creates a routing table containing a list of routes learned from other multicast routers' route report messages. Using these route report messages, the router builds a routing table and a shortest-path tree for each source. The DVMRP multicast router also keeps track of the following links:

Parent Link

The expected interface to receive multicast packets from a source (the interface that leads to the previous-hop router back to the source).

Child Link

The set of interfaces on which multicast packets are forwarded for each (source, group) pair. The EOS device uses the child link information to perform Reverse Path Broadcasting (RPB).

Leaf Link

a child link that no router uses to reach a source. For a given source, if no members of a particular group on the subnet are associated with a leaf link, DVMRP truncates the leaf link from the shortest path tree using the Truncated Reverse Path Broadcasting (TRPB) algorithm.

The DVMRP router also assigns the following router functions:

Designated Router

The router with the lowest IP address on a subnet becomes the designated router. The designated router is responsible for sending IGMP Host Membership Query datagrams on the subnet. When a multicast

router starts, it considers itself to be the designated router until it receives a Host Membership Query or Report datagram from a neighbor router with a lower IP address.

Dominant Router

To avoid duplicate multicast datagrams when more than one router exists on a virtual interface, one router is elected as the dominant router for a particular source. The dominant router is the router that is responsible for forwarding multicast datagrams on a subnet for a source (it has a route to the source with the lowest metric on that virtual interface).

Subordinate Router

A subordinate router for a virtual interface is the downstream router that considers this interface to be its parent link. Information from a subordinate router helps the DVMRP router decide whether to truncate the shortest path tree. For each route entry, the subordinate router helps decide if the subnet for that virtual interface is a leaf subnet.

Forwarding Table

In conjunction with the routing table, DVMRP creates a forwarding table. The forwarding table contains group information (source and group pairs) that is applied to the routing table's shortest-path tree. The forwarding table helps the EOS device forward multicast datagrams to each member of the group using the routing table's shortest-path tree.

The DVMRP router can receive Prune messages from downstream routers in the shortest-path tree if the attached subnet contains no group members for the particular (source, group) pair. In this way, the EOS device can prune the shortest-path delivery tree, allowing datagrams to be forwarded only to the subnets in which the specified group is located.

The DVMRP leaf router also prunes the shortest-path delivery tree if it no longer receives IGMP Host Membership Report messages or if all members left a group. The forwarding table maintains an entry in its cache until the timeout period is reached. During the timeout period, if the DVMRP router learns that members rejoined a group, it sends a Graft message to the upstream routers indicating that a member rejoined and allows the branches of the shortest-path tree to reattach.

12.6

Multicast Open Shortest Path First Protocol

To propagate routing information among multicast routers, you can use a multicast routing protocol such as MOSPF. MOSPF is an extension of the base version 2 OSPF Protocol and is backward compatible with OSPF (devices running OSPF interoperate with MOSPF devices). The introduction of multicast extensions does not impact unicast IP traffic. MOSPF devices identify other MOSPF-capable devices for forwarding multicast IP packets. Unlike DVMRP, where separate routing protocols for unicast and multicast packets are run, OSPF and MOSPF run a single copy of the protocol. But like DVMRP, MOSPF forwards multicast traffic based on both the source and destination address; this process is known as *source and destination routing*.

The MOSPF Protocol does not provide the ability to tunnel through non-MOSPF capable devices. MOSPF devices must be directly interconnected with each other. If you do not set up direct connections between MOSPF devices, multicast packets may not be delivered, even though unicast connectivity is maintained.

While forwarding multicast packets, MOSPF may replicate packets along the way. The replication is performed only at tree branches where replication is absolutely necessary. Although multiple copies may be forwarded, the packet is not modified (except the TTL field, where it is decremented by 1 at each hop). No IP-over-IP encapsulation is performed. The destination address is always listed as Class D multicast address. To avoid packet duplicates, equal-cost multiple path forwarding in MOSPF is not possible.

When sending multicast IP packets, MOSPF conforms to link-layer encapsulation. Over Ethernet interfaces, the mapping between IP multicast and datalink multicast address is used. Over other kinds of LAN interfaces, link-level multicast or broadcast is used. Over WAN media, IP multicast packets are encapsulated as unicast packets.

OSPF partitions the network topology into a number of routing domains, with ASBRs interconnecting routing domains. Within a routing domain, OSPF allows multiple areas to be interconnected by ABRs. Areas may be transit, stub, or backbone. MOSPF partitions the network topology in the same way as OSPF; you can use the same topology for both OSPF and MOSPF. For more information on the OSPF topology, see [IP Routing Configuration on page 119](#).

12.6.1

Group Membership Learning

MOSPF uses the IGMP Protocol to monitor multicast group membership on directly attached LANs. MOSPF periodically sends IGMP queries and listens to IGMP replies. The membership information learned is then used to build group-membership link state advertisements (LSAs).

On a LAN, only the designated MOSPF router (usually the device with the highest router priority) sends queries at the interval specified by the `-MIP QueryInterval` parameter to the all hosts address (224.0.0.1) and listens to IGMP replies. The MOSPF designated router (DR) processes IGMP replies and performs the IGMP maintenance work on the network. The DR is responsible for flooding group membership information throughout the routing domain by issuing group membership LSAs. When a new group is learned, MOSPF sends a new group membership LSA. When a group is aged out, MOSPF flushes the corresponding group membership LSA. When the MOSPF router resigns as the DR, it flushes all locally generated group membership LSAs.



NOTE: In a mixed environment in which MOSPF and OSPF devices reside on the same LAN, an MOSPF device must become the DR to monitor group membership, generate group-membership LSAs, and forward multicast packets onto the LAN. Therefore, you should assign OSPF devices a router priority of 0 to prevent them from becoming the DR, thereby allowing an MOSPF device to become the DR.

12.6.2

Shortest Path First Tree

MOSPF uses the group membership LSA with the OSPF database, which provides complete topology information about the area and routing domain. The group-membership LSAs describe the location and address of all multicast groups in an area and routing domain. The group membership database is built by the IGMP Protocol and enables delivery of multicast packets.

MOSPF devices use the group membership LSA information to compute the shortest path first (SPF) tree, which enables delivery of multicast packets to remote destinations. The SPF tree is rooted at the packet's source address toward all destination group members and describes the intermediate hops from the source to all possible destinations belonging to the same group. Different sources are likely to have different trees. The SPF tree is pruned only toward the intended destination; all paths and devices that do not lead to group members are pruned from the tree. A separate tree is built for each source and destination pair.

The SPF tree is computed on demand (when a packet is received). A cache entry is created with the source and destination pair, and the upstream node and downstream interface information is recorded. The SPF tree is then discarded, freeing all resources along with it. The newly created cache entry is used for forwarding decisions, and the entry is stored in the forwarding database. Future received packets with the same source and destination pair can locate the forwarding decision from the database without resorting to another SPF computation.

12.6.3

Forwarding Cache

Each MOSPF device in the path of a multicast packet selects elements for forwarding based on the contents of its forwarding cache. The forwarding cache consists of the local group database and the SPF tree. Each cache entry contains information about received multicast packets from the neighboring node (upstream router or LAN) and where multicast packets should be forwarded (downstream interfaces or

MOSPF neighbors). Each downstream interface has a time-to-live (TTL) value associated with it. The TTL value indicates the number of hops a datagram can travel to reach the nearest multicast destination or be discarded. The hop count prevents packets from redundant forwarding and conserves bandwidth. You can restrict the hop count further using the `-MIP THreshold` parameter.

The cached information is not aged or periodically refreshed; the information is kept as long as enough system resource are available, or until the next topology change. However, the forwarding cache may need to be flushed under the following circumstances:

- OSPF topology changes
- Group membership LSA changes with identical multicast destination
- Local group database changes with identical multicast destination

12.6.4

Interarea Multicasting

When multicast routing occurs between areas (interarea multicasting), source and destination addresses may not reside in the same area, the ABR must have multiple copies of the OSPF link databases (one for each area), and the MOSPF device must build separate SPF trees for each area.

ABRs are responsible for interconnecting areas (transit or stub) to the backbone and other areas. The backbone area is considered a transit area, with area number 0 reserved for it. All ABRs must be connected to the backbone area, either directly or through virtual links. The ABRs are responsible for summarizing reachability information from the backbone to other areas, and from other areas to the backbone. These summaries take the form of summary LSAs.

When running the MOSPF Protocol, you must use the `-MOSPF MABR` parameter to configure a portion of the OSPF ABRs as interarea multicast forwarders, which are ABRs with multicast extensions enabled. An interarea multicast forwarder must be an ABR, but not all ABRs need to be interarea multicast forwarders.

Interarea multicast forwarders calculate all the reachable group addresses from their areas. They convey group membership information to other areas by summarizing the group membership LSAs from their attached areas into the backbone. They do not summarize group membership information from the backbone to other areas. All interarea multicast forwarders concurrently and independently perform this action.

After the EOS device summarizes group membership LSAs into the backbone, the backbone area has complete information regarding all the reachable group memberships. The backbone area may not know the exact location of group members subnets (because that requires the detailed topology information from within the area), but it knows which area is interested in which group address. Non-backbone areas have only group membership information for their area and do not know that some group members exist in other areas.

For multicast packets to flow between areas, all interarea multicast forwarders announce wildcard multicast receiver status (equivalent to the default route for unicast traffic) into attached areas. A wildcard multicast receiver is a device to which all multicast packets should be forwarded regardless of the multicast destination. With sufficient routing information in a backbone area, a wildcard multicast receiver is not needed. Interarea multicast forwarders do not announce the wildcard multicast receiver into the backbone.

Wildcard multicast receiver status is automatic; no user configuration is required. In nonbackbone areas, all interarea multicast forwarders are wildcard multicast receivers. Backbone areas do not require these receivers.

When MOSPF devices are used between areas, they perform one SPF computation for the source and destination per attached area. Each area has its own link state database, and the SPF computation exclusively uses the LSAs within the area. The backbone area is treated the same as other areas.

12.6.5

Interautonomous System Multicasting

When multicast routing occurs between autonomous systems (interautonomous system multicasting), some MOSPF devices must be configured as inter-AS multicast forwarders. These inter-AS multicast forwarders have additional routing information for forwarding multicast packets outside the routing domain. They can concurrently run another inter-AS multicast protocol in the same device or be configured with static external routes. The current implementation does not support static routes. You can configure inter-AS multicast devices through the `-MOSPF Dvmrp` and `PolicyControl` parameters.

The MOSPF Protocol guarantees that all inter-AS multicast forwarders receive all multicast packets. When multicast packets are received from outside the MOSPF domain, MOSPF assumes those packets reach the inter-AS multicast forwarder through a Reverse Path Forwarding algorithm. The DVMRP also uses a Reverse Path Forwarding algorithm.

All inter-AS multicast forwarders declare themselves as wildcard multicast receivers in the backbone area. After reaching the backbone area, all multicast packets are required to reach all inter-AS multicast forwarders regardless of destination.

12.7

Multicast Client

The Multicast Client is a generic EOS component, implemented as part of MIP, that enables EOS to source/receive multicast data packets from any internal EOS application. The Multicast Client sets up the data path for applications to send/receive multicast data packets to and from the EOS IP layer, enabling the Multicast Client to source or receive multicast data packets from internal applications within EOS.

This allows any application to join/leave on a multicast group address and to receive multicast packets on that group address. The application is not aware of the underlying multicast protocol used for sending joins/leaves and for receiving multicast streams.

The CCGW Call Manager and Netcom Manager applications use the Multicast Client. The Multicast Client sets up and builds the multicast streams to be used for analog voice and Netcom traffic transmission.

The Multicast Client maintains an application table that stores the information for each registered application. When an application registers with the Multicast Client, the table is filled in with the following information:

- Application name – given by the application.
- Application ID – assigned by the Multicast Client.
- Services flag – given by the application.
- Forwarding cache entry type – given by the application.
- Port ID – assigned by the Multicast Client
- Receiving routine – given by the application.

The interested applications call an application registration routine to register with the Multicast Client for its services.

If the registration is successful, the Multicast Client:

- Adds an entry to the application table.
- Assigns a unique application ID.
- Returns the application ID back to the application.
- Enables the appropriate multicast routing protocol on the internal port(s) assigned to the application.

If the registration is unsuccessful, the Multicast Client returns an error value.

After an application registered for the Multicast Client's services, it calls a signaling join routine to receive multicast data packets on the specified group address and a signaling leave routine to stop receiving multicast data packets on the specified group address. The Multicast Client signals the configured multicast routing protocol about the joins and leaves received for the specified multicast group.

12.8

Multicast Group Limitation

The multicast group limitation feature was implemented in the EOS MIP and PIM services to limit the number of multicast groups that can be handled by an EOS device. Multicast group limitation can also be configured on device in the core. This feature limits the damage a compromised site can cause to the rest of the network. If a device using this feature is attacked by a hacker, the larger multicast domain can continue to function, even though the device itself may have denial of service.

The multicast group limitation feature limits the number of IGMP or PIM joins that a site can transmit into the core. This prevents the transmitting site from causing devices in the core to allow for a large number of calls, eventually either overstressing the CPU or overusing device memory and bringing PIM down.

The multicast group limitation feature allows you to independently configure the maximum number of entries allowed in the MIP local groups table and the PIM multicast routing table. After the configured number of entries was received, further entries are denied, and a syslog message and trap are generated. For details on the `mnrMultiCastGroupsExceeded` trap, see [Traps Supported on EOS Devices on page 871](#) and the descriptions of syslog messages 1700 and 1701 in [Syslog Messages on page 997](#).

The MIP (IGMP) group limit is configurable on a per-port basis, while the PIM multicast group limit is configurable on a per-device basis.

12.8.1

IGMP (MIP) Multicast Group Limitation

When the IGMP multicast group limitation feature is implemented on an EOS device, the device limits the number of IGMP entries in the MIP local group table. Any IGMP report that causes the table size limit to be met or exceeded is discarded without informing the source. The device accepts and processes IGMP reports for groups already in the MIP local group table regardless of whether or not the threshold was met or exceeded.

MIP Multicast Group Limit Configuration

To configure the MIP multicast group limit with the MIP service `LocalGroupsLim` parameter on a per-port basis, you must use:

```
SETDefault !<portlist> -MIP LocalGroupsLim {<limit> | Off | Default}
```

where `<limit>` is the maximum number of multicast group entries allowed in the MIP local group table (1-65536).

For details, see the description of the `LocalGroupsLim` parameter in "MIP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.



IMPORTANT: To avoid issues with device memory, we recommend a limit no higher than 2000 entries for an MNR S2500 router or GGM 8000 gateway and 15000 entries for an S6000 router.

MIP Multicast Group Limitation Statistics

The `SHOW -SYS STATISTICS -MIP` command displays the number of discarded IGMP packets that were discarded because the multicast group limit was met or exceeded.

MIP (IGMP) Multicast Group Limitation Report

EOS generates and increments an IGMP group limitation report each time the configured MIP (IGMP) multicast group limit is met or exceeded.

You can use the MIP service `LocalGroupsLimRp` parameter to view and flush this report. For details, see the description of the `LocalGroupsLimRp` parameter in “MIP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*

12.8.2

PIM Multicast Group Limitation

When the PIM multicast group limitation feature is implemented on an EOS device, the device limits the number of entries in the PIM multicast routing table. Any packet that causes the table size limit to be met or exceeded is discarded, and further entries are denied. The device accepts and processes messages for groups already in the PIM multicast routing table regardless of whether or not the threshold was met or exceeded.

PIM Multicast Group Limit Configuration

You can configure the PIM multicast group limit with the PIM service `mcastGroupsLim` parameter on a per-device basis by entering:

```
SETDefault -PIM mcastGroupsLim {<limit> | Off | Default}
```

where `<limit>` is the maximum number of multicast group entries allowed in the PIM multicast routing table (1-1048576).

For details, see the description of the `mcastGroupsLim` parameter in the *Enterprise OS Software Version 18.2 Reference Guide*.

To avoid issues with device memory, we recommend that you set the limit no higher than 2000 entries for an MNR S2500 router or GGM 8000 gateway and 15000 entries for an S6000 router.

PIM Multicast Group Limitation Statistics

The `SHow -SYS STATistics -PIM` command displays the number of PIM multicast groups that were denied because the multicast group limit was met or exceeded. The number of IGMP joins rejected by PIM, PIM joins and prunes, registers, and multicast data packets that were discarded.

PIM Multicast Group Limitation Report

EOS generates/increments an PIM group limitation report each time the configured PIM multicast group limit is met or exceeded. Use the PIM service `mcastGroupsLimRp` parameter to view and flush this report. For details, see the description of the `mcastGroupsLimRp` parameter in “PIM Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*

12.9

Configuring IP Multicast Routing for Local Area Networks and Point-to-Point Links


You can use the information in this section to configure basic IP multicast routing over LAN ports and Point-to-Point Protocol (PPP) links. You can enable both DVMRP and MOSPF if you want to perform route exchanges between the two domains.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Your ports and paths are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- You set up OSPF for IP unicast routing, if you plan to use MOSPF as the multicast routing protocol. See [IP Routing Configuration on page 119](#).
- You are familiar with the protocols supported by the EOS device.

Figure 78: Basic IP Multicasting Configuration



Procedure:

1. On the multicast router, assign an IP address to each port that performs IP multicasting by entering:

```
SETDefault !<port> -IP NETaddr = <IP_address> [<subnet_mask> [Ones | Zeros [MTU]]]
| UnNumbered
```

You must assign the IP addresses to LAN ports and to WAN ports using PPP as the serial line protocol. PPP does not require that you assign an IP address to each wide area port. If you do not want to assign an IP address to a wide area port, you must set the value of the `-IP NETaddr` parameter to `UnNumbered`. An advantage of not assigning an IP address to each wide area port is that you conserve valuable network and subnet numbers.

Step example:

To assign an IP address to port 1 and port 2 of the multicast router in Figure 12-1, enter:

```
SETDefault !1 -IP NETaddr = 192.10.10.1 255.255.255.0
SETDefault !2 -IP NETaddr = 129.10.2.1 255.255.0.0
```

2. Enable the MIP service by entering:
SETDefault -MIP CONTROL = Enable
3. Determine which multicast routing protocol you want to use by performing one of the following actions:

If...	Then...
If you want to attach to the Internet Multicast backBONE (MBONE), or if you do not OSPF as the routing protocol,	enable DVMRP on each interface using multicast routing by entering: SETDefault !<port> -DVMRP CONTROL = Enable

If...	Then...
If you run OSPF on your LAN and want multicasting support within an autonomous system,	enable MOSPF on each interface participating in multi-cast routing by entering: SETDefault !<port> -MOSPF CONTrol = Enable
If you want to use a sparse mode multicast routing protocol,	enable PIM-SM on each interface participating in multi-cast routing by entering: SETDefault !<port> -PIM CONTrol = sm PIM-SM requires just any one kind of unicast routing protocol.
If you connect to the MBONE,	perform the following actions: - Enable DVMRP on each interface using multicast routing by entering: SETDefault !<port> -DVMRP CONTrol = Enable - Enable MOSPF on each interface participating in multi-cast routing by entering: SETDefault !<port> -MOSPF CONTrol = Enable

The MOSPF Protocol depends on the OSPF Protocol for proper operation. In order to use MOSPF, you must first ensure OSPF is operating correctly. For more information on how to enable OSPF, see [IP Routing Configuration on page 119](#).

You may need to set additional parameters to complete the configuration for PPP. For more information, see [Wide Area Networking Configuration with PPP on page 341](#).

To configure multicast routing using DVMRP over Frame Relay, see [Configuring DVMRP Multicasting over Frame Relay on page 280](#).

12.10

Verifying the IP Multicast Routing Configuration Overall Status

Procedure:

- Display the parameter settings in the MIP, DVMRP, MOSPF, and PIM-SM Services by entering the following commands:

```
SHoW [!<port>] -MIP CONFIguration
SHoW [!<port>] -DVMRP CONFIguration
SHoW [!<port>] -MOSPF CONFIguration
SHoW [!<port>] -PIM CONFIguration
```
- Verify that the MIP, DVMRP, or MOSPF Services are enabled.
- If you use DVMRP as the multicast routing protocol, verify the entries in the routing table, forwarding table, and the neighboring router table by performing the following actions:
 - Display routing and forwarding table entries by entering the following commands:

```
SHoW -DVMRP RouteTable [<subnet>[/<mask>]] [Long]
SHoW -DVMRP ForwardTable [<subnet>[/<mask>]] [<group>]
```
 - In the routing table, check each source subnet and verify that the status is Up.

- c. In the forwarding table, for multicast datagrams that was sent, ensure that there are entries that correspond to the source (source subnet) and destination (multicast group).

For additional information about the displays, see [Controlling the Routing Table on page 283](#) and [Controlling the Forwarding Table on page 284](#).

- d. Display neighboring router information by entering:

```
SHoW -DVMRP NeighborRouter
```

- e. Verify that the addresses of neighboring routers appear in the list.

For more information about this display, see “DVMRP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

4. If you use MOSPF as the multicast routing protocol, verify the entries in the forwarding table by entering:

```
SHoW -MOSPF ForwardTable
```

For multicast datagrams that were sent, make sure there are entries in the forwarding table that correspond to the source and destination.

The MOSPF forwarding table is built only when the device attempts to forward IP multicast packets. The table shows packets the device recently processed including those successfully forwarded or discarded. The forwarding table varies from device to device because not all devices forwarded multicast packets. devices may periodically flush the forwarding table when topology changes are made.

5. Examine the local group membership table by entering:

```
SHoW [!<port> | !*] -MIP LocalGroups [<Group_addr>]
```

You must verify that group memberships of local hosts are displayed in the table.

If you run MOSPF, only the designated router (DR) and backup designated router (BDR) collect local group information; the table may be empty for non-DR and BDRs.

6. If you use PIM-SM as the multicast routing protocol, check the PIM-SM RP-Set, multicast routing table, and PIM-SM neighbor list by performing the following actions:

- a. Display the bootstrap router and RP Set by entering:

```
SHoW -PIM RPSet
```

- b. Display the multicast routing table by entering:

```
SHoW -PIM MRT
```

- c. Display the PIM-SM neighbor list by entering:

```
SHoW -PIM NR
```

- d. Verify that the RP-Set is not null.

If the RP-Set is Null, then no multicast routing entries can be created.

7. Examine the MIP forward cache by entering:

```
SHoW -MIP ForwardCache
```

If data flows for any multicast stream, this table shows the multicast group, source, and ports associated with that stream.

12.11

Viewing IP Multicasting Statistics

Statistics for DVMRP and MOSPF reflect data that is forwarded, not control messages. Some control statistics are provided in the MIP Service statistics.

You can collect statistics for a specific period by using the `SampleTime` and `STATistics` parameters. For more information on these parameters, see “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. For information on interpreting the statistics, see [Statistics Displays on page 967](#).

Procedure:

To view IP multicasting statistics, enter the following commands:


```
SHoW -SYS STATistics -MIP
SHoW -SYS STATistics -DVMRP
SHoW -SYS STATistics -MOSPF
```

12.12

DVMRP Configuration Troubleshooting

If you are unable to make connections to other networks after setting up the device, this procedure can help correct problems in making connections involving more than one multicast routing device. If the device continues to operate improperly after completing this procedure, contact your network supplier for assistance.

Figure 79: Multicast Routing Topologies Troubleshooting



You can find neighboring multicast-capable routers and retrieve status by using:

```
MRInfo <target_IP> [!<port>] [<timeout (0-120)>]
```

For example, if you are unable to connect to network 128.60.0.0 on Multicast Router 4, you must use:

```
MRInfo 128.50.0.1
```

The `MRInfo` command sends an `AskNeighbors` packet to request neighboring router information. The display provides the addresses of the neighbor, the neighbor's neighbors, the metric, threshold, and status. The status indicates whether the link is down or disabled. It also indicates whether a tunnel link is down, which would prevent Multicast Router 4 from sending packets to Multicast Router 1 or to the host on network 128.60.0.0 that listens to multicast address 224.10.10.10.

If no reply packet is received from 128.50.0.1, you can try entering the `MRInfo 128.40.0.1` command to see if the tunnel between Multicast Router 1 and Multicast Router 2 is up.

You can also discover the multicast tree from a specified receiver to the source by using:

```
MTraceRoute <source> <destination> [G <group>] [H <reports>] [!<port>] [T <timeout>] [W  
<gateway>] [R <Resp_addr>] [L <Resp_ttl>]
```

For example, if you want to see the multicast tree from the Host on network 128.60.0.0 to Multicast Router 4 on network 128.20.0.0, you must use:

```
MTraceRoute 128.60.0.1 128.20.0.1
```

Each device in the tree that receives a multicast trace route packet adds its forwarding information associated with the request to the request packet and forwards the packet to the upstream device. When the request packet reaches Multicast Router 1, Multicast Router 1 sends a multicast trace route response back to Multicast Router 4 because the source address is on one of its subnets. The display shows the route from the source to the destination, including hop count, IP subnet, the multicast routing protocol used, threshold, delay time, and error flags.

For more information about the `MRInfo` and `MTraceRoute` commands, see "Commands" in the *Enterprise OS Software Version 18.2 Reference Guide*

12.13

Multicast Routing Configuration Customization

After you set up and check the basic multicast routing configuration, an EOS device configured as a multicast router begins multicast packet routing among group members.

You can further customize your multicast routing configuration as follows:

- Control local group membership queries.
- Adjust the threshold on multicast datagrams.
- Configure by using the DVMRP Protocol.
 - Configure a multicast tunnel for DVMRP routers separated by a nonmulticast device.
 - Configure scoping (filtering) to prevent traffic from forwarding beyond a boundary router to a set of addresses.
 - Configure multicasting over Frame Relay.
 - Configure a metric.
 - Control the bandwidth (rate limit) allocated for multicast datagram traffic.
 - Configure routing policies.
 - Configure forwarding policies.
 - Configure route aggregation.
 - Control the DVMRP routing and forwarding tables.
- Configure by using the MOSPF Protocol.
 - Configure interarea multicast routing.
 - Configure interautonomous (AS) multicast routing.
 - Configure forwarding policies.
 - Display the forwarding table.

12.13.1

Controlling Local Group Membership Queries

You can control how often Internet Group Management Protocol (IGMP) query messages are sent by the designated IGMP querier to request local group membership information.



NOTE: The designated IGMP querier is the device with the lowest IP address.

Procedure:

To control local group membership queries, enter:

```
SETDefault !<port> -MIP QueryInterval = <seconds>(5–5400)
```

The default setting of this parameter is 120 seconds.

Adjusting the setting of the `QueryInterval` parameter affects the `MembershipExpirationTime`, the length of time a local group membership is valid without confirmation. The `MembershipExpirationTime` is set to two times the value of the `QueryInterval` parameter plus 20 seconds.

By adjusting the `QueryInterval` parameter, you control how long entries remain in the local group membership table.

12.13.2

Configuring IGMP Proxy

EOS supports IGMP proxy (IGMP-based multicast forwarding). When a device is configured as an IGMP proxy agent, it does not participate fully in multicast routing; instead it forwards multicast data solely on the basis of IGMP membership information. A device acting as an IGMP proxy agent has a single upstream interface, where IGMP proxy is configured, and all other interfaces are downstream interfaces. The device performs the host portion of the IGMP protocol on its upstream interface and the routing portion of the IGMP protocol on its downstream interfaces.

Procedure:

To configure IGMP proxy, enter:

```
SETDefault !<port> -MIP IgmpControl = ([Disable|Enable|Proxy] [V2|V1])
```



NOTE: EOS devices running EOS version 11.4 or higher use IGMP version 2, and the `IgmpControl` parameter is set to V2 by default. An IGMP version 2 device may be placed on a subnet where a device on the subnet was not upgraded to IGMP version 2. However, EOS maintains compatibility with IGMP version 1 (you can set the `IgmpControl` parameter to V1) to support hosts and local LANs that are running IGMP version 1.

On the downstream interface, you must set `IgmpControl` to `Enable`

On the upstream interface, you must set `IgmpControl` to `Proxy`

12.13.3

Configuring IGMP Proxy with Source Address Translation


EOS also supports an enhanced implementation of IGMP proxy: IGMP proxy with source address translation (SAT). This implementation enables a device configured as an IGMP proxy agent to forward multicast data packets received on any downstream interface (where multicast routing is enabled) to the interface where the IGMP proxy is configured.



NOTE: IGMP proxy with SAT can be enabled only when IGMP proxy is enabled on any port.

When IGMP proxy with SAT is enabled, the device translates the source IP address of the multicast UDP data packets received on any downstream interface to the IP address of the interface where IGMP proxy is configured. The device then forwards the packet to the upstream (proxy) interface. The device also checks the multicast IP (MIP) proxy group table to see if any other interface received IGMP join information for that multicast group. If so, then the data packets are also forwarded to any downstream interface with subscription information.

Figure 80: IGMP Proxy Configuration Example



The Site Router acts as an IGMP proxy agent for the Site Controller.

On port 1 (!1), the Site Router is an IGMP querier, that sends out IGMP membership queries at a regular interval specified by the MIP service `QueryInterval` parameter.

On virtual port 1 (!V1), the Site Router is an IGMP proxy agent, that sends IGMP membership reports to Core Router 1 (which runs a multicast protocol) in response to the IGMP membership queries.

On virtual port 2 (!V2), IGMP control is disabled; any multicast data packets or IGMP packets received on this port are dropped.

When the Site Router receives a multicast data packet from the Site Controller, the Site Router changes the source IP address to the IP address of its !V1 (proxy) interface, then sends the packet toward Core Router 1. Before the packet is forwarded, the IP and UDP checksum are recalculated since the IP Header field (source IP address) changed. The source address translation is needed because the PIM protocol is running on the core routers. Since the source IP address of the packet is now that of an interface directly connected to Core Router 1, Core Router 1 accepts the packet and sends a Register pack to the Rendezvous Point (RP). If IGMP proxy with SAT was not enabled, the PIM protocol on Core Router 1 would drop the packet.

Procedure:

1. Disable any multicast routing protocol, such as PIM-SM, running on the device.
2. Enable multicast routing by entering:

```
SETDefault -MIP CONTROL = Enable
```

3. Configure the IGMP control on the device by performing the following actions:
 - a. Configure the upstream interface to be the IGMP proxy interface by entering:

```
SETDefault !<port> -MIP IgmpControl = Proxy
```

Step example:

```
SETDefault !V1 -MIP IgmpControl = Proxy
```

- b. Enable IGMP control on the downstream interfaces on which you want received multicast packets to be forwarded to the IGMP proxy interface by entering:

```
SETDefault !<port> -MIP IgmpControl = Enable
```

Step example:

```
SETDefault !1 -MIP IgmpControl = Enable
```

Since IGMP control is not enabled on virtual port 2 (!V2) (the IgmpControl parameter is left at the default Disable value), any multicast data packets or IGMP packets received on virtual port 2 are dropped.

4. Enable IGMP proxy source address translation by entering:

```
SETDefault -MIP ProxySAT = Enable
```

12.13.4

LAN QoS Configuration to Support Geographically Redundant IP Simulcast Prime Sites

In a geographically redundant IP simulcast prime site topology, the prime site is split into two redundant and geographically separate prime sites connected by extended prime site LAN links. LAN QoS is supported across the extended prime site links to help ensure that IGMP packets sent by a CCGW located in one prime site are received by devices in the redundant (and geographically separate) prime site.

EOS implements the following features to support LAN QoS across the extended prime site links:

- When EOS sends out an IGMPv2 Membership Report, the IGMP join packet is replicated and sent multiple times at regular intervals to increase the probability that it is successfully transmitted across the extended prime site LAN link. The MIP service `MultReports` parameter configures how many times the replicated IGMP packets are sent and at what interval. For more information about the `MultReports` parameter, see “MIP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: The current IGMP Proxy implementation sends multiple Proxy Joins, one after another with no gaps. The number of Proxy Joins sent is configured by the MIP service `ProxyReportCnt` parameter. This implementation is unchanged.

- TOS marks IGMP join packets with a high priority value so that they are prioritized over other lower priority packets.

These features reduce the risk of an upstream neighbor in a geographically redundant prime site topology missing a host's request to join a call.

TOS Marking for IGMP Packets

EOS uses the TOS byte in the IP header to define the TOS value. To increase the probability that IGMPv2 Membership Reports are received successfully across an extended LAN link connecting geographically redundant IP simulcast prime sites, EOS marks IGMP packets as described in the following table.

Table 27: IGMP TOS Values by Packet Type


Packet Type	TOS Value	IP Header TOS Byte Value
IGMPv2 Membership Reports	5	0xA0
IGMPv2 Leaves	4	0x80
IGMP Queries (general or group specific)	4	0x80

12.13.5

Multicast Datagram Threshold Adjustment

You can adjust the threshold on an EOS device to prevent multicast packets whose time-to-live (TTL) value is less than threshold from forwarding to the given interface. By adjusting the default value, you can provide scope control and prevent certain multicast datagrams from forwarding out of your network.

Figure 81: Threshold Adjustment



To adjust the threshold value, you must use:

```
SETDefault {!<port> | !<tunnel_ID>} -MIP THreshold = <value> (1-255)
```

By default, this value is set to 1.

For example, on Multicast Router B, you can use:

```
SETDefault !3 -MIP THreshold = 191
```

Suppose the host in China sends a packet to multicast address 224.10.13.2 with an IP TTL of 192. As the packet is forwarded by each multicast device in its path, the IP TTL value is decremented. When the packet reaches Multicast Router B, the IP TTL value is 190. The packet is not forwarded to Multicast Router A because packets with a TTL value less than the configured threshold of 191 are prevented from reaching the United States.

12.13.6

Customization with DVMRP Protocol


You can further customize your multicast routing configuration if you use the DVMRP Protocol as the multicast routing protocol.

12.13.6.1

Configuring a DVMRP Multicast Tunnel

When two DVMRP devices are separated by a nonmulticast device, you can use a multicast tunnel to build a virtual link between the two devices. Packets to be tunneled over the virtual link are encapsulated by IP-over-IP (IP protocol number is set to 4). No group is associated with a tunnel, and the only neighbor on a tunnel is the remote-end device (nonmulticast devices are not considered to be neighbors).

Figure 82: Multicast Tunnel Configuration



Procedure:

1. Create a virtual point-to-point link between the pair of multicast devices by entering:

```
SETDefault !< tunnel_ID> -DVMRP TUnnel = <local-end_IP> <remote-end_IP> [<tTl> (1-255)]
```

Step example:

On Multicast Router A, you can enter:

```
SETDefault !T1 -DVMRP TUnnel = 128.1.10.10 192.1.1.3 3
```

On Multicast Router B, you can enter:

```
SETDefault !T1 -DVMRP TUnnel = 192.1.1.3 128.1.10.10 3
```

You can configure up to 32 tunnels.

The local-end IP address can be any IP address assigned to the system. The remote-end IP address must be unique; you cannot assign tunnels with different local IP addresses and the same remote IP address. The remote-end IP address cannot belong to one of the directly connected subnets if the underlying subnet has broadcast or multicast capability. By default, the TTL is set to 64.

You must set the TTL value to a value greater than or equal to the number of unicast devices in between the multicast devices plus the value of the -MIP Threshold parameter on the remote device interface.

If IP (unicast) routing is not enabled, you must configure a static route for the remote end of the tunnel.

2. Enable DVMRP routing on the specified tunnel interface by entering:


```
SETDefault !< tunnel_ID> -DVMRP CONTrol = Enable
```

12.13.6.2

Configuring DVMRP Scoping

To configure scoping (filtering), you must configure a set of multicast destinations that are not reachable through the boundary router port or tunnel.

Figure 83: DVMRP Scoping



Procedure:

To configure scoping (filtering), enter:

```
ADD {! <port> | ! <tunnel_ID>} -DVMRP BoundaryAddr <IP_addr> [<subnet_mask>]
```

Step example:

To configure Multicast Router B with a boundary address so that packets destined to the group multicast addresses 239.1.0.1 through 239.1.255.1 are dropped, enter:

```
ADD !1 -DVMRP BoundaryAddr 239.1.1.1 255.255.0.255
```


Packets from area 1 destined to the above address ranges do not reach area 2; packets from area 2 destined to one of the blocked address also do not reach area 1.

You can block a single address by not specifying the subnet mask (the default subnet mask is 255.255.255.255).

12.13.6.3

Configuring DVMRP Multicasting over Frame Relay

Figure 84: Multicasting over Frame Relay



Procedure:

1. Set up the Frame Relay Service. See [Wide Area Networking Configuration with Frame Relay on page 308](#).
2. Assign an IP address to each device wide area port that performs IP multicasting by entering:
`SETDefault !<port> -IP NETaddr = <IP_address> [<subnet_mask> [Ones | Zeros [MTU]]]`
3. Add a neighbor address over the Frame Relay network by entering:

```
ADD !<port> -DVMRP NEighbor <FR_DLCI>
```

You must also specify the Frame Relay data link connection identifier (DLCI) address associated with the permanent virtual circuit.

Step example:

Only Multicast Router A and Multicast Router B are participating in multicast routing. To configure the neighbor address, on Multicast Router A specify the DLCI address of Multicast Router B by entering:

```
ADD !3 -DVMRP NEighbor @44
```

On Multicast Router B, enter the same command and specify the DLCI address of Multicast Router A.

4. Enable the DVMRP routing protocol on each wide area port by entering:

```
SETDefault !<port> -DVMRP CONTrol = Enable
```

5. Display neighboring router information by entering:

```
SHow !<port> -DVMRP NEighborRouter [<IP_addr>]
```

If you specify an <IP_addr>, only neighboring router information for this IP address is displayed. For more information about elements in the display, see “DVMRP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

12.13.6.4 Configuring a DVMRP Metric

Procedure:

To configure a metric, or administrative cost on a device, enter:

```
SETDefault {!<port> | !<tunnel_ID>} -DVMRP Metric = <value> (1-31)
```


The default metric is 1.

You may want to adjust the metric if you have multiple routes to the same source and want one route selected over the other. For example, suppose that DVMRP learns about two routes to the same source. Route 1 has an administrative cost of 25; Route 2 has an administrative cost of 3. The DVMRP Protocol selects Route 2 because it is the route with the lowest metric.

12.13.6.5 Controlling the DVMRP Rate Limit for Multicast Traffic

The DVMRP rate limit is the bandwidth measured in kilobits per second.

Figure 85: Controlling the Rate Limit



To control your multicast traffic, you must configure the rate limit if you are connected to the MBONE, which anticipates traffic at a rate of 500 kbps.

When Multicast Router A receives two 125-byte packets, it queues the packets into the ports transmit queues (because of the rate limit settings) instead of immediately forwarding them. Multicast Router A controls packet forwarding as follows:

- After 1 millisecond, port 2 assigns tokens at a rate limit of 2,000 kilobits per second (kbps) (2,000 bits per millisecond or 250 bytes per millisecond). The router extracts both packets from port 2's transmit queue and forwards them to port 2's attached LAN.
- Port 3 assigns tokens at a rate limit of 1,600 kbps (1,600 bits per millisecond or 200 bytes per millisecond). The router can forward only the first 125-byte packet on the attached LAN after 1 millisecond. After the next millisecond, port 3 receives another 200 byte token and can transmit up to 275 bytes (200 - 125 + 200); therefore, the router forwards packet 2 (175 bytes) on port 3's attached LAN.
- The tunnel interface (!T1) assigns tokens at a rate limit of 500 kbps (500 bits per millisecond or 62.5 bytes per millisecond). After 2 milliseconds, the router forwards packet 1 (125 bytes) on the tunnel interface. After 4 milliseconds, the router forwards packet 2 (175 bytes) on the tunnel interface.

Procedure:

To control the rate limit that is allocated for multicast datagram traffic, enter:

```
SETDefault {!<port> | !<tunnel_ID>} -DVMRP RateLimit = <Kbits/second> (0-100000)
```

The default is 0, which means that no limit is applied to the given interface, and the interface uses its full bandwidth.

Step example:

To set the rate limit on Multicast Router A, enter the following commands:


```
SETDefault !2 -DVMRP RateLimit = 2000  
SETDefault !3 -DVMRP RateLimit = 1600  
SETDefault !T1 -DVMRP RateLimit = 500
```

12.13.6.6

Configuring DVMRP Route Aggregation

With DVMRP route aggregation, you can combine the characteristics of several different routes so that a single route can be advertised. By combining several networks into one supernet, you reduce the number of route report messages and the size of the routing table.

Figure 86: DVMRP Route Aggregation



Procedure:

1. Specify a list of networks that DVMRP advertises as a single supernet route by entering:

```
ADD -DVMRP AggregateRange <subnet>/<mask> [<metric>]
```

Step example:

To combine the routes to sites 1, 2, 3, and 4 into a range so that only a single route is advertised on Router A, enter:

```
ADD -DVMRP AggregateRange 192.1.0.0/16
```

2. Specify a list of routes that DVMRP explicitly advertises by entering:

```
ADD -DVMRP AggregateRange <subnet>/<mask> [<metric>]
```

Step example:

If you do not want site 4 included in the aggregation range, enter:

```
ADD -DVMRP AggregateExcept 192.1.5.0/24
```

3. Enable route aggregation and the DVMRP routing protocol by entering:

```
SETDefault !4 -DVMRP CONTROL = (Enable, Aggregate)
```

Router A advertises a single network (192.1.0.0/16) that summarizes each of the three connected sites and also explicitly advertises the exception route (192.1.5.0) to Router B. Without the use of aggregation, Router A advertises each route with a separate entry as shown in the Router C routing table, which grows in size. With route aggregation, the Router B routing table has an entry for 192.1.0.0 and 192.1.5.0.

12.13.6.7

Controlling the Routing Table

You can control how often an EOS device sends route report messages, delete entries in the routing table, and display the routing table.

Procedure:

- To control how often the device sends route report messages containing the complete routing table, enter:

```
SETDefault -DVMRP UpdateTime = <seconds>(5-5400)
```

By default, DVMRP updates the routing table every 60 seconds. By changing this setting, you affect how long a route is considered valid (RouteExpirationTime) and how long a route exists without confirmation (GarbageCollectionTime). The RouteExpirationTime is equal to three times the value of this parameter, and the GarbageCollectionTime is equal to five times the value of this parameter.

By increasing the value of the UpdateTime parameter, you can reduce the amount of route report traffic but you may also increase the size of the routing table. The UpdateTime parameter determines how long a neighbor is considered "up" without confirmation (NeighborExpireTime) and when to consider the associated virtual interface as a leaf link (LeafConfirmationTime). The NeighborExpireTime is set to two times the value of this parameter plus 20 seconds, and the LeafConfirmationTime is set to three times the value of this parameter plus 20 seconds.

- To flush entries in the routing table learned from DVMRP, enter:

```
FLush -DVMRP RouteTable
```

- To display the routing table, enter:

```
SHow -DVMRP RouteTable [ <subnet>[ / <mask>]] [Long]
```

If you specify the <subnet> or <mask> syntax, the routing table for the range of specified subnets is displayed. If you specify Long, the display shows a lists of ports that connect to child subtrees and leaf subnets.

Step example:

To display the following table, enter:

```
SHow -DVMRP RouteTable Long
```

SourceSubnet	SubnetMask	FromGateway	Metric	Status	TTL	InPort	OutPorts
20.0.0.0	255.0.0.0	11.11.11.11	3	Up	200	1	2, 3*, 4, 5
30.0.0.0	255.0.0.0	----	0	Up	--	2	1, 3*, 4, 5
40.0.0.0	255.0.0.0	11.11.11.11	6	GC	100	1	2, 3*, 4

The display consists of the following items:

Table 28: Routing Table Items

Item	Description
SourceSubnet	The original subnet from which the multicast datagram originated.
SubnetMask	The subnet mask of the source subnet.
FromGateway	The previous hop device that leads back to the source. If no gateway is specified, the subnet is directly connected.
Metric	The routing metric of the path back to the source subnet.
Status	The status of this route entry: Up, GC (Garbage-Collect), HD (Hold-Down), and Down.
TTL	The amount of time (in seconds) remaining before an entry is removed from the routing table.
InPort	The incoming port for the multicast datagrams from that source.
OutPorts	A list of ports on which multicast datagrams originated from this source are forwarded. An asterisk (*) indicates that the outgoing port connects to a leaf of the multicast delivery tree rooted at this source.

12.13.6.8

Controlling the Forwarding Table

You can specify how long you want to keep a (source, group) pair in the forwarding table and display the contents of this table.

Procedure:

- To control how long entries remain in the forwarding table, enter:

```
SETDefault -DVMRP CacheTime = <seconds> (300-86400)
```

The default value of this parameter is 300 seconds. You can adjust the setting up to 1 day (86,400 seconds). By adjusting the `CacheTime` parameter, you can control the size of the forwarding table.

- To display entries in the forwarding table, enter:

```
SHoW -DVMRP ForwardTable [ <subnet>[ / <mask>] ] [ <group>]
```

You can display the current table for each (source, group) pair. If you specify only the subnet, all group entries associated with that subnet are displayed. If you specify only the group, all source subnets associated with that group are displayed. If you specify both the subnet and group, only that particular entry is displayed.

Step example:

To display the following table, enter:

```
SHoW -DVMRP ForwardTable
```

SourceSubnet	MulticastGroup	TTL	InPort	OutPorts
20.0.0.0	224.1.1.1	200	1 Pr	2p 3p 4p
	224.2.2.2	100	1	2p 3 4
	224.3.3.3	250	1	2 4b
30.0.0.0	224.1.1.1	300	1	2 3 4
	239.4.4.4	100	1	Sc

The display consists of the following items:

Table 29: Forwarding Table Items

Item	Description	
SourceSubnet	The original subnet of multicast datagrams.	
MulticastGroup	The group address to which multicast datagrams are destined from the origin.	
TTL	Time-to-live indicates how much time (in seconds) are left before removing a source and group entry from the table.	
InPort	Indicates the incoming port for the multicast datagrams from that source.	
	Pr	A Prune message is sent to the upstream device.
	Sc	The multicast group address is configured as a boundary address, and no traffic for this group address is forwarded from that port.
OutPorts	Indicates the ports that multicast datagrams belonging to this group are forwarded.	
	p	The port receives all the Prune messages of the downstream neighboring routers, and no multicast datagrams are forwarded to this port.
	b	The multicast group address is configured as a boundary address, and no traffic for this group address is forwarded to this port.

12.13.7

Customization with MOSPF Protocol

You can further customize your multicast routing configuration if you use the MOSPF Protocol as the multicast routing protocol.


12.13.7.1

Configuring Interarea Multicasting

To perform interarea multicasting when running the MOSPF Protocol, you must configure the Area Border Router (ABR) as an interarea multicast forwarder, which is an ABR with multicast extensions enabled.

As shown in the following figure, the ABR connects two areas to the backbone. You must configure the ABR as an interarea multicast forwarder so that it can summarize group membership information from attached nonbackbone areas into the backbone and to forward multicast packets between areas.

Figure 87: Interarea Multicasting



Procedure:

To allow multicasting between areas, enter:

```
SETDefault -MOSPF MABR = Enable
```

The device must be an OSPF ABR for the MABR parameter to take effect. By default, this parameter is enabled.

For more information, see [Interarea Multicasting on page 265](#).

12.13.7.2

Configuring MOSPF Routing Policies

Using the routing policies supported by MOSPF, you can control the reporting of routes learned from other sources for interautonomous system multicasting.

Procedure:

1. Enable DVMRP/PIM-SM routing information to be advertised into the MOSPF domain by entering:

```
ADD -MIP RoutePolicy from <DVMRP / MOSPF / PIM> to <DVMRP / MOSPF / PIM> <Subnet>/  
<mask> [Aggregate | Individual | Reject] [<metric>] [Type1 | Type2]
```

Supply the subnet and mask of the address range of the DVMRP route to be advertised. The mask value is the number of leading 1s in the mask and ranges from 0 to 32.

The <Subnet>/<mask> describes a range of addresses. For example:

- 10.0.0.0/8 describes all the subnets within network 10.
- 10.1.0.0/16 describes all the subnets within 10.1.0.0.
- 0.0.0.0/0 describes all subnets.

An address can fall into multiple <Subnet>/<mask> ranges. In this situation, the range with the highest mask bits is chosen. The range 0.0.0.0/0 is always the lowest priority.

The **Aggregate** option specifies that MOSPF advertises a single subnet/mask route, which can summarize multiple networks into a single network. The **Individual** option specifies that all individual source subnets are accepted and advertised as learned into the MOSPF domain. The **Reject** option specifies that the specified source network is rejected (not advertised).

You can optionally supply a metric value from 0 to 65,535.

You can select either **Type1** or **Type2**. **Type1** advertises the routes as a type 1 external LSA, which is always preferred over a type 2 external LSA for the same destination.

Step example:

To accept and aggregate routes from 192.10.10.0 advertised as a single route into the MOSPF domain, enter:

```
ADD -MIP RoutePolicy from DVMRP 192.10.10.0/24 Aggregate
```

To accept and advertise all routes learned from 129.213.0.0 sourced from the DVMRP domain, enter:

```
ADD -MIP RoutePolicy from DVMRP 129.213.0.0/16 Individual
```

To reject all other routes, including transmissions from the MBONE, enter:

```
ADD -MIP RoutePolicy from DVMRP 0.0.0.0/0 Reject
```

2. Enable the MOSPF device to perform interautonomous system multicast forwarding by entering:

```
SETDefault -MIP ForwardControl from MOSPF to DVMRP Enable
```

When this command is executed, the MOSPF device declares itself as a wildcard multicast receiver to all its attached areas to attract multicast packets to all destinations. It imports specified routes sourced from DVMRP into the MOSPF routing domain as external LSAs.


This parameter enables the MOSPF domain to accept only DVMRP-sourced multicast packets. For information about how to configure the DVMRP domain to accept MOSPF-sourced multicast packets, see [Customization with DVMRP Protocol on page 277](#). Failure to configure the DVMRP routing policies results in half-duplex communication.

12.13.7.3

Configuring MOSPF Forwarding Policies

By using MOSPF forwarding policies, you can filter destination groups and control data packet forwarding between MOSPF and DVMRP domains.

Figure 88: MOSPF Destination Group Filtering



Procedure:

1. Configure a list of destination group addresses whose data packets are accepted and forwarded, or rejected and dropped, by entering:

```
ADD -MIP ForwardPolicy from <DVMRP | MOSPF | PIM> to <DVMRP | MOSPF | PIM>  
<SubnetPrefix>/<mask> [Accept | Reject]
```

The *<subnet>/<mask>* syntax describes a range of addresses to either be accepted or rejected by MOSPF. For example:

- 239.0.0.0/8 describes all the addresses within network 239.

- 239.1.0.0/16 describes all the addresses within network 239.1.
- 239.1.10.0/24 describes all the addresses within network 239.1.10.

The **Accept** option specifies that the multicast router do the following:

- If the multicast router receives a packet from the DVMRP domain with a destination address that matches this destination group filter, then the multicast router accepts it and forwards it into the MOSPF domain.
- If the multicast router receives a packet from the MOSPF domain with a destination address that matches this destination group filter, then the multicast router accepts it and forwards it into the DVMRP domain.

The **Reject** option specifies that the multicast router do the following:

- If the multicast router receives a packet from the DVMRP domain with a destination address that matches this destination group filter, then the multicast router rejects it and drops the packet and never forwards it into the MOSPF domain.
- If the multicast router receives a packet from the MOSPF domain with a destination address that matches this destination group filter, then the multicast router rejects it and drops the packet and never forwards it into the DVMRP domain.

Step example:

To configure data packets to the destination group 224.1.0.1 to be rejected on the multicast router, enter:

```
ADD -MIP ForwardPolicy from DVMRP to MOSPF 224.1.0.0/16 Reject
```

2. Enable the policy by entering:

```
SETDefault -MIP ForwardControl from <DVMRP | MOSPF | PIM> to <DVMRP | MOSPF | PIM>  
<DestGroup | NoDestGroup>
```

Data packets to destination group addresses 224.1.0.0 to 224.1.255.255 between DVMRP and MOSPF domains are rejected and dropped by the multicast router.

12.14

Configuring Multicast Border Router Functionality

The Multicast Border Router (MBR) functionality provides for the efficient interoperation among independent multicast routing protocols, such as DVMRP, MOSPF, and PIM-SIM. The MBR allows sources and receivers inside multiple autonomous multicast routing domains (each running a different multicast routing protocol) to communicate. These domains must be connected by MBRs where the primary role of the MBR is to forward the traffic from one domain to another.

An MBR consists of two or more active multicast routing interfaces, each running an instance of some multicast routing protocol, such as DVMRP or PIM. The EOS device is configured to forward packets between two or more independent multicast domains. Only one multicast protocol can be active on each interface.

The **ForwardControl** parameter in the **-MIP Service** controls the forwarding of packets between the multicast routing domains on the MBR.

Procedure:

- To control the forwarding between multicast routing domains, enter:

```
SETDefault -MIP ForwardControl from <DVMRP | MOSPF | PIM> to <DVMRP | MOSPF | PIM>  
Enable | Disable
```

When you set the **ForwardControl** parameter to **DestGroup**, data packets are forwarded between the “from” and “to” domains according to the lists established by the **-MIP ForwardFilter** parameter.

Step example:

When a device has DVMRP, MOSPF, and PIM interfaces to forward data packets between the DVMRP and PIM domains, enter:

```
SETDefault -MIP ForwardControl from DVMRP to PIM Enable
```

- To control the routing information exchange between the multiple protocols, enter:

```
Add -MIP RouteFilter <DVMRP / MOSPF / PIM> to <DVMRP / MOSPF / PIM> <GroupPrefix>/  
<mask> [Aggregate | Individual | Reject] [<metric>] [Type 1 | Type 2]
```

Step example:

To accept routes from 192.10.10.0 advertised as a single route into the MOSPF domain from the DVMRP domain, enter:

```
ADD -MIP RouteFileter from DVMRP to MOSPF 192.10.10.0/24 aggregate
```

Chapter 13

Protocol Independent Multicast-Sparse Mode (PIM) Configuration

Multicast routing mechanisms such as Distance Vector Multicast Routing Protocol (DVMRP) and Multicast Open Shortest Path First (MOSPF) send data packets (in the case of DVMRP) or membership reports (in the case of MOSPF) periodically on many links that do not lead to receivers and senders. The periodic broadcasting of information by these protocols is required to identify the location of the interested receivers for a specific multicast session. This approach is useful in networks where bandwidth is plentiful, or when there are large number of senders and receivers for a multicast session. However, when senders and receivers to multicast sessions are distributed sparsely across a wide area, such schemes are not efficient; they lead to wasted bandwidth on expensive WAN links and require the maintenance of “routing-state” on devices that are not on the forwarding tree for the multicast session.

Protocol Independent Multicast-Sparse Mode (PIM-SM) is a multicast routing protocol designed to resolve the above mentioned inadequacies with the “broadcast-and-prune” protocols such as DVMRP and MOSPF. The term “protocol independent” comes from the fact that PIM can work with any unicast routing protocol. DVMRP, on the other hand, maintains its own unicast routing table, and MOSPF depends on the OSPF link-state-database to build source-based trees.

PIM-SM builds a per-group (or per multicast session) shared multicast distribution tree centered at a rendezvous point and requires receivers to explicitly join to this shared distribution tree before they receive data traffic. This approach ensures that bandwidth does not get wasted on links that do not have any downstream receivers. Since a “shared-tree” mechanism could result in suboptimal paths for data traffic from the sender to receivers, PIM-SM also supports the ability to switch to a source-specific tree if the data traffic warrants it.

13.1

Rendezvous Points

Rendezvous points (or RPs) in PIM-SM provide the mechanism for receivers of a multicast group to encounter senders to the same group. Receivers can be detectable through RPs and use them to receive information about senders to the group. PIM-SM builds a per-group (or per-multicast-session) shared multicast distribution tree rooted on the RPs, and requires receivers to explicitly join to this shared distribution tree before they can receive data traffic. Designated routers (DRs) with directly connected sources encapsulate multicast data traffic in PIM control messages and unicast it to the rendezvous points until the rendezvous point signals to the DR to start sending multicast natively.

13.2

Bootstrap Router

To obtain information about the rendezvous points, all devices within a PIM domain (a contiguous set of devices that all implement PIM-SM) collect bootstrap messages. Bootstrap messages contain the list of active rendezvous points and the group prefixes served by each. The domain's bootstrap router (BSR) is responsible for originating the bootstrap messages, which are then forwarded hop-by-hop by PIM devices. A PIM-SM domain must have at least one device which is capable of functioning as a bootstrap router. When multiple candidate bootstrap routers exist in the PIM-SM domain, the bootstrap messages are used to carry

out a dynamic bootstrap router election to determine the PIM domain's bootstrap router. Unlike in DVMRP or MOSPF, in PIM-SM the bootstrap message is the only message that is flooded in the entire domain.



NOTE: If you configure RP sets for a PIM domain using the `StaticRP.cfg` file, it is recommended that you disable BSR functionality (you must set the PIM service `BSRControl` parameter to `Disable`) on all the PIM devices so that the devices do not accept BSR packets.

13.3

Group Operations

Joining a Group

When a receiver wants to join a multicast session “G,” it conveys its membership information through the Internet Group Management Protocol (IGMP). This IGMP message is called an IGMP join (or IGMP report). The PIM-designated device on the LAN receives this IGMP report and looks up the associated rendezvous point for this group. After creating the appropriate multicast route entry, the DR sends a join message via multicast (to the ALL-PIM-ROUTERS group, 224.0.0.13) on the RPF interface towards the rendezvous point. As this join message propagates hop-by-hop to the rendezvous point, each intermediate upstream router towards the rendezvous point creates or updates its multicast route entry for this multicast group. Thus a shared tree rooted at the rendezvous point gets created, for forwarding data traffic destined to the group G.

Sending Data to a Group

When a sender (S) starts sending multicast data packets to a group, its DR initially delivers each packet to the rendezvous point for distribution down the rendezvous point-tree. The sender's DR encapsulates each data packet in a PIM-Register message and unicasts it to the rendezvous point for that group. The rendezvous point decapsulates each register message and forwards the enclosed data packet natively to downstream members on the shared rendezvous point-tree.

Switching from a Shared Tree to a Shortest Path Tree

If the data rate of the source warrants the use of a source-specific shortest path tree (SPT), the rendezvous point or PIM devices with local IGMP members may construct a new multicast route entry that is specific to the source. The entry is referred to as (S,G) state. The devices can send periodic join messages toward the source S (instead of the rendezvous point). When this (S,G) join message reaches the DR directly connected to the source, all devices between the DR of the source and the device that initiated the (S,G) join process cause (S,G) state to forward data packets destined for group G using an SPT. After data traffic for group G starts arriving on the SPT, the device that initiates the (S,G) join removes itself from the shared tree (since it now has a better path to the source). When PIM-SM is implemented on EOS devices, the RP always switches to a shortest path tree after seeing the first packet. Furthermore, if a device switches to the shortest path tree, it never returns to the shared path tree.

Leaving a Group

When the group has no more directly connected listeners (or receivers), the DR on the LAN gets notified via IGMP. If the DR has neither local members nor downstream receivers, it initiates the process to prune itself off the forwarding tree for this group. This pruning process is performed by sending PIM-Prune messages to the source or the rendezvous point, depending on whether the DR had (S,G) or (*,G) state for this multicast group.

13.4

Designated Router Election

PIM devices periodically transmit Hello messages to discover neighboring PIM devices. To elect the DR, each PIM device on the network examines the Hello messages received on an interface and performs the DR

election. The PIM devices can perform the DR election using one of the following two methods, depending on whether or not DR election by priority is enabled on the link:

- DR election by priority
- DR election by IP address

You enable and disable DR election by priority and set DR priority values using the `DRPriority` option of the `HelloOption` parameter. For more information about the `HelloOption` parameter, see “PIM Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

DR Election by Priority

If all PIM devices on the link advertise a non-zero DR priority value, the DR is elected by priority. Each PIM device examines the received Hello messages and compares its own DR priority value with those of its neighbors. The device with the highest DR priority value is elected as the DR for that network. If two devices share the high DR priority value, the tie is broken by IP address: of the devices with the highest DR priority value, the device with the highest IP address is elected as the DR.

DR election by priority gives you convenience and flexibility in configuring the DR in a PIM-SM-enabled link. The advantage of DR election by priority is that you can make a particular PIM device the DR for a network simply by assigning it the highest DR priority value – you do not have to change the device’s IP address.



NOTE: If a PIM device receives Hello packet with no DR priority value, it assumes that the sender has disabled the DR priority option and stores DR priority value=0 for that sender. In such a case, a DR election by IP address is performed as DR election by priority is not enabled on all PIM devices.

DR Election by IP Address

If DR election by priority is disabled on any of the PIM devices on the shared link, DR election must be performed by IP address. Each PIM device examines the received Hello messages and compares its own IP address with those of its neighbors. The device with the highest IP address is elected as the DR.

The disadvantage of DR election by IP address is that in order to make a particular PIM device the DR for a network you must assign it the highest IP address on the link – and it is not always possible to renumber the IP addresses of devices in an existing network.

DR Election After a Device Shuts Down

If any PIM device goes down, all remaining devices check the validity of the election by DR priority. This is done because it is possible that the device that went down was the only device that did not have DR election by priority enabled.

- If all the remaining devices have DR election by priority enabled (all devices are advertising a non-zero DR priority value), a DR election by priority is performed, and the device with the highest DR priority value is elected as the DR.
- If any one of the remaining devices has DR election by priority disabled (it is not advertising a DR priority value, and the other devices are storing DR priority=0 for this device), then no DR election is performed, and the device with the highest IP address remains as the DR.

DR Election After a New Device is Added

As soon as a new device (unknown to its neighboring devices) comes up, it receives the Bootstrap and Hello packet from the current DR. The DR priority option for all devices is then checked and a DR election is performed, depending on the DR priority options of all the devices.

- If the new device has DR election by priority enabled but any other device on the network has DR election by priority disabled, the DR election is performed on the basis of IP address.
- If the new device has DR election by priority enabled and all other devices on the network also have DR election by priority enabled, the DR election is performed on the basis of DR priority value.

- If the new device has DR election by priority disabled, the DR election is performed on the basis of IP address—regardless of whether DR election by priority is enabled or disabled on the other devices.

13.5

Delayed PIM Hellos Disabling

When a PIM device comes up, before it sends out a hello message on any port, it waits first for all the PIM ports to come up and then for a timer to allow unicast routing to converge. You can set the value of the unicast convergence timer using the `UnicastWaitTime` parameter. To disable the timer, set the `UnicastWaitTime` to 0, in which case PIM hellos are sent out immediately after all PIM ports are up. For details, see “PIM Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

13.6

GenID Option

A PIM device is identified as a new neighbor in any of the following situations:

- PIM functionality is enabled on the device interface for the first time, and the device enters the PIM network for the first time.
- The device reboots in the PIM network with PIM-enabled configurations.
- The device’s interface towards the network comes up.

When a new neighbor is identified on a PIM-SM network, as described above, the network converges slowly and suffers traffic outages until the next periodic Bootstrap message is received. The Generation Identifier (GenID) option alleviates this problem by triggering Bootstrap and Join/Prune messages to the new neighbor.

When you enable the GenID option on a PIM device, a different GenID number is randomly assigned to each interface of the device when it boots up. The GenID numbers are sent periodically as a Hello option inside the PIM Hello packet. They remain the same for each interface as long as that interface is up. When the GenID option is included in the Hello packets of the devices on a PIM-SM network, a neighbor reboot or interface up state can be detected easily by the change in the GenID numbers.

When a PIM device receives a Hello packet, the device checks if the GenID value it receives is the same as the value it stored for that neighbor. If the value is different, the local device defaults to reading the Hello packet as coming from a new neighbor device and performs the following functions:

- The local device sends a Hello packet to the new neighbor.
- If the local device is the current DR, it sends a Bootstrap packet to the new neighbor.
- If the new neighbor is the current DR – this can happen if the DR reboots before the local device times it out – and the local device is the runner-up candidate for DR election on that network, the local device sends a Bootstrap packet to the DR (the new neighbor).
- If the new neighbor is the upstream RPF neighbor for any of the entries in the PIM-SM routing table of a local device, the local device sends Join/Prune messages to the new neighbor.
- For all multicast routing table entries, the local device checks to see if the new neighbor is the better RPF neighbor for that entry.

For details about configuring the GenID option, see the description of the GenID parameter in “PIM Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: Motorola Solutions recommends that you enable the GenID option on all PIM devices on the network. If no GenID option is specified in a Hello message, the receiving device is not able to detect the quick reboot, or interface up state from the neighbor and thus it may take longer for that PIM network to converge.

13.7

Triggered Bootstrap Messages

Bootstrap messages are multicast periodically by the bootstrap router to all the other PIM devices in the domain. A unicast bootstrap message is also triggered by the designated router (DR) on a link when the DR detects a new neighbor on that link.

Whenever a BSR receives the candidate rendezvous point (CRP) advertisement from a new RP (for which the BSR does not have any existing state), the BSR triggers a multicast bootstrap message on all the interfaces. This bootstrap message contains all the RPs in the current RP set.

An elected BSR triggers a bootstrap message in the following scenarios:

- Whenever the BSR receives a CRP advertisement, it checks to determine if it is an advertisement from a new RP or a periodic refresh from the RPs existing in the current RP set. If the CRP advertisement is from a new RP, the BSR triggers a bootstrap message to all the PIM devices in the domain.
- Whenever the elected BSR receives a bootstrap packet from a loser BSR, the elected BSR triggers a bootstrap message to avoid having two elected BSRs.
- Whenever an elected BSR reboots and receives a unicast bootstrap packet from its neighbor, the elected BSR triggers a multicast packet so that all entries in all PIM devices in the domain are refreshed.

BSR Advertisements Disabled on Leaf Interfaces

In the EOS implementation of PIM-SM, BSR advertisements are disabled on PIM leaf interfaces for the following reasons:

In a PIM domain, the BSR sends a BSR advertisement packet at periodic intervals, as defined by the `BSRPeriod` parameter (the default is every 60 seconds). This packet contains the list of all the rendezvous points (RPs) in the PIM domain and is the only packet that is flooded in the entire PIM domain. When a PIM device receives this packet, it:

- Checks to see if the packet is received on the correct RPF interface towards the BSR; if not, the device drops the packet.
- Updates/refreshes its own RP set.
- Forwards the packet to all PIM interfaces other than the one on which the packet was received.

When a BSR advertisement packet is sent down a PIM interface, it is meant for any other PIM device(s) on that interface. If the PIM interface is a leaf interface (a PIM interface that has no other PIM neighbors on it). However, the BSR advertisement packet is wasted.

For this reason, before a PIM device (BSR or non-BSR) sends out a multicast bootstrap packet, it checks for the existence of any PIM neighbor on that interface. If the interface has no PIM neighbors (if it is a leaf interface), the device does not send out the BSR advertisement. This check is not required for unicast bootstrap messages because PIM devices send these messages in response to recognizing a new neighbor. Thus the interface cannot be a leaf interface.

13.8

Triggered Candidate RP Advertisements

Candidate RP advertisements are periodic PIM messages unicast to the BSR by those devices that are configured as CRPs. The BSR aggregates RP information from all the RPs and sends the RP set information on all directly-connected interfaces.

When the CRP receives a bootstrap packet in the following scenarios, it triggers a CRP advertisement toward the BSR in the following scenarios:

- The CRP does not have any BSR information.
- There is a change in the elected BSR.

13.9

Bootstrap Segmentation

The size of a bootstrap message depends on the number of group prefixes and the number of CRPs associated with each group prefix.

Bootstrap segmentation was introduced into the Enterprise OS software to address the following problems that arise when a single bootstrap packet is sent to advertise all the group prefixes and all the CRPs associated with each group prefix:

- If the bootstrap packet size is larger than 1480 bytes, an incomplete RP set is advertised.
- If the MTU (maximum transmission unit) of the network is less than 1480 bytes, the bootstrap packet is fragmented at the IP level.
- If the bootstrap packet is lost, the PIM devices do not receive the RP set for any group prefix until the next periodic bootstrap message.
- If the number of group prefixes reaches a certain level, the number of CRPs allowed for each group prefix is limited. For example, if there are 50 group prefixes to be advertised, for two CRPs in each group prefix, the size of the bootstrap packet would become: $14 + 50 * (12 + 10 * 2) = 1614$ bytes.

The effect of bootstrap segmentation is to segment a single bootstrap message into a number of different bootstrap segments, with each segment containing the CRPs for one or more group prefixes, depending on the maximum IP packet size allowed. These segments are known as bootstrap fragments. For a single group prefix, if there are many CRPs, such that the packet size exceeds the maximum allowed limit, the group prefix can extend over more than one bootstrap fragment. You configure the maximum size for bootstrap fragments using the `BSFragLimit` parameter. For more information about `BSFragLimit` parameter, see "PIM Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

A bootstrap message is originated in the following scenarios:

- Periodically by the current BSR of the domain.
- Triggered by the BSR upon receiving a new CRP advertisement.
- Triggered by the designated router (DR) when a new `GenID` value is received from a non-DR router.
- Triggered by the runner-up DR when a new `GenID` value is received from the DR.

If bootstrap segmentation is necessary, it is performed in all these scenarios. The sending device forms a bootstrap packet or multiple bootstrap fragments, depending on the bootstrap fragmentation limit (`BSFragLimit`) configured at the originator. All the bootstrap fragments belonging to a single bootstrap message contain the same fragment tag—a 16-bit randomly-generated number included in the header and used to distinguish fragments belonging to different bootstrap messages.

In addition to the fragment tag, bootstrap fragments include the following fields:

RP-Count

The total number of CRP addresses included in the entire bootstrap message for the corresponding group prefix. A device does not replace its old RP set for a given group prefix until or unless it receives "RP-Count" addresses for that prefix; the addresses could be carried over several fragments. If only part of the RP set for a given group prefix is received, the device waits until all the information for that group prefix is received. If a group prefix extends over more than one fragment, the RP-Count field is present (with the same value) in all the fragments containing that group prefix.

Frag RP-Cnt

The number of CRP addresses included in this fragment of the bootstrap message, for the corresponding group prefix.

The PIM bootstrap message/fragments is/are flooded to the entire PIM-SM domain so all PIM devices are aware of the current RP set serving different group prefixes. Upon receiving a bootstrap message/fragments from a preferred BSR, the PIM device performs one of the following actions:

- Forwards the bootstrap fragment as is on all of its PIM-enabled interfaces except the one over which the message arrived.



NOTE: The receiver does not aggregate the received fragments, even if the `BSFragLimit` configured for it allows it to do so.

- If necessary, segments the fragment further, depending on the configured `BSFragLimit`. If the receiver segments the fragment further it forwards the new fragments all with the same fragment tag.

In addition, the receiving device looks at the “RP-Count” field for each group prefix. If the received fragment contains a number of CRP addresses equal to the “RP-Count,” the receiving device accepts the RP set for that group. If the received fragment does not contain a number of CRP addresses equal to the “RP-Count,” the receiving device waits for subsequent fragments containing the remaining CRP addresses. If the receiving device does not receive a number of CRP addresses equal to the “RP-Count” before it either receives a bootstrap packet with a new fragment ID or the bootstrap fragment holdtime period expires, the receiving device discards the incomplete RP set.



NOTE: The bootstrap fragment holdtime = 30 seconds.

13.10

Join/Prune Aggregation

Join/prune aggregation is a method by which a device can skip sending the first periodic join/prune message upstream toward the RP or Source in order to preserve network bandwidth.

When a receiver joins a PIM-SM domain, it sends an IGMP join toward the DR. The DR in turn sends a triggered join toward the RP. The (*, G) join travels hop-by-hop toward the RP, and in each device it passes through, the multicast tree state for group G is instantiated. Eventually, the join either reaches the RP or reaches a device that already has (*, G) join state for that group.

After the DR sends the triggered join, it sends periodic join/prune messages upstream toward the RP after every JoinPrune period (the default is 60 seconds). The JoinPrune timer may expire anytime after the triggered join was sent. As a result, the DR may be sending redundant information toward the RP within the span of the JoinPrune timer. This information is redundant, because the upstream devices already created the state as a result of the triggered join message. It unnecessarily creates a larger packet size for the join/prune message and wastes network bandwidth.

Join/prune aggregation addresses this inefficiency. Because upstream devices keep the state for (3.5 * JoinPrune period) seconds, the DR can skip sending the first periodic join/prune message for all groups. This reduces the size of the periodic join/prune messages and saves bandwidth. You can enable and disable join/prune aggregation with the `JoinPruneAgg` parameter. For details, see the description of the `JoinPruneAgg` parameter in “PIM Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

13.11

Multiple Joins

Multiple joins are a method by which a PIM device sends multiple joins to provide higher reliability in a network in which there is a high probability of dropped packets.

PIM-SM, as an IP-based protocol, does not provide reliable joins between devices. When a receiver joins a PIM-SM domain, it sends an IGMP join toward the DR, which in turn sends a triggered join toward the RP. If this first triggered join is dropped, a device upstream to the DR does not know about the newly joined receiver until it receives the next periodic join from the DR. The periodic join is not sent until the JoinPrune period expires (the default is 60 seconds). Because of that a call with a duration less than the JoinPrune period may not be placed at all if the first triggered join is dropped.

The multiple joins feature addresses this potential problem. PIM-SM devices emit multiple (*, G) and (S,G) joins when an internal multicast state changes. By having each device send the first triggered join multiple

times, paced at a calculated interval, the probability of a join never making it to the upstream device is greatly reduced.

For details about how to configure multiple joins, see the description of the `MultipleJoins` parameter in “PIM Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

13.12

Static RP Sets

In addition to the dynamic RP set configuration (dynamic RP) implemented in previous versions, EOS version 13.0 and higher supports distribution of RP set information (RP -> multicast address binding) via a static RP set configuration (static RP).

In dynamic RP, the RP set information is propagated through bootstrap messages originated by the designated bootstrap router (BSR). This can cause problems in large networks, where any error in a small part of the network can cause an error in the bootstrap packet, making the entire network unusable.

The static RP feature allows you to configure the RPs for specified multicast address ranges and download the static RP set information to EOS devices in a separate ASCII file (the `StaticRP.cfg` file) that is downloaded with the device’s `boot.cfg` file. You configure static RPS using the PIM service `ADD StaticRP` command; each `ADD` command configures an RP for a specified multicast address range.

The `StaticRP` parameter allows you to:

- Add and delete static RPs.
- Show all configured static RPs.
- Show all static multicast groups.
- Display a static entry audit summary.

For details about the `StaticRP` parameter, see “PIM Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: If you are configure RP sets for a PIM domain by using the `StaticRP.cfg` file, we recommend that you disable BSR functionality (set the PIM service `BSRControl` parameter to `Disable`) on all the PIM devices so that the devices do not accept BSR packets.

Static Entry Audit Summary

You must issue the `SHoW -PIM StaticRP Summary` command to display the static entry audit summary. This report provides a list of the static RP configuration errors detected by the software, including the following information for each error:

- `Multicast Range` – the multicast IP address range to which the static RP configuration applies.
- `Dynamic RP` – the RP IP address propagated dynamically by the BSR.
- `Static RP` – the RP IP address configured statically with the `ADD -PIM StaticRP` command.
- `Error Code` – one of the following error names: `Missing Entry`, `Mismatched RP`, or `Mismatched CRP`.

Table 30: Static RP Errors Reported by the EOS `SHoW -PIM StaticRP Summary` Command

Error (Error Code #)	Description	Corrective Action
Missing Entry (1)	The bootstrap packet contains a multicast range that has no corresponding entry in the <code>StaticRP.cfg</code> file.	Add this entry to the <code>StaticRP.cfg</code> file.

Error (Error Code #)	Description	Corrective Action
Mismatched RP (2)	The RP for the multicast address in the <code>StaticRP.cfg</code> file is different than the one advertised in the bootstrap packet.	Correct the configuration in the <code>StaticRP.cfg</code> file.
Mismatched CRP (3)	This device is the candidate RP (CRP) for this multicast address, but a different RP for this address is configured in the <code>StaticRP.cfg</code> file.	Correct the configuration in the <code>StaticRP.cfg</code> file.

13.13

Internal IP Address Support

The PIM-SM service supports System IP, a software-only interface that emulates an IP interface. The system IP interface is always up (active and functional); it allows IP operations to take place as long as at least one interface on the device is up.

Currently PIM-SM allows one `CandidateBSR` / `CandidateRP` per system. With System IP support, you can assign the System IP address to these two parameters. When you assign the System IP address to the `CandidateBSR` and/or `CandidateRP`, the configured System IP address is sent out to the other devices, and all PIM-SM configurations reflect the configured System IP address as the `CandidateBSR` and/or `CandidateRP`. The PIM multicast routing table also reflects the System IP address on the device on which it was configured as well as on the neighboring devices.

If you change the System IP address, PIM-SM is notified. If the `CandidateBSR` and/or `CandidateRP` are configured, the IP address is automatically updated and advertised.

13.14

Event Driven/Triggered PIM-SM

PIM-SM makes use of the system's unicast routing database (RTM and BGP) to perform reverse path forwarding (RPF) checks on the source of data traffic. PIM-SM organizes its multicast routing table (MRT) into a hierarchical set of lists:

- The top level list is a list of all the possible RPs.
- Associated with each RP is a list of group prefixes served by this RP.
- Associated with each group prefix is a list of active group addresses.
- Associated with each active group address is a list of active sources sending on this multicast address.

RPF checks update the routes to the group multicast addresses on which active sources are sending and the RPs associated with those addresses. You can use the `RpfRefresh` parameter to choose between a periodic polling scheme (in which PIM-SM performs RPF updates via periodic polling) or a change notification scheme (in which PIM-SM is notified immediately when the underlying routing subsystem detects a change). The change notification scheme (`RpfRefresh = Triggered`) reduces the latency encountered in propagating routing changes to PIM-SM. RPF updating is performed when route change notifications are received, rather than on a strictly periodic basis.

13.15

PIM Packet Formats

In the current version of PIM Sparse Mode (2A II PIM), control messages use the IP protocol number 103. PIM control messages are either unicast (such as Register messages from the PIM designated router to the rendezvous point), or multicast hop-by-hop (such as Join/Prune and Assert to the ALL-PIM-ROUTERS multicast group (224.0.0.13)).

PIM-SIM TOS Values

TOS values are included in PIM-SM network control packets to allow for prioritization of PIM-SM control packets over a network infrastructure. The TOS values assigned to the delay-critical and non-delay-critical PIM-SM packets are as follows:

Delay-Critical Packets (TOS = 6)

- Triggered Join/Prune
- Assert
- Null Register (For Register packets which contain data, the TOS value is copied from the TOS value set in the encapsulated packet.)

Non-Delay-Critical Packets (TOS = 4)

- Periodic Join/Prune
- Bootstrap (periodic multicast or triggered unicast)
- Candidate for RP Advertisement
- Hello

13.16

PIM SNMP Traps

EOS supports the following PIM SNMP traps:

- pimExcessiveRegistersRP
- pimExcessiveRegistersDR
- pimExcessiveRouteEntries
- pimNoReceivers


You can use the `SnmpTraps` parameter to enable or disable generation of the PIM SNMP traps on the device and to configure the values above which the supported traps are generated.

For more information about the `SnmpTraps` parameter, see “PIM Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

13.17

Configuring PIM-SM

Perform this procedure to configure PIM-SM as demonstrated on the following example of a sample PIM-SM configuration.

Figure 89: PIM-SM Sample Configuration**Procedure:****1. If you want to configure the pim1 device, perform the following actions:**

- a. Set the baud rate and enable the paths by entering:

```
SETDefault !3a,!3b -PATH BAud=256 CONTrol=e
```
- b. Configure PPP links by entering:

```
SETDefault !3a,!3b -PORT OWNer=ppp
```
- c. Enable IP routing by entering:

```
SETDefault -IP CONTrol=ro
```
- d. Configure the console prompt by entering:

```
SETDefault -SYS NMPrompt="pim1 #"
```
- e. Configure IP addresses by entering the following commands:

```
SETDefault !3a -IP NETAddr=192.168.1.1
```

```
SETDefault !3b -IP NETAddr=192.168.4.1
```
- f. Enable MIP control by entering:

```
SETDefault -MIP CONTrol=e
```
- g. Enable OSPF by entering:

```
SETDefault !3a,!3b -OSPF CONTrol=e
```
- h. Enable PIM-SMv2 by entering:

```
SETDefault !3a,!3b -PIM CONTrol=sm
```
- i. Enable CandidateRP by entering:


```
SETDefault -PIM SMControl=crp
```

- j. Add a CandidateRP address by entering:

```
ADD -PIM CandidateRP 192.168.1.1 AddressList CRP1
```



NOTE: The AddressList you specify in the ADD -PIM CandidateRP command (CRP1) must be previously created by using the ADD -IP AddressList command.

2. If you want to configure the **pim2** device, perform the following actions:

- a. Set the baud rate and enable the paths by entering:

```
SETDefault !3a,!3b -PATH BAud=256 CONTROL=e
```

- b. Configure PPP links by entering:

```
SETDefault !3a,!3b -PORT OWNeR=ppp
```

- c. Enable IP routing by entering:

```
SETDefault -IP CONTROL=ro
```

- d. Configure the console prompt by entering:

```
SETDefault -SYS NMPrompt="pim2 #"
```

- e. Configure IP addresses by entering the following commands:

```
SETDefault !3a -IP NETAddr=192.168.1.2
```

```
SETDefault !3b -IP NETAddr=192.168.2.2
```

- f. Enable MIP control by entering:

```
SETDefault -MIP CONTROL=e
```

- g. Enable OSPF by entering:

```
SETDefault !3a,!3b -OSPF CONTROL=e
```

- h. Enable PIM-SMv2 by entering:

```
SETDefault !3a,!3b -PIM CONTROL=sm
```

- i. Configure the PIM-SM candidate bootstrap router by entering:

```
SETDefault -PIM SMControl=cbsr
```

- j. Add a candidate bootstrap router address by entering:

```
ADD -PIM CandidateBSR 192.168.2.2
```

3. If you want to configure the **pim3** device, perform the following actions:

- a. Set the baud rate and enable the paths by entering:

```
SETDefault !3a,!3b -PATH BAud=256 CONTROL=e
```

- b. Configure PPP links by entering:

```
SETDefault !3a,!3b -PORT OWNeR=ppp
```

- c. Enable IP routing by entering:

```
SETDefault -IP CONTROL=ro
```

- d. Configure the console prompt by entering:

```
SETDefault -SYS NMPrompt="pim3 #"
```

- e. Configure IP addresses by entering the following commands:

```
SETDefault !1 -IP NETAddr=192.168.30.1
```

```
SETDefault !3a -IP NETAddr=192.168.3.3
```

```
SETDefault !3b -IP NETAddr=192.168.2.3
```

- f. Enable MIP control by entering:**

```
SETDefault -MIP CONTrol=e
```

- g. Enable OSPF by entering:**

```
SETDefault !1,!3a,!3b -OSPF CONTrol=e
```

- h. Enable PIM-SMv2 by entering:**

```
SETDefault !1,!3a,!3b -PIM CONTrol=sm
```

- 4. If you want to configure the **pim4** device, perform the following actions:**

- a. Set the baud rate and enable the paths by entering:**

```
SETDefault !3a,!3b -PATH BAud=256 CONTrol=e
```

- b. Configure PPP links by entering:**

```
SETDefault !3a,!3b -PORT OWNer=ppp
```

- c. Enable IP routing by entering:**

```
SETDefault -IP CONTrol=ro
```

- d. Configure the console prompt by entering:**

```
SETDefault -SYS NMPrompt="pim4 #"
```

- e. Configure IP addresses by entering the following commands:**

```
SETDefault !1 -IP NETAddr=192.168.40.1
```

```
SETDefault !3a -IP NETAddr=192.168.3.4
```

```
SETDefault !3b -IP NETAddr=192.168.4.4
```

- f. Enable MIP control by entering:**

```
SETDefault -MIP CONTrol=e
```

- g. Enable OSPF by entering:**

```
SETDefault !1,!3a,!3b -OSPF CONTrol=e
```

- h. Enable PIM-SMv2 by entering:**

```
SETDefault !1,!3a,!3b -PIM CONTrol=sm
```

Chapter 14

Router Discovery Protocol Configuration

This chapter describes how to configure your system to use Internet Control Message Protocol (ICMP) Router Discovery Protocol (RDP) messages to dynamically generate a list of active neighbor router addresses. RDP is an extension of ICMP, and assists hosts in discovering neighboring routers. It is also a requirement per RFC 1812 for IPv4 devices on all connected networks that support either IP multicasting or IP broadcast addressing.

RDP is a process defined by RFC 1256 that allows a device to use two messages, router advertisements and router solicitations, to discover the addresses of neighboring routers. The RDP process works only on devices enabled for the process that are in the same subnetwork. The device listens for router advertisements, or can solicit an address by sending a router solicitation message.

The discovery process is dynamic - the list collected contains only the addresses of active devices. Devices that are not actively sending router advertisements are dropped from the list, and are reinserted in the list only when they come back up and begin sending messages again.

RDP is not a routing protocol; it allows hosts only to keep track of neighboring routers, not to determine which device is best to reach a particular destination. This protocol uses two ICMP messages to provide a simple router discovery method. The method provides a list of device addresses on a multicast link without manual configuration and that is independent of the used routing protocol.

Although RDP cannot make routing decisions, if a host makes a poor choice for a first-hop router for a destination, RDP receives an ICMP Redirect message identifying a better route.

RDP Features

The router advertisement message includes a preference level for each advertised device address. When a host system must choose a default device, it is expected to select device addresses with a high preference level. You set this preference level when you specify the list of devices that participates in the RDP process.

To make sure that hosts ignore devices that go down, the router advertisement message also includes a lifetime field that specifies the maximum length of time that advertised addresses are to be considered valid by hosts. The default advertising rate is every 10 minutes, and the default lifetime span is 30 minutes. The defaults minimize the load imposed on the links by the periodic transmission of the messages, but you can change the defaults as needed. For example, you may want to decrease the lifetime value so that you can become aware of devices that go down before the 30-minute period is up.

Other Timer Considerations

When an interface enabled for advertising the device address becomes active, the device begins transmission of periodic router advertisement messages. The interval between the first three messages cannot be greater than 16 seconds. After these first three messages, the interval is randomly chosen from the values configured for each interface. A new, random interval is chosen for each transmission, to reduce the possibility of all devices transmitting packets at the same time.

Periodic advertisements are either multicasted to the all-host address (244.0.0.1) or broadcasted to the limited-broadcast address (255.255.255.255), depending on the system configuration. The device also transmits advertisements in response to host solicitations. When the source IP address is not set to 0, the reply is unicasted to the host; otherwise, it is multicasted to the configured address.

The host transmits up to three RDP router solicitation messages when an interface becomes active. The first message is transmitted within one second and then retransmitted at 3-second intervals. Transmissions stop when a host receives a valid router advertisement message. You can configure router solicitations to be either multicast to the all-routers address (244.0.0.2) or broadcast to the limited-broadcast address (255.255.255.255).

RDP Terms

Multicast link

A link over which the IP multicast or IP broadcast services are supported, including media such as satellite and point-to-point.

Neighbor

A device with an IP address belonging to the same subnet.

Default device


The device address that has the highest preference level. Unless a host is redirected or is configured to use a specific device address, it must choose a default device address for a particular destination. You can set the preference level to encourage or discourage use of a particular device.

14.1

Setting Up RDP

You can use this process to enable your system for RDP. Depending on your network requirements, you can use the default values of the parameters in the various services, or you can configure the EOS device with custom values.

Figure 90: Routers Participating in RDP



Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Your ports and paths are configured. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- IP routing is configured. See [IP Routing Configuration on page 119](#) and [IP Multicast Routing Configuration on page 258](#).

Procedure:

1. Define participating devices by performing the following actions:

- a. Define the list of devices that participate in the router discovery process by entering:

```
ADD RouterList <IP_address> [NoAdvertise] [<preference_level>|Infinity]
```

You must enter the IP address for each device to be advertised. The length of this list is limited only by available system memory.

You can also enter the IP address for each device, and assign a preference level to specify the default devices that learn from router advertisement messages by using:

```
ADD RouterList <IP_address> <preference_level>
```

The `<preference_level>` option indicates the preferences for selecting the default device.

The `<preference_level>` option is a 32-bit, signed, twos-complement integer, which allows you to enter a definitive number to specify the selection criteria for the default device. The higher the value you assign, the more preference the device address has.

Step example:

To list the routers for participation in RDP, enter:

```
ADD -RDP RouterList 10.0.0.1
ADD -RDP RouterList 10.0.0.2
ADD -RDP RouterList 10.0.0.3
ADD -RDP RouterList 10.0.0.4
```

To designate a device as the default, enter:

```
ADD -RDP RouterList 10.0.0.1 100
```

- b. To indicate that an address is not to be used as a default device address, enter:

```
ADD -RDP RouterList 10.0.0.4 Infinity
```

The `Infinity` option indicates a minimum value (0x80000000) that prevents the device from being picked up by hosts as a default device.

- c. To block a device from participating in the discovery process, enter its IP address with the `NoAdvertise` option.

Step example:

```
ADD -RDP RouterList 10.0.0.6 NoAdvertise
```

2. Configure the RDP timers by performing the following actions:

- a. Specify a value for the `LifeTime` parameter in router advertisement messages by entering:

```
SETDefault !<port> -RDP LifeTime = <seconds>(4-9000) | Default
```

The default is 30 minutes (1800 seconds).

- b. Specify a value for the maximum interval between two router advertisement messages by entering:

```
SETDefault !<port> -RDP MaxInterval = <seconds>(4-1800) | Default
```

The default is 10 minutes (600 seconds), which must be less than the value of the `LifeTime` parameter.

- c. Specify a value for the minimum interval allowed between two router advertisement messages by entering:

```
SETDefault !<port> -RDP MinInterval = <seconds>(3-1800) | Default
```

The default is 75 percent of the `MaxInterval` value (nine minutes or 450 seconds), which must be less than the value set for `MaxInterval`.

Step example: The default lifetime of the router advertisement message is 30 minutes, and the default interval between the messages is ten minutes. To change these values to a lifetime value of 12 minutes and an interval of 6 minutes, you must use the following commands:

```
SETD !1 -RDP LifeTime = 720
SETD !1 -RDP MaxInterval = 360
SETD !1 -RDP MinInterval = 300
```

You must always set the value of `MinInterval` to be less than the value you set for `MaxInterval`.

3. Enable/disable RDP by performing the following actions:

- a. To enable or disable RDP globally, enter:

```
SETDefault !<port> -RDP CONTROL = ([Auto | Enable | RespondOnly | Disable],
[Multicast | Broadcast])
```

The default is Auto, which enables RDP on local area networks, but not wide area networks.

- b. To enable RDP on a wide area network, enter:

```
SETDefault !1 -RDP Control = Enable
```

- c. To specify that packets are multicasted, enter:

```
SETDefault !1 -RDP Control = Multicast
```

When the device is set in host mode, the Multicast option sends router solicitation messages out with the IP destination set to the all-routers address (244.0.0.2).

When the device is set in router mode, the Multicast option sends router advertisement messages out with the IP destination address set to the all-host IP address (244.0.0.1). This is the default mode.

- d. To specify IP broadcasting when it is enabled, enter:

```
SETDefault !1 -RDP Control = Broadcast
```

Both router solicitation and router advertisement messages are sent out with the IP destination set to the limited-broadcast IP address (255.255.255.255).

- e. To only respond to solicitation messages, enter:

```
SETDefault !1 -RDP Control = RespondOnly
```

The device does not send out periodic advertisements, but it responds to solicitation messages from the host.

4. Discover neighboring RDP routers by having the system send out router solicitations by entering:

```
DiscRouteRs [!<port> | <source_IP>] [Broadcast]
[<timeout (1-30 seconds)>]
```

To enable systems in host mode (!0), to discover neighboring routers, specify either an outgoing port number or one of the system source IP addresses to be sent with the router solicitations. After the command is entered, the system transmits router solicitations every second for the length of time specified with the timeout option. During the timeout period, any router advertisements that are received are displayed.

14.2

Verifying the RDP Configuration

Procedure:

To verify that RDP is recognized by the IP network, display the values associated with RDP by entering:

```
Show [!<port>] -RDP CONFIGuration
```

14.3

Troubleshooting the RDP Configuration

You can troubleshoot the RDP operation by using one or more of the following commands.

Procedure:

- To display the set of interfaces enabled or disabled for RDP, enter:

```
Show [!<port>] -RDP CONTROL
```

- To display the router list, enter:
`SHow -RDP RouterList`
- To delete one or all of the interfaces enabled or disabled for RDP, enter:
`DELeTe -RDP RouterList {<IP_address>|ALL}`
- To flush the router list and allow the routes to be relearned, enter:
`FLush -RDP RouterList`
This command only works when the device is in host mode.
- To display the value of the router advertisement lifetime field, enter:
`SHow [!<port>] -RDP LifeTime`
- To display the value of the maximum interval between router advertisement messages, enter:
`SHow [!<port>] -RDP MAxInterval`
- To display the value of the minimum interval between router advertisement messages, enter:
`SHow [!<port>] -RDP MInInterval`

Chapter 15

Wide Area Networking Configuration with Frame Relay

This chapter describes how to configure your EOS device to establish serial line connectivity through Frame Relay.

15.1

Setting Up The Frame Relay Service

You can use this process to configure your device to transmit and receive data over a Frame Relay interface. You can configure congestion control for individual virtual circuits on a logical or virtual port you are configuring for Frame Relay.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- WAN ports on your device are configured.

Determine if your Frame Relay network supports the Local Management Interface (LMI) Protocol. For information about LMI, see [Local Management Interface \(LMI\) Protocol on page 330](#).

Determine if you have a partially meshed or nonmeshed topology:

- If you have a partially meshed or nonmeshed topology, and you plan to enable the Internet Protocol (IP) or the Routing Information Protocol (RIP), ensure that the next-hop split horizon feature is enabled.
- If you have a partially meshed or nonmeshed topology and you plan to enable Open Shortest Path First (OSPF), ensure that you set `-OSPF CONTro1` to `NonMesh` to enable the point-to-multipoint interface.

Process:

1. Configure the Frame Relay service. See [Configuring the Frame Relay Service on page 308](#).
2. Configure the congestion control by performing one of the following actions:
 - Configure congestion control for device ports configured for NTTLMI. See [Configuring Congestion Control for NTTLMI Protocol Users on page 309](#).
 - Configure congestion control for ports configured to use an LMI protocol other than NTTLMI. See [Configuring Congestion Control for Other LMI Protocol Users on page 310](#).
3. Configure PVCs and SVCs. See [Configuring Frame Relay SVC Connections on page 311](#).
4. Verify the Frame Relay configuration. See [Verifying the Frame Relay Configuration on page 313](#).

15.1.1

Configuring the Frame Relay Service

Use this procedure to configure your device to transmit and receive data over a Frame Relay network.

Procedure:

1. Enable the Frame Relay service by setting the owner of the serial interface to Frame Relay by entering:


```
SETDefault !<portlist> -PORT OWNeR = FrameRelay
```



NOTE: If PORT OWNeR is set to Auto, Frame Relay is detected and configured automatically and this step may not be necessary.

For networks running RIP with the port up and the -RIPIP CONTroL parameter set to TAlk, the CONTroL parameter DynamicNbr option is automatically enabled. With the DynamicNbr enabled, neighbors are automatically added. If DynamicNbr is not enabled, you must add neighbors manually.

2. Perform one of the following actions:

- If your Frame Relay network supports the LMI Protocol, ensure that the appropriate LMI Protocol is enabled.
- If your Frame Relay network does not support the LMI Protocol, disable this protocol.

The EOS software includes four types of LMI:

Consortium LMI
Annex-D LMI
NTT LMI
ITU LMI

You must configure the software with the type of LMI that the switching equipment supports:

- Consortium LMI to be configured by specifying LMI
- Annex-D LMI to be configured by specifying ANSiLMI
- NTT LMI to be configured by specifying NTTLMI
- ITU LMI to be configured by specifying ITULMI.

If the port is configured for auto detect, the type of LMI is determined dynamically. To manually enable the specific LMI or to completely disable the LMI Protocol, you must use:

```
SETDefault !<portlist> -FR CONTroL = [NoLMI | LMI | ANSiLMI | NTTLMI | ITULMI]
```

15.1.2

Configuring Congestion Control for NTTLMI Protocol Users

Procedure:

To configure congestion control for device ports configured for NTTLMI, enter:

```
SETDefault !<portlist> -FR DLCIR = <vcid> <cir>
```

where:

<vcid> (virtual circuit identifier) value is one of the following:

- For a permanent virtual circuit (PVC), the <vcid> value is the data link circuit identifier (DLCI) assigned by your Frame Relay service provider from range of 16 through 991.
- For a switched virtual circuit (SVC), the <vcid> value is a unique virtual circuit identifier number that you assign to an SVC from the 16 through 991 DLCI range of numbers that was not assigned by your service provider for a PVC. For more information on the -FR DLCIR parameter, see "FR Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

<cir> value specifies the rate of device data (in kbps) that the Frame Relay network commits to transfer under normal conditions.



CAUTION: Failure to specify <cir> values for the DLCIs causes unpredictable results.

15.1.3

Configuring Congestion Control for Other LMI Protocol Users

For information on the parameters described in this procedure, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Procedure:

1. To activate congestion control for a particular port and individual `<vcid>` connection, and to specify how many consecutive BECN=1 frames (`<step>`) Frame Relay sends to the device port, before the maximum throughput rate is reduced to a level below cir, enter:

```
SETDefault !<portlist> -FR CongestControl = <vcid> <No / Yes> <step> (1-999)
```

2. To specify the throughput parameters for data coming in and going out a specified virtual circuit (`<vcid>`) connection on a specified port, enter:

```
SETDefault !<portlist> -FR CIRbothdir = <vcid> <cir> <mincir> <Bc> <Be>
```

where:

`<vcid>` identifies the Frame Relay virtual circuit used to establish the connection between the remote user and the local port.

`<cir>` specifies the rate of device data (in kbps) that the Frame Relay network commits to transfer under normal conditions.

`<mincir>` specifies the minimum rate of device data throughput (in kbps) that the calling user is committed to accept for the call. Normally you specify `<mincir>` for SVCs, not PVCs. Specify the `<mincir>` value for PVCs only if the PVC `<cir>` is set to zero and you want to use congestion control.

`<Bc>` specifies the maximum of device data bits (in Kilobits) that the network commits to transfer under normal conditions during the time interval (Tc) measured in seconds. Tc (Committed Rate Measurement Interval) is the time interval during which the device can send only Bc committed amount of data and Be excess amount of data. According to the standards, the duration of Tc is proportional to the “burstiness” of the traffic. Tc is computed (from the subscription parameters of CIR and Bc) as $Tc = Bc/CIR$. Tc is not a periodic time interval; rather it is used only to measure incoming data, acting as a sliding window. Incoming data triggers the Tc interval, which continues until it completes its computed duration.

`<Be>` specifies the maximum of uncommitted device data bits (in Kilobits) in excess of Bc that the network attempts to deliver during the time interval (Tc) measured in seconds. If `<Be>` is nonzero, the device does not transmit more than `<Be>` number of data bits over the Tc time interval.




NOTE: When you configure a PVC, the values you enter for `<cir>`, `<Bc>`, and `<Be>` must match the values for these variables provided by your service provider.

15.1.4

Frame Relay Congestion Control Configuration Example

Figure 91: Congestion Control Configuration Example



This example describes how to configure the logical ports of an EOS device for Frame Relay Congestion Control. The following values are used in the configuration:

- Virtual port 1 on physical path 2 over DLCI (vcid) 21 with a *<cir>* of 32, *<mincir>* *<Bc>* of 32, and a *<Be>* of 32.
- Virtual port 2 on physical path 2 over DLCI (vcid) 22 with a *<cir>* of 16, *<mincir>* of 16, *<Bc>* of 16, and a *<Be>* of 16.
- Virtual port 3 on physical path 3 over DLCI (vcid) 22 with a *<cir>* of 16, *<mincir>* of 16, *<Bc>* of 16, and a *<Be>* of 32.
- All ports are set to a *<step>* value of 4.

To activate congestion control for all virtual circuits, and to specify that Frame Relay send four consecutive BECN=1 frames (*<step>*) to the device port to notify it of congestion on the specific virtual circuit before the maximum throughput traffic is reduced below *<cir>*, you must use the following commands:

```
SETDefault !V1 -FR CongestControl 21 Yes 4
SETDefault !V2 -FR CongestControl 22 Yes 4
SETDefault !V3 -FR CongestControl 21 Yes 4
```

To specify the *<cir>*, *<mincir>*, *<Bc>*, and *<Be>* values for all the logical ports, you must use the following commands:

```
SETDefault !V1 -FR CIRbothdir 21 32 32 32 32
SETDefault !V2 -FR CIRbothdir 22 16 16 16 16
SETDefault !V3 -FR CIRbothdir 21 64 64 64 0
```

15.1.5

Configuring Frame Relay SVC Connections

You can use either PVCs or SVCs to connect remote users through the Frame Relay network to your device port. PVCs are identified by Frame Relay virtual circuit IDs (*<vcid>*), which are DLCI numbers assigned by your service provider to make the connection. The range of available DLCIs that can be assigned by the service provider is determined by the LMI protocol configured for your device port.

SVCs are identified by Frame Relay virtual circuit IDs, which are numbers that match DLCI numbers available from the range of DLCIs that were not assigned by the service provider for PVC connections.

Unlike PVCs, which use the DLCIs to make their connections, SVCs must be configured for the following addresses (telephone numbers) to establish their connections:

- The customer premises equipment (CPE) link address – the telephone number assigned by your network service provider for your local device port. This address is used by all SVCs configured for this port.
- The destination address (telephone number) of the remote user associated with the local virtual circuit ID.
- The address of the local SVC – telephone number associated with the local virtual circuit ID (the DLCI you assigned it). The CPE link address is used to identify the local virtual circuit if the address of the local SVC is not set.

You can configure SVCs as dynamic connections that go down after the data transfer stopped for the time interval specified with the `-FR SvcIdleTimer` parameter, or you can configure them like PVCs – as static connections that remain up after the data transfer stopped.

For more information on the parameters used in this procedure, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Procedure:

1. Identify the local customer premises equipment (CPE) address (telephone number) assigned by your Frame Relay network service provider for your device to identify the local device port to the network by entering:

```
SETDefault !<portlist> -FR LinkAddress = [ <E.164 address> (1-15 digits) | <X.121 address> (1-15 digits)]
```

2. Identify the SVC destination address (telephone number) of the remote user associated with the local virtual circuit ID (<vcid>) by entering:

```
ADD !<portlist> -FR SvcDestAddress = <vcid> [ <E.164 address> (1-15 digits) | <X.121 address> (1-15 digits)]
```

3. Optional: Identify the local SVC telephone number associated with a specified virtual circuit ID (<vcid>) to the remote user by entering:

```
SETDefault !<portlist> -FR SvcLocalAddress = <vcid> [ <E.164 address> (1-15 digits) | <X.121 address> (1-15 digits)]
```

where <vcid> specifies a virtual circuit identifier number from the range of DLCI numbers available and that were not assigned by your service provider for a PVC connection. The local address configured for this <vcid> number is used to identify the local virtual circuit used to establish the connection between the remote user and the device port. If you do not set the address of the local SVC, the CPE link address is used to identify the virtual circuit.

An SVC connection is automatically established depending on whether the port is busy with another caller on another <vcid> assigned to the port.

4. Optional: Manually activate an SVC connection that ensures a connection for your <vcid> to that port (optional) by entering:

```
ADD !<portlist> -FR SvcConnection <vcid>
```

5. Optional: Manually disconnect an SVC connection for your <vcid> to that port (optional) by entering:

```
DELeTe !<portlist> -FR SvcConnection <vcid>
```

6. Optional: Set your `SvcIdleTimer` parameter for how long an SVC remains idle before it shuts down by entering:

```
SETDefault!<portlist> -FR SvcIdleTimer <vcid> [None | (0-3600 seconds)]
```

The default setting is 80 seconds.

7. Optional: Enable or disable SVC operation on the specified port by entering:

```
SETDefault !<portlist> -FR SvcMode = [0 = disabled | 1 = enabled]
```

The default setting is 0 (disabled).

15.1.6

Verifying the Frame Relay Configuration

Procedure:

To verify the Frame Relay configuration, enter:

```
SHoW -FR CONFIguration
```

The device displays current Frame Relay configuration information.

15.2

Configuring Transparent Bridging over Frame Relay

Prerequisites:

Ensure that:

- Ports and paths in your LAN are configured. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Frame Relay service is set up. See [Setting Up The Frame Relay Service on page 308](#).

If your Frame Relay network does not support the LMI Protocol, and you want to bridge to specific DLCI neighbors, obtain DLCI neighbor addresses to add to the DLCI table. For more information on DLCI addresses, see [Frame Relay Networks on page 324](#).

Procedure:

1. If your Frame Relay network does not support the LMI Protocol or if you want to bridge to specific DLCI neighbors, add DLCI neighbors to the static DLCI neighbor table by entering:

```
ADD !<port> -BRidge DlcINeighbor = <dIci>
```

Even if the LMI Protocol is enabled, you can add DLCI neighbors to the static DLCI neighbor table to bridge to specific DLCI neighbors. Static DLCI neighbors take precedence over neighbors learned dynamically with the LMI Protocol.

If LMI protocol is running consortium LMI, the valid range for subscriber numbers is 16 to 1022. For other LMI protocols, the range is 16 to 991.

2. Verify if transparent bridging is enabled for the appropriate wide area port or virtual port by entering:

```
SHoW -BRidge TransparenTBridge
```

By default, transparent bridging is enabled on all ports. If transparent bridging is disabled for the wide area port, you must enter:

```
SETDefault !<port> -BRidge TransparenTBridge = TransparenTBridge
```

3. Verify that bridging is enabled by entering:

```
SHoW -BRidge CONFIguration
```

4. If bridging is disabled, enable it for the system by entering:

```
SETDefault -BRidge CONTrol = Bridge
```


15.3

Configuring IP over Frame Relay

You can use this procedure to configure your device to transmit and receive data over a Frame Relay interface using the IP routing protocol.

You can configure an EOS device to simultaneously route multiple protocols over Frame Relay to one or more remote network connections. In the topology illustrated in the following figure, the local network supports TCP/IP traffic and routes information through a single Frame Relay connection to both types of remote networks.

Figure 92: Routing Multiple Protocols over Frame Relay Network



If your network is small and relatively stable, Motorola Solutions recommends that you configure the `-ARP CONTROL` parameter with the `NoInArp` value. This static address resolution reduces network overhead during initialization.

If your network is large and requires frequent reconfiguration, Motorola Solutions suggests that you configure the `-ARP CONTROL` parameter with the `InARP` value. This dynamic configuration can save you some network administration work.

InARP entries in the IP address table are learned when:


- IP addresses are configured.
- A new DLCI is available.

InARP entries in the IP address table are deleted when:

- The IP address table is flushed. After this occurs, InARP immediately sends out InARP requests and discovers new entries.
- An existing DLCI becomes unavailable.

To minimize the network overhead, after an IP address associated with a specific DLCI is discovered, it is treated as a static entry and is not aged out.

Figure 93: Configuring IP over Frame Relay



Prerequisites:

Ensure that:

- Ports and paths in your LAN are configured. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Frame Relay service is set up. See [Setting Up The Frame Relay Service on page 308](#).

Determine the IP addresses for each wide area port of your device that is attached to the Frame Relay network. If you use virtual ports, you must assign each virtual port a unique IP address on a separate IP subnet.

Obtain:

- Frame Relay DLCI addresses to create IP-to-FR mappings for static routes (static configurations only)
- IP addresses of neighbors that should send or receive RIP update packets



NOTE: If you configure your network for static address resolution over a Frame Relay network, make sure you set the `-ARP CONTROL` parameter to `NoInArp` using the `SETDefault !<port -ARP CONTROL = NoInArp` syntax.

Procedure:

1. Assign an IP address to each port or virtual port on each device that is directly attached to the Frame Relay network.

Step example:

To assign the address 128.1.0.1 to port 3 on router A, enter:

```
SETDefault !3 -IP NETaddr = 128.1.0.1
```

2. Perform one of the following actions:

- For static address resolution over a Frame Relay network, specify the IP-to-Frame Relay DLCI address mapping information.

Example:

To specify IP-to-Frame Relay DLCI mapping information for the devices directly attached to the Frame Relay network, you can use the following commands for Router A (IP address 128.1.0.1):

```
ADD -IP ADDRESS 128.1.0.2 @22
```

```
ADD -IP ADDRESS 128.1.0.3 @22
```

The valid range for Frame Relay DLCIs is 16 through 991 for user permanent virtual circuits. In this example, you can use DLCI values in place of @.



NOTE: You must specify this information for DLCIs associated with ports as well as virtual ports.

You can use similar commands on Router B (IP address 128.1.0.2) and Router C (IP address 128.1.0.3) to specify the IP address and DLCI mapping information.

- For dynamic address resolution over a Frame Relay network, enable the ARP Service to automatically discover IP addresses for the DLCIs.

Example:

To enable InArp on port 3, you must enter:

```
SETDefault !3 -ARP CONTro1 = InArp
```

3. Add each device on a Frame Relay network with which the system wants to communicate as a neighbor by performing the following actions:



NOTE: You can skip this step if the `DynamicNbr` option of the `-RIPIP` and `-OSPF CONTro1` parameters are enabled. If `DynamicNbr` is disabled, you must specify this information for ports as well as virtual ports.

- a. Specify a list of neighbor addresses to which RIP sends update packets.

Step example:

To transmit RIP packets from Router B, which runs RIP, to Router C, enter:

```
ADD !3 -RIPIP AdvToNeighbor 128.1.0.3
```

- b. Add IP addresses of neighbors on each device port participating in RIP.
- c. Do not change the default for all neighbors.
- d. Specify a list of neighbor addresses to which OSPF send update packets.

Step example:

To transmit OSPF packets from Router B, which runs OSPF, to Router C, enter:

```
ADD !3 -OSPF Neighbor 128.1.0.3
```

- e. Add IP addresses of neighbors on each device port that participates in OSPF.

4. Enable the dynamic routing protocols for IP by using RIPIP, OSPF, or Integrated IS-IS (IISIS) for each port and/or virtual port.

Step example:

To learn routes dynamically on port 3 by using RIPIP, you must determine if the Frame Relay network is fully meshed or nonmeshed.

If the Frame Relay network is fully meshed, enter:

```
SETDefault !3 -RIPIP CONTro1 = (TAlk, Listen, FullMesh)
```

If the Frame Relay network is non-meshed, enter:

```
SETDefault !3 -RIPIP CONTro1 = (TAlk, Listen, NonMesh)
```

Setting the `CONTro1` parameter to `TAlk` and `Listen` enables the device to send and receive routing information with other devices using RIP.

To enable routes dynamically on port 3 using OSPF, you must determine if the Frame Relay network is fully meshed or non-meshed.

If the Frame Relay network is fully meshed, enter:

```
SETDefault !3 -OSPF CONTro1 = (Enable, FullMesh)
```

If the Frame Relay network is non-meshed, enter:

```
SETDefault !3 -OSPF CONTro1 = (Enable, NonMesh)
```

You must configure all the OSPF neighboring devices with the same mode: `FullMesh` or `NonMesh`. Both modes apply to ports as well as virtual ports.

After OSPF operation is enabled, the device exchanges routing information with other devices using OSPF.

To enable routes dynamically by using IISIS, see [IP Routing Configuration on page 119](#).

5. Verify that IP routing is enabled on each device that is attached to the Frame Relay network by entering:

```
SHoW -IP CoNFIGuration
```

6. If IP routing was disabled, enable it by entering:

```
SEtDefault -IP CoNTRol = ROute
```

15.4

Multilink Frame Relay Configuration

Multilink Frame Relay (MFR) allows multiple virtual links (constituent virtual circuits or CVCs) of the same CIR to be multiplexed together under one virtual port (aggregated virtual circuit or AVC), thus providing an increased bandwidth for one port.

You can create an AVC by using the MFR option in the POrt service VirtualPort parameter, and you can define the attributes of the AVC by using the Frame Relay service MfrMinToTrunk, MfrFrag and MfrFrmLossTmr parameters.

To configure AVC, see [Adding an Aggregated Virtual Circuit on page 317](#).

For details about the POrt service VirtualPort parameter and the Frame Relay service MfrMinToTrunk, MfrFrag and MfrFrmLossTmr parameters, see the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: MFR is supported on HSSI, T1/E1, and X.21 interfaces. However the CVCs of an AVC must be of the same physical interface type (all HSSI, all T1, all E1, or all X.21) and each CVC in an AVC also must have the same physical speed.

Testing indicates that when MFR is enabled on an S2500 virtual port, device performance drops. Specifically, the S2500's CPU is not able to process MFR traffic efficiently and thus is only able to achieve approximately 90% of bandwidth on 64-byte packets using two T1 links and approximately 79% of bandwidth using E1 links.

15.4.1

Adding an Aggregated Virtual Circuit

Procedure:

1. Create the AVC by entering:

```
ADD !v<virtual_port> -Port VirtualPort MFR <P1@D1> [, <P2@D2> ... <P8@D8>]
```

where <P1@D1> ... <P8@D8> are the CVCs (defined by path number and DLCI address) that can be attached to the AVC.



NOTE: All CVCs attached to an AVC must have the same CIR.

Step example:

To attach path 3, DLCI 16; path 4, DLCI 17; and path 5, DCLI 18 to an AVC on virtual port 10, enter:

```
ADD !V10 -Port VirtualPort MFR 3@16, 4@17, 5@18
```

2. Define the minimum CIR value for the AVC by entering the following:

```
SEtDefault !<portlist> -FR MfrMinToTrunk = <CIR>
```

where <CIR> is the minimum CIR value, in kilobits, required to keep the MFR AVC active.

If the total CIR for the AVC falls below this value, the AVC is declared down.

3. Enable or disable fragmentation for MFR packets and configure the fragment size per AVC by entering:

```
SETDefault !v<virtual_port> -FR MfrFrag = ([Enable <bytes> (64-1500) | Disable])
```

The default setting is Disable.

4. Configure how long MFR waits before declaring packet(s) lost by entering:

```
SETDefault !v<virtual_port> -FR MfrFrmLossTmr = <milliseconds> (5-100)
```

The default setting is 20 milliseconds.

15.5

Priority Queuing and FRF.12 Fragmentation

To ensure that high-priority traffic is processed and delivered in a timely fashion, the EOS software implements priority queuing: traffic of different types is placed in different priority queues, depending on the requirements of each application, and the highest priority traffic is handled first.

The EOS software supports the following types of priority queuing:

- Realtime priority queuing delivers voice packets and certain types of packets supporting emergency applications in real time. Lower priority packets are not permitted to block the transmit path. The data path is always available should realtime packets appear.
- With non-realtime priority queuing via the priority queue handler (PQH), lower priority packets are permitted to go down the transmit queue when there are no high-priority or realtime packets present at the time.

Related to realtime priority queuing is Frame Relay fragmentation. Packets exceeding a specified size (FragPktSize) are fragmented to ensure that lower-priority packets do not cause undesirable delay to realtime packets. You configure fragmentation on a per-VC basis. When fragmentation is enabled on a VC, fragmented segments are interleaved with non-fragmented packets over the same VC as well as across all other VCs. This allows realtime packets to be scheduled in a timely fashion.

The Frame Relay service supports the following parameters related to realtime priority queuing and Frame Relay fragmentation:

- RealTimeQueue – Indicates the type of priority queuing to be used on the Frame Relay interface: realtime or non-realtime.
- QueueMethod – Indicates whether priority queuing is implemented on a per-VC basis or on a per-interface basis.
- RtgDEbitLevel – Indicates the priority level at and below which a packet is marked for discard (the DE bit is set to 1) if the Frame Relay network experiences congestion.
- FragPktSize – Specifies the maximum packet size, on a per-VC basis. Packets larger than the specified maximum are fragmented.
- FragType – Specified the type of fragmentation control or disables fragmentation on a specified VC.
- RtgFragLevel – Indicates the priority level at and below which packets are candidates for fragmentation.

For details about these parameters, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

15.5.1

Configuring Realtime Priority Queueing

Procedure:

1. If you want to configure a central site (core) device realtime priority queueing, perform the following actions:
 - a. Enable realtime priority queueing on the interface by entering:

```
SETD !<path> -FR RealTimeQueue = RealTimeQueue
```

- b. Specify that priority queuing is implemented on a per-VC basis by entering:

```
SETD !<path> -FR QueueMethod = VcBasedQueue
```

- c. Specify the priority level at or below which packets have their DE bits set to 1 (making them the first packets to be discarded if the Frame Relay network experiences congestion) by entering:

```
SETD !<path> -FR RtqDEbitLevel = <priority_level>
```

where *<priority_level>* is PR3, PR4, PR5, or PR6.



NOTE: The RtqDEbitLevel can be different than the RtqFragLevel.

2. If you want to configure a remote (edge) device realtime priority queuing, perform the following actions:

- a. Enable realtime priority queuing on the interface by entering:

```
SETD !<path> -FR RealTimeQueue = RealTimeQueue
```

- b. Specify that priority queuing is implemented on a per-path basis by entering:

```
SETD !<path> -FR QueueMethod = PathBasedQueue
```

- c. Specify the priority level at or below which packets have their DE bits set to 1 (making them the first packets to be discarded if the Frame Relay network experiences congestion) by entering:

```
SETD !<path> -FR RtqDEbitLevel = <priority_level>
```

where *<priority_level>* is PR3, PR4, PR5, or PR6.



NOTE: The RtqDEbitLevel can be different than the RtqFragLevel.

15.5.2

Configuring FRF.12 Fragmentation

To activate fragmentation, you must specify a type of fragmentation control (FragType) and set a maximum fragment size (FragPktSize).



NOTE: You can configure fragmentation independently for each VC.

Procedure:

1. Enable fragmentation by selecting the type of fragmentation control:

- For standard FRF.12 fragmentation, enter:

```
SETDefault !<portlist> -FR FragType = <vcid> Static
```

- For priority-based FRF.12 fragmentation, enter:

```
SETDefault !<portlist> -FR FragType = <vcid> PriorityBasedFRF12
```

Priority-based FRF.12 fragmentation allows the device to fragment one packet of each priority level in the same direction at the same time. This prevents large fragmentable packets from being backed up behind large fragmentable lower-priority packets. For details, see [FRF.12 Fragmentation on page 332](#).

2. Specify the maximum fragment size (30-1500 bytes) by entering:

```
SETDefault !<portlist> -FR FragPktSize = <vcid> <FragPktSize>
```



NOTE: When fragmentation is enabled on a VC, any packets exceeding the size you specify are fragmented.

3. Optional: Specify the priority level by entering:

```
SETDefault !<portlist> -FR RtqFragLevel = <vcid> [PR3|PR4|PR5|PR6]
```

The available priority levels are PR3, PR4, PR5, and PR6. This configuration is useful when multiple data streams with different priority levels are carried within a VC.

The priority levels map to TOS values at and below which traffic is subject to fragmentation. For the mapping between priority levels and TOS values, see [Priority Queue-to-TOS Mapping on page 320](#).

You can use the <vcid> option to specify the DLCI number corresponding to the virtual port for which you want to set the priority level.



NOTE: The RtqFragLevel can be different than the RtqDEbitLevel.

15.5.3

Priority Queue-to-TOS Mapping

Depending on the priority level that you specify for the RtqFragLevel or RtqDEbitLevel parameter, packets with TOS values at or below a certain level are fragmented.

Table 31: Priority Queue-to-TOS Mapping

Priority Level	TOS Value
7	7
6	6
6	5
5	4
5	3
4	2
4	1
3	0

The value in the Priority Level column indicates the priority level at which you are setting the RtqFragLevel or RtqDEbitLevel parameter. The highest corresponding value in the TOS Value column indicates the TOS value at and below which packets are fragmented (RtqFragLevel) or the DE bit is set to 1 (RtqDEbitLevel) if you specify this priority level. For example, if you set the RtqFragLevel parameter to 6, packets with TOS values of 6 or less are fragmented; if you set the RtqFragLevel parameter to 4, packets with TOS values of 2 or less are fragmented; and so on.

When you set the RtqFragLevel or RtqDEbitLevel parameter, an informational message indicates the TOS level at and below which packets are fragmented (RtqFragLevel) or the DE bit is set to 1 (RtqDEbitLevel).

15.6

FRIPHC Negotiation

EOS is interoperable and is compliant with FRF.Iphc, a Frame Relay Forum (FRF) standard that proposes two peers negotiate various parameters before starting to compress IP packets. EOS implementations of CRTP over Frame Relay prior to release 11.6.025 did not employ any negotiation between peers and used proprietary Frame Relay control byte and NLPID values. Beginning with release 11.6.025, EOS employs negotiation between peers and uses the control byte and NLPID values proposed in FRF.Iphc.

You can enable or disable FRIPHC (Frame Relay IP header compression control) negotiation using the FrfIPhc parameter.

For details, see “Frame Relay Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: The EOS implementation of FRF.Iphc is not backwards-compatible with EOS version 11.4.

15.6.1

Parameters Negotiated Between Peers

In order to establish compression of IP datagrams sent over a Frame Relay VC, each end of the VC must agree on a set of configuration parameters to be used for the compression. This negotiation takes place as soon as the VC comes up; the “PORT_UP” event on the virtual port triggers exchange of control frames between the peers.

Negotiation consists of a simple handshake to enable the IP header compression algorithms and to agree upon the header compression parameters for each direction of the VC. Each direction is negotiated independently, so it is possible to have header compression enabled in one direction only, or to have different algorithms and/or parameters enabled in each direction.

Packets are sent uncompressed over the VC until both sides agree or completed negotiation of the parameters. Negotiation begins as soon as the virtual port comes up. Any (UDP/RTP data) traffic sent while the negotiation is taking place is transmitted uncompressed.

When negotiation is complete, console messages provide the status of the FRF.Iphc negotiation:

- If the negotiation is successful, a message similar to the following appears:
Thu Oct 11 09:38:24 2001 Port V3 Frf.Iphc in OPERATIONAL state
- If FRF.Iphc is disabled on the remote node, a message similar to the following appears:

Thu Oct 11 09:40:43 2001 Port V3 Frf.Iphc in DISABLED state



NOTE: If the `FrFiphc` parameter on the local port is set to `NoNegotiate`, no console messages are displayed; however, the `SHOW -FR FrfIPhc` command displays a status of `FORCED OPERATIONAL`, meaning that the port did not negotiate but uses the compliant control byte and NLPID values.

Table 32: Values Used by EOS in IP Header Compression Negotiation

Field	Value use in EOS
IP type ID	4 (IPv4)
TCP_SPACE	0 (does not compress TCP packets)
NON_TCP_SPACE	255 (supports UDP_8 and RTP_8)
F_MAX_PERIOD	0
F_MAX_TIME	0
MAX_HEADER	168 (no limit; suggested value)
Suboptions	0 (no suboptions included)

During negotiation, the state machine uses a *handshake timer* to dictate the interval after which a request is repeated and a *counter* to determine the maximum number of times to try a request. The following values are hard-coded in the EOS implementation:

- Handshake timer – 3 seconds (range 1-10)
- Counter – 4 (range 1-10)

15.7

Configuring Disaster Recovery in a Frame Relay Environment


Disaster recovery is a mechanism that allows a user to maintain connectivity between central and remote sites in the event of failure of a physical line or the Frame Relay network. This feature provides a way to recover from the loss of a primary permanent virtual circuit (PVC) in a Frame Relay network by triggering a backup PVC.

If the primary PVC becomes unavailable, as determined by the LMI Protocol, the traffic destined for the primary PVC is forwarded over the backup PVC, maintaining connectivity between nodes. Upon recovery of the primary PVC, the backup PVC is deactivated, and traffic is again forwarded over the primary PVC. You can configure the backup PVC on a separate link to provide redundancy, and the backup PVC can be a leased or a dial-up link.

You can use this process to configure a primary line, a backup PVC, a backup PVC on a separate link, and a backup PVC on a separate dial-up link in the event of the failure of a primary line. You must configure these lines on both ends of the link.

The following figure shows a configuration that is fully redundant between site A, the central site, and remote sites B and C. At site A, the primary PVCs are on link 5, and the backup PVCs are on link 6. Both of these lines are leased lines. At site B, the primary PVC is on link 2, and the backup link is on link 3. Link 3 at site B is a dial-up line. At site C, the primary PVC is on link 3, and the backup PVC is on a dial-up link 2.

Figure 94: Fully Redundant Network Among All Sites



Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Your wide area interfaces are configured.
- Virtual ports are configured.
- Primary link is configured. See [Setting Up The Frame Relay Service on page 308](#).
- You acquired services from a Frame Relay service provider.

Process:

1. Configure a primary PVC. See [Configuring a Primary PVC on page 323](#).
2. Configure a backup PVC. See [Configuring a Backup PVC on page 323](#).
3. Configure a backup link. See [Configuring a Backup Link on page 324](#).

15.7.1

Configuring a Primary PVC

To configure a primary PVC, you must first set up a virtual port with that PVC. Before setting up virtual ports, you must ensure that the owner of the wide area port associated with the path through which the virtual ports are defined is set to Frame Relay. You must create a virtual port for each remote network that is attached to a Frame Relay cloud. You also must configure the primary link according to [Setting Up The Frame Relay Service on page 308](#).

Procedure:

1. Configure the primary PVC between site A and site B by performing the following actions:

- a. Set up a virtual port for site A by entering:

```
ADD !V1 -PORT VirtualPort 5@30
```

This command designates 5@30 as a primary PVC.

- b. Set up a virtual port for site B by entering:

```
ADD !V1 -PORT VirtualPort 2@30
```

This command designates 2@30 as a primary PVC.

When creating virtual ports, you must designate the same PVC on both ends of the connection as primary.

2. Configure the primary PVC between site A and site C by performing the following actions:

- a. Set up a virtual port for site A by entering:

```
ADD !V2 -PORT VirtualPort 5@40
```

This command designates 5@40 as a primary PVC.

- b. Set up a virtual port for site B by entering:

```
ADD !V1 -PORT VirtualPort 3@40
```

This command designates 3@40 as a primary PVC.

When creating virtual ports, you must designate the same PVC on both ends of the connection as primary.

15.7.2

Configuring a Backup PVC

Procedure:

1. Configure the backup PVC from site A to site B by entering:

```
ADD !V1 -FR BackupPVC 6@45
```

2. Configure the backup PVC from site B to site A by entering:

```
ADD !V1 -FR BackupPVC 3@45
```

You must designate the same PVC as the backup on both ends of the connection between sites A and B.

3. Configure the backup PVC from site A to site C by entering:

```
ADD !V2 -FR BackupPVC 6@55
```

4. Configure the backup PVC from site C to site A by entering:

```
ADD !V1 -FR BackupPVC 2@55
```

You must designate the same PVC as the backup on both ends of the connection between sites A and C. By default, the port is brought down when the primary PVC fails, even when the backup PVC is available.

15.7.3

Configuring a Backup Link

You can add a backup PVC to a previously configured virtual port to provide redundancy. If the backup PVC is on a separate path, you must attach this path to a separate port.

Procedure:

1. Enable the Frame Relay Service by setting the owner of the serial interface to Frame Relay by entering:

```
SETDefault !6 -PORT OWNEr = FrameRelay
```

2. If your Frame Relay network supports the LMI Protocol, ensure that the appropriate LMI Protocol is enabled.

For more information about enabling the LMI Protocol, see [Setting Up The Frame Relay Service on page 308](#).

Step example:

To manually enable the LMI Protocol, enter:


```
SETDefault !6 -FR CONTro1 = LMI
```

15.8

Frame Relay Networks

The wide area device supports both bridging and routing of multiple protocols over Frame Relay. In most cases, only one wide area protocol runs over a path or serial line in a Frame Relay environment; however, you can configure dual permanent virtual circuits (PVCs) bound to a single logical port on a boundary router.

Figure 95: Configuring Dual PVCs Over a Single Logical Port



To add a backup PVC on path 3 for logical port 2, you must perform the following actions:

1. Configure port 2 of Router A for Frame Relay over a backup PVC 3 at DLCI 21 by using:

```
ADD !2 -FR BackupPVC 3@21
```


2. Configure port 2 of Router B for Frame Relay over a backup PVC 2 at DLCI 21 by using:

```
ADD !2 -FR BackupPVC 2@21
```

Frame Relay allows your device to transmit and receive data from any other device on the Frame Relay network over a PVC and over an SVC, which provide virtual connections to all other nodes on the network.

15.8.1

PVC and SVC Connections

PVCs are permanent circuits between the remote user and the local device. SVCs are dynamic circuits set up between the remote user and the local device and are disconnected after data activity stopped. You can also configure an SVC as a static connection that works like a PVC by setting the `-FR SvcIdleTimer` setting to 0 or None. For details, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

PVCs use the data link circuit identifier (DLCI) numbers assigned by the Frame Relay service provider to establish the connection between the remote user and the local device logical or virtual port. The DLCI numbers assigned by the service provider can range from 0 to 1023. The service provider assigns a subset from this range and reserves the rest. The range assigned depends on the protocol for which the device port was configured:

- For ports configured for ANSI or NTTLMI, the service provider assigns DLCI numbers ranging from 16 to 991. Numbers 0 to 15 and 992 to 1023 are reserved and not assigned.
- For ports configured for other LMI protocols, the service provider assigns numbers ranging from 16 to 998. Numbers 0 to 15 and 999 to 1023 are reserved and not assigned.

When you configure SVCs, you assign a virtual circuit identifier `<vcid>` number for the SVC that is mapped to a number from the numbers in the DLCI range that were not assigned for PVCs by the service provider. For example, if the service provider assigned DLCI numbers 16 through 50, 85, 87, and 89 for PVCs from the range of 16 to 991, you can assign vcid numbers for SVCs from the other numbers still available within the range of 16 to 991. Although an SVC is assigned a virtual circuit identifier number that maps to a specific DLCI number, because SVCs are dynamic, any virtual circuit available (other than those assigned for PVCs) when a connection is needed is used by the SVC to establish the connection.

PVC Connection

For PVCs, the connection between the remote user and the local device port through the Frame Relay network is established when you enter the DLCI number assigned by the service provider. You must enter the DLCI number when you configure the individual protocol that transmits through the Frame Relay network.

SVC Connection


Unlike when you establish a PVC connection, to establish an SVC connection between a remote user and the local device port through the Frame Relay network, you must perform [Establishing an SVC Connection on page 326](#). In this procedure, you enter a virtual circuit identifier (vcid) number that identifies the SVC. The vcid number must match a number from the range of DLCI numbers available and not assigned by the service provider for PVCs.

The following figure shows device port 2 (in Santa Clara), over Frame Relay, connecting a remote user:

- In New York on an SVC assigned vcid 30. The `-FR SvcLocalAddress` parameter is used to configure vcid 30 with telephone number 1-408-764-5734, and the `-FR SvcDestAddress` parameter is used to configure the vcid 30 with telephone number 1-718-858-0859.
- In Boston on an SVC assigned vcid 35. This vcid was configured with a telephone number 1-408-764-6848 at the local site and 1-617-648-7652 at the remote site.
- In Atlanta on an SVC assigned vcid 40. This vcid was configured with a telephone number 1-408-764-8532 at the local site and 1-404-596-8598 at the remote site.

- In Miami on a PVC through DLCI 25. Since this DLCI is a PVC, a configuration with a telephone number is not required.

Figure 96: SVC and PVC Connections over Frame Relay



Annotation	Description
1	Use SvcLocalAddress to convert to: 1-408-764-5734
2	Use SvcDestAddress to convert to: 1-718-858-085

15.8.1.1

Establishing an SVC Connection

Procedure:

1. When configuring the individual protocol that transmits through Frame Relay, enter the vcid number that identifies the SVC used for the connection.
2. Identify the local customer premises equipment (CPE) link address (telephone number) assigned by your Frame Relay network service provider for your device to identify the local device port to the network by entering:

```
SETDefault !<portlist> -FR LinkAddress = [ <E.164 address> (1-15 digits) | <X.121 address> (1-15 digits)
```

The -FR LinkAddress parameter identifies the local device port for all SVCs configured for that port.

3. Enter the destination address (the telephone number) of the remote user associated with the device virtual circuit ID:

```
ADD !<portlist> -FR SvcDestAddress = <vcid> [ <E.164 address> (1-15 digits) | <X.121 address> (1-15 digits)
```

where <vcid> specifies a virtual circuit identifier number that is mapped to a DLCI number from the range of numbers available that were not assigned by your service provider for a PVC connection.

4. Identify the local SVC telephone number associated with a specified virtual circuit ID (<vcid>) to the remote user (optional) by entering:

```
SETDefault !<portlist> -FR SvcLocalAddress = <vcid> [ <E.164 address> (1-15 digits) | <X.121 address> (1-15 digits)]
```

where *<vcid>* specifies a virtual circuit identifier number that is mapped to a DLCI number from the range of numbers available that were not assigned by your service provider for a PVC connection. The CPE link address is used to identify the local virtual circuit if the address of the local SVC is not set.


15.8.2

Frame Relay Topologies

Fully Meshed Frame Relay Topology

A fully meshed Frame Relay topology, is a topology in which each node on a network is directly connected to all other nodes on the network. Each node is connected to the other nodes through a PVC, and each PVC has a DLCI associated with it. This DLCI may appear as a different number to each end of the PVC.

Figure 97: Fully Meshed Frame Relay Topology




This topology is composed of four routers (A, B, C, and D). Through the established PVCs, Router A is connected to Router B, Router C, and Router D; Router B is connected to Router A, Router C, and Router D; and so on.

Non-meshed Frame Relay Topology

A non-meshed Frame Relay topology, is a topology where each node on a network is not necessarily connected to all other nodes on the network.


Figure 98: Non-meshed Frame Relay Topology



This topology is composed of four routers (A, B, C, and D). Through the established PVCs, Router A is connected to Router B, Router C, and Router D. Router B, Router C, and Router D are connected to Router A only, not to one another.

 **NOTE:** Transparent bridging does not operate correctly in some non-meshed topologies. For example, in the following figure, the transparent bridge properly forwards traffic received on !v1 to !v2. However, traffic received from one of the transparent bridge's remote connections on !v3 is not properly forwarded to the other two remote connections on !v3. Therefore, do not configure transparent bridging in this type of non-meshed topology. The flooding algorithm floods packets on a per-port basis, not on a neighbor-per-port basis.


Figure 99: Transparent Bridging Unsuitable in some Non-meshed Topologies



Partially Meshed Frame Relay Topology

A partially meshed Frame Relay topology is a topology in which some nodes on a network are directly connected to nodes on the network (as in a fully meshed topology) and other nodes are not (as in a non-meshed topology).

Figure 100: Partially Meshed Frame Relay Topology



This topology is composed of four routers (A, B, C, and D). Through the established PVCs, Router A, Router B, and Router C are connected to one another, but Router D is connected to Router A only.

You can use the following solutions to work around the lack of connectivity between Router B, Router C, and Router D in nonmeshed and partially meshed topologies:

- Next-hop split horizon

If you route IP-RIP, this protocol offers the next-hop split horizon feature.

In IP-RIP, you can enable next-hop split horizon by setting `-RIP IP CONTro1` to `NonMesh`. In the partially meshed topology, if you route IP-RIP and you set `-RIP IP CONTro1` to `NonMesh`, a list of neighbors containing Router B, Router C, and Router D is generated by the system, or you can configure them as neighbors using the `-RIP IP AdvToNeighbor` parameter.

For more information, see “RIP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

- Virtual ports

Another solution for the lack of connectivity in non-meshed and partially meshed Frame Relay topologies is to create virtual ports. Virtual ports are supported by bridging and all routing protocols over a Frame Relay network.

For more information on virtual ports, including the number of virtual ports supported per platform, see [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).

15.8.3


Frame Relay Addresses

Before attaching your device to a Frame Relay network, you must obtain one or more virtual circuit identifiers, called DLCIs, from the Frame Relay service provider. A DLCI identifies a circuit between two devices from the end users’ perspective. Each end of the circuit can have a different DLCI number for the link.

The DLCI number can range from 0 to 1023, but the service provider only assigns subscriber numbers ranging from 16 to 991. For ANSI and NTT LMI, 0 to 15 and 992 to 1023 are reserved. For LMI, 0 to 15 and 999 to 1023 are reserved.

In the following figure, Router A, Router B, Router C, and Router D are assigned the DLCI numbers 36, 38, 40, and 41, respectively.

Figure 101: Frame Relay Addressing Example



The following are examples of what occurs when packets are sent from one device to another:

- When Router A sends a packet to Router B, it uses DLCI 38. When the packet arrives at Router B, the network changes the DLCI to 36 to indicate to Router B that the packet originated at Router A.
- When Router A sends a packet to Router C, it uses DLCI 40. When the packet arrives at Router C, the network changes the DLCI to 36 to indicate to Router C that the packet originated at Router A.
- When Router B sends a packet to Router C, it also uses DLCI 40. When the packet arrives at Router C, the Frame Relay network changes the DLCI to 38 to indicate to Router C that the packet originated at Router B.

EOS devices can operate in both local and global addressing schemes used by the Frame Relay network.

- In the standard (local) addressing convention, the DLCI number has only local significance; a duplicate number can be used by other devices.
- In the global addressing convention, identifiers used throughout the Frame Relay network are unique, and all traffic to a node has the same destination DLCI number.

15.8.4

Local Management Interface (LMI) Protocol

The LMI Protocol runs between the device data terminal equipment (DTE) and the Frame Relay network switching equipment data communications equipment (DCE). The LMI Protocol provides information about all devices that are accessible on the Frame Relay network by listing all DLCIs connecting the local system with the remote devices. The LMI Protocol improves reliability between the DTE and DCE devices by exchanging keepalive packets that are sent every 5 to 30 seconds, depending on the configuration. If the LMI Protocol is disabled, the device operates under the assumption that all the DLCIs are active whether they are up and running or not.

Signaling that is used in LMI can include bidirectional signaling or unidirectional signaling (a subset of bidirectional signaling).

- In unidirectional signaling, Status Enquiry messages are sent in one direction, and Status messages are sent in the other direction in response to the Status Enquiry messages.

- In bidirectional signaling, both Status Enquiry and Status messages are sent in both directions.

You can use several FR service parameters to configure how LMI is run on a port. These parameters consist of:

- Timer Timeout Values
 - T391 (implemented as `KATime`)
 - T392 (implemented as `PVTime`)
- Count Limit Values
 - N391 (implemented as `FSEInterval`)
 - N392 (implemented as `DCEErrThreshold` and `DTEErrThreshold`)
 - N393 (implemented as `DCEMonEvent` and `DTEMonEvent`)

The T391 value determines how often Status Enquiry messages are expected to be received on a port. Every N391th Status Enquiry message that is sent out of a port requests Full Status information. N392 and N393 values are used in error processing; when N392 errors are encountered in the last N393 events, the port is considered failed.

You can use all of the parameters listed above if you are running bidirectional signaling; you can use only some if you are running unidirectional signaling.

- On every port where you are running DTE LMI, you need to configure T391 (`KATime`), N391 (`FSEInterval`), N392_DTE (`DTEErrThreshold`), and N393_DTE (`DTEMonEvent`) properly.
- On every port where you are running DCE LMI, you need to configure T392 (`PVTime`), N392_DCE (`DCEErrThreshold`) and N393_DCE (`DCEMonEvent`) properly.
- Since bidirectional LMI consists of running both DTE LMI and DCE LMI, on every port where you are running bidirectional LMI, you need to configure all the parameters listed above properly.

For details about LMI parameters, see “Frame Relay Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: The LMI Protocol is enabled by default on your EOS device. Some Frame Relay switches, however, do not run the LMI Protocol. In this situation, set the `-FR CONTro1` parameter to `NoLMI`.

15.8.5

C RTP Twice Algorithm

The C RTP Twice Algorithm is a feature designed to minimize the number of packets dropped by C RTP when a single voice packet is dropped due to a CRC error.

C RTP is used to minimize bandwidth usage by compressing the IP, UDP, and RTP headers of a voice packet. With the C RTP Twice Algorithm, the C RTP decompressor handles incoming packets as follows:

- Compares the four-bit sequence number for the received packet to the expected value (the sequence number of the current packet +1). If the difference is nonzero, the decompressor sends a context state packet to the compressor but does **not** invalidate the context.
- Assumes that lost packets had the same deltas (as the deltas for ID, RTP sequence number and RTP timestamp) calculated from the current out-of-order packet. If the current packet does not have all the deltas, the decompressor uses the deltas stored in the context to decompress.
- If a non-zero UDP checksum exists in the current out-of-order packet, calculates the UDP checksum of the reconstructed packet and drops the packet only if the checksums do not match.

EOS supports both 8-bit and 16-bit context IDs (CIDs) in C RTP.

- 8-bit CIDs allow for 256 sessions per direction (receive/transmit) per port.
- 16-bit CIDs allow for 512 sessions per direction per port for non-MFR ports and 2048 sessions per direction per port for MFR AVC ports.

If a port running 8-bit or 16-bit CID support occupies all the possible sessions (256 for 8-bit and 512/2048 for 16-bit), further streams are forwarded on without compression.

You can use the IP service `RtphcCidBits` parameter to configure 8-bit or 16-bit CID support. For details, see “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

15.8.6

FRF.12 Fragmentation

Each VC is an independent leased circuit with its own speed and configured CIR that supports a specific application. For VCs supporting realtime applications, such as voice, long packets are fragmented to minimize the time that the transit path is occupied. Fragmentation and reassembly is turned on only if the VC is configured for fragmentation. If fragmentation is configured for a specific Frame Relay VC, all packets to be transmitted to the specified VC are examined to determine if they are candidates for fragmentation and reassembly per FRF.12.

When FRF.12 fragmentation is enabled on a VC, non-delay-critical packets exceeding a specified MTU are fragmented to prevent these long packets from potentially blocking the transmit path. Fragmentation is configured on a per-VC basis, and the fragmentation procedure is transparent to the Frame Relay network interconnecting the DTEs. Each VC manages its own fragmentation requirements during transmission, and the reassembly module at the peer collects all fragments of a packet for delivery to Frame Relay applications. To properly support realtime data transfer with multiple priority levels, FRF.12 fragmentation is performed only when a frame is permitted to be transmitted.

If a packet consists of many fragments, only the first n fragments may be sent directly to the device driver queue. The value of n depends on the CIR for the VC and the latency requirement. In general, a one-fragment delay is considered to be acceptable for realtime applications such as voice. The remaining fragments of the original long packet are placed in the fragment queue for later transmission.

Fragments are transmitted as permitted by the available bandwidth and the delay tolerance so that non-fragmented packets of higher priority levels can interleave with fragments. Fragments belonging to different VCs may be interleaved in transmission, but only the non-fragmented packets of a priority level higher than the highest level that supports fragmentation (as defined by the `RtqFragLevel` parameter) can interleave with fragmented packets. Once a packet is fragmented, it is either transmitted immediately or it is sent to the fragment queue.

FRF.12 fragmentation supports two logical data streams within a VC: one for delay-critical traffic and one for non-delay-critical traffic. The EOS implementation of priority queueing, on the other hand, supports five priority levels and thus five logical data streams. To support various applications using the priority scheme without exceeding the (two data stream) capability of FRF.12, fragmented packets belonging to different priority queues temporarily share one logical fragment data stream. This ensures that packets are transmitted and received in order and prevents out of sequence transmission.

A Frame Relay priority queue list for all Frame Relay ports is established when the Frame Relay service becomes operational. Each port addresses all the active VCs on the interface. For each VC, five queues are supported - one for each priority level. In addition, a special fragment queue is supported for fragments of packets in transit, for a total of six queues for each VC. When fragmentation takes place, fragments belonging to a lower priority queue are promoted to the highest priority queue that supports fragmentation (as defined by the `RtqFragLevel` parameter). For example, consider a scenario in which priority level 5 is the highest level that supports fragmentation (`RtqFragLevel` is set to PR5). When a priority level 3 long packet is fragmented, the first n fragments - as permitted by the delay bandwidth requirement - are transmitted to the device driver, and the remaining fragments are placed in the priority 5 queue.

FRF.12 Packet Interleaving (Priority-Based FRF.12 Fragmentation)

FRF.12 packet interleaving (priority-based FRF.12) is a custom implementation of FRF.12 fragmentation. In standard FRF.12, the device can fragment only one packet at a time in each direction. Priority-based FRF.12 allows the device to fragment one packet of each priority level in the same direction at the same time. This prevents large fragmentable high-priority packets from being backed up behind large fragmentable

lower-priority packets. If a fragmentable high-priority packet is received while the device is sending fragments from a lower-priority packet, the device fragments and sends the high-priority packet immediately, then finishes sending the fragments from the lower-priority packet.

The following EOS user interface changes from version 11.8 to version 12.0 support the addition of priority-based FRF.12:

- The available types of FRF.12 fragmentation are modified. In previous implementations, you could specify the following fragmentation types:

FCP – Fragmentation Control Protocol (a modified form of standard FRF.12 fragmentation)

Static – Standard FRF.12 fragmentation

With priority-based FRF.12, you can now specify the following fragmentation types:

Static – Standard FRF.12 fragmentation

PriorityBasedFRF12 – Priority-based FRF.12 fragmentation

- The minimum fragmentation packet size (configurable with the FR service `FragPktSize` parameter) was decreased from 64 bytes to 30 bytes.

For details about these parameters, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE:

Because priority-based FRF.12 is not an industry standard form of FRF.12, it is supported only on Motorola Solutions MNR routers and GGM 8000 gateways and not on any other company's communication devices.

FRF.12 is only run for packets that have a priority level up to the level configured with the `RtqFragLevel` parameter.

15.8.7

Frame Relay Congestion Control

Frame Relay network data traffic running over an individual virtual circuit of an device port exceeds network capacity at times. You can configure two types of Frame Relay congestion control for individual virtual circuits to monitor and control the throughput of outgoing data when network capacity is exceeded.

- One for device ports configured for NTTLMI protocol
- One for device ports configured for other LMI protocols.

Congestion control, for device ports configured to use an LMI protocol other than NTTLMI, optimizes the traffic throughput of a virtual circuit when the traffic exceeds the network capacity by having the device reduce its traffic load. After the congestion disappears, the device increases its traffic load again.

When a Frame Relay network is congested, it sets the frames with the Backward Explicit Congestion Notification (BECN) bit set to 1 to notify the device of the congestion. When the first BECN=1 frame arrives at the device port to activate congestion control, the current throughput rate on the virtual circuit is reduced to a maximum rate equal to the value of the committed information rate (`<cir>`).

To turn on congestion control and to set up how many consecutive BECN=1 frames must be sent before the current throughput rate on the virtual circuit is reduced to a maximum rate below the value of `<cir>` for the `<vcid>`, you must use:

```
SETDefault !<portlist> -FR CongestControl <vcid> <No / Yes> <step_(1-999)>
```

How the congestion control is executed is determined by the values you enter for the `- FR CIRbothdir` parameter, which uses the following syntax:

```
SETDefault !<portlist> -FR CIRbothdir <vcid> <cir> <mincir> <Bc> <Be>
```

For more information about how to configure Frame Relay Congestion Control, see:

[Configuring Congestion Control for NTTLMI Protocol Users on page 309](#)

[Configuring Congestion Control for Other LMI Protocol Users on page 310](#)

[Frame Relay Congestion Control Configuration Example on page 311](#)

When configuring congestion control, you must observe the following rules:

- Setting *<step>* to too large a number means that sufficient BECN=1 frames may not reach the device port in time to activate congestion control while the traffic overload is occurring.
- You should configure the *<cir>* and *<Bc>* values for a reasonable time interval (*Tc*) because $Tc = \frac{Bc}{cir}$. Typically, you should configure these values for a *Tc* of one second.
- If you set the value of *<Be>* to be nonzero, the device switches to a continuous monitoring throughput mode. This means that the device continuously monitors its throughput even if there are no BECN=1 frames coming in, which lowers device performance.

When the first BECN=1 frame arrives at the device port to activate congestion control, the current throughput rate on the virtual circuit is reduced to a maximum rate equal to the value of *cir*. When the configured number of consecutive BECN=1 frames arrive at the device port, the maximum throughput rate lowers 1/4 to 0.750 times the value of *cir* as listed in the following table.

Table 33: Congestion Control Throughput Rate Reduction

Current Device Transmission Throughput Rate	Throughput Rate Reduction
$cir \geq \text{Current Rate} > (0.750) \times cir$	$(0.750) \times cir$
$(0.750) \times cir \geq \text{Current Rate} > (0.500) \times cir$	$(0.500) \times cir$
$(0.500) \times cir \geq \text{Current Rate} > (0.250) \times cir$	$(0.250) \times cir$
$(0.250) \times cir \geq \text{Current Rate}$	No Action

Further throughput rate reduction occurs only when the next number of consecutive BECN=1 frames arrive at the device port. After the reduction, a throughput increase of 1/8 the existing maximum throughput occurs when the port receives a number of consecutive BECN=0 frames equal to 1/2 of the number of consecutive BECN=1 frames configured for *<step>*.

For example, if you configured *<step>* for 5 consecutive BECN=1 frames for a particular port, 3 BECN=0 frames are required to increase the throughput by 1/8 of the existing maximum throughput because 2.5, which is half of 5, is rounded up to 3 frames.

The following example shows how congestion control reduces maximum throughput rates for a port as it receives BECN=1 frames, and increases maximum throughput rates as it receives BECN=0 frames.

In this example, *<cir>* is equal to 100 kbps, and *<step>* is 5:

1. The port receives one BECN=1 frame and the maximum transmit rate is 100 kbps or *<cir>*.
2. The port receives 5 consecutive BECN=1 frames, which limits the maximum transmit rate to 75 kbps.
3. The port receives 5 more consecutive BECN=1 frames, which limits the maximum transmit rate to 50 kbps.
4. The port receives 3 consecutive BECN=0 frames, which increases the maximum transmit rate to 50 times 1.125 = 56.25 kbps.
5. The port receives 3 more consecutive BECN=0 frames, which increases the maximum transmit rate to 56.25 times 1.125 = 63.28 kbps.
6. The port receives 5 more consecutive BECN=1 frames, which increases the limits to the maximum transmit rate equal to 50 kbps.

15.8.8

Disaster Recovery over Frame Relay

Disaster recovery is a mechanism that allows you to maintain connectivity between your central and remote sites in the event of a primary line failure or the Frame Relay network failure. This section describes how to use virtual ports and explains possible points of failure in a Frame Relay network.

Virtual Ports Used for Disaster Recovery

To configure disaster recovery, you must create virtual ports. A PVC attached to the virtual port is designated as the primary PVC. After you create a virtual port, you can add a backup PVC to that virtual port to support disaster recovery. For additional information about configuring virtual ports over Frame Relay, see [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).

You must configure virtual ports on both ends of the connection, and you must designate the same PVC on both ends of the connection as a primary PVC. When configuring a backup PVC, you must designate the same PVC on both ends of the connection as the backup PVC, as shown in [Figure 94: Fully Redundant Network Among All Sites on page 322](#). The PVCs 5@30 at site A and 2@30 at site B are both designated as primary PVCs. 6@45 at site A and 3@45 at site B are both designated as backup PVCs. If you designate a PVC as primary on one end and backup on the other end, all packets are dropped.

You can configure the backup PVC on a separate link to provide redundancy. If the backup PVC is on a separate link, this link can be a leased line or a dial-up line. You can use a dial-up line only on one end of the connection. In most cases, a permanent connection is established between the device at the central site and the Frame Relay switch. At the remote site, the device establishes a connection with the switch using a dial-up link.

When using a dial-up link on a Frame Relay network, only the EOS device can dial out by establishing a connection with the Frame Relay switch. The Frame Relay switch cannot dial out and can accept only incoming connections. When the switch accepts these incoming connections, it activates the PVC associated with that link. Since the switch cannot dial out to establish an end-to-end connection with the EOS device, you must establish a permanent connection between the device and the Frame Relay switch on one side of your configuration. Establishing this connection enables the other side to dial out if the primary PVC fails.

The following points of failure are possible on a Frame Relay network:

- Failure of a local path
- Failure of a local switch
- Failure of a remote path
- Failure of a remote switch
- Failure in the Frame Relay network

A fully redundant network is one on which there is no single point of failure. In a partially redundant network, at least one point of possible failure exists. Data is not transmitted across a backup PVC when the primary PVC is active, which ensures correct sequencing of packets.

The triggering of the backup PVC is kept transparent to the network layer protocols by using DLCI substitution. In the case of a primary PVC failure, Frame Relay sends and receives data using the backup DLCI by substituting the primary DLCI for the backup DLCI in the packet before passing data to the network layer protocols. As long as a PVC exists that can carry the traffic on a virtual port, there is no change in port status, and the network layer protocols are unaware that the backup DLCI is used.

In a Frame Relay disaster recovery environment, sessions may need to be restarted if a failure occurs on the primary PVC. By default LMI and Annex-D LMI provide full status information for the DLCIs every 60 seconds. If a failure occurs at one end of the primary PVC, it may take up to 60 seconds to inform the other end of the PVC. If you are using a dial-up line for the backup PVC, additional time may be necessary to establish the connection. These delays can cause session timeout.

Partially Redundant Networks

The following examples show the locations of redundant links and possible points of failure in partially redundant networks.

Partially Redundant Network Point of Failure – Example 1

In the following figure, central site A is connected to remote site B. A redundant link (shown as the dotted line) is configured at site B, but not at site A. In this configuration, a failure of the link at site A or a failure of the Frame Relay network can bring down the connection between site A and site B.

Figure 102: Partially Redundant Network with Redundant Link at Site B




Table 34: Covered and Uncovered Links – Partially Redundant Network with Redundant Link at Site B

Links Covered for Failure	Links Not Covered for Failure
Link failure at site B	Link failure at site A
	Frame Relay network failure

Partially Redundant Network Point of Failure – Example 2

The following figure shows a partially redundant configuration in which the redundant link is located at site A. At site B, a link failure or a Frame Relay network failure can bring down the connection between the two sites.

Figure 103: Partially Redundant Network with Redundant Link at Site A




Table 35: Covered and Uncovered Links – Partially Redundant Network with Redundant Link at Site A

Links Covered for Failure	Links Not Covered for Failure
Link failure at site A	Link failure at site B
	Frame Relay network failure

Partially Redundant Network Point of Failure – Example 3

The following figure, the network configuration has redundant links at both site A and site B. The point of failure in this configuration is the Frame Relay network

Figure 104: Partially Redundant Network: Redundant Links at Site A and Site B




Table 36: Covered and Uncovered Links – Partially Redundant Network: Redundant Links at Site A and Site B


Links Covered for Failure	Links Not Covered for Failure
Link failure at either end	Frame Relay network failure

Fully Redundant Networks

The following example shows the location of a redundant link between two sites.

The following figure shows a Frame Relay configuration that is fully redundant between site A and site B. There is no single point of failure between sites A and B, because redundancy is provided for all possible points of failure between these nodes.

Figure 105: Fully Redundant Network Between Site A and Site B



Sites A and B have redundant links, and both sites A and B are connected to two different Frame Relay networks. DLCI 40 is the primary PVC going between site A and site B. DLCI 45 is the backup PVC between site A and site B. Each of these PVCs belongs to a different Frame Relay network accessible through different

device interfaces. The primary and backup PVCs are on different interfaces. The network layer protocols believe they are talking to DLCI 40, even when the backup DLCI 45 took over. In this configuration, there is no redundancy between site A and site C.

15.9

Frame Relay Queue-Depth Exceeded Traps and Variables

The Frame Relay MIB (RFC 1315) includes the following queue-depth-exceeded traps and variables:

Queue-depth-exceeded traps:

- `Caution level` – Sent when the Frame Relay WAN link exceeds the user-configured caution threshold (a percentage of the queue limit, specified with the Frame Relay service `CautionTrap` parameter or via SNMP using the variables in the Circuit Statistics table). For more information, about the `CautionTrap` parameter, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.
- `Queue-depth exceeded` – Sent when the Frame Relay WAN link exceeds the queue limit.
- `Total queue-depth exceeded` – Sent when the total number of queued packets in all priority queues for the Frame Relay WAN link exceeded the sum of the queue limits for all priority queues.

For details about using these and other traps supported by the EOS SNMP agent, see [Traps and Trap Profiles on page 664](#).

Global variable:

Implements a time delay between traps. You set the length of the time delay with the FR service `TrapTimeInterval` parameter. The default value for `TrapTimeInterval` is 300 seconds. The delay is implemented as follows:

- The first time the Frame Relay WAN link exceeds the configured queue depth caution threshold or queue limit, EOS drops any packets above the queue limit, sends an `mnrWANIIfQueueDepthCaution` or `mnrWANIIfQueueDepthExceeded` trap, and sets a time stamp.
- If the Frame Relay WAN link exceeds the configured queue depth caution threshold or queue limit again, EOS checks to see if the interval specified by the Frame Relay service `TrapTimeInterval` parameter (300 seconds by default) expired since the last trap was sent.
 - If the `TrapTimeInterval` did not expire, no additional trap is sent, but any packets above the queue limit are dropped.
 - If the `TrapTimeInterval` expired, a second trap is sent, its time stamp becomes the new starting point for the `TrapTimeInterval`, and any packets above the queue limit are dropped.

For more information about the `TrapTimeInterval` parameter, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Additional variables in the Circuit Statistics table:

Two variables are statistics; the other three support the caution level trap.

- `frCircuitStatsPRQXmitByte` – The total number of bytes transmitted from this priority queue.
- `frCircuitStatsTotalPacketsDropped` – The number of packets dropped from this priority queue because the queue was full.
- `frCircuitStatsCautionEnabled` – A Boolean control variable that enables or disables the generation of caution traps.
- `frCircuitStatsCautionPercent` – Sets the queue depth at which a caution trap is generated, as a percentage of the total queue depth.

- `frCircuitStatsMaxQueueDepth` – The size of the priority queue, in packets. This is a read-only variable. Frame Relay calculates the queue size based on the transmission rate of the connection and the CIR of the DLCI.

For details about the priority queuing statistics you can retrieve remotely via SNMP, see [Statistics Retrieved Through SNMP on page 687](#).

15.10

Frame Relay Autostartup

Autostartup does not determine the type of data network to which you subscribe. After autostartup, you can set this value using the `PDNtype` parameter in the Frame Relay Service. This is a tuning feature.

For example, the following command configures port 2 to FR Service with the Sprint public data network:

```
SETDefault !2 -FR PDNtype = SPRint
```

For further information about the `PDNtype` parameter, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

15.11

MFR Overview

MFR allows multiple virtual links (constituent virtual circuits or CVCs) of the same CIR to be multiplexed together under one virtual port (aggregated virtual circuit or AVC), thus providing an increased bandwidth for one port. Incoming packets can be sent across any CVC of the AVC. EOS employs a round robin scheme to decide which packets traverse through which CVCs and uses sequence numbers to preserve the in-order delivery of frames required by Frame Relay. Frames are received from the constituent links and either saved in the MFR AVC structure or forwarded, depending on the sequence number of the MFR packet received.

When a packet is transmitted out of the AVC, EOS determines how the packet are sent as follows:

- If the incoming packet size is greater than the minimum size configured with the `FR MfrFrag` parameter, EOS:
 - Splits the packet so that, when possible, the segments are equal to the size configured with the `MfrFrag` parameter.
 - Prepends the MFR header.
 - Updates the sequence number.
 - Transmits the segments on the next location of the round robin.
- If the incoming packet size is less than the minimum size configured with the `FR MfrFrag` parameter, EOS:
 - Prepends the MFR header.
 - Updates the sequence number.
 - Transmits the packet on the next location of the round robin.



NOTE: With MFR implemented, CVCs use LMI to report status. However, in order for LMI to report status, the parent port of the CVC must be configured as such (the `FR CONTrol` parameter set to `ANSiLMI`). If the `FR CONTrol` parameter is set to `NoLMI`, the CVC does not use LMI for reporting status.

Frame Relay Forum Specification for MFR

The Frame Relay Forum specifies two standards for MFR. EOS implements the End-to-End Multilink Frame Relay Implementation Agreement - FRF.15.

Salient aspects of the specification are as follows:

- Aggregation takes place at the VC data transfer level only.
- Frames may be transmitted over different CVCs.
- Frame reassembly procedures accept fragments of a frame from any CVC within an AVC.
- The sequence number is 12 bits. EOS uses these bits as follows:
 - The lower 8 bits plus the “R” (reserved) bit in the standard FRF.15 header are used for the sequence number (which is 9 bits long).
 - The higher 4 bits are used to let the receiver know that a modified MFR header is used. They signal the peer about the resetting of the sequence numbering scheme (to zero) and to carry the priority level value in every packet to aid the receiver in reassembly according to priority level.
- FRF.15 proposes one sequence numbering scheme per AVC; however, EOS uses five schemes, one for every priority level supported by Frame Relay in EOS. This helps the receiver buffer packets of similar priority values and avoids a condition in which higher-priority packets have to wait for a lower-priority packet to arrive before they can be forwarded on to the upper layers for further processing.
- The status of an AVC is identified as follows:

An AVC is active if the total aggregated CIR available to the AVC is equal to or greater than the configured minimum trunk CIR. The minimum trunk CIR is configured with the `MfrMinToTrunk` parameter. For details, see “Frame Relay Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Chapter 16

Wide Area Networking Configuration with PPP

This chapter describes how to configure wide area networking using the Point-to-Point Protocol (PPP). PPP is a standard protocol that provides serial line connectivity between two Motorola Solutions Enterprise OS devices or between a Motorola Solutions Enterprise OS device and another vendor's device running PPP.

After you set up and check the serial lines, PPP performs the following functions:

- Negotiates the maximum size of a packet that can be received over a serial line
- Manages a serial line
- Maintains serial line quality

16.1

Packet Size Negotiation

During device startup, the two devices connected by the serial line negotiate the maximum size of a packet that each can receive. Once the devices agree on the size of the packet, the negotiation is complete and bridging or routing can begin.

To change the size of the packet, you must use:

```
SETDefault !<port> -PPP MaxRcvUnit = <bytes>(1-4500)
```

For example, if you want to change the size of the packet from the default of 1,500 to 1,000 bytes on port 3, you must use:

```
SETDefault !3 -PPP MaxRcvUnit = 1000
```

In this example, even if the packet size is negotiated between both devices to be 1,000 bytes, the device continues to receive packets up to 1,500 bytes in size. However, it discards packets greater than 1,500 bytes in size.


16.2

Serial Line Management

A common problem experienced with T1 lines is the loopback of packets. If a T1 line goes down, packets transmitted are looped back as shown in the following figure. Because the packets are looped back, the devices may perceive that they are still sent and received and may not realize that the line is down. To detect a physical loopback problem, each link management packet sent from a device includes a magic number that is checked upon receipt.

The magic number is an option that is negotiated during device startup. If both devices at the ends of the serial line connection support this option, then unique numbers are negotiated. These unique numbers are used later in the line quality management packets for loop detection.

Figure 106: Detecting Loopback of Packets with Magic Numbers



In this configuration, imagine that Router 1 sends out a Link Control Protocol (LCP) packet with a magic number (Abra). When the packet is transmitted to Router 2, Router 2 checks the magic number in the received packet against its unique magic number that it negotiated earlier. If the magic numbers are different, it concludes that a loopback is not occurring.

If, upon receipt, Router 1 and Router 2 check the magic numbers in their respective returned packets and detect their own magic numbers, both devices conclude that a loopback is occurring and declare the line down.

When a loopback is detected, the following system message appears on the system console:

```
Path <n> loopback
```

After the loopback is removed, the line comes back up. When PPP completes handshaking, the following message appears on the console:

```
Path x is up
```

This message indicates that PPP is back in the normal state.

16.3

Serial Line Quality Maintenance

PPP attempts to maintain the quality of a serial line. If data packets are not received and no echo reply is received in response to echo requests for approximately eight seconds, PPP brings the line down. Once the line goes down, PPP continuously attempts to negotiate until it successfully brings the line up.



NOTE: The actual time to bring down a line varies depending on the baud rate on the link.

16.4

Authentication Methods

PAP uses a two-way handshake method to establish the identity of the peer before a link is established, thereby providing greater security than basic PPP settings. With PAP, a peer determines when the Authentication Request is made. The PeerID/Password pair is sent over the wire in a plain text form.

CHAP uses a three-way handshake sequence to establish the identity of a peer before establishing a link. You may repeat this sequence at any time if the `AuthReptIntvl` parameter is set, to ensure that the identity of the peer did not change. CHAP relies on the equivalent of userid and password, a name/secret pair, to produce a challenge value used for authentication. Because the name/secret pair is never sent directly on the circuit, CHAP provides a higher level of security than PAP authentication. With CHAP, the authenticator controls when the authentication request is made.

16.5

Load Sharing and Load Balancing

When multiple serial links are assigned to a port running PPP, load sharing or load balancing can be used on that port to make more efficient use of the bandwidth of the links.

With load sharing, the fastest link mapped to each port is selected as the primary link, and the other links are considered secondary. When data needs to be sent on that port, it is sent on the primary link until the link is saturated, then the balance of the data is sent over secondary links.

The advantage of load sharing is that if you configure the secondary link for bandwidth-on-demand, it is used only when the data exceeds the bandwidth of the primary link. The disadvantage of load sharing is that it can cause packet misordering, which may be undesirable for some network protocols.

You can use load sharing only when data traffic on the port is not sensitive to packet misordering. If the data traffic has a combination of network protocols, some of which are sensitive to packet misordering, and some that are not, you need to select the sequencing feature using the mnemonic filtering scheme to ensure that packets are not misordered for sequencing sensitive protocols. Remember that sequencing works for bridged packets only.

With load balancing, data is split over parallel serial links while preserving sequencing. If the link speeds are the same, the load is split evenly. However, if the link speeds differ, the data is split in proportion to the difference between the speed of the links. For example, if you have two links, and the speed of one is 60% greater than the speed of the other, the faster link receives approximately 60% more data traffic. Through load balancing, all active links can be used to their full capacity.

Load balancing is accomplished only by using the PPP MultiLink Protocol (MLP) (RFC 1717). With MLP, each packet is assigned a sequence number to guarantee in-sequence delivery. In addition, packets may be divided into fragments, which are also assigned sequence numbers.

Packets and packet fragments are sent over the available serial links. The receiving station, also running MLP, reassembles the fragments into packets based on the sequence numbers. Once a packet is completely reassembled, it is released to the client protocol. Packets that cannot be completely reassembled due to lost fragments are discarded.

Whether packets are fragmented depends on their size and the traffic on the links. For example, some packets are too small to benefit from fragmentation. Moreover, if you have only two links, and one of them is saturated, you do not gain a performance advantage by sending packet fragments to this link.

If you use bandwidth-on-demand to back up a single primary link, and you also use load balancing, packets are sent unfragmented on the primary link. As soon as the data traffic exceeds the bandwidth of the primary link, the secondary link is brought up, which in turn enables load balancing. Once the data traffic drops below a user-specified threshold, the secondary link is brought down and the packets are once again sent unfragmented on the primary link.

If you use bandwidth-on-demand to back up multiple primary links, and you also use load balancing, packets can be fragmented over the multiple primary links. As soon as the data traffic exceeds the bandwidth of all the primary links, the secondary link is brought up, and load balancing is extended to the primary and secondary links. When traffic drops below a specified threshold, the secondary link is brought down and packets are again load-balanced, but only across primary links.

When you use load balancing, paths manually assigned to the same port are referred to as a bundle. The following is a list of alternate ways to create a bundle:

- Dynamically binding paths to a port through dial pooling. You do this using the `SysCallerID` port parameter, which allows you to tell the remote device the port number it must assign to the dial pooling path. You can use dial pooling with either load sharing or load balancing.

For more information about configuring the `-SYS SysCallerID` parameter, see “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. For more information about dial pooling paths, see [Protocol Reservation Configuration on page 383](#).

If MLP is turned off (`-PPP MlpCONTRol = Disabled`), then load sharing is automatically used.

- Standard bundling using Endpoint Discriminator (ED) and authentication. EOS devices provide class 3, MAC address for the ED. Authentication is achieved through PAP or CHAP if either is configured or negotiated. The bundle ID used by EOS devices to identify links belonging to the same bundle is defined as a combination of ED and authentication.

For ED support you use the `TxEndpointDisc` parameter in the PPP service. This parameter indicates whether to send ED for a call on the link.

16.6

Point-to-Point Protocol Communication

Only one wide area protocol is allowed to run over one port, regardless of the number of paths assigned to the port.

Figure 107: One Path Per Port Configuration

Only one path was assigned to one port. Router 1 is running Point-to-Point Protocol (PPP) over ports 3 and 4.





Figure 108: Two Paths Per Port Configuration

Two paths are assigned to one port. Router 1 is running PPP over port 3, which has paths 3 and 4 assigned to it.



If you assign multiple paths to one port, the load sharing feature is enabled. For more information on load sharing, see [Load Sharing and Load Balancing on page 343](#).

By default, PPP is enabled on serial interfaces. Serial lines running PPP can bridge or route all protocols supported by the EOS device.

Enabling Point-to-Point Protocol

While enabling Point-to-Point Protocol (PPP), keep in mind the following considerations:

- Before configuring PPP, you must be logged on with Network Manager privilege.
- When you enable PPP on the device, you must also enable PPP on the device at the other end of the serial connection.

If the owner of a port is not PPP, you must use:

```
SETDefault !<port> -PORT OWner = PPP
```

Setting an Authentication Protocol

You can configure PPP to prevent unauthorized access, especially for dial-up lines over the Public Switched Telephone Network (PSTN), and also to administer multiple remote users. PPP handles authentication using either the Password Authentication Protocol (PAP), the Challenge Handshake Authentication Protocol (CHAP), or the Extensible Authentication Protocol (EAP). EAP is used to authenticate remote clients only. MD5-Challenge, which is identical to CHAP, is also supported.

When you are setting up authentication, you must specify a user ID and password pair as part of the `-PPP AuthLocalUser` and `AuthRemoteUser` parameters. How the pair is used depends on whether PAP or CHAP is configured as the authentication protocol with the `AuthProTocol` parameter.

With CHAP, you can also specify an interval value in minutes to repeat authentication and to ensure that the identity of a peer is not changed after a link is established. The `AuthReptIntvl` parameter sets the interval value.

16.6.1

Setting Up PAP

In the configuration shown in [Figure 108: Two Paths Per Port Configuration on page 344](#), Router 2 (the remote device) makes a dial-on-demand call for a user on the network.

When Router 2 makes a dial-on-demand call, no SCID is sent. The `Config_Req` packet includes the endpoint discriminator (ED), which is the MAC address of the port configured to make the call. When Router 1 receives the `Config_Req` packet request, it queries Router 2 for PAP or CHAP information. Router 2 responds with the configured userid and password pair. Router 1 looks up the userid/ password pair and binds the link to the specified port.

If additional bandwidth is required, path 4 comes up connecting to port 3. Path 4 also included the ED and is bundled with path 3.



NOTE: The password that you use for both `AuthLocalUser` and `AuthRemoteUser` must be the same as the password used by the remote device. You must also specify PAP as the `AuthProTocol`.

Procedure:

1. Configure the local device (Router 1) by performing the following actions:

- a. Specify the `AuthLocalUser` parameter with a user ID and password by entering:

```
SETDefault !<port> -PPP AuthLocalUser = ([ "<userid>" | None], "<password>")
```

You must remember to enclose the password in double quotes. The user ID and password are case-sensitive and can be up to 16 ASCII characters long.



NOTE: You can set up PAP so that either end of the link may initiate authentication for a session.

- b. Configure your remote user identification information by entering:

```
ADD !<port> -PPP AuthRemoteUser ["<userid>" | None], "<password>"
```

You must specify both the user ID and the password so that multiple remote users have unique user identification information.

- c. Enable PAP as your authentication protocol by entering:

```
SETDefault !<port> -PPP AuthProTocol = Pap
```

- d. Enable all the settings you have specified in the previous steps by entering:

```
SETDefault !<port> -Port CONTrol = Enabled
```

2. Configure the remote device (Router 2) by performing the following actions:

- a. Enable the end point discriminator by entering:

```
SETDefault !<port> -PPP TxEndpointDisc = Enabled
```

When a device is initiating a call, as it is in this example, the device sends the end point discriminator.

- b. Specify the AuthLocalUser parameter with a user ID and password by entering:

```
SETDefault !<port> -PPP AuthLocalUser = ([ "<userid>" | None], "<password>")
```

Remember that the user ID and password must match a user ID and password configured on the device that receives the call.

- c. Enable PAP as your authentication protocol by entering:

```
SETDefault !<port> -PPP AuthProTocol = Pap
```

16.6.2

Setting Up CHAP



NOTE: When you set up CHAP you must always specify a userid.

Procedure:

1. Specify the AuthLocalUser parameter with a userid and password by entering:

```
SETDefault !<port> -PPP AuthLocalUser = "<userid>", "<password>"
```

The password and userid are case-sensitive and can be up to 16 printable ASCII characters long.

2. Configure your remote user identification information, specify the userid and password so that multiple remote users have unique user identification information by entering:

```
ADD !<port> -PPP AuthRemoteUser "<userid>", "<password>"
```

If you are setting up CHAP for remote use, make sure that the remote userid and password pairs are added as AuthRemoteUser entries at the local end, and that the local userid and password pair is added at the remote end.

3. Optional: Specify how often CHAP repeats authentication to verify the identity of the remote user by entering:

```
SETDefault !<port> -PPP AuthReptIntvl = <minutes> (0-255)
```

If you specify 0, repeat authentication is disabled.

4. Enable CHAP as your authentication protocol by entering:

```
SETDefault !<port> -PPP AuthProTocol = Chap
```

5. Enable all the settings you specified in [step 4](#) by entering:

```
SETDefault !<port> -PORT CONTrol = Enabled
```

16.6.3

Verifying Your PAP or MPPE Configuration

Procedure:

1. Verify your settings for the AuthLocalUser parameter by entering:

```
ShowDefault !<port> -PPP AuthLocalUser
```

The AuthLocalUsers are displayed for each port you configured.

2. Verify your settings for AuthRemoteUser by entering:

```
ShowDefault !<port> -PPP AuthRemoteUser
```

The userids are displayed for remote users you added.

3. Verify the protocol that was configured by entering:

```
ShowDefault !<port> -PPP AuthProtocol
```

The authentication protocol you configured is displayed.

4. Verify the supported MPPE encryption keys by entering:

```
ShowDefault !<port> -ppp ENcryptCONtrol
```

16.6.4

Setting Up EAP

You can use EAP only for RAS clients on unbound paths. EAP is an option that is disabled by default using the PPP Service DefaultAptCtl parameter.

When an unbound path begins LCP negotiations, the device sends a Configure Request with the Authentication Protocol option set to EAP. By default, EAP is set as the device's most preferred authentication protocol.

The following authentication scenarios are supported:

- EAP MD5-Challenge authentication relay to RADIUS servers with EAP extension support
- EAP MD5-Challenge authentication relay to RADIUS servers without EAP extension support

Procedure:

1. Set up the RADIUS server. See [Remote Access Services Configuration on page 467](#).
2. Enable EAP by entering:

```
SETDefault -PPP DefaultAptCtl = EAP
```

16.7

NCP Configuration

The network control protocol (NCP) configuration list is in the PPP component and is system-wide. When the EOS software establishes an LCP connection with a RAS client, it sends out the NCP request packets specified in the NCP configuration list. You can configure the NCP configuration list to contain only NCPs that are supported by the RAS client. This way only configured NCPs are negotiated by the RAS client.



NOTE: You can enable or disable an NCP by specifying the appropriate element in the list. If an NCP is disabled, no Configuration Request packet for that NCP is sent, and the incoming Configuration Request for that NCP is ignored.

The EOS software supports the following configuration request packets:

- IPCP
- OSI

To configure NCP, you must specify the list of negotiable NCPs by using:

```
SETDefault -PPP NcpProtocolId = (IPcp, Osi)
```

16.8

Activating LAPB to Reduce Noisy Lines

Typically, when you use the PPP Protocol, you do not need to activate the LAPB service. If there is a problem with noisy lines when using PPP, you need to activate the LAPB Service.

The device has noisy lines in the following situations:

- A temporary lack of response occurs, even though the other end is active.
- LAPB assumes that a line is down, even when it is not, because the clock went down temporarily.
- Frames need to be retransmitted frequently.

While LAPB provides a reliable data link, it does add some protocol overhead. Consequently, you need to evaluate the need for LAPB by monitoring the error rates on your lines. To activate LAPB, see [Enabling LAPB for a PPP Link on page 396](#).

Procedure:

1. Increase the T1 parameter to lengthen the amount of time that LAPB waits for an acknowledgment.

This ensures that LAPB does not retransmit a frame unnecessarily.



NOTE: When selecting your value for the T1 parameter, be aware that the value you enter is internally divided by 250 milliseconds. As a result, any value you enter less than 250 actually equals zero.

Step example:

To increase the amount of time to 4,000 milliseconds on path 3, enter:

```
SETDefault !3 -LAPB T1 = 4000
```

2. Increase the N2 parameter. This increases the maximum number of times a frame is sent after a timeout.

Step example:

To increase the N2 parameter value to 12 on path 3, enter:

```
SETDefault !3 -LAPB N2 = 12
```

3. Decrease the size of the FrameSeq parameter.

When the line is noisy, keep the window size low to keep the number of unacknowledged frames low. In addition, a low window size reduces the number of retransmittals required when the software encounters a corrupted packet.

Step example:

To decrease the FrameSeq parameter to basic sequencing on path 3, enter:

```
SETDefault !3 -LAPB FrameSeq = Basic
```


Chapter 17

Port Bandwidth Management

Bandwidth management provides two operating modes: system bandwidth management and manual bandwidth management. You can enable these modes with the `-PORT DialInitState` parameter; the `DialOnDemand` option enables system bandwidth management mode, and the `ManualDial` option enables manual bandwidth management mode. For information about the manual bandwidth management mode, see [Configuring System Bandwidth Management Mode \(DOD\) on page 366](#).

System bandwidth management provides you with automated bandwidth management features. You provide bandwidth allocation guidelines, and the system automatically manages the virtual pipe for you by monitoring traffic rates on the line. When traffic increases, bandwidth management may automatically allocate additional bandwidth; when traffic decreases, it may automatically decrease the bandwidth.

Changes to bandwidth and line characteristics take effect immediately. The system also uses a phone number list that you define to automatically obtain additional bandwidth. If bandwidth requirements consume more than one path, the system picks additional paths to form the virtual pipe and meet the requirements.




NOTE: Port bandwidth management is applicable to PPP WAN paths only.

Virtual Pipe

Port *bandwidth management* is a process that applies static bandwidth, dynamic bandwidth, or a combination of these to provide a port with the bandwidth it needs to meet current requirements. At each port, a set of serial path resources are configured to provide a bandwidth bundle called a *virtual pipe*.

Figure 109: Use of Resources through the Virtual Pipe



Bandwidth

Static bandwidth is provided by a configuration of one or more leased lines or dial paths to a port. The static resources are dedicated to a single port. Leased lines can provide continuous dedicated bandwidth to the port. Static dial paths can also provide incremental bandwidth, or bandwidth that becomes available only when a decision is made to dial them up.

Dynamic bandwidth is provided by dial-up line paths, which are allocated from a dial pool. Incremental bandwidth is provided by dial paths. The port can use either analog or digital lines that are allocated to it, and additional dial path resources can be added incrementally.

Bandwidth management operates on a port-by-port basis. It monitors line use based on rate of traffic and increases or decreases bandwidth based on limits that you specify. Network protocols that use the port are unaware of the underlying physical links, which bandwidth management bundles together into the virtual pipe to meet the port bandwidth requirements.

Bandwidth Aggregation

The main function of bandwidth management is to determine the aggregate bandwidth that is provided to the set of protocols passing through the port. However, a WAN operates most efficiently when it can allow for variations in the type and amount of traffic passing through it. In addition to bandwidth management, the software provides the protocol reservation feature, which allocates portions of the virtual pipe to specified traffic such as the IP or OSI. As traffic passes through the pipe, you can also enable the Point-to-Point (PPP) Multilink Protocol (MLP) to distribute packets more evenly over the virtual pipe.

You can reserve bandwidth for the protocols traversing the WAN using the `-PORT ADD ProtocolRsrv` parameter. For more information, see [Protocol Reservation Configuration on page 383](#).

Load balancing equalizes traffic flow and ensures that packets that may be fragmented over the links arrive at their destination in the correct sequence. Load balancing is accomplished using the PPP Multilink Protocol as described in RFC 1717; you enable load balancing using the `-PPP MlpCONTROL` parameter.

17.1

Dial-on-Demand

Dial-on-demand (DOD) is an economical way to use phone lines when communicating between devices. It is supported only under system bandwidth management mode. You can enable DOD with the `DialOnDemand` option of the `-PORT DialInitState` parameter. DOD is triggered on when there is traffic on a port, and triggered off when the port is idle or experiences a decrease in traffic congestion. The careful monitoring of traffic provided by system bandwidth management allows more cost effective-use of DOD lines.

A connection is established when the system automatically dials a phone number specified with the `-PORT DialNoList` parameter, or the phone number configured in the DTR modem. The line stays up as long as transmit traffic is present. When there is no more data, the call is terminated. The call is automatically reestablished without any intervention when there is data to be sent across the line. Connections that are no longer in use are temporarily terminated until new demand occurs. The `-PORT DialIdleTime` parameter determines how long a connection must be idle before the call is terminated.

In general, DOD limits background traffic to routed packets (IP-routed packets) and other network protocol packets that are absolutely necessary to maintain the functionality and integrity of the overall network. The software feature set provides you with the necessary parameters for controlling the traffic over that DOD link. For phone lines and the connections associated with those lines to operate properly in the DOD state, the network layer protocols running over those connections must use statically defined routes or Open Shortest Path First (OSPF) demand circuits (per RFC 1793). Currently IP is the only network layer protocols supported with DOD. For procedures and configuration examples of IP routing over a DOD link, see [Configuring Routing over DOD Links on page 377](#).

17.2

Bandwidth-on-Demand

Bandwidth-on-demand (BOD) is triggered on when the system detects traffic congestion on a port configured for system management mode. You specify the bandwidth at which a port should operate normally, then define the maximum amount of bandwidth above this setting that the port can have. Together these settings define the maximum width of the virtual pipe.

The BOD allocation strategy provides a flexible approach for configuring WAN dial-up lines. For example, the normal operating bandwidth of a WAN with two 64 kbps ISDN lines could be configured together for a total bandwidth of 128kbps, or be configured at 64 kbps bandwidth with incremental increases up to 128 kbps, as traffic needs required. Depending upon traffic across your network, you can choose to configure one wide virtual pipe to handle the traffic, or configure a narrower virtual pipe that expands and contracts as traffic increases or decreases.

You define bandwidth allocation using the `-PORT NORMALBandwidth`, `BODThreshold`, `BODIncrLimit`, and `DialSamplPeriod` parameters, which specify bandwidth settings and the conditions that trigger BOD. Bandwidth management monitors the incoming and outgoing rate of traffic and uses the settings to prevent dropped packets by changing the size of the virtual pipe (allocating or removing lines and bandwidth) as required by traffic demands.

You can also configure a line as a general purpose line that can be allocated for any purpose, including disaster recovery, using the `UnReStricted` option of the `-PATH DialCONTRol` parameter. When a line failure causes the port bandwidth to drop below the level specified with the `-PORT NORMALBandwidth` parameter, the DOD strategy and bandwidth management work to restore the specified bandwidth. If traffic conditions warrant additional bandwidth, then BOD increases the bandwidth accordingly.

17.3

System Bandwidth Disaster Recovery

System Bandwidth disaster recovery is the disaster recovery threshold, which is defined as the minimum of the normal bandwidth threshold, as defined by the `-PORT NORMALBandwidth` parameter, and the total amount of configured leased line bandwidth that is assigned to the port, excluding disabled paths.

When the total active bandwidth from the leased line paths falls below the disaster recovery threshold, bandwidth management tries to recover the port bandwidth to the target set with the `-PORT NORMALBandwidth` parameter using dial paths. In this event, a path configured for disaster recovery is given preference. You configure a line specifically for disaster recovery by using the `DisasterRcvry` option of the `-PATH DialCONTRol` parameter.

17.4

System Bandwidth Path Configurations

Table 37: Path Configurations Available with System Bandwidth Management

Strategy	Path Configurations Available
Dial-on-demand	Single or multiple dial paths acting as a bundle
Bandwidth-on-demand	Single or multiple dial paths for bandwidth-on-demand bandwidth aggregation
Disaster recovery	Single or multiple dial paths for disaster recovery on all dial PPP ports

17.5

Bandwidth Management Terms

The following terms are used to explain concepts such as dial pools and the ports that can use them, and the bandwidth management strategies:

Bandwidth management

Process that applies static bandwidth, dynamic bandwidth, or a combination of these to provide a port with the bandwidth it needs to meet current requirements.

Dial pool

The pool of dial paths that can be dynamically bound to any properly configured port.

Disaster recovery threshold

The minimum of the normal bandwidth threshold and the total amount of configured leased line bandwidth that is assigned to the port.

Dynamic binding

The association of a path in the dial pool to a port when it is needed.

Dynamic path

A path that can be used by more than one port. You create a dynamic path by unbinding it from its port. A dynamic path is stored in the dial pool.

Initially, dynamic paths are not bound to any port, but dynamically bind to make an outgoing call. To receive an incoming call, a dynamic path receives the call while still in the dial pool, and then is bound to a port.

Dynamic paths can bind to different ports without user action, but only one port can bind at a time.

Once the path becomes inactive, it unbinds from the port and becomes available for other ports.

Normal bandwidth threshold

Bandwidth threshold defined by the `-PORT NORMALBandwidth`, `BODThreshold`, `BODIncrLimit`, and `DialSamplePeriod` parameters.

Port

A port is a logical interface used by the software to represent a connection to a network.

When the DOD path is up, the device routes the packets as expected in the normal NCP and SPX1 connection processes.

Static binding

The association or binding of a path to a port as defined at system initialization time or by user configuration.

Static path

A path is assigned (bound) to one port and can be used only by one port; a static path cannot be shared. By default, all paths are static at system initialization time.

Virtual path

A path used by a device to represent multiple logical paths multiplexed over a single interface. ISDN B, ISDN PRI, and DS0 channels delivered by channelized T1/E1 or switched-56 are presented as distinct virtual paths. A virtual path can be used as a static or dynamic resource.

Virtual pipe

A term that describes a port of variable bandwidth.

Virtual port

A port that is not associated with a physical interface. Virtual ports allow you to configure multiple destinations through a single interface.

17.6

Supported Communication Resources

You can use port bandwidth management to allocate a broad range of communication resources, from public telephone lines using inexpensive analog modems to digital circuits providing throughput at rates up to 45 Mbps.

Table 38: Communication Resources Supported by Enterprise OS Port Bandwidth Management Feature

Resource	Throughput
Telephone line; analog line using the public telephone system	Up to 56 kbps
Leased line; dedicated higher quality analog line	Up to 56 or 64 kbps
ISDN BRI; dedicated narrowband digital service	56/64 or 112/128 kbps
ISDN PRI; dedicated narrowband digital service	23 or 30 channels at 64 kbps/56kbps
Switched-56; nondedicated digital service	Up to 56 kbps
Fractional T1; dedicated digital circuit	64 kbps to 1.544 mbps in increments of 64 kbps
T1 or E1 line; dedicated digital circuit	24 or 30 channels at 64 kbps each
T1 or E1 channel; dedicated digital circuit	1.544 Mbps or 2048 kbps
T3 or E3 line; dedicated digital circuit	Up to 45 Mbps

For information about how to configure these resources, see [Configuring WAN Resources on page 362](#).

Sync Dial Lines

When you use sync dial lines (serial DTE dial-up lines), the EOS software supports the V.25 bis standard, which allows a DTE device to communicate with a Data Communication Equipment (DCE) modem or terminal adapter (TA). With V.25 bis, you can configure and store phone numbers in software for the modem or TA on the EOS device. The phone number is sent to the device modem or TA when dialing occurs.

The software can also activate modems or TAs that store a phone number in the firmware of a data terminal ready (DTR) dialed modem or TA. You can use DTR modems only in static configurations; dynamic paths selected from a resource pool rely upon telephone numbers stored in the device software.

17.7

Paths to Ports Associations

The EOS software uses the concepts of *ports* and *paths* to address interface connections. The basic interface connection is a port. A port is a logical interface that represents a connection to a network. The next logical connection is the path, which is the physical interface that connects the device to a physical medium such as an Ethernet local-area network (LAN) or a serial line. In an ISDN environment, a path additionally represents the channel over which data is transmitted. This section describes the paths that you can configure on ports under bandwidth management.

Static Versus Dynamic Paths

You can unbind static paths from their ports and save them in a dial pool to be shared by more than one port. The paths in the dial pool are called *dynamic paths*. A path in the dial pool can be *dynamically bound* to a port or PPP virtual port when the path is needed to transfer data on a dial-up line. After the path is bound to the port, port-based dialing occurs. When the traffic is no longer present and the line is idle for a specified period, the path is unbound from the port and returned to the dial pool.

A port can bind to multiple dial paths. For each port, bandwidth can be dynamically allocated by bundling the multiple dial paths into the virtual pipe. In this way, bandwidth can be allocated when and where it is most needed.

Dynamic Dial Path Pooling

You can use a dial pool to increase the reliability of your network configuration, achieve multideestination dialing by using dial phone number lists and modem pooling, and provide dynamic backup for leased or dial-up lines.

With multideestination dialing, you can allocate a small number of paths that are unbound from their ports to wait in the dial pool for an incoming call. You can create a PPP virtual port on the central device for each remote site and have all the virtual ports use the dial pool for path resources.

When the system receives an incoming call from a remote site, the dynamic path that answers is bound to a virtual port, which is standing by with the appropriate configuration information for the calling network. For the binding to occur, the identification of the remote site is transmitted to the central device.

In network configurations where dynamic paths are bound to ports for all different types of WAN services, the system caller ID (SCID) identifies the remote site device to a central site device. You configure the SCID with the `-SYS SysCallerID` parameter. Using the SCID in this way works only for EOS devices at the remote site. If you use SCID for a router from another vendor, the acceptance of an incoming call from this router can occur only with a static path and port configuration.

In network configurations where dynamic paths are bound or static paths are configured to ports for use with ISDN switched-circuit services, you can also specify the remote site caller ID with the Calling Line Identification Presentation (CLIP) dial string. The dial string is the ISDN phone number of the calling device.

Standard bundling on PPP links uses both the text endpoint discriminator (ED), and authentication. EOS devices can provide a class 3, MAC address value for ED, but they also accept classes 1 through 5 for ED from other systems. Authentication is achieved through PAP or CHAP. The bundle ID used by the device to identify links belonging to the same bundle is defined as a combination of ED and authentication. You configure the ED with the `-PPP TxEndpointDisc` parameter.

Binding dynamic paths to ports to establish a connection between a remote site and a central site using CLIP enables you to identify and connect EOS devices and routers from other vendors at the remote site to the central site. You can configure a port to identify remote site devices with only SCID or with only CLIP. If you configured a port for CLIP and SCID, CLIP is used.

When you configure CLIP for a port, you add the remote site ISDN telephone number to the CLIP database using the `-PORT CLIPList` parameter and set the `-PORT DialRcvrState` parameter to AnswerCLI. For details, see the descriptions of the `DialRcvrState` and the `CLIPList` parameters in “Port Service Parameters” in the *Enterprise OS Software Version 18.2 User Guide*.

Because not all sites that use a dial pool call the central site at the same time, it is possible to share a small group of paths with a larger group of sites. Each site that can potentially call into the dial pool has its own virtual port predefined, so there can be more virtual ports configured to use the dial pool than there are dynamic paths assigned to the dial pool. However, if all the remote sites dial the central device at the same time and only a small number of paths exist in the dial pool, some of the call attempts may fail due to a lack of path resources. These calls can be redialed at a later time.

For a summary of the terms used in this section, see [Bandwidth Management Terms on page 351](#). For more information about the dial pool, see [Resource Aggregation on page 355](#).

Table 39: Valid Port and Path Configurations

Type of Path	Default Ports	Virtual Ports
Static leased line	1 or more	Not supported
Static dial path	1 or more	Not supported

Type of Path	Default Ports	Virtual Ports
Dynamic dial physical path	1 or more	1 or more

17.8

Resource Aggregation

Dial Number List

The software can select a phone number from a list of destination phone numbers associated with a port. In addition, the device can automatically select alternate phone numbers as backup if a previously dialed phone number does not result in a connection. The software knows whether the dial-up line is a static or dynamic path. If it is a static path, the software obtains a phone number from the dial number list specified with the `-PORT DialNoList` parameter and brings up the path. If DTR dialing is used, the software uses the phone number stored in the modem.

If the dial-up line is a dynamic path, the software searches a dial pool to obtain an additional path. The software obtains a path from the dial pool, and binds it to the port. The software obtains a phone number for the dial number list specified with the `-PORT DialNoList` parameter. By using the dial number list, the software can try other phone numbers if the first number is unavailable. When searching the dial number list, the software seeks a path that matches the baud rate set for the path and chooses the path that matches or most closely matches that rate. By using the dial pool, the software can select another dynamic path if the first path is not working. When traffic conditions return to normal, the dial-up path is unbound from the port and returned to the dial pool.

Prioritized Path Preferences

You can prioritize dynamic and static dial paths for use by specific ports by using the `-PORT PathPreference` parameter. No other port can use these reserved paths, unless these paths are reserved by more than one port. The software tries the preferred list of paths first before using path resources in the rest of the dial pool, and seeks static paths first. By specifying your path preference, you can reserve path resources for your dial-up lines, and ensure that a path is always available.

When you specify the selection sequence for a static path, you actually specify the priority sequence of selection for use by a port. A static path is inserted at the end of the path preference list by default. A static dial path with either the V.25 bis or DTR dial mode can be inserted in the list, but DTR does not use the dial number in the `DialNoList` parameter to dial out; instead, it uses the dial number stored in the modem.

Leased line paths cannot be included in the path preference list because bandwidth management cannot bring a leased line up or down; leased line paths are brought up when the port is enabled.

17.9

Manual Bandwidth Management

You can enable manual bandwidth management with the `-PORT DialInitState` parameter; the `ManualDial` option enables manual bandwidth management mode. Manual bandwidth management mode requires user intervention to control bandwidth on the line. You can set a fixed amount of bandwidth by using the `-PORT NORMalBandwidth` parameter, and then issue the `DIAL` command to bring up the line. Bandwidth management tries to meet the bandwidth specification, but does not monitor traffic or make any dynamic changes based on traffic rates. Bandwidth changes are only made according to what you specify with the `-PORT NORMalBandwidth` parameter.

For more information on how to configure manual bandwidth management, see [Configuring Manual Bandwidth Management Mode on page 374](#).

Manual Dial

Under manual bandwidth management, you bring up a line manually by using the `Dial` command. The call remains connected until a timer expires or until you end the connection by using the `HangUp` command. After the call was disconnected, it can be reestablished only by issuing another `Dial` command.

The `Dial` command has a path mode and a port mode, and operation of this command differs depending upon the mode specified. The path-based `Dial` command is used mainly for testing and in event-based macros that automate line backup processes. For the description of the `Dial` command, see “Commands” in the *Enterprise OS Software Reference Guide*.

The port-based `Dial` command manually dials on the specified port. The command accepts static or dynamic port numbers and an optional dial string. If you enter a dial string, the number must be listed in the `-PORT DialNoList` parameter in order for the call to be placed. The call is placed on the available highest-priority phone number specified for the port by using the `-PORT DialNoList` parameter. If the highest prioritized phone number is not available, the software tries to use the next phone number specified for the port, providing more than one phone number is configured.

You must insert a telephone number in the `-PORT DialNoList` parameter for dialing to occur, (although you can temporarily override the phone numbers in the `DialNoList` parameter by specifying a port and dial string with the port-based `Dial` command). To complete the call, the software automatically finds a path by first checking if a path is available in the path preference list. If a path is available, that path is used. If no path is available, the software determines whether the port can use the dial pool. If the port can use the dial pool, the software checks for an available path in the dial pool, binds it to the port, and makes the call.

Bandwidth management manages `Dial` command calls and makes bandwidth evaluations based on the `-PORT NormalBandwidth` setting. All ports configured for `Dial` must have a positive bandwidth setting. Bandwidth management aggregates bandwidth resources as needed to meet the `NormalBandwidth` setting, dialing for more resources if additional bandwidth is needed, or hanging up or substituting resources if less bandwidth is needed. You specify which phone number and path to select with the `-PORT DialNoList` and `PathPreference` parameters.

Manual Hangup

The port-based `HangUp` command manually brings down all dial path resources and is the default. Under manual bandwidth management mode, the port is brought down when `HangUp` is issued unless there are leased lines active.

You can hang up calls on dynamic paths or static paths. If the path is dynamic and is currently bound to a port, the `HangUp` command disconnects the call, unbinds the path, and places the path back into the dial pool.

Manual Bandwidth Management Disaster Recovery

Under manual bandwidth management, the software only maintains the normal operating bandwidth of a port; there is no monitoring and automatic allocation of additional resources if the line goes down. You can specifically configure disaster recovery for a leased line by using the `DisasterRcvry` option of the `-PORT DialControl` parameter. When disaster recovery is enabled, manual bandwidth management restores the line bandwidth back to its original setting when a leased line fails and the port bandwidth falls below what was specified for normal operation.

17.10

Bandwidth Management Status and Statistics Displays

You can monitor the dial path status, including the state of the ports under bandwidth management control, by using the `-PORT DialStatus` command. The display produced by this command shows total

port bandwidth, messages indicating congestion levels, and the intentions of the bandwidth manager for allocating additional resources.

Bandwidth management DOD statistics displays enable you to control costs by monitoring connection charges. For details, see [Statistics Displays on page 967](#).

17.11

Configuring Wide Area Networking with Async PPP

Perform this process to set up and configure wide area networking using Async PPP.

Async PPP is used on paths whose `LineType` is `Dial` and whose `DialMode` is `ATdial`.

You perform most of the configuration for Async PPP by using commands in the `PATH` service; the only `PORT` service configuration required for use of AsyncPPP is the specification of PPP as `PORT OWNER`, and inclusion of Async types in the `DialNoList`.



NOTE: LAPB does not operate over an AsyncPPP or AT dial path. As a result, you cannot use history-based compression (`COMPRESSType` set to `History` on a port) with Async PPP.

Procedure:

1. Set up a path for Async PPP and AT Dial. See [Setting Up a Path for Async PPP and AT Dial on page 357](#).
2. Initialize the modem. See [Initializing the Modem on page 358](#).
3. Add a number to the phone number list. See [Adding a Number to the Phone Number List on page 359](#).

17.11.1

Setting Up a Path for Async PPP and AT Dial

When you configure the path for AT dial operation, the path attempts to initialize the modem whenever the path is enabled or the device notices that the modem was (re)attached or power cycled (determined by a drop then rise of the CTS input signal).

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).

Procedure:

1. Configure the path for dial with AT technology by entering the following commands:


```
SETD !<path> -PATH LineType = Dialup
SETD !<path> -PATH DialMode = ATdial
```
2. Configure the external device type as Async so the port can distinguish it from non-async modem paths by entering:


```
SETD !<path> -PATH ExDevType = Async
```
3. Enable CTS flow-control so the modem can avoid buffer overflow by preventing the device from transmitting by entering:


```
SETD !<path> -PATH AsyFlowCtrl = CTS
```
4. Configure the baud rate for the path by entering:


```
SETD !<path> -PATH BAud = <baud_rate>
```

Speeds supported for Async PPP paths are limited to the range of 1.2k bps through 38.4k bps; if you configure a higher speed, the system uses 38.4k, although no error message is displayed.

The baud rate specifies the speed of data between the modem and the EOS device, not the modem connect (carrier) speed. In general, this speed should be higher than the carrier speed (assuming the modem supports CTS flow-control); some modems enforce this. The baud rate you choose must be supported by the modem and should probably be the lowest baud rate just above whatever the actual maximum modem (carrier) throughput is expected to be.

5. Put the parameters you set in the through into effect by (re)enabling the path by entering:

```
SETD !<path> -PATH CONTrol = Enable
```

17.11.2

Initializing the Modem

You can use this procedure to define and assign a modem initialization string for a path.

To initialize the modem, the EOS device sends an initialization string, which is a configurable set of AT commands intended to configure the modem appropriately for use on the port. You define the modem initialization strings with the `AsyMdmInitStr` parameter, and assigned the strings to a path with the `AsyMdmName` parameter.



NOTE: The default initialization string is simply “AT” which presumes that the modem defaults (or is preconfigured) to the appropriate settings for operation with the EOS device.

Procedure:

1. Define an initialization string by entering:

```
ADD -PATH AsyMdmInitStr <name> " <string>"
```

where:

`<name>` is a token used to identify the string for use on specific paths. It consists of sixteen (16) characters, which can be letters, digits, underscore (_), hyphen (-), period (.), or asterisk (*); the first character must be a letter.

`<string>` is the AT command string used to configure the modem. It must be enclosed in double quotes, and should begin with AT.



NOTE: Although the parameter accepts up to 80 characters of printable ASCII as the string definition, EOS software allows a maximum of 50; if you define a larger initialization string, it is not used and the path does not come up. In addition, the size of the modem command buffer may further reduce the available length of the initialization string.

The initialization string should configure the modem as needed; configuration requirements for a modem to operate with a Motorola Solutions EOS platform include:

- No command echo
- Numeric result codes
- Minimal result codes set
- Auto-answer disabled
- Character format 8/N/1
- Software flow-control (XON/XOFF) disabled
- Hardware flow-control (CTS) enabled
- Fixed-speed DTE link (doesn't follow carrier)
- CD output signal asserted according to carrier
- DTR drop during connection causes hang-up

- CTS asserted when in command mode

For more information about these configuration settings, and some information on determining an appropriate initialization string to meet EOS device and modem compatibility requirements, see [Modem Initialization Strings on page 359](#).

Step example:

```
ADD -PATH AMIS courier "AT&F1 E0V0Q0X0&A0 &C1&D2&R1 S0=0 &C1&D2"
```

2. Specify which predefined initialization string to use on the path by entering:

```
SETD !<path> -PATH AsyMdmName = <name>
```

where:

<name> is one of the names in the AsyMdmInitStr table, defined in [step 1](#). If <name> is not specified or if ? is used, the EOS software help displays the currently defined names. You can also display AsyMdmInitStr table with the SHow command. (The reserved name none removes any assignment, so the path uses the default initialization string.)

If an undefined <name> is assigned to the path, a warning message is issued but the assignment is allowed to take effect. The device uses the default initialization string on that path until a definition for <name> exists. Also, <name> is case-insensitive. For example, the following command refers to the name "courier" defined in [step 1](#) above:

```
SETD !3 -PATH AMN = Courier
```

3. To put the new assignment into effect by (re)enabling the path, enter:

```
SETD !<path> -PATH CONTroL = Enable
```

17.11.3

Adding a Number to the Phone Number List

When a path was configured for AT dial and Async PPP operation, it is available to originate or answer calls for a PPP port. For a PPP port to use the AT dial path to make a call, it must have a phone number in its DialNoList with an assigned type of Async.

For more information, see the description of the DialNoList parameter in "Port Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Procedure:

To add an async number to the phone number list, enter:

```
ADD !<port> -Port DialNoList "<phone-no>" [Baud=<baud>] Type=Async [Pos=<1-xx>]
```

The default baud rate attribute for an Async path is 9.6.

The string specified as the phone number **must not** include the AT prefix or the D (dial) command; when attempting to dial, the device adds the prefix ATDT (or ATDP) to the specified string and sends it to the modem. The first character of the string can be used to specify tone (T) or pulse (P) dialing.

17.12

Modem Initialization Strings

For a device to make use of an AT dial modem, a number of conditions are imposed on modem operation. The commands required for the modem to operate this way can be used as an "initialization string" (assigned

by using the `PATH AsyMdmInitStr` and `AsyMdmName` parameters) that the device sends to the modem whenever the path is enabled or the modem is (re)attached or power cycled.



NOTE: Some modem initialization may be required offline to put the modem into a state where the device sends it commands and it accepts them; in particular, CTS must be active and CD inactive for the device to send the initialization string.

The modem must accept an incoming call and dial an outgoing call using the A and D commands, respectively. In addition, the modem must be configured as follows:

- No command echo
- Numeric result codes
- Minimal result codes set
- Auto-answer disabled
- Character format 8/N/1
- Software flow-control (XON/XOFF) disabled
- Hardware flow-control (CTS) enabled
- Fixed-speed DTE link (does not follow carrier)
- CD output signal asserted according to carrier
- DTR drop during connection causes hang-up
- CTS asserted when in command mode

These requirements are grouped roughly into three categories for more detailed discussion below. In some cases, specific modem commands may be referenced; you may need to consult your modem technical documentation to determine if a particular command is supported (and/or required) for your modem.

Recommended initialization strings for some modems are listed at the end of this section.

17.12.1

Command Interaction

The EOS device sends commands to the modem in order to configure it (using the initialization string), place a call, or accept a call. Commands should not be echoed; the only command-mode output expected from the modem is a single-digit response indicating the result of an operation:

- 0 – OK (command successful)
- 1 – CONNECT (carrier established)
- 2 – RING (incoming call indication)
- 3 – NO CARRIER (exiting call failed)
- 4 – ERROR (command failed)

You must configure the modem so that automatic answer mode is disabled. When an incoming call arrives, the modem should signal the device using a code 2 (RING); if the device chooses to answer the call, it responds with an ATA (answer) command.

You can achieve this configuration using basic AT commands which are fairly standard across a number of modems, although defaults tend to vary:

- E0: Do not echo commands
- Q0: Display result codes
- V0: Use numeric result codes
- X0: Minimal result code subset

- S0=0: Disable auto-answer mode



NOTE: Some modems may have additional result code commands that you may need to use to restrict the result code subset. You may want to use other command settings relating to modem command interaction, for example disabling local/remote escape sequences.

17.12.2

Signal Interaction

When an AT dial path is enabled, the EOS device asserts DTR and RTS. When connected, the device indicates a hang-up request by toggling the DTR signal. (The RTS signal is toggled only when switching between command mode and data mode, and can be ignored.)

The device uses CTS for modem presence in command mode. Commands are only sent if CTS is asserted; a CTS toggle restarts the sequence. In data transfer mode, however, CTS may be used for flow-control. The device expects CD (modem carrier) only in response to an answer or dial command; it should drop to indicate hang-up. The only restriction on DSR is that it not drop in command mode (but it may follow CD).

Signal control commands are somewhat less consistent across different modems, but the following commands apply to several modems:

- &C1: CD according to carrier
- &D2: DTR drop initiates hang-up
- &S0: DSR remains asserted

These commands may vary slightly on a per-modem basis, and some modems may have several commands that interact to control a single signal. Flow-control options are even more varied, see [Online Operation on page 361](#).

17.12.3

Online Operation

Most online modem operation is independent of the EOS device. Choices may depend on the quality of your telephone service and the type of modem(s) with which you communicate. These are some of the constraints you should keep in mind:

- DTE speed must not follow carrier speed
The device path speed is determined by the `BAud` parameter. Some modems automatically change their DTE speed based on carrier speed, particularly when answering a call. If the modem and EOS device use different speeds, the connection fails; it is essential that the modem maintain a constant DTE speed -- for example, based on the last AT command.
- Character format of 8/N/1.
The EOS device uses 8 databits and no parity on an AT dial path; if the modem does not automatically pass this transparent, you must configure it to do so. (Most modems learn the character format from the AT characters themselves.)
- Software flow-control (XON/XOFF) should be disabled.
The EOS device does not use XON/XOFF, and during LCP option negotiation it always proposes an ACCM of zeros. Thus any XON/XOFF characters in data must be able to traverse the modem link as is.
- Hardware flow-control (CTS) may be enabled.
If the `AsyFlowCtrl` parameter is set to `CTS`, the EOS device does not transmit data when CTS is low; this may improve operation by preventing overrun of the modem buffer. Most modems allow some form of hardware flow-control (some with restrictions) to take advantage of this. (The modem need not honor RTS flow-control; the device does not use it, and simply keeps RTS asserted.)
- Hang-up delay disabled.

The EOS device assumes the modem is available again after dropping DTR for a short time. With an error-control in use, some modems delay a hang-up initiated by DTR drop if there is buffered data to deliver. In the EOS device environment, much of the value of this extra delay is lost, and it should be minimized if possible. Many modems use S-Register 38 to configure this delay.

- Escape sequences disabled.
The EOS device does not use escapes to enter modem command mode. Many modems require “guard time” in escape sequences; Async PPP data from the device is explicitly framed, making chances of an inadvertent escape negligible. But if a modem allows escape sequences without a guard time, this may be a concern.
- Lower bound on carrier speed.
Some modems offer the capability to negotiate carrier speed no lower than a specified “floor”, or reject connection if signal quality is too low. This may be desirable to prevent excessively slow links when a connection is of unusually poor quality. (If a connection is too slow the PPP link may not come up at all.)

In general, the configuration commands for these options vary widely between different modem manufacturers and models.

17.12.4

Initialization Strings Examples

The following table lists initialization strings for several modems successfully used with EOS devices. These strings generally show the minimum configuration required; additional configuration may be desirable (or even required) for a particular environment.

Table 40: Modem Initialization Strings


String	Modems
"AT&FQ0V0E0&C1&D2&S0X0S0=0"	Hayes (Accura 33.6)
"AT&FQ0V0E0&C1&D2&S0X0S0=0S2=128S38=0"	Motorola Solutions (ModemSURFR 33,600)
"AT&F1E0V0Q0X0&C1&D2S0=0&R1&A0"	3Com/USR (Courier, Sportster)
"AT&F8&W0&FE0V0Q0X0&Q1&C1&D2&E4&E1&E12 %E0S0=0S13=0"	Multitech (MT1932ZDX)

17.13

Configuring WAN Resources

Perform this process to prepare your WAN resources for use with bandwidth management.

Figure 110: Basic Bandwidth Management Port and Path Configurations



Process:

1. Configure dial-up lines with a modem or TA. See [Configuring Dial-Up Lines with A Modem Or TA on page 363](#).
2. Configure leased lines. See [Configuring Leased Lines on page 365](#).
3. Configure ASPC leased lines. See [Configuring ASPC Leased Lines on page 366](#).

17.13.1

Configuring Dial-Up Lines with A Modem Or TA

You can use digital ISDN and analog serial lines to establish connectivity with remote sites so that these sites can send updates to a central location.



NOTE: You must complete this procedure at both ends of the link.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).

Procedure:

1. Set the LineType to Dialup by entering:

```
SETDefault !<path> -PATH LineType = Dialup
```

2. Specify the baud rate for the device by performing the following actions:

- a. Set the baud rate for a serial line by entering:

```
SETDefault !<path> -PATH Baud = <kbps> (1.2-52000)
```

The default baud rate for a serial line is 64 kbps. The Autostartup feature automatically detects modem connections on some devices; it does not sense the baud rate for ISDN paths with external TAs attached.

- b. Set the baud rate for an ISDN line connected with a TA by entering:

```
SETDefault !<connectorID> -PATH Baud = <kbps> (1.2-52000)
```

The Autostartup feature automatically detects modem and TA connections on devices with external TAs attached.

3. If you configure a device without FlexWAN ports, select the connector type by entering:

```
SETDefault !<path> -PATH CONNector = V35 | RS232 | RS449 | G703 |  
HSSI | X21
```

4. If you configure a device without FlexWAN ports, set the transmit clock for the EOS device by entering:

```
SETDefault !<path> -PATH CLock = TestMode | External | Internal
```

The **External** value allows the device to derive the transmit clock from either the send or receive timing clock supplied by the digital service unit/channel service unit (DSU/CSU) or by the attached modem.

5. Perform one of the following actions:

If...	Then...
If you want to select DTR dialing mode,	enter: SETDefault !<path> -PATH DialMode = DTRdial You must select DTR dial to configure a line using a modem that uses the DTR signal to initiate a call.
If you want to select V.25 bis standard,	perform the following actions: a. Enter: SETDefault !<path> -PATH DialMode = V25bis You must select V.25 bis to configure a DTE serial line using a V.25 bis-compatible modem. b. Specify the telephone number of the dialed remote site by entering: ADD !<port> -PORT DialNoList "<phone_no>"

For more information about the dial list configuration, see [Configuring the Dial List on page 367](#).

6. Specify the external device type attached to the DTE path by entering:

```
SETDefault !<path> -PATH ExDevType = [Modem | Bri | Sw56]
```

Port-based dialing that uses phone numbers from the dial-number list always looks at the setting of the **ExDevType** parameter to select an appropriate path for the phone number and phone technology. The default setting of the **ExDevType** parameter is **Modem**; the default setting for the **Type** attribute of the **DialNoList** parameter is also **Modem**.

7. Set the path characteristics for the line by entering:

```
SETDefault !<path> -PATH DialCONTRol = ([DYNamic | STATic], [DisasterRcvry |  
NoDisasterRcvry | UnReSTricted])
```

The **-PATH DialCONTRol** parameter provides several options for setting the line.

The **STATic** value allows the selected path to be statically bound to its corresponding port and is the default.

The **DYNamic** setting unbinds a path from its corresponding port and adds the path to the dial pool. A static path is not part of the dial pool. Placing a path in the dial pool allows the path to be used by any dial port. To make a dial path a dynamic dial path, you must set the **-PATH LineType** parameter to **Dialup**.

You can also choose to set the line specifically for disaster recovery, or as unrestricted to allow it to be used for any purpose including disaster recovery. The **NoDisasterRcvry** option prevents the line

from being used for disaster recovery and is usually assigned to the slowest or least reliable line on the network.

Step example:

To configure the analog serial line on path 3 for no disaster recovery, enter the following commands:

```
SETDefault !2 -PATH LineType = Dialup
SETDefault !2 -PATH Baud = 28.8
SETDefault !2 -PATH CONNector = RS232
SETDefault !2 -PATH CLock = External
SETDefault !2 -PATH DialMode = DTRdial
SETDefault !2 -PATH ExDevType = Modem
SETDefault !2 -PATH DialCONTRol = NoDisasterRcvry
```

8. Enable the line and make sure all settings on the path take effect by entering:

```
SETDefault !<path> -PATH CONTRol = Enabled
```

17.13.2

Configuring Leased Lines

You can use this procedure to configure leased lines to establish connectivity with remote sites so that these sites can send updates to a central location.



NOTE: You must complete this procedure at both ends of the link.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).

Procedure:

1. Set the LineType to Leased by entering:

```
SETDefault !<path> -PATH LineType = Leased
```

2. Specify the baud rate for the device by entering:

```
SETDefault !<path> -PATH Baud = <kbps> (1.2-52000)
```

3. If you configure a device without FlexWAN ports, select the connector type by entering:

```
SETDefault !<path> -PATH CONNector = V35 | RS232 | RS449 | G703 |  
HSSI | X21
```

4. If you configure a device without FlexWAN ports, set the transmit clock for the EOS device by entering:

```
SETDefault !<path> -PATH CLock = TestMode | External | Internal
```

The External value allows the device to derive the transmit clock from either the send or receive timing clock supplied by the digital service unit/channel service unit (DSU/CSU) or by the attached modem.

5. Enable the line and make sure all settings on the path take effect by entering:

```
SETDefault !<path> -PATH CONTRol = Enabled
```

17.13.3

Configuring ASPC Leased Lines

You can use this procedure to configure Austel Semi Permanent Circuit (ASPC) leased lines to use with bandwidth management. You can configure an ISDN interface for the ASPC line type when this line type is supported by the ISDN switch in use. ASPC is a leased line type that is controlled like a dial line but operated as a leased line. Unlike other leased lines, a phone number is associated with the ASPC line.



NOTE: You must complete this procedure at both ends of the link.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).

Procedure:

1. Set the line type to ASPC by entering:

```
SETDefault !<path> -PATH LineType = ASPC
```

2. Set the ASPC number by entering:

```
SETDefault !<path> -PATH ASPCNumber "dial string"
```

The "dial string" may contain up to 52 alphanumeric characters, and it must be enclosed in double quotation marks. Dial string usually contains the ISDN phone number of the ASPC switch. If the dial string is empty (only two double quotation marks are present), the existing number is deleted. The maximum number of entries per path is one.

3. Enable the line and make sure all settings on the path take effect by entering:

```
SETDefault !<path> -PATH CONTrol = Enabled
```

17.14

Configuring System Bandwidth Management Mode (DOD)

You can use this procedure to enable system bandwidth management mode.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Your WAN communication resources are enabled. See [Configuring WAN Resources on page 362](#).

Procedure:

To enable system bandwidth management mode, enter:

```
SETDefault !<port> -PORT DialInitState = DialOnDemand
```

This command also enables dial-on-demand.

17.15

Configuring Bandwidth-on-Demand

Perform this procedure to enable bandwidth-on-demand for the system bandwidth management mode.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Your WAN communication resources are enabled. See [Configuring WAN Resources on page 362](#).

Procedure:

1. Specify the amount of bandwidth the port should bring up when enabled by entering:

```
SETDefault !<port> -PORT NORMAlBandwidth = <kbps>
```

2. Enable BOD and specify the amount of additional bandwidth that bandwidth management can allocate for a port by entering:

```
SETDefault !<port> -PORT BODIncrLimit = <kbps>
```

This syntax specifies incremental bandwidth levels that can be allocated for the port.

3. Specify the conditions that trigger additional path resources for BOD by entering:

```
SETDefault !<port> -PORT BODTHreshold = <%>(0-100)
```

4. Specify the amount of time bandwidth management should wait to take action to bring a port up or down by entering:

```
SETDefault !<port> -PORT DialSamplPeriod = <seconds>(0-300), (0-300)
```

Step example:

To configure a DOD line with a normal port bandwidth of 64 kbpsm, enter the following commands:

```
SETDefault !3 -PORT DialInitState = DialOnDemand
```

```
SETDefault !3 -PORT NORMAlBandwidth = 64
```

```
SETDefault !3 -PORT BODIncrLimit = 64
```

```
SETDefault !3 -PORT BODTHreshold = 50
```

```
SETDefault !3 -PORT DialSamplPeriod = 30, 60
```

Traffic must exceed 32 kbps for 30 seconds before bandwidth management brings up additional bandwidth using BOD. The additional dial path is taken down when the rate of traffic is less than 32 kbps for longer than 60 seconds.

17.16

Configuring the Dial List

You can use this procedure to configure a dial list that allows the software to select a phone number from a list of destination phone numbers associated with a port. Numbers in the dial list are selected sequentially, but you can use the EOS software to order the numbers in the dial list.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.

- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Your WAN communication resources are enabled. See [Configuring WAN Resources on page 362](#).

Procedure:

1. To allow your EOS device to dial out, configure the dial number list by entering:

```
ADD !<port> -PORT DialNoList "<phone_no>" [Baud = <rate> (1.2-16000)] [Type = Modem  
| Bri | Sw56 | WE | WEH0] | [Pos = <number>]
```

You can enter this command more than once to append a phone number or profile to the dial list.

For ISDN, the phone number includes a dial prefix, country code, and area code and possibly a subaddress. You can use hyphens (-) to separate the prefix from the country code from the phone number. If you specify a subaddress, you must separate the phone number from the subaddress with a semicolon (;). You can configure up to 16 phone numbers per port.

For V.25 bis dialing, the phone number can include a dial prefix, country code, and area code.



NOTE: The EOS software uses the Baud and Type keywords to make a path match. It is important to enter the same device type as you entered in the -PATH ExDevType parameter; the software makes a best match for the baud rate. See [Configuring WAN Resources on page 362](#) for further details.

2. Optional: Specify the number of time the software attempts to redial the remote system if the call attempt fails by entering:

```
SETDefault !<port> -PORT DialRetryCount = <number> (0-20)
```

If dialing is based on a static port and path binding, the software first tries to make the call. If the attempt fails to bring the path up, the software tries the call again using the same or different path. The call attempts continue until the dial retry count is reached.

You can also append phone numbers to the end of the list, insert a phone number into a specific position in the list, edit an existing phone number, or delete an existing phone number. For details, see [Phone Number Operations on page 369](#).

17.16.1

Dial List Configuration Examples

Dial List Configuration Example 1

To enter a Los Angeles phone number for port 2 that consists of a long-distance dial prefix 1 (assume that the configured device is located in Santa Clara), the phone number 213-456-7000, and the subaddress 101, you must use:

```
ADD !2 -PORT DialNoList "1-213-456-7000;101" Baud = 56 Type = Bri
```

Dial List Configuration Example 2

To add a London phone number for port 2, at the end of a dial-up phone list, that consists of the international dialing code 011, the U.K. country code 44, and the phone number 213-456-7000, you must use:

```
ADD !2 -PORT DialNoList "011 44 213 456 7000"
```

This entry ignores the Baud rate, Type, and Pos (position on the list).

For DTR dialing, the phone number is irrelevant, because the outgoing telephone number is stored in the modem.

Dial List Configuration Example 3

You can configure the dial number list to dial the same number repeatedly by adding multiple copies of the number. You can prefix the phone number with a variable number of periods to distinguish the duplicate entries by using the following commands:

```
ADD !V1 -PORT DialNoList "123 4567"
ADD !V1 -PORT DialNoList ".123 4567"
ADD !V1 -PORT DialNoList "..123 4567"
```

The EOS device dials 123-4567 three times. This technique works for both ISDN and analog phone numbers.

17.17

Phone Number Operations

Adding a Phone Number

To insert a phone number into a specific position in the dial number list, you must use the `Pos` (position) keyword with a non-zero number after the dial string.

For example, to insert a phone number for port 4 into position 2 of the dial number list that contains 10 phone number, you must use:

```
ADD !4 -PORT DialNoList "510 555 7000" Pos = 2
```

The software inserts the new phone number into position 2. The phone number that was previously in position 2 is now in position 3. If the phone already exists in the dial number list, it is moved to position 2. You can insert the same phone number twice by using blanks or other non-dialing characters. You also can include the `Baud` and `Type` keywords in any order when inserting phone numbers into the dial number list.

Editing an Existing Phone Number

To edit an existing phone number in the dial number list, you can change the position in the list, change the baud rate, and change the device type.

For example, if port 3 was assigned 612-345-3989 in position 2 with a baud rate of 64 kbps, you can change the baud rate by using:

```
ADD !3 -PORT DialNoList "612 345 3989" Pos = 2 Baud = 14.4
```

Because the dial string is case-sensitive, make sure to match it exactly to successfully edit an existing string using characters other than numbers.

Deleting a Phone Number

To remove a phone number or profile from the dial number list, you must use:

```
DELeTe !<port> -PORT DialNoList "<phone_no>"
```

The profile name is case-sensitive and must be matched exactly to be deleted.

17.18

Paths to Ports Binding

EOS software uses the concept of port and path bindings to pair a logical interface (port) to a physical network resource (path) such as an ISDN line. A port can bind to multiple dial paths. For each port, the bandwidth can be dynamically allocated by bundling the multiple dial paths into the virtual pipe. This concept allows bandwidth to be allocated when and where it is most needed.

The software also allows you to create dynamic paths by unbinding static paths from their ports and saving them in a dial pool to be shared by more than one port. A path in the dial pool can be bound to a port when the path is needed for data transfer events associated with dial-up.

17.18.1

Converting a Static Path to a Dynamic Path

Paths are static by default.

Procedure:

1. Unbind a path from its corresponding port and add the port to the dial pool by entering:

```
SETDefault !1 -PATH DialCONtrol = DYNAmic
```

2. Enable the path by entering:

```
SETDefault !1 -PATH CONtrol = Enabled
```

17.18.2

Changing a Dynamic Path to a Static Path

You can use this procedure to change a dynamic path to a static path and remove it from the dial pool.

Procedure:

1. Convert a dynamic path to static by entering:

```
SETDefault !1 -PATH DialCONtrol = STATic
```

2. Reassign a path or multiple paths to a port by entering:

```
ADD !<port> -PORT Paths <path>
```

When a static dial path is added to a port, it is automatically inserted at the end of the path preference list; see the [Identifying Remote Sites with SCID and CLIP on page 370](#) for further information about the path preference list.

3. Enable the path by entering:

```
SETDefault !1 -PATH CONtrol = Enabled
```

17.18.3

Identifying Remote Sites with SCID and CLIP

To bind a dynamic path to a port to connect a remote site device to a central site, you must identify the remote site device.

If you connect a remote site device to a central site device over ISDN or non-ISDN lines, you can use the system caller ID (SCID) to identify the remote site device to the central site device only if the remote site device is an EOS device. The SCID for a remote site can be a telephone number or a text string.

If you connect a remote site device to a central site device over ISDN lines, you can use a CLIP dial string to identify Motorola Solutions and other-vendor devices at the remote site to a central site device. The dial string is an ISDN phone number plus an optional subaddress. You configure this number with the `-PORT CLIPlist` parameter.

If you configure a device port for multiple uses such as an ISDN dial-up path, a modem dial-up path, and a leased path, you should create a virtual port and identify the remote user using SCID.

You can configure a port to identify remote site devices with only SCID or with only CLIP. When you configure a virtual port for CLIP, you must use SCID to identify the remote site. The SCID identifier is ignored when the dynamic dial-up path binds to the virtual port.

If you configured a port for SCID and CLIP, CLIP is used for incoming ISDN dial-up path calls. If the port is configured for something other than ISDN dial-up paths, SCID is used.

Procedure:

1. If you want to configure a local device port to use SCID to identify a remote device, perform the following actions:
 - a. Identify the central site to the remote sites by entering:


```
SETDefault -SYS SysCallerID = "<string>"
```

where *<string>* is a text string (within quotes) up to 31 characters long and unique to the system that identifies the central site (for example, a city name such as "Sunnyvale") to the remote sites.
 - b. Create a virtual port for the central site device that uses SCID to identify to the remote user by selecting SCID when creating the -PORT VirtualPort parameter and entering:


```
ADD !<port> -PORT VirtualPort {<connectorID.channelID> {<FRDLCI>} | {SCID "<SysCallerID>"}}
```

The SCID is a text string of characters within quotes that identifies the remote site, such as "San Diego," to the remote site.
2. If you want to configure a local device port to use CLIP to identify a remote device, perform the following actions:
 - a. Select SCID for the -PORT VirtualPort parameter, by entering:


```
ADD !<port> -PORT VirtualPort {<connectorID.channelID> {<FRDLCI>} | {SCID "<SysCallerID>"}}
```

The SCID is a text string of characters within quotes that describes the remote site, such as "San Diego." The SCID is ignored when a dynamic path is binding to the virtual port configured for CLIP.
 - b. Add the dial string to the CLIList database that identifies the remote site device to the central site virtual port by entering:


```
ADD !<port> -PORT CLIList "dial string"
```

The dial string can be up to 50 alphanumeric characters in length and must be enclosed in quotes. The dial string is the remote site ISDN phone number. With ISDN, the phone number usually includes the dial prefix, country code, area code, and possibly a subaddress assigned to your ISDN interface. You can use hyphens (-) to separate the prefix, country code, and area code. If you specify a subaddress, you must separate the phone number from the subaddress with a semicolon (;).
 - c. Set the central site device virtual port to check the identity of incoming ISDN calls with the CLIList database by setting the -PORT DialRcvrState parameter to AnswerCLI by entering:


```
SETDefault !<port> -PORT DialRcvrState = NoAnswer | Answer | AnswerCLI
```

If the identity of the incoming call matches the dial string in the CLIList database exactly, the dial-up path is bound to the port to make the connection.

17.19

Configuring the Path Preference List

A path preference list reserves a path for use by a group of ports and sets the order of line use. Prioritization is accomplished by position in the path preference list. (Leased line paths cannot be included in the path

preference list because bandwidth management cannot bring a leased line up or down; leased line paths are brought up when the port is enabled.) A path can be reserved by more than one port.

You can append one or more dial-up paths to the end of the path preference list, insert one or more paths into a specific position in the list, or delete an existing path in the list. For more information, see [Path Operations on page 373](#).

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Your WAN communication resources are enabled. See [Configuring WAN Resources on page 362](#).

Procedure:

1. Specify the paths added to the path preference list as dial-up.

Step example:

To specify path 1 as dial-up, enter:

```
SETDefault !1 -PATH LineType = Dialup
```

2. Unbind a path from its assigned port and add it to the dial pool by entering:

```
SETDefault !1 -PATH DialCONTRol = DYnamic
```

3. Enable changes to the path by entering:

```
SETDefault !1 -PATH CONTRol = Enabled
```

In this command, you must use the *<path>* syntax if you specify a DTE serial path and use the *<connectorID.channelID>* syntax if you specify an ISDN path.

4. If you use ports (as opposed to PPP virtual ports) with dynamic lines, enable the port or virtual port to use dial pool resources and map the remote system caller ID to a specific port by performing the following actions:

a. Enter:

```
ADD !<port> -PORT PATHs SCID" <SysCallerId>"
```

Step example:

To allow port 4 to use the dial pool for a path resource for outgoing calls and to map incoming calls with the caller ID of "London" to port 4, enter:

```
ADD !4 -PORT PATHs SCID"London"
```

The string you enter for the caller ID is case-sensitive and can contain up to 31 characters.

- b. Ensure that each remote site is configured with unique caller ID by entering:

```
SETDefault -SYS SysCallerID = "<string>"
```

The SysCallerID string is limited to 31 characters. The string should be administratively assigned and be unique across the network.

5. If you use PPP virtual ports, enable the port or virtual port to use dial pool resources and map the remote system caller ID to a specific port by performing the following actions:

a. Enter:

```
ADD !<port> -PORT VirtualPort {SCID" <SysCallerId>"}
```

The caller ID string is case-sensitive and can contain up to 31 characters.

Unlike Frame Relay virtual ports, which are always associated with a particular path, PPP virtual ports can potentially use any path in the dial pool.

Step example:

To create PPP virtual port V3 and allow it to use the dial pool for its path resources for outgoing calls and to map incoming calls with the caller ID of "NewYork" to virtual port V3, enter:

```
ADD !V3 -PORT VirtualPort SCID"NewYork"
```

- b. Ensure that each remote site is configured with unique caller ID by entering:

```
SETDefault -SYS SysCallerID = "<string>"
```

The SysCallerID string is limited to 31 characters. The string should be administratively assigned and be unique across the network.

6. Reserve the paths and define their priority by entering:

```
ADD !<port> -PORT PathPreference [<path>] [, ...] [Pos = <1- number>]
```

Step example:

To specify paths 2.1 and 2.2 for use by port 5, enter:

```
ADD !5 -PORT PathPreference 2.1, 2.2
```

7. Enable the previous port changes by entering:

```
SETDefault !1 -PORT CONTrol = Enabled
```

8. Check the current configuration of the path preference list by entering:

```
SHow -PATH DialPool
```

By default, the software adds dial paths to the end of the list if you do not specify the position.

After you configure paths using the -PORT PathPreference parameter, no other ports can use the reserved paths except the designated ones. The software tries the preferred list of paths first before using path resources in the rest of the dial pool.



NOTE: A dynamic path can appear in the path preference list for more than one port. A static path can appear in the path preference list for only the port to which it is bound.

17.20

Path Operations

Appending a Path

To append one or more dial paths to the end of the path preference list, you must use:

```
ADD !<port> -PORT PathPreference [<path>] [, ...] [Pos = <1- number>]
```

For example, assume the path preference list for port 2 includes dial-up paths 3 and 2.1, and you want to append dial paths 5 and 6 to the end of the path preference list. You must use:

```
ADD !2 -PORT PathPreference 5, 6
```

After this command is executed, when port 2 needs a path resource, the software uses the preferred paths first. The order of their use is 3, 2.1, 5, and 6.

Adding a Path

To add one or more dial paths into a specific position in the path preference list, you must use the Pos (position) keyword with the desired position number.

The paths are added into the list as following:

- The software deletes any duplicate paths from the list.

- The software then adds the dial paths to the path preference list by inserting them starting at the specified position.

If you want to add more than one path, you must list the paths in the intended order. For example, assume the path preference list for port 3 includes dial-up paths 3, 4, and 5, and you want to insert dial path 2.1 into position 2. You must use:

```
ADD !3 -PORT PathPreference 2.1 Pos = 2
```

After this command is executed, the path preference list includes paths 3, 2.1, 4, and 5.

If you want to insert more than one dial path, you must list the paths in the intended order. For example, assume the path preference list for port 3 includes dial-up paths 2.1 and 4. To insert dial path 3 and 5 into position 2 and 3, you must use:

```
ADD !3 -PORT PathPreference 3, 5 Pos = 2
```

After this command is executed, the path preference list includes paths 3, 2.1, 4, and 5.

You can also use the `ADD -PORT PathPreference` command to change the position of a path that already exists in the path preference list. For example, assume the path preference list for port 3 includes dial-up paths 4, 6, 7, and 5. To reposition path 6 into position 3, you must use:

```
ADD !3 -PORT PathPreference 6 Pos = 3
```

The software inserts path 6 at position 3; path 7 that was originally in position 3 is now in position 2. If the position specified is larger than the existing list, the path is appended to the end of the list by default.

Deleting a Path

To remove one or more dial-up paths from the path preference list, you must use:

```
DELEte !<port> -PORT PathPreference <path> [, ...]
```

In this syntax, you must use `<path>` if you specify a DTE serial path. If you specify an ISDN path, substitute `<connectorID.channelID>` for `<path>`.

The paths to be deleted can be listed in any order. If you try to delete a path that does not exist in the list, an error message is displayed.

17.21

Configuring Manual Bandwidth Management Mode

With manual bandwidth management mode, you can control the connect sequence and bandwidth settings for a one-time call on a line. You use port-based dialing and the `DIAL` command to manually dial on the specified port. You can use this procedure to configure manual bandwidth management mode and manually connect a line, and to enable disaster recovery.



NOTE: When you enable disaster recovery, you configure the software to bring up bandwidth to meet the target when a leased line goes down and bandwidth on the port falls below that specified with the `-PORT NORMALBandwidth` parameter.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your device are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Your WAN communication resources are enabled. See [Configuring WAN Resources on page 362](#).

Procedure:

1. If you want to place a call under manual bandwidth management mode, perform the following actions:

- a. Enable system bandwidth management mode by entering:

```
SETDefault !<port> -PORT DialInitState = ManualDial
```
 - b. Specify the amount of bandwidth that bandwidth management should bring up when this port is enabled by entering:

```
SET !<port> -PORT NORMAlBandwidth = <kbps>
```
 - c. Place the call by entering:

```
Dial !<port> [-PORT] ["<dial-string>"]
```

This command accepts a static or dynamic port number. If you enter a telephone number in the optional dial string, it must also be listed in the dial number list; however, you can enter a port and telephone number to temporarily override entries in the dial number list. For more information about setting up a dial list, see [Configuring the Dial List on page 367](#).
 - d. Add the telephone number to dial by entering:

```
ADD !<port> -PORT DialNoList "<phone_no>" [Baud = <rate> (1.2-16000)] [Type = Modem | Bri | Sw56 | WE | WEH0] | [Pos = <number>]
```
 - e. Disconnect the call by entering:

```
HangUp !<port> [-PORT]
```

HangUp brings down all dial path resources unless there are leased lines active.
2. If you want to configure disaster recovery using manual bandwidth management mode, perform the following actions:
 - a. Enable manual bandwidth management by entering:

```
SETDefault !4 -PORT DialInitState = ManualDial
```
 - b. Enable disaster recovery on the port by entering:

```
SETDefault !4 -PORT DialCONTRol = DisasterRcvry
```
 - c. Specify the normal bandwidth for this port entering:

```
SET !<port> -PORT NORMAlBandwidth = <kbps>
```

17.22

Verifying the Bandwidth Management Configuration

Procedure:

1. Display dial-up configuration information by entering:

```
SHow [!<port> | !*] -PORT DialCONFig
```

This display shows the port state (up or down), port function (disaster recovery or BOD), the paths that are in use, the path state, and the dial string for active outgoing calls. The path in use can be a static path or a dynamic path from the dial pool. If the path is from the dial pool, the information is displayed similar to a static path.
2. Display path configurations by entering:

```
SHow -PATH CONFiguration
```

You must verify that the paths are enabled and their status is up.
3. Display port configurations by entering:

```
SHow -PORT CONFiguration
```

You must verify that the ports are enabled and their status is up. Also make sure that dynamic ports are selecting path resources from the dial pool.

17.23

Troubleshooting the Bandwidth Management Configuration

Procedure:

1. Display the status of the lines under bandwidth management control by entering:

```
SHOW [!<port> | !*] -PORT DialStatus
```

The display shows the state of the ports under bandwidth management. Displays include total port bandwidth utilization, and messages indicating congestion levels and the intentions of the bandwidth manager for allocating additional resources. If the port is to be used for an outgoing call, the dial string (phone number) is also displayed.

2. Display a time-stamped dial history for the specified port or for all ports by entering:

```
SHow [! <port> | !*] -PORT DialHistory
```


17.24

Load Balancing over Multiple Dial-Up Links

Load Balancing equalizes traffic flow and ensures that packets that may be fragmented over the links arrive at their destination in the correct sequence. Load Balancing is accomplished using the PPP Multilink Protocol (MLP) as described in RFC 1717. The `-PPP MlpCONTRol` parameter enables this protocol.

The example configuration depicted in the following figure uses system bandwidth management mode. Bandwidth is set to 64 Kbps and BOD is triggered on when traffic exceeds 32 Kbps; bandwidth management adds paths to increase bandwidth to the 128 Kbps limit. The paths are taken down once traffic returns to 32 Kbps or less for longer than 60 seconds.

Figure 111: PPP MLP Load Balancing



To enable load balancing over two ISDN BRI channels, you must use the following commands:

```
SETDefault !2 -Port name="ToBetty"
SETDefault !2.1 -Path LineType=Dialup
SETDefault !2.2 -Path LineType=Dialup
ADD !ToBetty -Port DialNoList "4085551212" t=bri baud=64
ADD !ToBetty -Port DialNoList "4085551313" t=bri baud=64
SETDefault !MLPPPath1 -Path LocalDialNo="1234567"
SETDefault !MLPPPath2 -Path LocalDialNo="2345678"
```

```
SETDefault !MLPPath1 -PAth DialCONTRol=(STAtic,UnReSTRicted)
SETDefault !MLPPath2 -PAth DialCONTRol=(STAtic,UnReSTRicted)
ADD !ToBetty -POrt PAth 2.1,2.2
SETDefault !2.1 -PAth CONTRol=Enabled
SETDefault !2.2 -PAth CONTRol=Enabled
SETDefault !ToBetty -POrt DialInitState=DialOnDemand
SETDefault !ToBetty -POrt NORMAlBandwidth=64
SETDefault !ToBetty -POrt BODIncrLimiT=64
SETDefault !ToBetty -POrt BODTHreshold=50
SETDefault !ToBetty -POrt DialSamplPeriod=0,30
SETDefault !ToBetty -POrt DialIdleTime=300
SETDefault !ToBetty -ppp MlpCONTRol=Enabled
SETDefault !ToBetty -POrt CONTRol=Enabled
```

 **NOTE:** You must configure the parameters on both devices.


You can verify that both paths are up and that MLP is enabled by using:

```
SHoW -PATH CONFIguration
SHoW -PPP STATUS
SHoW -PPP MlpSTATIstics
```

17.25

Configuring Routing over DOD Links

Figure 112: IP Network Design Example: IP over a DOD Link



Procedure:

1. If you want to configure IP over a DOD link, perform the following actions:
 - a. Configure the network address for port 1 and port 4 on Router 2 by entering the following commands:

```
SETDefault !1 -IP NETaddr = 139.87.20.1
SETDefault !4 -IP NETaddr = 139.88.20.1
```


Task	Command or Parameter	Description	Applies to
	-PATH LocalDialNo	Associates a phone number to your ISDN path.	ISDN
	-PATH LocalSubAddr	Specifies a subaddress to the phone number you specified for your ISDN path.	ISDN
	-PATH SPIDdn1 and SPIDdn2	Specifies the Service Profile Identifiers (SPIDs) and directory numbers (DNs) for North American BRI ISDN dial-up modes.	ISDN
	-PATH RateAdaption	Specifies a method that determines the data rate to be used on an ISDN path.	ISDN
Enable manual bandwidth management mode (manual dial mode)	-PORT DialInitState	ManualDial option allows you to manually dial calls. You specify bandwidth settings rather than let bandwidth management monitor the line and adjust settings as needed.	DTE, ISDN
Connect	Dial	Manually connects a dial-up path or port.	DTE, ISDN
	-PORT CLIList	Adds (or deletes) a "dial string" (the ISDN phone number and subaddress) to a list of ISDN numbers to be used by the called party to map the incoming call to the appropriate port and to bind an ISDN dynamic dial-up path to the port to complete the call. This parameter is also used to screen out any calls that do not have a match in the CLIList database. The CLIList entries can take effect only if -PORT DialRcvrState was set to AnswerCLI. See the -PORT DialRcvrState parameter in this table for more details. The binding of a path to a port with a CLI number supersedes and ignores the binding between path and port set up by a system caller ID (SCID) number.	ISDN
	-PORT DialNoList	Adds, deletes, edits, and displays a list of phone numbers with their associated attributes (baud rate, phone number, and	DTE, ISDN

Task	Command or Parameter	Description	Applies to
		position in the list) for a device port. The device port tries to match an incoming call with this list of phone numbers. If there is a match, the port tries to find a path to make the connection.	
	-PORT AutoDial	Connects the dial-up path assigned to a port as soon as the path is enabled.	DTE, ISDN
Disconnect	HangUp	Manually disconnects a dial-up path or port.	DTE, ISDN
	-PATH DialCarrierTime	Defines the number of seconds the system must wait for carrier signals on the line that connected.	DTE, ISDN
	-PORT DialIdleTime	Sets the idle timer in seconds for a dial-up line before the line is disconnected if it is not in use.	DTE, ISDN
Retry a dial-up connection	-PORT DialRetryCount	Specifies the number of times to retry the call if the call attempt fails.	DTE, ISDN
	-PORT DialRetryTime	Sets the initial value in seconds to wait before attempting to reconnect after a connection failed because the carrier was not detected or for any other reason that the path did not come up.	DTE, ISDN
Configure port attributes for answer-only line	-PORT DialRcvrState	Sets the call receiver dial control state. Also, by selecting AnswerCLI you specify that the device tries to match an incoming call to the port specified and to bind an ISDN dynamic dial-up path to that port if the ISDN number (and subaddress) matches exactly the ISDN number in the CLIList database. See the -PORT CLIList parameter for more details.	DTE, ISDN
	-PORT DialInitState	NoDialOut option prevents outgoing calls when -PORT DialRcvrState is set to Answer .	DTE, ISDN
Enable system bandwidth management mode (dial-on-demand)	-PORT DialOnDemand	DialOnDemand option enables system bandwidth management for bandwidth-on-demand and	DTE, ISDN

Task	Command or Parameter	Description	Applies to
		dial-on-demand modes, and monitors the line; additional lines are brought up or down based on traffic demand.	
Configure for bandwidth-on-demand	-PATH DialCONTRol	Assigns the dial path unrestricted use as an additional resource for adding bandwidth.	DTE, ISDN
	-PORT NORMalBandwidth	Specifies the port bandwidth setting.	DTE, ISDN
	-PORT BODTHreshold	Configures the threshold that triggers the BOD line up and down.	DTE, ISDN
	-PORT BODIncrLimit	Configures the maximum incremental bandwidth that can be allocated using BOD.	DTE, ISDN
	-PORT DialSamplPeriod	Sets the time to sample threshold conditions before taking an action to bring a path up or down, based on transmit traffic load for BOD.	DTE, ISDN
Configure disaster recovery	-PORT DialCONTRol	Restricts the dial path for use as disaster recovery only.	DTE, ISDN
Configure a port to use dynamic dial path resources	-PORT PathS	Assigns dial pool resources to a specified port and identifies the remote system caller ID. For dial-up lines of any kind, the remote caller ID is a text string (such as a city name).	DTE, ISDN
	-PORT VirtualPort	Creates a virtual port that uses path resources from the dial pool and identifies the path and circuit types attached to the port. It also identifies the remote site associated with the virtual port. To configure a device port for multiple uses (dial-ups of any kind), identify the remote site with a system caller ID (SCID) text string. The SCID for remote sites can be a telephone number or a text string. When you configure ISDN dial-up lines, you can also identify remote sites with a Calling Line Identification Presentation (CLIP) dial string, which is the re-	DTE, ISDN

Task	Command or Parameter	Description	Applies to
		mote site ISDN telephone number. Remote sites identified with SCID can only connect Motorola Solutions remote sites to the central site, while remote sites identified with CLIP can connect Motorola Solutions and other-vendor devices at the remote sites with the central site.	
Configure path attributes for a dial-up path	-PATH DialCONTRol	Sets path attributes for static and dynamic dial-up paths.	DTE, ISDN
	-PORT PathPreference	Configures the dial path usage preference.	DTE, ISDN
Display port and path status	-PORT DialSTATUS	Provides the status of the dial-up service and the state of bandwidth management for the specified dial ports.	DTE, ISDN
Display port dial history	-PORT DialHistory	Provides the dial history for the specified port.	DTE, ISDN
Display the dial pool status	-PATH DialPool	Provides the status and configuration of the paths in the dial pool.	DTE, ISDN

Chapter 18

Protocol Reservation Configuration

The protocol reservation feature enables you to assign a percentage of bandwidth to designated packets transmitting out of a WAN port that meet certain conditions. You can set protocol reservation for all WAN ports on the EOS device or for specified WAN ports. However, configuring protocol reservation for specific ports is not recommended because it can affect network performance.

Protocol reservation affects only traffic transmitted from the local device (the transmit direction). You cannot configure protocol reservation for traffic received by the local device (the receive direction). If you want protocol reservation for traffic in both directions, you must configure protocol reservation on both devices (local and remote) that are sending traffic to each other.


Protocol reservation enables you to reserve bandwidth for lower bandwidth usage, interactive, response-time-sensitive, or transaction-oriented network applications, which are normally crowded out by heavy bandwidth usage applications such as file transfer or mail.

For example, in a multiprotocol environment that includes response-time-sensitive and mission-critical Telnet packets mixed in with other protocol traffic (such as IP), Telnet devices throttle back the data transmission rate to the end station when they sense available bandwidth decreasing. If other network protocols increase this bandwidth consumption, Telnet devices throttle back the data transmission rate further, which slows the response time of Telnet packets even more.

To avoid this situation, use protocol reservation to provide a percentage of bandwidth for the Telnet packets and to restrict the percentage of bandwidth to the other more aggressive protocol packets. This ensures that the response-time-sensitive Telnet packets can pass through the port in a timely manner.

The following figure illustrates the situation described in the preceding paragraphs. In the first diagram, only Telnet traffic is travelling over the port, using up the available bandwidth that provides the end user adequate response time. However, if you decide to also send IP packet traffic over that same port (as shown in the second diagram) then the IP packets continually use as much bandwidth as possible until the Telnet traffic is “crowded out” of the bandwidth, which greatly reduces the response time of the Telnet devices on the network. This crowding out is due to the connectionless nature of how IP works versus the connection-oriented nature of Telnet traffic.


Figure 113: IP and Telnet Traffic Contention (No Protocol Reservation)



To deal with this situation, you can use protocol reservation to reserve a percentage of the port bandwidth to each protocol. The following figure shows an example in which Telnet allocated 55 percent of bandwidth and IP traffic allocated 35 percent of traffic (the other 10 percent is the default allocation for other traffic).

The allocation of the bandwidth configured with protocol reservation occurs only when the different packet types actually contend for the bandwidth of the configured port.

Figure 114: IP and Telnet Traffic Contention (With Protocol Reservation)



Telnet traffic and IP packets dividing percentage of bandwidth based on protocol reservation percentages

You can use protocol reservation to allocate recommended bandwidth for other protocols besides IP and Telnet traffic. You also have wide flexibility in determining the protocols for which you want to reserve bandwidth, and how much bandwidth you want to reserve for these protocols. For more information, see [Protocol Reservation Functionality on page 391](#).

18.1

Protocol Reservation Procedural Overview

This section provides an overview of how the protocol reservation for different traffic types can be configured. Because procedures for each of the traffic types vary, read this section to determine the proper procedure for your configuration.

For specific step-by-step configuration procedures, see:

[Configuring Protocol Reservation for Bridged Traffic on page 387](#)

[Configuring Protocol Reservation for IP-Routed Packets on page 388](#)

[Protocol Reservation on Virtual Ports on page 390](#)

You can configure protocol reservation using a variety of procedures, depending on the type of packet traffic you configured. The following table lists the traffic types that you can configure for protocol reservation and the procedure used to configure each.

Table 42: Packet Types and Configuration Procedures for Protocol Reservation

Traffic Types	Configuration Procedure	Mask	Tag	See Configuration Examples in This Chapter
All Bridged traffic including IP	Filter Service ¹	Built-in masks or user-defined masks NetBIOS requires user-defined bridged masks	User-defined	Configuring Protocol Reservation for Bridged Traffic on page 387

Traffic Types	Configuration Procedure	Mask	Tag	See Configuration Examples in This Chapter
IP-Routed traffic such as FTP, IP, IPDATA, ICMP, SMTP, TCP, TELNET, and UDP	-IP FilterAddr parameter ²	Built-in protocol masks	User-defined	Configuring Protocol Reservation for IP-Routed Packets on page 388

¹ See [Mnemonic Filtering Configuration on page 433](#) and "Filter Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

² For more information on IP filtering, see [Configuring Packet Filtering on page 131](#) and "IP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

When you configure a `Filter POLICY` for use with protocol reservation, the `Filter POLICY` should not specify a port number. Not specifying a port number ensures that the protocol reservation valve controls bandwidth as defined. The examples in this chapter follow this recommendation and operate correctly regardless of the configuration of the WAN port (Frame Relay with or without virtual ports, or PPP).

Protocol reservation uses bandwidth allocation rules to determine how to allocate bandwidth of one traffic type compared to the bandwidth of other traffic types. For more information, see [Protocol Reservation Functionality on page 391](#).



NOTE: When you enter the `-PORT PROTOCOLRsrv` command, you must specify a physical WAN port, not a virtual port. This rule applies for all device port configurations.

More specifically, you configure protocol reservation using the following procedures for each traffic type:

- Procedure for mnemonic filtering

To configure protocol reservation by using the mnemonic filtering procedure, you must perform the following actions:

 - Using various Filter Service parameters, assign a built-in or user-defined mask, create a filter policy, and designate the type of packet filtering that is performed.
For mnemonic filtering descriptions, lists of built-in masks, and instructions on how to create user-defined masks, see [Mnemonic Filtering Configuration on page 433](#).
For syntax and descriptions of Filter Service parameters, see "Filter Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.
 - Set the `-PORT QueueCONTROL` parameter to `PROTOCOLRsrv` and use the `- PORT PROTOCOLRsrv` parameter to assign bandwidth percentage and a tag for the packet traffic type.
For syntax and descriptions of the `-PORT PROTOCOLRsrv` parameter, see "Port Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.
- Procedure for IP filtering

To configure protocol reservation by using the IP filtering procedure, you must perform the following actions:

 - Use IP Service parameters to create a filter and enable filtering.
For IP filtering descriptions and examples, see [Configuring Packet Filtering on page 131](#).
For syntax and descriptions of the IP Service parameters, see "IP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.
 - Set the `-PORT QueueCONTROL` parameter to `PROTOCOLRsrv`, and use the `- PORT PROTOCOLRsrv` parameter to assign bandwidth percentage and a tag for the packet traffic type.
For syntax and descriptions of the `-PORT PROTOCOLRsrv` parameter, see "Port Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

18.2

Protocol Reservation with Frame Relay Virtual Ports

When you configure protocol reservation for traffic sent over Frame Relay virtual ports, you must configure protocol reservation on the physical port. You can set up the filter to tag packets on virtual ports, but if you configure the filters using this method, you must configure filters for all virtual ports assigned to the physical port. If you have a large number of virtual ports and you configure protocol reservation filters for each virtual port, system performance is negatively impacted. It is recommended that you configure the protocol reservation filters to apply to the device instead of individual ports, then configure the protocol reservation percentages to apply to individual ports.


18.3

Configuring Protocol Reservation for Bridged Traffic

This procedure shows how to configure protocol reservation for all bridged protocol packets by using the mnemonic filtering.

In the following figure, a router is bridging IP traffic. You want to reserve 40 percent of the bandwidth for IP-RIP traffic, 35 percent for IP traffic, and 20 percent for bridged OSI traffic transmitting from WAN port 4. Five percent of the bandwidth is automatically set aside as a default for untagged traffic.

Figure 115: Hardware Configuration for Bridged Packets Example



For more information and examples of how to use the mnemonic filtering procedure to set up protocol reservation, see [Mnemonic Filtering Configuration on page 433](#).

Procedure:

1. Add a filter policy for each protocol with built-in IP-RIP, IP, and bridged OSI filter masks by entering the following commands:

```
ADD -Filter POLICY POLICY1 PROTOcolRsrv ANY_IPRIP IPRIP
ADD -Filter POLICY POLICY2 PROTOcolRsrv ANY_IP IP
ADD -Filter POLICY POLICY3 PROTOcolRsrv BRIDGED_OSI OSI
```

When you configure filters to control protocol reservation, the filters affect all ports on the device. You can configure filters for specific ports, but this configuration is not recommended for use with protocol reservation.

2. Apply the bridge filtering policies by entering:

```
SETDefault -Filter SElection = BRidge
```

3. Enable the Filter Service by entering:

```
SETDefault -Filter CONTrol = Enable
```

4. Define the percentage of bandwidth to be reserved for each protocol and enter name tags that match those entered in the `-Filter POLiCy` commands in [step 1](#) by entering the following commands:

```
ADD !4 -PORT PROToColRsrv ANY_IPRIP 40
```

```
ADD !4 -PORT PROToColRsrv ANY_IP 35
```

```
ADD !4 -PORT PROToColRsrv BRIDGED_OSI 20
```

5. Specify the `PROToColRsrv` option for the `-PORT Service QueueCONTrol` parameter by entering:


```
SETDefault !4 -PORT QueueCONTrol = PROToColRsrv
```

18.4

Configuring Protocol Reservation for IP-Routed Packets

This procedure provides an example of how to use the IP filtering procedure to configure protocol reservation for IP-routed packets.

Figure 116: Hardware Configuration for IP-Routed Packets Example



For more information and examples of how to use the IP filtering procedure to set up protocol reservation, see [Configuring Packet Filtering on page 131](#).

Prerequisites:

Ensure that:

- You added an IP filter that assigns 20 percent of reserved bandwidth for all Telnet sessions, 30 percent of reserved bandwidth for all FTP packets, and 25 percent for all other IP packets transmitted through port 2.
- You set the `-IP FilterDefAction` parameter so that all packets that do not meet the filtering conditions are forwarded.

Procedure:

1. Set up IP routing. See [IP Routing Configuration on page 119](#).
2. Add IP filters by entering the following commands:

```
ADD -IP FilterAddrs ALL ALL PROTOcolRsrv = TELNETTAG Telnet
ADD -IP FilterAddrs ALL ALL PROTOcolRsrv = FTPTAG FTP
ADD -IP FilterAddrs ALL ALL PROTOcolRsrv = ALLOther-IP IP
```

The filters are added to:

- Assign 20 percent of reserved bandwidth for all Telnet packets and designates a tag name of Telnet to identify the packets.
- Assign 30 percent of reserved bandwidth for all FTP packets and designates a tag name of FTP to identify the packets.
- Assign 25 percent of reserved bandwidth for all other IP packets and designates a tag name of ALLOther-IP to identify the packets.

When you configure filters to control protocol reservation, the filters affect all ports on the device. You can configure filters for specific ports, but this configuration is not recommended for use with protocol reservation.

3. Add an IP filter default action that forwards any packets that do not satisfy the filter requirements by entering:

```
SETDefault -IP FilterDefAction = Forward
```

4. Enable the IP filtering feature by entering:

```
SETDefault -IP CONTROL = Filtering
```

5. Assign 20 percent of bandwidth to the PROTOcolRsrv name tag Telnet, 30 percent of the bandwidth to the PROTOcolRsrv name tag FTP, and 25 percent of bandwidth to the PROTOcolRsrv name tag ALLOther-IP for port 2 by entering the following commands:

```
ADD !2 -PORT PROTOcolRsrv TELNETTAG 20
ADD !2 -PORT PROTOcolRsrv FTPTAG 30
ADD !2 -PORT PROTOcolRsrv ALLOther-IP 25
```

6. Set PROTOcolRsrv as the QueueCONTROL option for port 2 by entering:

```
SETDefault !2 -PORT QueueCONTROL = PROTOcolRsrv
```

After you entered these commands, any packets sent out by the system through port 2 that have the name tag Telnet are allocated 20 percent of the bandwidth, packets with the name tag FTP are allocated 30 percent of the bandwidth, and IP packets with the name tag ALLOther-IP are allocated 25 percent.

Five percent of the bandwidth is allocated for all untagged traffic, and the remaining 20 percent of the bandwidth is added to the default to be used by the configured protocols or by the untagged traffic on a first-come first-serve basis.

18.5


Protocol Reservation Allocating IP Protocol Types

Using IP filtering, how you define the tags for IP packets or other protocols in the TCP/IP protocol suite determines the percentage of bandwidth that is used for each protocol. If you configure a percentage of bandwidth for a specific protocol, such as UDP, TCP, or Telnet, then those packets are removed from the

percentage allocated to IP. However, if you define a percentage of bandwidth for IP only, then all the IP-related protocols such as UDP, TCP, and Telnet are included within that percentage.


The following figure illustrates how this allocation works. In the example, 60 percent of the bandwidth is allocated to IP.

Figure 117: IP Protocols Allocation (UDP Included)



The following figure shows the same configuration but with 60 percent of the bandwidth allocated to IP and 20 percent allocated to UDP. While UDP traffic is no longer included in the 60 percent of bandwidth allocated to IP, TCP and Telnet are still allocated as a subset of the IP bandwidth percentage. The bandwidth allocated to UDP in this case is exclusive of bandwidth allocated to IP.

Figure 118: IP Protocols Allocation (UDP Excluded)




18.6

Protocol Reservation on Virtual Ports

In this example, when you configure protocol reservation for Frame Relay virtual ports, you configure the protocol reservation on the parent port, not on the virtual port. No special configuration for virtual ports is necessary if you do not specify port numbers when you set up the filter masks.

The following figure shows a configuration with virtual ports. In this configuration, on IP Router A, you configure the specific filter masks without specifying a port number. You then configure the filter percentages by using the `-PORT PROTOCOLsrv` parameter on port 2, the parent port to virtual ports 1 and 2.

Figure 119: Protocol Reservation on Virtual Ports



18.7

Protocol Reservation Functionality

Protocol reservation allows you to assign a percentage of bandwidth to designated packets transmitting through a specified WAN logical port (no virtual ports) and meeting certain conditions. The conditions can include protocol type, packet length, and packets destined for a specified address, among others.

Protocol Reservation Controlling Bandwidth for Traffic Types

Protocol reservation provides a valve that sits above the transmit queue and controls the amount of bandwidth reserved for specific types of data (identified through filtering schemes). During times of stress when bandwidth of the WAN link is consumed beyond a threshold, the protocol reservation valve engages and works to normalize ratios of traffic types down to the configured percentages. If the bandwidth is not utilized to threshold, the device does not attempt to achieve the configured percentages so all packets can be serviced.

You identify the traffic through the use of mnemonic or manually configured filters. OSI traffic is treated as "default". You cannot assign percentages to these types of protocols; however, you can allocate percentages if these protocols are bridged. When the protocol reservation valve is engaged, each packet is checked against the configured filter(s), which introduces some latency when compared with traffic flow over an uncongested link. However, reserving minimum percentages for a protocol helps to prevent session loss, which may occur during traffic bursts.

Some protocols reduce their transmit rate when congestion is sensed, and thus may not use all of their allocated bandwidth. Protocol reservation automatically allocates any unused bandwidth to other protocols, and the desired effect may not be achieved. If you are using a protocol that reduces its transmit rate and does not utilize its configured bandwidth, you may want to use the priority queuing feature, which allows you to control the order in which packets are serviced between high, medium, and low priority queues. For more information, see [Filtering Process on page 436](#).

Protocol reservation supports PPP virtual ports, but for Frame Relay protocol reservation is supported only on the parent ports at the physical port level (specific percentages are not applied to individual virtual ports).

Tuning

The protocol reservation valve normalizes bandwidth to configured percentages over time intervals. If a large packet is encountered, the packet must be passed to the driver transmit queue in its entirety; it is not fragmented into smaller sizes. Other packets are passed to the queue behind it. Percentages are maintained

over time but it is still possible for some traffic to experience latency in extremely busy environments if larger packets fill up the driver transmit queue.

Do not allocate more bandwidth for a protocol than you can use. If a protocol cannot reach the percentage of bandwidth allocated to it, then the bandwidth not used by the protocol is used by other protocols.

Bandwidth Allocation Process Rules

Protocol reservation uses bandwidth allocation process rules for allocating bandwidth to WAN ports.

Bandwidth Normalization

You can use the protocol reservation features for IP routing and all bridging protocol traffic. To allow any traffic that is not tagged to go through, at least 5 percent of the bandwidth is reserved for this untagged traffic. This reservation is called the default queue.

If the total configured bandwidth percentages for the port exceed 95 percent, the values are balanced by the system so that the default queue still has its default allotment of approximately 5 percent of the available bandwidth. The rest of the bandwidth is distributed among the entries configured for the port in a ratio to the percentages that were configured for each.

This process of distributing the ratio is called normalization. Since the distribution only uses whole numbers for a percentage, the fraction remainders of each protocol are added to the default queue. As a result, the default queue can, at times, have a percentage greater than 5.

The following table provides an example of the traffic type, configured bandwidth, and the normalized numbers that occur when the configured assignments of bandwidth exceed 95 percent.

Table 43: Traffic Type, Configured Bandwidth, and Normalized Bandwidth

Traffic Type	Configured Bandwidth	Normalized Bandwidth
Default (untagged traffic)	5 percent	7 percent
IP traffic	95 percent	37 percent

Distribution of Non-Allocated Bandwidth

If the total configured bandwidth percentages are less than 95 percent, the non-allocated bandwidth is added to the default to be given to the configured protocols or for untagged traffic on a first-come first-served basis.

For example, consider a scenario in which you configure protocol reservation for a WAN port with the following bandwidth allocations:

- 50 percent of the bandwidth for IP traffic
- 20 percent of the bandwidth for IP-RIP traffic
- 5 percent automatically set aside as default bandwidth for untagged traffic

The remaining 25 percent of the bandwidth is added to the default to be used for IP-RIP traffic, or for untagged traffic, whatever traffic needs it first.

Chapter 19

Data Compression Configuration

This chapter describes how to configure data compression, an optional feature that you can use to enhance the effective throughput on Point-to-Point Protocol (PPP) and Frame Relay connections.



NOTE: If you use a modem that already performs compression, Motorola Solutions suggests that you do not configure data compression.

Data compression performs additional processing on the contents of each packet to look for repetitive patterns. Consequently, it is most effective when there is sufficient CPU cycles available to handle the additional processing. Data compression is most effective on slow lines.

19.1

Supported Data Compression Types

The EOS software supports the following types of data compression:

Tinygram compression

Packet-level compression that can be configured for PPP links only.

History-based compression

Link-level compression that you can configure for PPP links. History-based link-level compression requires a reliable link for proper operation. When you enable history-based compression over a PPP link, you must run Link Access Procedure, Balanced (LAPB) to provide the reliable link (for details, see [Supported Data Compression Types on page 393](#)). Frame Relay does not provide a reliable link and therefore does not support history-based compression.

Per-packet compression

Link-level compression that you can configure for PPP and Frame Relay links.

All three types of compression operate and are configured independently.

19.2

Tinygram Compression

Tinygram compression is packet-level compression. Tinygram compression is performed on packets with a length of less than 64 bytes. An increase in effective throughput is achieved by suppressing the transmission of trailing nulls, or hexadecimal zeros, in packets that are encapsulated in the Ethernet frame format. This type of compression is also referred to as local area transport (LAT) compression (since LAT packets are typically small in size and are padded up to 64 bytes with trailing nulls). The receiving end of a compressed packet can easily recreate the original packet by adding the trailing nulls. Because the CPU cycles involved in the stripping and adding of the trailing nulls are significantly less than the time it takes to transfer those nulls across a slow-speed line, the effective throughput of the system is increased.

The effectiveness of tinygram compression depends on the characteristics of your system. The following recommendations are based on line speeds of 64 kbps.

You can consider using tinygram compression in the following situations:

- In environments with small data packets using null character padding (for example, in LAT environments)

You should avoid using tinygram compression in the following situations:

- In environments where packets are generally transmitted with enough data to create 64-byte packets requiring no padding

- In environments where small data packets use random data for padding (for example in some Telnet environments).

19.3

Link-Level Compression

Link-level compression is performed over all packets sent on a specified link. The effective throughput is increased by sending fewer bytes across the link, as with tinygram compression. The algorithm used for link-level compression looks for repetitive data patterns in packets and replaces them with shorter length codes.

The EOS software supports the following types of link-level compression algorithms:

- History-based (the default when you select link-level compression)
- Per-packet

The algorithm used for history-based link-level compression looks for repetitive data patterns across multiple packets and replaces them with shorter length codes. The sending and receiving ends both build up a history buffer, and encode and decode the data in the packet according to that buffer. The history buffer has the last 2 KB of data. For proper encoding and decoding, the history buffer at each end must always be synchronized. Because the history information is transferred along with compressed data, the sending side must be assured that the receiving side reliably gets the data. As a result, history-based compression can operate **only** over a reliable data link. History-based compression requires that the LAPB Service be configured and operational over all the links on which this type of compression is desired.

The algorithm used for per-packet link-level compression looks for repetitive patterns within a packet and replaces them with shorter length codes. With per-packet compression, the sending and receiving ends do not preserve the history between packets. As a result, per-packet compression does not need to operate over a reliable data link, and you do not have to configure the LAPB Service over all links on which this type of compression is desired.

Because history-based compression looks for repetitive data across multiple packets, it is more effective in shrinking a packet size, which includes the line throughput. When considering history-based compression, however, you must consider the memory required to maintain a history buffer (approximately 26 KB of memory per interface), particularly if history-based compression is enabled on several links. Because a history buffer is not maintained in per-packet mode, the memory requirement for per-packet compression is considerably less than for history-based compression.

The effectiveness of link-level compression depends on the characteristics of your system. The following recommendations are based on line speeds of 64 kbps.

You can consider using link-level compression in the following situations:

- In environments with repetitive patterns in the bit stream transferred (with file transfers or email, for example)
- In environments where slow lines (64 kbps or lower) are used

You should avoid using link-level compression in the following situations:

- In environments with patterns in the bit stream that are *not* repetitive (in image files, for example)
- In environments where high-speed lines are used
- When the overall throughput of a system is already below normal

If you decide to use link-level compression, you must further decide which type to use: history-based or per-packet. The following recommendations, are based on line speeds of 64 kbps.

You can consider using history-based link-level compression in the following situations:

- In remote sites with 1–2 LAN ports and 1–2 WAN ports or in central sites with minimal LAN traffic and 3–6 WAN ports

- In sites where the wide area links are considered reliable

You can consider using per-packet link-level compression in the following situations:

- In central sites with more than 6 WAN ports
- In any site where there is significant LAN-to-LAN traffic

You should use the preceding recommendations as general guidelines. In cases where link compression is used, verify the effectiveness of the algorithm with the `-PORT LinkCompStat` parameter. If you do not see a significant difference between the raw and compressed bytes, the serial line throughput increase may not be enough to offset the overhead of applying the algorithm. If your system displays overall performance degradation, you should reevaluate the continued use of link compression.

19.4

Configuring Tinygram Compression

To enable tinygram compression, perform this procedure at both ends of the link.

The `-PATH TinyGramcomp` parameter allows you to compress all bridged Ethernet packets that are 64 bytes and are padded with trailing zeros. When the packet is sent on a serial line, the receiving side reinserts the zeros before forwarding the packet to an Ethernet LAN. This compression method is effective only on serial lines and is normally used in Digital Equipment Corporation/local area transport (DEC/LAT) terminal-to-host environments.

Procedure:

1. Enable tinygram compression on a specific path by entering:
`SETDefault !<path> -PATH TinyGramcomp = Enabled`
2. Verify the PATH configuration by entering:
`SHoWDefault !<path> -PATH TinyGramcomp`
3. Activate tinygram compression by entering:
`SETDefault !<path> -PATH CONTroL = Enabled`

19.5

Configuring History-Based or Per-Packet Compression

To enable history-based or per-packet link-level compression, perform this procedure at both ends of the link.

Procedure:

1. Enable link-level history based compression on a particular port by entering:
`SETDefault !<port> -PORT COMPResType = HIStory`
 **NOTE:** Frame Relay does not support history-based data compression.
2. Select the per-packet link-level compression algorithm by entering:
`SETDefault !<port> -PORT COMPResType = PerPacket`
3. For the configured compression types to take effect, re-enable the port by entering:
`SETDefault !<port> -PORT CONTroL = Enabled`

Before you re-enable the port with this command, you must apply optional configuration. See [Enabling LAPB for a PPP Link on page 396](#) and [Frame Relay Configuration Options on page 396](#).

19.6

Enabling LAPB for a PPP Link

If you plan to run history-based compression on a PPP link, you can use this procedure to set up LAPB.

Procedure:

1. Enable LAPB at both ends of the link by entering:
`SETDefault !<path> -LAPB CONTrol = Enable`
2. Configure one end of the serial link as data terminal equipment (DTE) by entering:
`SETDefault !<path> -LAPB InterfaceType = DTE`
3. Configure the other end of the serial link as data communications equipment (DCE) by using:
`SETDefault !<path> -LAPB InterfaceType = DCE`

19.7

Frame Relay Configuration Options

When you set up data compression for Frame Relay links, you can specify data compression for each individual data link connection identifier (DLCI). The `-FR COMPresSType` parameter overrides the `-PORT COMPresSType` parameter.

Data compression is not negotiated over Frame Relay connections. You must configure data compression appropriately at both ends of the link.

- To enable per-packet data compression on a specific DLCI in an Frame Relay link, you must use:
`SETDefault !<port> -FR COMPresSType = <dlci> PerPacket`
- To disable data compression on a specified DLCI, you must use:
`SETDefault !<port> -FR COMPresSType = <dlci> NONE`
- To use the compression type configured by the `-PORT COMPresSType` parameter, you must use:
`SETDefault !<port> -FR COMPresSType = <dlci> DEFault`
The `DEFault` value for the `-FR COMPresSType` parameter in this command requires that the DLCI use the compression type configured on the PORT Service.

19.8

Verifying the Link-Level Compression Effectiveness

In multiprotocol and mixed-application environments, it may be difficult to achieve a consistent high level of compression. The effectiveness of compression is measured by the ratio of the uncompressed data to the compressed data, also known as the compression ratio. The software can display the number of raw and compressed bytes, which you can use to measure the effectiveness of link-level compression for your environment.

After you configure link-level compression, you may need to determine its effectiveness for your network environment. To decide whether link-level compression is beneficial on your network, use the following procedure.

Procedure:

1. Display the link compression status by entering:
`SHow !<port> -Port LinkCompStat`
The system displays statistics accumulated since link compression was configured.
Example:


```
-----Compression Statistics for Port = 4 LCN = 1-----
Owner      = FrameRelay
CompType   = PerPacket
TX_Raw     = 742
TX_Comp    = 200
TX_Ratio   = 3.71
RX_Raw     = 452
RX_Comp    = 162
RX_Ratio   = 2.79
TX_Fail    = 0
RX_Err     = 0
```

2. If you want to recalculate link compression performance, remove the old statistics by entering:

```
FLush !<port> -Port LinkCompStat
```

3. Compare the number of raw bytes to compressed bytes for both the transmit and receive sides and check all ports that use compression.

Compression is very CPU-intensive. You may want to disable compression if you do not get a favorable compression ratio, which depends on the nature of the data, or if the overall system performance suffers because of CPU overloading.

Chapter 20

Multiprotocol Data Prioritization

This chapter describes how to use the data prioritization feature to assign a priority (high, medium, or low) to most packets that are forwarded over a wide area network using Point-to-Point Protocol (PPP) or Frame Relay.

You can achieve the following benefits by using the data prioritization feature:

- Control data traffic on heavily used wide area networks.
- Allow the following types of packets to have a higher priority over other data traffic on a wide area network:
 - Network-critical traffic; for example, bridge spanning tree packets (system-configured)
 - Mission-critical traffic; for example, voice packets (user-configured)
 - Time-critical traffic; for example, Telnet packets (system-configured)
 - Specific protocol packet; for example, Internet Protocol (IP)-routed packets (user-configured)
- Allow sessions to be prioritized according to the session characteristics.

20.1

Data Prioritization Overview

The EOS software implements a prioritization scheme that assigns a priority to each packet then forwards the packet to an urgent-, high-, medium-, or low- priority queue. The packets are forwarded from the queues onto a wide area network using PPP or Frame Relay in an order controlled by a queue arbitration algorithm.

The system assigns a priority to some packets. The following table lists these packet types and the priority assigned to them by the system. You cannot reconfigure the priority of these packets. The system assigns a priority to Telnet packets that originate from the box; it does not assign a priority to routed Telnet packets.

Table 44: Packets with System-Assigned Priorities

Packet Type	Priority Level
Bridge Spanning Tree	Urgent
OSPF	High
Telnet	High

You can assign a priority to all other types of packets in the following ways:

- For IP-routed packets, you can assign a high, medium, or low priority using the `SETDefault -IP QueuePriority` command.
- For all packets other than IP-routed packets, you can set up a mask that determines what packets should be prioritized and a policy that determines what priority the packets should be assigned.

Any packet not specifically assigned a priority receives its priority from the setting of the `-PORT DefaultPriority` parameter. The default setting of this parameter is medium.

When multiple ports are attached to one path, no one particular port receives a higher priority over another port. All ports attached to one path receive the same priority.

For complete information on assigning priorities to packets, see [Configuring Data Prioritization on page 402](#). For more information, see “PORT Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

20.1.1

Priority Assigned to Packets

All packets are assigned either an urgent, high, medium, or low priority before they are transmitted over a wide area interface.

The system assigns a priority to a packet using the following process:

1. The system assigns a priority to certain packet types. You cannot change the priority assigned to these packets.
2. The system assigns a priority to IP-routed packets with the setting of the `-IP QueuePriority` parameter. You can change the default settings of these parameters.
3. The system assigns a priority to all packets other than IP-routed packets through a prioritization filter. You set up a prioritization filter by configuring masks that determine the types of packets that should be prioritized using the `ADD -Filter MASK` command. You should assign a configuration policy that determines the priority each type of packet using the `ADD -Filter POLICY` command.
4. The system assigns a priority to all other packets that do not receive their priority through any of the previously discussed methods via the setting of the `-PORT DefaultPriority` parameter. You can change the default setting of this parameter.

After the system assigns a priority to a packet using one of methods described above, the packet is forwarded to a queue for transmission to a wide area network. For example, if a packet is assigned a priority by using the `-IP QueuePriority` parameter, the packet does not undergo the following methods. The system handles packet priority assignment in this way to reduce the number of packets that are sent through the prioritization filter because filtering can impact system performance.

20.1.2

Queue Types


After the system assigns a priority to a packet automatically or through configured or default parameter settings, the packet is forwarded to one of four types of queues: urgent, high, medium, or low. The following figure shows how particular types of packets are assigned priorities and how they are forwarded to the various queues.

The following figure assumes that you configured:

- A medium priority for IP-routed packets.
- The value of low for the `-PORT DefaultPriority` parameter.

The other packets are assigned their priorities by the system.

Figure 120: Prioritizing and Forwarding Packets to Queues



The system forwards all packets in the urgent-priority queue before packets in the high-, medium-, and low-priority queues. You cannot assign the urgent priority to a packet. The system automatically assigns an urgent priority to bridging spanning tree packets only.

After all packets in the urgent-priority queue are forwarded, the system forwards packets from the high-, medium-, and low-priority queues according to a queue arbitration algorithm. For more information on this

algorithm, see [Queue Arbitration Algorithm on page 401](#). The system automatically assigns a high priority to open shortest path first (OSPF) and Telnet packets and forwards these types of packets to the high-priority queue.

20.1.3

Queue Arbitration Algorithm

Instead of forwarding all packets from high-priority queues, then all packets from medium-priority queues, and so on, the system uses a queue arbitration algorithm, which ensures that the high-, medium-, and low-priority queues are serviced according to an interleave factor.

You can configure the interleave factor by entering:

```
SETDefault -PORT QueueInterLeave
```

This command allows you to set up the forwarding ratio of high- to medium- priority packets and of medium- to low-priority packets.

The algorithm implements an 8-slot queue that is composed of a variable number of high-, medium-, and low-priority slots. Urgent packets bypass this queue and are immediately forwarded.

The number of high-, medium-, and low-priority slots in the queue are based on the setting of the `-PORT QueueInterLeave` parameter. The system chooses the closest of five possible 8-slot patterns.

Table 45: -PORT QueueInterLeave Parameter Values and 8-Slot Patterns

-PORT QueueInterLeave Value	Corresponding 8-Slot Pattern
6, 1	HHMHHLHH
3, 2 ¹	HHMHMHLH
2, 2	HHMHHLHM
2, 1	HMHLHMHL
1, 1	HMHLHMLM
¹ Default value	

The actual transmission rate of high, medium, and low packets on a WAN link may not exactly match the 8-slot pattern because of the packet receive rate and receive pattern. For example, only medium and low packets may arrive in one window of time, and if no high packets are available to send at this time, then the medium and low packets are sent out on the serial line. The low and medium transmission queue can get flushed to favor high-priority packets. For example, in a situation where the packet receive rate is much higher than the WAN link can handle, many packets are dropped, and with the default queue arbitration pattern of 5 high, 2 medium, and 1 low, the actual transmission ratio may be 6 high, 1 medium, and 1 low.

To display the 8-slot pattern selected by the system, you must use:

```
Show -PORT QueuePATtern
```

A ratio based on the pattern also is displayed. For example, for the default value of the `-PORT QueueInterLeave` parameter (3, 2), the following display appears:

```
HHMHMHLH (5:2:1)
```

This display indicates that the first and second packets are sent from the high-priority queue, the third packet is sent from the medium-priority queue, and so on. Once the eighth packet is sent, the algorithm wraps to the beginning of the pattern again and the first and second packets are sent from the high-priority queue, and so on.

The contents of this display are based on the setting of the `-PORT QueueInterLeave` parameter and are the results of a translation algorithm. The contents of this display are calculated by the system and cannot be reconfigured.

If a packet is due to be sent from the high-priority queue but that particular queue is empty, a packet from the medium-priority queue is sent instead; if the medium-priority queue is also empty, a packet from the low-priority queue is sent instead. If a packet is due to be sent from the medium-priority queue but that particular queue is empty, a packet from the high-priority queue is sent instead. If a packet is due to be sent from the low-priority queue but that particular queue is empty, a high-priority packet is sent instead.

For more information, see the “POrt Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

20.2

Configuring Data Prioritization

This procedure describes how to assign a medium priority to IP-routed packets and a low priority to all other packets.

For more information on the parameters used in this procedure, see “Filter Service Parameters”, “IP Service Parameters”, and “POrt Service Parameters” chapters in the *Enterprise OS Software Version 18.2 Reference Guide*.

Prerequisites:

Ensure that:

- You determined what types of packets you want to prioritize and what priority you want to assign to each type.
For example, you may want to assign IP-routed packets a high priority, OSI packets a medium priority, and all other packets a low priority. You also may want to prioritize types of packets within the IP protocol itself.
- You familiarized yourself with the different ways you can assign a priority to a type of packet and determined which option you want to use for each type of packet you want to prioritize.
For example, you may want to assign a high priority to IP-routed packets, a medium priority to bridged OSI packets by setting up a mask and policy, and a low priority to all other packets.
- You determined the interleave factor, which is defined as the ratio of packets that you want forwarded from high- to medium-priority queues and the ratio of packets you want forwarded from medium- to low-priority queues. For more information on the interleave factor and queue arbitration, see [Data Prioritization Overview on page 398](#).

Procedure:

1. To assign a priority to IP-routed packets, enter:

```
SETDefault -IP QueuePriority = <H / M / L / DEFault>
```

If you retain the default setting of `DEFault` for any of the commands listed above, the system uses the setting of the `-PORT DefaultPriority` parameter.

Step example:

To assign a medium priority to IP-routed packets, you must enter:

```
SETDefault -IP QueuePriority = M
```

2. If you want to assign a priority to packets other than IP-routed packets, perform the following actions:
 - a. Set up a mask that determines the types of packets that should be prioritized.
You can set up either a built-in or a user-defined mask. Since using a built-in mask requires less configuration, Motorola Solutions recommends this option. To determine if a built-in mask exists

for the type of packet you want to prioritize, see [Mnemonic Filtering Configuration on page 433](#). If a built-in mask that suits your purposes exists, you can skip this step.

For more examples of various masks configuration, see [Mnemonic Filtering Configuration on page 433](#).

Step example:

Since a built-in mask for OSI packets does not exist, if you want to assign a medium priority to OSI packets, you need to configure one by using:

```
ADD -Filter MASK <maskname> <location> [ <operation>] <pattern>
```

To configure a mask that prioritizes packets with a value greater than %45 at the first byte of data, you must enter:

```
ADD -Filter MASK some_data dl.data+%0>%45
```

- b. Set up a policy that determines the priority each type of packet should be assigned by entering:

```
ADD -Filter POLicy <polycyname> <action> <masks> [ <context>]
```

3. If you want to change the default priority, enter:

```
SETDefault -PORT DefaultPriority = <H / M / L>
```

The default priority for packets to which you do not specifically assign a priority (as in [step 1](#) and [step 2](#)) is medium.

Step example:

To change the default priority for all packets other than IP, enter:

```
SETDefault -PORT DefaultPriority = L
```

4. If you want to reconfigure the default ratio of forwarded packets, enter:

```
SETDefault !<port> -PORT QueueInterLeave = <ratio1> <ratio2> (1-10)
```

The default ratio of packets forwarded from high- to medium-priority queues is 3 and the ratio of packets forwarded from medium- to low-priority queues is 2.

For more information on forwarding ratios and queue arbitration, see [Data Prioritization Overview on page 398](#).

Step example:

To change the default values to 6 and 3 on port 3, enter:

```
SETDefault !3 -PORT QueueInterLeave = 6 3
```

5. If you anticipate that high-priority packets may be slowed down by many medium- and low-priority packets, (especially large medium- and low-priority packets) in the queue, adjust the number of medium- and low-priority packets forwarded to the queue by entering:

```
SETDefault !<port> -PORT QueueThrottle = <1-40>
```

By default, the serial line driver accepts as many medium- and low-priority packets as it can possibly handle.

Motorola Solutions recommends setting the value between 1 and 20 if you are using a slow-speed line (64K or below) and between 21 and 40 if you use a high-speed line (for example, T1).

20.3

Data Prioritization Parameters

The following table summarizes the available prioritization parameters and the services that support these parameters.

Table 46: Data Prioritization Parameters

Service	Parameter	Description
Filter	MASK	Sets up a mask that determines which packets should be prioritized.
	POLICY	Sets up a policy that determines the priority a particular type of packet is assigned.
IP PORT	QueuePriority	The <code>-IP QueuePriority</code> parameter sets the priority of IP-routed packets, respectively; if these parameters are set to <code>DEFAULT</code>, the system uses the setting of the <code>-PORT DefaultPriority</code> parameter. The <code>-PORT QueuePriority</code> parameter displays the settings of the following parameters: <code>-IP QueuePriority</code> <code>-PORT DefaultPriority</code>
IP	FilterAddrs	The <code>QPriority</code> option for the <code>FilterAddrs</code> parameter specifies a queue priority value of high, medium, or low.
PORT	DefaultPriority	Sets the default priority of packets if one of the following conditions apply: - The <code>-IP QueuePriority</code> parameter is set to <code>DEFAULT</code>. - A mask and prioritization policy is not configured for a particular type of packet.
	QueueInterLeave	Sets and displays the interleave factor, which is defined as the forwarding ratio of high- to medium-priority and medium- to low-priority packets.
	QueuePATtern	Displays the interleave factor configured by the <code>-PORT QueueInterLeave</code> parameter translated by the system into a high, medium, and low pattern. A ratio based on the high, medium, and low pattern also displays.
	QueueThrottle	Controls the number of medium- and low-priority packets that are forwarded to the driver each time packets from the priority queue are forwarded on to the wide area network.

For more information on the parameters, see the “Filter Service Parameters”, “IP Service Parameters”, and “PORT Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Chapter 21

UDP/LAPB Packet Relay Configuration

The EOS implementation of UDP/LAPB packet relay adds a state machine definition and the capacity for processing UDP messages.

When a FlexWAN serial port on an S6000 router is configured as a LAPB relay port (Port Owner is set to LAPBR), the following occurs:

- Any UDP packets with port number 2100 are forwarded to the LAPB relay port.
- Any LAPB frames received on the LAPB port are converted to UDP packets with port number 2100 and are transmitted to the Ethernet port receiving those particular UDP packets.

Windows of 3 UDP packets and 7 LAPB frames are provided to relay to the LAPB port and to the UDP port respectively.



NOTE: LAPB relay is supported only on S6000 FlexWAN serial ports.

21.1

Configuring UDP/LAPB Packet Relay on MNR S6000 Routers

Procedure:

1. Start the LAPB relay software by setting the port owner to LAPB relay by entering:

```
SETDefault !<port> -Port OWNeR = LAPBR
```

Step example:

To start the LAPB relay software on !3a, enter:

```
SETDefault !3a -Port OWNeR = LAPBR
```

2. Configure the LAPB interface type by entering:

```
SETDefault !<port> -LAPB InterfaceType = DTE | DCE
```

Step example:

To configure the LAPB interface type to DTE for port 3a, enter:

```
SETDefault !3a -LAPB InterfaceType = DTE
```

3. Set the T1 value by entering:

```
SETDefault !<port> -LAPB T1 = <milliseconds> (100-1000000)
```

T1 value means the maximum time that the LAPB protocol waits for an acknowledgement once a frame is transmitted.

Step example:

To set the T1 value to 500 milliseconds for port 3a, enter:

```
SETDefault !3a -LAPB T1 = 500
```

4. Set the bandwidth of the LAPB relay port by entering:

```
SETDefault !<port> -PAth BAud = <kbps> (1.2-16000)
```

Step example:

To set the bandwidth of port 3a to 256 kbps, enter:

```
SETDefault !3a -PAth BAud = 256
```

Chapter 22

Cooperative WAN Routing Configuration

The Motorola Solutions cooperative WAN routing (CWR) solution provides redundant WAN failover capabilities, allowing core and exit routers to interface directly with site and inter-zone links via a relay panel.

The CWR solution requires the following:

- 2 MNR S6000 routers with one or two 12-port T1/E1 modules installed in each router, depending on how many T1/E1 ports are required for the specific implementation.
- 1 relay panel
- 2 or 4 12-port relay cables (part number DKN6128A (3-foot cable) or DKN6129A (6-foot cable)).
 - 2 12-port relay cables for CWR applications requiring 1-12 T1/E1 ports.
 - 4 12-port relay cables for CWR applications requiring 13-24 T1/E1 ports.
- EOS software version 14.0 or higher

The CWR solution implements S6000 routers deployed in pairs for redundancy and switchover protection. S6000 routers configured and deployed in this way are referred to as CWR peers. The CWR peers control the passive relays on the relay panel, which switch all the links on the interface to one of the routers. The routers detect a cable or CWR peer failure and switch the relays for the site or inter-zone links to the inactive router.

22.1

Configuring MNR S6000 Routers as CWR Peers

You can use this procedure to configure CWR functionality on S6000 routers. For details about the EOS CWR service parameters, see “CWR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: CWR functionality is supported only on MNR S6000 routers with a 12-port T1/E1 module installed, running EOS software version 14.0 or higher, and connected to a relay panel.

Prerequisites:

Ensure that:

- Routers are connected to the relay panel.
- RJ-48C cables are attached between the T1/E1 interfaces on the relay panel and the appropriate device.

Procedure:

1. Log on to the system with Network Manager privileges.
2. On each router, perform the following actions:

This configuration is the same as the one used when configuring non-CWR T1/E1 ports and paths, except that the connector range for the 12-port T1/E1 module is 4A through 4L and 5A through 5L (as opposed to 4A through 4D and 5A through 5D when the S6000 is equipped with a 4-port UltraWAN module).

Physical T1/E1 connectors 1-24 on the relay panel correspond to logical paths on the S6000 with the following mapping:

Table 47: Physical T1/E1 Connectors and S6000 Logical Paths

T1/E1 Ports	Logical Paths
T1/E1 port 1	Logical path 4A
T1/E1 port 2	Logical path 4B
T1/E1 port 3	Logical path 4C
T1/E1 port 4	Logical path 4D
T1/E1 port 5	Logical path 4E
T1/E1 port 6	Logical path 4F
T1/E1 port 7	Logical path 4G
T1/E1 port 8	Logical path 4H
T1/E1 port 9	Logical path 4I
T1/E1 port 10	Logical path 4J
T1/E1 port 11	Logical path 4K
T1/E1 port 12	Logical path 4L
T1/E1 port 13	Logical path 5A
T1/E1 port 14	Logical path 5B
T1/E1 port 15	Logical path 5C
T1/E1 port 16	Logical path 5D
T1/E1 port 17	Logical path 5E
T1/E1 port 18	Logical path 5F
T1/E1 port 19	Logical path 5G
T1/E1 port 20	Logical path 5H
T1/E1 port 21	Logical path 5I
T1/E1 port 22	Logical path 5J
T1/E1 port 23	Logical path 5K
T1/E1 port 24	Logical path 5L

- a. Configure the Ultra WAN (T1/E1) connectors.



NOTE: T1/E1 connector configuration templates are available for your reference from the *Enterprise OS Software and User Documentation CD*. To access the templates, extract the .tar file from the CD-ROM onto a directory on your PC and navigate to:

`<X>:/<directory>/ 16.0.0.<xx>/u_guides/html/hw/pathcfg.htm`

where:

`<X:/>` is the drive letter for the directory

`<directory>` is the name of the directory into which you extracted the .tar file

`<xx>` is the specific build of EOS software provided on the CD-ROM.

Step example:

If you extracted the .tar file into directory `D:/EOS_SW`, and the CD-ROM contains EOS version 16.0.0.11, navigate to:

`D:/EOS_SW/16.0.0.11/u_guides/html/hw/pathcfg.htm`

- b. Disable all the paths on each T1/E1 connector you want to configure, and enable only the paths you want to use and map the DS0 timeslot channels to the paths by entering the following commands:

```
SETDefault !<port> -Port CONTROL = Disable
```

```
SETDefault !<port> -Port CONTROL = Enable
```

```
ADD !<port> -Path ChannelMap <channel numbers, separated by commas>
```

where *<port>* is the logical path number corresponding to the channel you want to configure.

Step example:

```
SETDefault !4A.1-!4A.24 -Port CONTROL = Disable
```

```
SETDefault !4A.1 -Port CONTROL = Enable
```

```
ADD !4A.1 -Path ChannelMap
```

```
1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24
```

- c. Configure the ports as Frame Relay ports and configure the Frame Relay parameters for these ports.

Step example:

```
SETDefault !<port> -Port OWNER = FrameRelay
```

```
SETDefault !<port> -FR CONTROL = ANSiLMI
```

```
SETDefault !<port> -FR LMIPROcedure = BIDIR
```

```
SETDefault !<port> -FR KATime = 3
```

```
SETDefault !<port> -FR FSEInterval = 3
```

```
SETDefault !<port> -FR DTEMonEvent = 4
```

```
SETDefault !<port> -FR PVTime = 3
```

```
SETDefault !<port> -FR DCEMonEvent = 4
```

- d. Create and enable virtual ports mapping Frame Relay DLCIs to the paths, by entering:

```
ADD !V<port> -Port VirtualPort MFR !<port>@<DLCI>
```

```
SETDefault !V<port> -Port CONTROL = Enable
```

where *<port>* is the virtual port number and *<DLCI>* is the DLCI number.

For more information, see [Configuring Frame Relay and PPP Virtual Ports on page 93](#).

- e. Configure additional Frame Relay and IP parameters as needed by your application.

For more information, see [Wide Area Networking Configuration with Frame Relay on page 308](#), "Frame Relay Service Parameters" and "IP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE:

If you configure large numbers of ports, you may find it helpful to use instance groups. For information about creating instance groups, see [Multiple Instance Lists on page 96](#).

3. On each router, enable CWR functionality on the 12-port T1/E1 module(s), by entering:

```
SETDefault !<port> -CWR CONTROL = Enabled
```

where *<port>* is the WAN port corresponding to the I/O slot in which the module is installed: 4 (I/O slot A) or 5 (I/O slot B).

Step example:

On Router A, enter the following commands:

```
SETDefault !4 -CWR CONTROL = Enabled
```

```
SETDefault !5 -CWR CONTROL = Enabled
```

On Router B, enter the following commands:

```
SETDefault !4 -CWR CONTROL = Enabled
```

```
SETDefault !5 -CWR CONTROL = Enabled
```

4. Set the default CWR state to active on one router and to inactive on the other router by entering:

```
SETDefault -CWR DefaultState = Active | Inactive
```

Step example:

On Router A, enter:

```
SETDefault -CWR DefaultState = Active
```

On Router B, enter:

```
SETDefault -CWR DefaultState = Inactive
```

5. For each router, configure the LAN IP addresses of the CWR peer, by entering:

```
ADD -CWR LanIpPair <Local_IP_Address> <Peer_IP_Address>
```

where:

<Local_IP_Address> is the IP address of either the physical LAN port or a non-tunneled VLAN port on the router you are configuring

<Peer_IP_Address> is the IP address of either the physical LAN port or a non-tunneled VLAN on the other router (the CWR peer).

Step example:

On Router A, enter:

```
ADD -CWR LanIpPair 10.1.1.1 10.1.1.2
```

On Router B, enter:

```
ADD -CWR LanIpPair 10.1.1.2 10.1.1.1
```

6. On each router, configure BGP, multicast routing, network management, OSPF, port, and NTP parameters as needed by your application.
7. Enable the ports and paths you are using and disable unused ports.

Chapter 23

Bridging Configuration

This chapter describes how to set up, customize, and troubleshoot a bridge.

23.1

Transparent Bridging

Transparent bridging is supported on Ethernet, and the following wide area networks: Frame Relay, PPP, and ISDN. When transparent bridging is enabled, the bridge forwards packets based on the destination address in the packets it receives. It also learns and records information about the location and addresses of devices on the surrounding networks, based on the source address in the received packets.

Maximum Transmission Unit

The maximum transmission unit (MTU) is the maximum packet size allowed; it varies according to the LAN media. You must configure applications that run in a multimedia bridged environment to use packet sizes that are equal to or smaller than the smallest of the MTU sizes in the extended LAN; otherwise, some media may drop oversized packets. If a particular application cannot accept smaller packets, using network layer routing instead of MAC layer bridging may provide a solution.

For IP packets bridged between interfaces that have mismatched MTU sizes, you can enable the IP fragmentation feature by setting the `-BRIDGE CONTrol` parameter to `IPFragment`. The bridge then fragments IP packets that are forwarded to ports with a smaller MTU size.

IP-Specific Issues

MAC-layer bridging typically does not bridge large frames between physical media that have dissimilar maximum frame sizes. To solve this problem, the EOS software implements fragmentation of bridged IP packets. IP fragmentation is supported between LAN media and also on WAN media. You can enable fragmentation by setting the `-BRIDGE CONTrol` parameter to `IPFragment`.



NOTE: Fragmentation may cause some deterioration in performance.

23.2


Spanning Tree Algorithm

The Spanning Tree algorithm detects loops and puts some bridge ports into blocking state, if necessary, so that only one route exists between any two stations. (A port in blocking state does not forward or receive packets.) Eliminating the extra paths creates a stable network configuration. When one or more bridges or ports in the stable topology fail, the algorithm automatically returns some ports from blocking state to forwarding state to ensure that all stations are connected.

For the Spanning Tree algorithm to be effective, all bridges in your extended network must run it.

An extended network without loops can be viewed as a Spanning Tree. A Spanning Tree is a topology in which one node is designated as the root, and any two nodes are connected to each other through one and only one route. The following figure depicts an example of the Spanning Tree structure in which one bridge represents the root, other bridges represent the nodes, and the communication lines represent the branches. The arrows illustrate the unique path that a packet from the San Francisco bridge takes when destined for the Los Angeles bridge. The topology would not be a Spanning Tree if there were also a line directly linking the San Francisco bridge and the San Jose bridge, or if the line between the San Jose bridge and the Santa Clara bridge were broken.

Figure 121: Spanning Tree Structure




When more than one bridge connects LANs, the network manager may inadvertently configure the network with loops, causing packets to circulate indefinitely. A loop exists if more than one path can be used to forward a packet from one end station to another.

For example, the dotted line in the following figure highlights a loop; packets from a station on LAN 1 can be forwarded to one or more stations on LAN 2 via either bridge A or bridge B. The destination stations receive duplicate packets because both bridge A and bridge B replicate the packet and then forward the packet to LAN 2. If the station sends out a broadcast packet, both bridges forward it to their attached networks, creating packets that circulate indefinitely.

The Spanning Tree algorithm detects and breaks loops that can form within a bridging topology.

Figure 122: Network with a Loop



Spanning Tree Algorithm Functionality

The Spanning Tree algorithm configures the network so that no loops exist in the extended network, and every two LANs can communicate with each other. This section lists the prerequisites required for the Spanning Tree algorithm to work and gives an example to explain how the algorithm arrives at a loop-free configuration.

For the Spanning Tree algorithm to configure the network correctly:

- Each bridge must be able to recognize a unique destination address.
- Each bridge must have a unique identifier (bridge ID) that contains a priority field and a data link address.
- Each port of a bridge must have a unique identifier (port ID) that contains a priority field and a port number.
- Each port must be associated with a path cost, which is determined by the speed of its network interface (the faster the speed, the smaller the cost).

Loop-Free Configuration with Spanning Tree Algorithm

To arrive at a loop-free configuration based on the bridge ID, port ID, and path costs, the Spanning Tree algorithm performs the following tasks:

- Selects a bridge that acts as the root of the Spanning Tree network. This is usually the bridge with the lowest bridge ID of all the bridges on the extended network.
- Selects a root port on each bridge (except the root bridge) that incurs the lowest root path cost when the bridge forwards a packet to the root bridge.
- Selects the designated bridge on each LAN that incurs the lowest path cost when forwarding a packet from that LAN to the root bridge. The port through which the designated bridge is attached to the LAN is called the designated port.
- Enables all root ports and designated ports so they can forward packets, and blocks all other ports.

The following example shows how the algorithm makes the selections, then eventually eliminates loops. The following figures illustrate an extended network. In these figures, the bridges are numbered from 1 to 5, where bridge 1 has the lowest data link address, and bridge 5 has the highest.

When the bridges are turned on, each assumes that it is the root bridge. Each bridge then transmits a packet called the Configuration Bridge Protocol Data Unit (CBPDU) through all its ports. A CBPDU contains information such as the ID of the bridge that the transmitting bridge considers the root bridge, the root path cost of the transmitting bridge, and the number of the source port.


When a bridge receives a CBPDU on one of its ports that contains information that is superior to that currently stored at that port, it stores the information at that port. If this CBPDU is received at the root port of the bridge, the bridge also forwards it with an updated message to all attached LANs for which it is the designated bridge.

If a bridge receives a CBPDU on one of its ports that contains information inferior to that currently stored at that port, it discards it. If the bridge is a designated bridge for the LAN from which the CBPDU is received, it sends that LAN a CBPDU containing the up-to-date information stored at that port. In this way, inferior information is discarded and superior information is propagated on the extended network.

Root bridge

Assume that each port in the following figure is equipped with an Ethernet interface that has a path cost of 100, and that the priority fields in the IDs of bridge 1 and bridge 3 are the same. Having the lowest bridge ID (because its data link address is the lowest), bridge 1 becomes the root bridge, and its CBPDU is superior to those of the other bridges. After exchanging a few CBPDUs and discarding the inferior ones, all bridges contain the same information that indicates that bridge 1 is the root bridge. Because a root bridge is automatically the designated bridge for all LANs to which it is attached, bridge 1 is also the designated bridge for LANs 1 and 2.

Figure 123: Root Bridge




Root port

Each bridge (except the root bridge) has to select a root port that incurs the least cost when the bridge forwards a packet to the root. The cost depends partly on the path cost of the port (determined by the speed of its network interface) and partly on the root path cost of the designated bridge for the LAN to which this port is attached.

For example, in the following figure, ports 1 and 2 of bridge 3 both have the same network interface type and the same path cost, while bridge 3 incurs less cost if it forwards a packet from port 1 than from port 2. Thus the Spanning Tree algorithm determines that port 1 should be the root port for bridge 3.


Figure 124: Root Port



Selecting a Designated Bridge – Single Bridge

If a LAN is attached to a single bridge, that bridge is the designated bridge of the LAN. For example, in the following figure, bridge 2 is the designated bridge for LAN 3, because bridge 2 is the only bridge attached to LAN 3.

Figure 125: Selecting a Designated Bridge – Single Bridge




Selecting a Designated Bridge – Multiple Bridges Attached to Network

For a LAN that is attached to more than one bridge, a designated bridge must be selected. For example, in the following figure, because LAN 4 is attached to bridge 3, bridge 4, and bridge 5, the algorithm must compare the root path costs of these bridges. In this case, their root path costs are the same. Having the lowest bridge ID, bridge 3 becomes the designated bridge for LAN 4. Because bridge 3 is attached to LAN 4 through port 2, port 2 is the designated port for LAN 4.

Bridge 1, which is the root bridge, is automatically the designated bridge for all attached LANs (LANs 1 and 2). Because bridge 2 is the only bridge attached to LAN 3, it becomes the designated bridge for LAN 3.

Figure 126: Selecting a Designated Bridge – Multiple Bridges Attached to Network



When a port is in forwarding state, it performs learning, filtering, and forwarding functions. When it is in blocking state, it performs none of these functions.

Only root ports and designated ports are put into forwarding state. Other ports, such as port 1 of bridge 4 and port 2 of bridge 5, are put into blocking state.

Because some ports are put into blocking state, none of the packets circulate on the extended network indefinitely.


23.2.1

Spanning Tree Algorithm with Wide Area Bridges

Local and wide area bridges participate in configuring loop-free networks by using the Spanning Tree algorithm.

In the following figure, two bridges connect two remote networks. On each bridge, one of the interfaces is a network interface, and the others are serial links connected to the other wide area bridge. To apply the Spanning Tree algorithm in such a network configuration, it is assumed that the serial links are attached to an imaginary backbone network on which no end stations exist. The only traffic on the backbone is the traffic between the bridges. With this assumption, all bridge interfaces operate as if they were network interfaces, and the same Spanning Tree principle described in the previous section applies.

Figure 127: Two Wide Area Bridges Connected to Imaginary Backbone



When wide area bridges with parallel links participate in the Spanning Tree algorithm, all remote links connected to the same wide area bridge are considered one network interface. The algorithm puts all links into either forwarding or blocking state. This ensures that the network topology can maximize the use of the bandwidth provided by parallel network links.

If you configure your wide area bridge with parallel links, make sure that both paths are assigned to the same port. If you use separate ports, the Spanning Tree algorithm considers each port to be a separate network. As a result, one port is put into blocking state. You can use parallel links on different ports as a backup. If the link in the forwarding state fails, the second link moves from the blocking state to the forwarding state.

23.2.2

Spanning Tree Protocol over PPP Configuration

When you connect two bridges over a PPP serial link, both bridges must operate in the same Spanning Tree domain. The EOS software supports the following configuration of the Spanning Tree protocol over PPP:

- Transparent bridge to transparent bridge (TB-TB)

The following configurations are not supported:

- Transparent bridge to source route transparent (TB-SRT)
- Source route to source route (SR-SR)
- Source route transparent to source route transparent (SRT-SRT)
- Source route to transparent bridge (SR-TB)
- Source route to source route transparent (SR-SRT)



NOTE: If you connect bridges in the unsupported configurations, the separate Spanning Tree domains are combined into a single domain.

To configure ports for transparent bridging, you must set the `TransparentBridge` parameter in the `BRIDGE` Service.

23.2.3

Spanning Tree Parameters Modification

The Spanning Tree Protocol (STP) Service controls parameters used by the Spanning Tree algorithm (for example, the priority field in the bridge identifier) to influence the final network configuration. For more information on setting STP parameters, see “STP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

23.2.4

Topology Reconfiguration

The Spanning Tree algorithm reconfigures the network topology when bridges are added or removed, or when the network manager changes the parameters.

Whenever a designated bridge for a LAN detects a topology change, it sends out a topology change notification Bridge Protocol Data Unit (BPDU) through its root port. This information is eventually relayed to the root bridge. The root bridge then sets the topology change flag in its CBPDU so that the information is broadcast to all bridges. It transmits this CBPDU for a fixed amount of time to ensure that all bridges are informed of the topology change.

If a port is changed from blocking state to forwarding state as a result of the topology change, the algorithm ensures that it propagates the topology information to all ports before that port starts forwarding data. This prevents temporary data loops.

If a bridge does not receive packets from an address within a fixed period of time, it removes that address from its routing table. After reconfiguration, the bridge removes these addresses faster to ensure that each active port still forwards packets to the right network after a topology change.

23.3

Load-Sharing Algorithm

When multiple paths are assigned to a port on an EOS device, a load-sharing algorithm is used. The load-sharing algorithm selects the highest bandwidth link as the primary link. Any outgoing data is transmitted through this link until a certain threshold (defined within software limits for that bandwidth) is reached. When the threshold is reached, packets are forwarded on the next highest bandwidth link. If the number of bytes queued on the primary link falls below the threshold, outgoing packets revert to the primary link.

23.4

Routing Tables

A bridge forwards packets according to information in the routing table. Each entry in this table lists an address, the network on which the station with that address can be found, and an indication of elapsed time since a packet was received from that node. For an interpretation of the routing table, see “BRIDGE Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Routing table entries:

Learned (dynamic) entries

Entries that the bridge learns from the network. The learned entries are subject to dynamic changes or deletion whenever the `-BRIDGE CONTrol` parameter is set to `Aging` and `LEarn`.

User-assigned (static or permanent) entries

Entries assigned by entering `ADD -BRIDGE ROUTe`. The user-assigned entries can be changed or deleted manually only through the `ADD` or `DELeTe` commands.

You can access the routing table of transparent bridges by using:

```
SHoW -BRIDGE AllRoutes
```

You can configure the size of the routing table on the transparent bridge by using the `-BRIDGE RouteTableSize` parameter. For complete information on parameters, see “BRIDGE Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.


23.5

Learning and Filtering

This section describes how a bridge learns the network configuration and adapts to the addition or removal of stations on the attached network segment in order to perform standard filtering. For information on Motorola Solutions mnemonic filtering and related filtering processes such as logging, sequencing, and packet prioritization, see [Mnemonic Filtering Configuration on page 433](#). For complete explanations of packet filtering parameters, see "Filter Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

The following figure shows two networks interconnected by a bridge. After the bridge receives a packet, it decides whether to forward it to the other network or discard it. To help make this decision, the bridge determines to which network the destination of the packet belongs.

Figure 128: Bridge Learning



When a bridge is operating, it receives packets from all attached networks. By looking at the source address of packets, the bridge learns the addresses of stations on each network and stores them in its routing table. For example, when the bridge receives a packet from network 1 with the address for host A as the source address, it learns that host A is on network 1. In the same way, it also learns that hosts B and C are on network 2.


If a packet is destined for the network where it originated, the bridge discards it. This is called standard filtering. For example, if the bridge receives a packet on network 2 from host C that is addressed to host B, and it determines from the learned entries in its routing table that host B is on the same network as host C, then it discards the packet.

In addition, the bridge uses the learned network configurations to forward packets destined for another network. For example, if the bridge receives a packet on network 2 from host C addressed to host A, it determines that host A is on another attached network, and forwards the packet to that network.

If the bridge receives a packet from a host on a network that is not learned yet, the bridge forwards the packet to all ports except the port on which the packet was received.

The bridge also can learn that a station was removed from one of its attached networks. For example, in the following figure, host C was moved from network 2 to network 1. The bridge no longer receives packets on network 2 with host C as the source address. The bridge record of the location of host C is no longer updated and is removed (aged) from the routing table. With host C attached to network 1, the bridge receives packets from network 1 with the address of host C as the source address, and learns that network 1 now includes host C.

Figure 129: Network Configuration After Host C is Moved



23.6

Configuring Basic Bridging

You can set up your transparent bridge to forward packets over the following types of wide area networks:

- Frame Relay
- PPP
- ISDN

Bridging over Frame Relay is supported over fully meshed, partially meshed, and nonmeshed topologies. If you plan to bridge Frame Relay over a partially meshed or nonmeshed topology, you must create a virtual port for each remote network that is attached to a Frame Relay cloud. For configuration information, including a discussion of fully meshed, partially meshed, and nonmeshed topologies and virtual ports, see [Wide Area Networking Configuration with Frame Relay on page 308](#). For information on the number of virtual ports supported on each EOS platform, see [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).


If you configure bridging over Frame Relay in a meshed topology with multiple data link connection identifier (DLCI) neighbors, the transmission of unknown unicast addresses and multicast (including broadcast) frames is processed separately from packets sent to known addresses. As a result, some frames may arrive at the destination nodes out of order.

To configure bridging over PPP, see [Wide Area Networking Configuration with PPP on page 341](#). For information about using the Spanning Tree Protocol over PPP, see [Spanning Tree Algorithm with Wide Area Bridges on page 416](#).

Transparent bridging is supported on Ethernet, Point-to-Point Protocol (PPP), Frame Relay, and Integrated Services Digital Network (ISDN).

In the following figure, the EOS device performs IP routing and bridging of other protocols between networks 1 and 2. If station 1 wants to route a packet to station 2 via the EOS device, the packet should be addressed to the EOS device. If for some reason an IP packet is unicast to station 2, the `FireWall` value ensures that the EOS device discards this packet.

Figure 130: Effects of FireWall on Bridging and Routing



If a device receives packets of a protocol type that was not configured on it, it forwards the packets as if it were a bridge.

 **NOTE:** Broadcast packets are not forwarded for any routed protocol, even if the `-BRIDGE CONTROL` parameter is set to `NoFireWall`. This procedure reduces unnecessary network traffic.

Prerequisites:

Ensure that:

- Ports and paths are set up according to [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- You are logged on to the system with Network Manager privilege.

Procedure:

1. To set a transparent bridge, enable bridging on the entire device by entering:

```
SETDefault -BRIDGE CONTROL = Bridge
```

If you do not want transparent bridging enabled on all ports, you can disable it on an individual port basis. For more information, see [Configuring Per-Port Transparent Bridging on page 425](#).

2. Configure a device for bridging and routing by performing the following actions:
 - a. Prepare your device for bridging.
 - b. Prepare your device for IP routing. See [IP Routing Configuration on page 119](#).
 - c. Decide whether `FireWall` should be configured in the `CONTROL` parameter of the `BRIDGE` Service.

`FireWall` causes the device to discard all packets of a configured protocol that are addressed to destinations other than the device. There is a performance cost when `FireWall` is enabled, because every bridged packet must be checked. `NoFireWall` bypasses this check, resulting in better bridging performance.

`FireWall` ensures that protocols configured for routing are never bridged by checking a type field in every packet to determine whether the packet should be routed or bridged. This level of checking slows down the performance of the device.

`NoFireWall` causes the device to skip the type field (except for broadcast packets, which are always checked). The device routes only data that is sent directly to the device at the MAC layer. All other packets are bridged. `NoFireWall` improves performance, but may forward incorrectly

addressed packets. For example, if the IP protocol is configured for routing, and an IP packet is received with an address different from the device address, that IP packet is bridged.

NoFireWall is selected by default. To change the CONTROL parameter, you must enter one of the following commands:

- SETDefault -BRidge CONTROL = FireWall
- SETDefault -BRidge CONTROL = NoFireWall

23.7

Verifying the Bridge Configuration

Table 48: SHow BRIDGE CONFIGURATION Display Example

Current Configuration Values						
CONTROL = (Aging,Bridge,NoFireWall,F0rward,LEarn)						
AgeTime = 300						
	Name	State	Type	Status	SRcSec	DStSec
	----	----	----	----	----	----
Port 1	Port_1	Forwarding	Ethernet	Reachable	None	None
Path 1	Path_1			Up (Wed Dec 31 16:00)		
Port 2	Port_2	Forwarding	Ethernet	Reachable	None	None
Path 2	Path_2			Up (Wed Dec 31 16:00)		
Port 3	Port_3	Disabled	Remote	Unreachable	Down	None
Path 3	Path_3		64 Kbps	Down (Wed Dec 31 16:00)		
Port 4	Port_4	Disabled	Remote	Unreachable	Down	None
Path 4	Path_4		64 Kbps	Down (Wed Dec 31 16:00)		

Procedure:

1. Verify the values assigned to the -BRIDGE CONTROL parameter by entering:

```
SHow -BRIDGE CONTROL
```

The current values of the -BRIDGE CONTROL parameter are displayed. Usually the following, default values are displayed: F0rward, LEarn, Aging, and NoFireWall. If the setting of any of these values is changed, the bridge performs one or more of the following processes:

- Not forward packets.
- Use only user-defined routes in its routing table to forward packets.
- Not check for addresses of nodes that appear to be dormant.
- Discard unicast packets of a protocol that is routed (except for unicast packets to the bridge itself).

2. Verify that transparent bridging is enabled on the appropriate ports by entering:

```
SHoW -BRidge TransparenTBRidge
```

The current values of the `-BRidge TransparenTBRidge` parameter are displayed. Make sure that `TransparenTBRidge` is selected on the appropriate ports.

3. Verify that the Spanning Tree Protocol is enabled by entering:

```
SHoW -STP CONTroL
```

The current values of the `-STP CONTroL` parameter are displayed. Make sure that the default value `Enabled` is selected, to ensure that the bridge participates in the Spanning Tree network configuration.

4. Check the configuration of the transparent bridge and the status of each port and path by entering:

```
SHoW -BRidge CONFIguration
```

The first two lines of the displayed table show the current values of the `-BRidge CONTroL` and `-BRidge AgeTime` parameters. The remaining lines show the status of each port and its associated paths, including the name (if any) of the type of line, the time at which the paths status last changed, and the source and destination security.

5. If the display indicates that a port or a path is down, perform the following actions:

- a. Check the configuration of each port by entering:

```
SHoW -PORT CONFIguration
```

- b. Check the configuration of each path by entering:

```
SHoW -PATH CONFIguration
```

6. Test the bridge by sending packets across it by making a connection from a device on one attached network to a host on another attached network.

For example, make a connection from a device on one attached network to a host on another attached network.

If you can successfully make a connection, the bridge is ready for normal operation.

If you cannot make a connection, see [Troubleshooting the Bridge Configuration on page 424](#).

23.8

Getting Statistics

You can collect statistics for a specified time period by using the `-SYS SampleTime` and `-SYS STATistics` parameters. For more information, see “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

For information on interpreting statistics displays, see [Statistics Displays on page 967](#).

Procedure:

1. Display statistics for bridged packets by entering:

```
SHoW -SYS STATistics -BRidge
```

2. Display statistics for all ports by entering:

```
SHoW -SYS STATistics -PORT
```

3. Display statistics for all paths by entering:

```
SHoW -SYS STATistics -PATH
```

4. If the display indicates that there are errors on the attached network (for example, cyclic redundancy check errors), ensure that:

- The transceiver cable is properly attached to the transceiver.

- The transceiver is properly attached to the network cable.
 - The network is properly terminated.
5. If the errors occur on a serial line, check the following items:
- Cable attachments
 - Channel service units (CSUs) and digital service units (DSUs)
 - Modems on each end of the serial line
- If the line is a leased line, request help from the company that leases you the line (for example, the telephone company).

23.9

Troubleshooting the Bridge Configuration

Procedure:

1. If one or more devices cannot communicate across the bridge, determine whether the filtering feature is enabled and if so, what types of filters were implemented, by performing the following actions:
 - a. Determine whether the filtering is enabled by entering:

```
SHoW -FIlter CONTroL
```
 - b. If the filtering is enabled, review the filters by entering the following commands:

```
SHoW -FIlter MASK
```

```
SHoW -FIlter POLIcy
```

A filter defined incorrectly can cause packets destined for certain addresses to be discarded.
 - c. If filtering is enabled, set up a counter to record the number of packets of a particular type that are forwarded by the bridge, by using the `-FIlter POLIcy` parameter.

For more information on bridge filtering, see [Mnemonic Filtering Configuration on page 433](#).
2. Determine if source and destination explicit forwarding were implemented by entering the following commands:

```
SHoW [!<port> | !*] -BRidge SRcSecurity
```

```
SHoW [!<port> | !*] -BRidge DStSecurity
```
3. Check the routing table by entering:

```
SHoW -BRidge AllRoutes
```

The address of the affected station should appear in the routing table, followed by the correct destination network number (port). If the address does not appear, make sure that the `-BRidge CONTroL` parameter settings include `LEarn` and `FORward`. If necessary, enter the `ADD -BRidge ROUte` command to add the address of the affected station.
4. Display bridge configuration information, check the status of each path and verify if each path is assigned to the appropriate network by entering:

```
SHoW -BRidge CONFIguration
```
5. Check the physical attachments of any network not listed as `REACHABLE` or any path not listed as `UP`.
6. Verify that the path is enabled by entering the following commands:

```
SHoWDefault -PORT CONFIguration
```

```
SHoWDefault -PATH CONFIguration
```
7. If the display indicates that a port or a path is down, perform the following actions:

- a. Check the configuration of each port by entering:
`SHow -PORT CONFIguration`
- b. Check the configuration of each path by entering:
`SHow -PATH CONFIguration`
8. Check for other activity on the bridge.
If there is no other activity, check the physical attachments of the bridge to its networks, including boards, back panel connectors, and transceiver or modem connectors. For lines to wide area bridges, check the CSU/DSU or modem and its configuration.
9. If a large number of errors occurs on the serial line of a wide area bridge to a remote network, check the physical lines.
For a detailed account of errors on a given path, you must enter:
`SHow -SYS STATistics -PATH`
It is possible to set some of the statistics to zero by entering:
`FLush -SYS STATistics [-<service>]`
10. If a pair of devices cannot communicate across the bridge, check to see whether another pair can communicate across it.
If the second pair communicates, the first pair of devices is the problem and not the bridge.
To determine whether a pair of bridges can communicate with each other, initiate the data link test using the `DLTest` command. This test allows the bridges to exchange test packets and displays the resulting statistics. For more information on the `DLTest` command, see "Commands" in the *Enterprise OS Software Version 18.2 Reference Guide*.
11. If possible, replace any bridge you suspect has a problem with another bridge or a repeater.
If the problem persists, the bridge is not the cause.

23.10

Bridge Customization

This section discusses only bridge customization with the `DStSecurity`, `FunctionalAddr`, `MultiCastAddr`, `ROUte`, `SRcSecurity`, and `TransparentBRidge` parameters. For information on other `BRidge Service` and `STP Service` parameters, see "BRidge Service Parameters" and "STP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

If you configure a logical network (port group), you can customize it by configuring global bridge properties, such as `AgeTime`, or individual properties of member ports, such as the ones described in this section. Configuring properties on the group port has no effect.

23.10.1

Configuring Per-Port Transparent Bridging

In addition to enabling transparent bridging on all ports (by setting the `-BRidge CONTRol` parameter to `Bridge`), you can enable or disable transparent bridging on specified ports.

For complete information on the `-BRidge TransparentBRidge` parameter, see "BRidge Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Procedure:

- To enable transparent bridging on specified ports only, enter:
`SETDefault !<port> -BRidge TransparentBRidge = TransparentBRidge`

- To disable transparent bridging, enter:

```
SETDefault !<port> -BRidge TransparentBRidge = NoTransparentBRidge
```

23.10.2

Adding or Deleting Static Entries



CAUTION: When you change the owner for any WAN port, you must delete all static routes that were configured for the previous owner and WAN type. You must use the `DELeTe -BRidge ROUte` command to delete these routes. Failing to delete the routes can cause a crash (fatal error) in the EOS software.

For more information on the `-BRidge ROUte` command, see “BRidge Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

For more information on routing tables, see [Routing Tables on page 418](#).

Procedure:

- To add a static entry in the routing table used by your bridge, enter:

```
ADD !<port> -BRidge ROUte <address>
```

- To delete a routing table entry, enter:

```
DELeTe [!<port>] -BRidge ROUte <address>
```

23.10.3

Bridge Security Configuration

You can use bridge security features to select certain stations whose packets are forwarded or blocked depending on their source or destination address. These features are applied only to packets traveling from one port of the bridge to another port. Packets addressed to the bridge are not affected. The security features include:

- Source explicit forwarding.
- Source explicit blocking.
- Destination explicit forwarding.
- Destination explicit blocking.
- Combined source and destination security.



CAUTION: Before you use the `SRcSecurity` and `DStSecurity` parameters, read the descriptions and examples in this chapter and see “BRidge Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. The `SRcSecurity` and `DStSecurity` parameters can affect bridge performance. Incorrect use of these parameters can cause the bridge to discard packets that you want to forward or to forward packets that you want to discard.

For information on restricting packet movement based on packet contents, see [Custom Filters Configuration on page 431](#).

23.10.3.1

Configuring Source Explicit Forwarding

The source explicit forwarding feature allows you to forward packets from specific source addresses, on a per-port basis, in conjunction with a routing table. To forward packets using the source explicit forwarding feature, you must enable forwarding on the port where the packet enters the bridge by using:


```
SETDefault !<port> -BRidge SRcSecurity = Fwd
```

For a packet to be forwarded, its source address must be a static entry in the routing table on the port where the packet enters the bridge. To add or delete static entries, use the `ADD` or `DELEte -BRidge ROute <address>` syntax. All other packets are discarded.

 **NOTE:** Some packets that meet forwarding conditions cannot be forwarded, because they are blocked by other constraints such as filtering or destination explicit blocking.

The following figure shows a bridge connecting two Ethernet networks, network A and network B. All stations on network B can communicate with all stations on network A. However, you can restrict packet forwarding so that only stations 1 and 2 on network A can communicate with the stations on network B. If stations 3 to 20 on network A send packets to any of the stations on network B, the packets are discarded.

Figure 131: Source Explicit Forwarding



Procedure:

1. Set the `-BRidge SRcSecurity` parameter to `Fwd` on the port where the packet enters the bridge.

Step example:

To set this parameter to `Fwd` on port 1 of the bridge, enter:

```
SETDefault !1 -BRidge SRcSecurity = Fwd
```

2. Add a static entry in the routing table for each source station from which you want packets to be forwarded. The static entry must be generated on the port where the packet enters the bridge.

Step example:

To add a static entry for station 1, enter:

```
ADD !1 -BRidge ROute %080002001234
```

23.10.3.2

Configuring Source Explicit Blocking

The source explicit blocking feature allows you to discard packets from specific source addresses on a per-port basis in conjunction with a routing table. The blocking feature is the reverse of forwarding. You must choose whichever feature (forwarding or blocking) allows you to enter fewer source addresses in the routing table.

 **NOTE:** On any port, only one of the two features (forwarding or blocking) can be turned on at a time.


To block packets using the source explicit blocking feature, you must enable the block feature on the port where the packet enters the bridge by using:

```
SETDefault !<port> -BRidge SRcSecurity = Blk
```

For a packet to be blocked, its source address must be a static entry in the routing table on the port where the packet enters the bridge. Add or delete static entries using the `ADD` or `DELEte -BRidge ROute <address>` syntax. All other packets are forwarded (subject to other constraints).

The following figure shows a bridge connecting two Ethernet networks, network A and network B. All stations on network B must be able to communicate with all stations on network A. However, you can restrict packet forwarding so that if stations 1 and 2 on network A send packets to the stations on network B, the packets are discarded. If stations 3 to 20 on network A send packets to the stations on network B, they are forwarded.

Figure 132: Source Explicit Blocking



Procedure:

1. Set the `-BRidge SRcSecurity` parameter to `Blk` on the port where the packet enters the bridge.

Step example:

To set this parameter to `Blk` on port 1, enter:

```
SETDefault !1 -BRidge SRcSecurity = Blk
```

2. Add a static entry in the routing table for each source station from which you want packets to be blocked. The static entry must be generated on the port where the packet enters the bridge.

Step example:

To add a static entry for station 1, enter:

```
ADD !1 -BRidge ROute %080002001234
```

If you used source explicit forwarding in this example, you would have to enter the addresses of stations 3 to 20 manually in the routing table.

23.10.3.3

Configuring Destination Explicit Forwarding

The destination explicit forwarding feature allows you to forward packets to specific destination addresses, on a per-port basis, in conjunction with a routing table.


To forward packets using this feature, you must enable the forward feature on the port where the packet enters the bridge by using:

```
SETDefault !<port> -BRidge DStSecurity = Fwd
```

For a packet to be forwarded, its destination address must be a static entry in the routing table on the port where the packet exits the bridge. You add or delete static entries using the `ADD` or `DELeTe -BRidge ROute <address>` syntax. All other packets are discarded.

The following figure shows a bridge connecting two Ethernet networks in a company, network A and network B. All stations on network B must be able to communicate with all stations on network A. However, you can set `DEF` so that network A stations can send packets only to two stations on network B: station 21 and station 22. If stations on network A send packets to stations other than 21 or 22 on network B, the packets are discarded.

Figure 133: Destination Explicit Forwarding



Procedure:

1. Set the `-BRIDGE DStSecurity` parameter to `Fwd` on the port where the packet enters the bridge.

Step example:

To set this parameter to `Fwd` on port 1 of the bridge, enter:

```
SETDefault !1 -BRIDGE DStSecurity = Fwd
```

2. Add a static entry in the routing table for each destination station to which you want packets to be forwarded.

The static entry must be generated on the port where the packet exits the bridge.



NOTE: Some packets that meet forwarding conditions may not be forwarded, because they are blocked by other constraints such as filtering or source explicit blocking.

Step example:

To add a static entry for station 21, enter:

```
ADD !2 -BRIDGE ROUTe %080002001234
```

23.10.3.4

Configuring Destination Explicit Blocking

The destination explicit blocking feature allows you to discard packets sent to specific destination addresses, on a per-port basis, in conjunction with a routing table. The blocking feature is the reverse of forwarding. You must choose whichever feature (forwarding or blocking) allows you to enter fewer source addresses in the routing table.



NOTE: On any port, only one of the two features (forwarding or blocking) can be turned on at a time.


To block packets using this feature, you must enable the blocking feature on the port where the packet enters the bridge by using:

```
SETDefault !<port> -BRIDGE DStSecurity = Blk
```

The destination address must be a static entry in the routing table on the port where the packet exits the bridge. Add or delete static entries using the `<ADD>` or `DELeTe -BRIDGE ROUTe <address>` syntax. All other packets are forwarded (subject to other constraints).

The following figure shows a bridge connecting two Ethernet networks, network A and network B. Any station on network B must be able to communicate with any station on network A. You can set Destination Explicit Blocking so that all stations on network A can communicate with the stations on network B except for stations 21 and 22. If stations on network A send packets to stations 21 or 22 on network B, the packets are discarded.

Figure 134: Destination Explicit Blocking



Procedure:

1. Set the `-BRIDGE DStSecurity` parameter to `Blk` on the port where the packet enters the bridge.

Step example:

To set this parameter to `Blk` on port 1, enter:

```
SETDefault !1 -BRIDGE DStSecurity = Blk
```

2. Add a static entry in the routing table for each destination station to which you do not want packets to be forwarded. The static entry must be generated on the port where the packet exits the bridge.

Step example:

To add a static entry for station 21, enter:

```
ADD !2 -BRIDGE ROUTE %080002001234
```

In this example, you can set the `DStSecurity` parameter to `Fwd`, but you must enter the addresses of stations 23 to 40 manually in the routing table, requiring more work for the network manager.

3. To discard all packets destined for a particular address, enter:

```
ADD -BRIDGE ROUTE <address> Off
```

23.10.3.5

Combined Source and Destiny Security

You can build a complex bridge security system by combining the source and destination explicit forwarding and blocking features.

You can turn on the `-BRIDGE SRcSecurity` and `-BRIDGE DStSecurity` commands at the same port. However, if both forwarding and blocking conditions are met, the blocking condition takes precedence and the packet is discarded.


For example, suppose that the `SRcSecurity` parameter is set to `Fwd` and the `DStSecurity` parameter is set to `Blk`. A packet originating from an address that is not a static entry does not meet the forwarding condition. This packet is discarded regardless of its destination address. A packet originating from an address that is a static entry does meet the forwarding condition. However, if the destination address of this packet is also a static entry, then this packet meets the blocking condition and the packet is discarded.



NOTE: Packets traveling on a network combining these features reach their destination only if all forwarding conditions are met and no blocking conditions are met.

The following figure shows a bridge connecting two Ethernet networks: network A and network B.

Figure 135: Combining Bridge Security Features



The network manager set up the following complex system of bridge security features that restrict communication between the two networks:

- Stations 1 and 2 on network A **cannot** communicate with stations 21 and 22 on network B, because DStSecurity on port 1 is set to Blk and stations 21 and 22 are static entries on port 2. Although SRcSecurity is set to Fwd, and stations 1 and 2 on network A are static entries on port 1, the setting is ignored because blocking takes precedence over forwarding.
- Stations 1 and 2 **can** communicate with stations 23 to 40 on network B, because:
 - SRcSecurity is set to Fwd on port 1.
 - Stations 1 and 2 are static entries on port 1.
 - DStSecurity is set to Blk on port 1, but stations 23 to 40 are not static entries on port 2.
- Stations 3 to 20 on network A **cannot** communicate with any station on network B because:
 - SRcSecurity is set to Fwd on port 1.
 - None of the stations 3 to 20 are listed as static entries on port 1.
- Stations 21 and 22 on network B **cannot** communicate with any station on network A because:
 - SRcSecurity is set to Blk on port 2.
 - Stations 21 and 22 are static entries on port 2.
- Stations 23 to 40 on network B **can** communicate with stations 1 and 2 on network A because:
 - SRcSecurity is set to Blk on port 2 but stations 23 to 40 are not static entries on port 2.
 - DStSecurity is set to Fwd on port 2 and stations 1 and 2 on network A are static entries on port 1.

23.10.3.6

Custom Filters Configuration

You can use the `-Filter MASK` and `-Filter POLICY` parameters to define a custom filter so that packets meeting the criteria specified in the filter are forwarded or discarded. You can also restrict forwarding of packets by protocol type or other packet contents.

For information on configuring filters and policies, see [Mnemonic Filtering Configuration on page 433](#).

23.11

Optimizing Bridge Performance



NOTE: Avoid configuring source and destination security features or filters.

Procedure:

1. Disable the firewall feature in mixed bridging and routing environments by entering:
2. After the bridge learned addresses, disable dynamic learning (if you do not need it) by entering:
3. If you do not need dynamic learning, increase the aging time for which entries remain in the routing table by entering:

```
SETDefault -BRidge CONTrol = NoFireWall
```

```
SETDefault -BRidge CONTrol = NoLEarn
```

```
SETDefault -BRidge AgeTime = <seconds> (10-1000000)
```

The default setting for this parameter is 300 seconds.

Chapter 24

Mnemonic Filtering Configuration

Filtering is an operation that determines whether specified packets are forwarded or discarded by an EOS device. The Filter Service controls these and other capabilities through the Filter POLicy parameter action options: Count, Discard, DodDiscard, Forward, PROToColRsrv <tag>, Sequence, Prioritization, and Trace. These action options are described in [Filtering Process on page 436](#).

You must configure prioritization separately. For complete information on the prioritization allocation, see [Multiprotocol Data Prioritization on page 398](#).

By using filtering in a bridged or IP-routed environment, you can:

- Achieve security and bandwidth protection by isolating specific segments of the network.
- Monitor network traffic by gathering statistics.
- Adjust the performance of your network to fit the traffic flow.
- Sequence packets so that they are received in the order they were sent.
- Reserve bandwidth for particular protocols, so that large-bandwidth user applications, such as file transfer and mail, share link capacity with lower bandwidth users such as interactive sessions and transaction-oriented applications.

Enterprise OS software includes the use of mnemonics and built-in masks for specific protocols in the configuration of filters. Through the use of built-in mnemonics, you can also create user-defined masks to meet more specialized needs.

For more information on the parameters used in creating filters and masks, see “Filter Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

24.1

Configuring Filters

When you configure filters, you perform the selection, qualification, and action steps using the ADD and DELete commands. The MASK parameter specifies the selection criteria, and the POLicy parameter specifies the context by qualifying the selection and associating the action.

You can use the same selection criteria (masks) in different contexts (policies). You can also combine different selection criteria while qualifying them and specifying the action.

Procedure:

1. Determine whether or not a built-in mask can be used.

You must identify the type of packet to be filtered. [Table 51: User-Defined Bridge Masks and DataLink Locations on page 440](#) identifies all types of packets for which built-in masks can be used.

2. If you can use a built-in mask, perform the following actions:

- a. Define the policy by using the ADD POLicy command.

Step example:

Suppose you want to discard all Internet Protocol (IP) multicast packets at port 2. To define the policy, enter:

```
ADD -Filter POLicy NoIPMC Discard IP MC AT !2
```

The message that policy NoIPMC is added appears on the screen.

- b. Continue using the ADD MASK and ADD POLicy commands for all types of packets to be filtered.

- c. Specify the action for packets that do not match any policy by setting the `DefaultAction` parameter:

```
SETDefault -Filter DefaultAction = [Forward | Discard]
```

When `DefaultAction` is set to `Discard`, all packets not matching a policy are discarded. All packets matching the policy are handled according to the policy.

- d. Enable filtering by entering:

```
SETDefault -Filter CONTROL = Enabled
```

3. If you cannot use a built-in mask, perform the following actions:

- a. Define your own mask by using the `ADD MASK` command.

For the information on the built-in mnemonics that you can use to construct user-defined masks for the BBridge services, see [Table 51: User-Defined Bridge Masks and DataLink Locations](#) on page 440.

Step example:

Suppose you want to define a pattern for a mask that is not built-in. For example, you may want to discard all packets that are longer than 512 bytes. Because you cannot represent this pattern as a built-in mask, you must enter the following command and the built-in mnemonics (`dl.length`) to define the mask:

```
ADD -Filter MASK longpkts dl.length>%0200
```

The message that `LONGPKTS` mask is added appears on the screen.



NOTE: The expected value must be an even number of digits.

- b. Define the policy by using the `ADD POLICY` command.

Step example:

Suppose you still want to discard all packets that are longer than 512 bytes at port 2. You defined the mask. To define the policy, enter:

```
ADD -Filter POLICY toolong Discard longpkts AT !2
```

The message that `TOOLONG` policy is added appears on the screen.

- c. Specify the action for packets that do not match any policy by setting the `DefaultAction` parameter:

```
SETDefault -Filter DefaultAction = [Forward | Discard]
```

When `DefaultAction` is set to `Discard`, all packets not matching a policy are discarded. All packets matching the policy are handled according to the policy.

- d. Enable filtering by entering:

```
SETDefault -Filter CONTROL = Enabled
```

24.2

Configuring a Filter Between Station Groups


To configure a filter for a group of logically related stations, you must use the `StationGroup` parameter.

When you use the `StationGroup` parameter, you must complete the following tasks:

- Assign a set of station addresses for easy reference.
- Give the group a name.
- Create a mask by referencing the station group name.

The following figure illustrates specifying a policy based on station groups.

Figure 136: Network Showing Station Groups



Stations belong to group A or to group B. Group A has stations on network 1 and network 2. Group B has stations on network 3 and network 4. After grouping the stations, you can create a policy that would, for example, prohibit a certain type of traffic between group A and group B.

Assuming that the media access control (MAC) address for station A1 is %0800020000a1 and the MAC address for station A2 is %0800020000a2, you can use the following procedure to configure a filter between group A and group B.

For more information on the `StationGroup` parameter, see “Filter Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Procedure:

1. Define a station group and add the MAC addresses of the stations belonging to the defined group.

Step example:

To create group A and group B and add appropriate addresses to them, enter:

```
ADD -Filter StationGroup group_a %0800020000a1
ADD -Filter StationGroup group_a %0800020000a2
ADD -Filter StationGroup group_a %0800020000a3
ADD -Filter StationGroup group_a %0800020000a4
ADD -Filter StationGroup group_a %0800020000a5
ADD -Filter StationGroup group_a %0800020000a6
ADD -Filter StationGroup group_b %0800020000b1
ADD -Filter StationGroup group_b %0800020000b2
ADD -Filter StationGroup group_b %0800020000b3
ADD -Filter StationGroup group_b %0800020000b4
ADD -Filter StationGroup group_b %0800020000b5
ADD -Filter StationGroup group_b %0800020000b6
```

2. Define masks by using the station groups.

Step example:

To create masks, enter:

```
ADD -Filter MASK from_group_a DataLink.SrcAddr = group_a
ADD -Filter MASK from_group_b DataLink.SrcAddr = group_b
ADD -Filter MASK to_group_a DataLink.DstAddr = group_a
ADD -Filter MASK to_group_b DataLink.DstAddr = group_b
```

3. Define policies by using the previously defined masks.

Step example:

To create policies, enter:

```
ADD -FI POLicy block_from_a Discard from_group_a, to_group_b, IP
ADD -FI POLicy block_from_b Discard from_group_b, to_group_a, IP
```

24.3

Filter Parameters

The following table lists and briefly describes the `Filter Service` parameters. For detailed descriptions of these parameters, see “Filter Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Table 49: Filter Service Parameters

Parameters	Descriptions
CONFIguration	Displays the overall configuration of the Filter Service.
CONTRol	Disables and enables the Filter Service. Must be enabled for any filter-related actions to be performed.
DefaultAction	Specifies the action applied to a packet if it does not match any of the policies configured. (If <code>DefaultAction</code> is set to <code>Discard</code> , and there are no forwarding policies defined, no packets are forwarded by the system.)
DIAGnostics	Shows the current decision tree that the system is using, including which <code>MASKs</code> are associated with which <code>POLicies</code> .
MASK	Defines the criteria used to select a packet for special handling.
MNEmonics	Displays all possible options for a location that can be used to construct a user-defined mask.
POLicy	Defines the system context within which the specified masks are applied and the action to be taken. Uses the <code>MASKs</code> that are defined, and applies specific operations to packets that match the <code>MASK</code> conditions of the <code>POLicy</code> .
SElection	Lists services for which the filter function can be invoked (<code>BRidge</code>).
StationGroup	Groups a set of station addresses for easy reference.

24.4

Filtering Process

A filter contains the following two components:

Mask

Defines the qualifications a packet must meet.

Policy

Defines which masks are to be applied and what action is to be taken for the packets that meet the criteria of the mask.

For packets using filters based on either user-defined masks or built-in mnemonic masks, the following Filter Service POLicy parameter action options are available: Count, Discard, DodDiscard, Forward, PROTOcolRsrv <tag>, Sequence, Prioritization, and Trace.

When you use filters with user-defined masks, you must determine location offsets and values to create the mask. Using built-in masks allows you to specify packet selection criteria without determining specific offsets, encapsulation, and frame formats. These built-in masks simplify filtering operations for the bridge and make filtering configurations transferable across interfaces of different types. Most built-in masks are defined for specific protocols. For details on these masks, see [Built-in Bridge Masks on page 439](#).

To support user-defined masks, Enterprise OS software has several built-in mnemonics that you can use to specify location and pattern. The locations and patterns are listed later in this chapter.


The filtering process involves selection, qualification, and action.

Selection

Selection identifies the packets on which filtering is performed. You can select packets for special action by specifying a particular pattern of data at a particular location. You can also specify other, more complicated, selection criteria. You can use the MASK parameter to select the packet.

The following figure illustrates an example of how the MASK parameter is used in the selection process. The location is typically specified as a string of hexadecimal numbers. In the following figure, the use of built-in BRidge mnemonics lets you specify the location at the DataLink layer. The offset for the same field within a packet can vary, depending on the encapsulation or frame format. For more information about the MASK parameter, see "Filter Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Figure 137: Filter Selection Process



Qualification


Qualification specifies the context of the filtering operation – the direction of travel and the ports affected.

After selecting a packet for special action using the MASK parameter, you can specify additional qualifications before the action is taken. For example, using Enterprise OS software, it is possible to select only those broadcast packets that arrive on a specified port, instead of all broadcast packets. You can use the POLicy parameter to specify qualifications for the packet.

Action

After the packet is selected and qualified, a specified action occurs. Use the POLicy parameter in the Filter Service to specify the desired action. The following figure illustrates the qualification and action processes using an example of the POLicy parameter. The subsections following the figure describe the POLicy actions supported by the Filter Service.

Figure 138: Filter Qualification and Action Processes



Policy actions supported by the Filter Service:

Count

When you use the `Count` option, the system counts packets that meet specified criteria.

For example, if you want to count all IP packets forwarded by the bridge before deciding how the bridge should handle them, you must use:

```
ADD -Filter POLicy IP_count Count ip
```

Discard

`Discard` option enables you to discard packets that match specific criteria.

DodDiscard

If you use the `DodDiscard` option for a dial-on-demand (DOD) port when the dial-up path is down, you can ensure that the packet is discarded and does not cause the dial-up path to be raised. If the path is up, the packet is forwarded, but is not considered as user traffic that keeps a dial-up path up.

Forward

Filters can prevent packets meeting certain criteria from being forwarded across the system or forward only those packets meeting specified criteria while blocking all others. When you use the `Forward` option, you forward packets that match specific criteria. For more information on forwarding, see “Filter Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

PROTocolRsrv <tag>

Protocol reservation assigns a specified percentage of bandwidth to designated packets passing through a specified port and meeting specified conditions. The specified conditions can be protocol type, packet length, packets destined for specified address, and more.

You set up protocol reservation with different procedures depending on the packet types configured for protocol reservation. The mnemonic filtering procedure applies to all bridged packets. The IP filtering procedure applies only to IP-routed packets. IP-routed packets are also filtered using the IP firewall feature. For information about the IP firewall feature, see [Internet Firewalls Configuration on page 551](#).

For a detailed description of the protocol reservation procedures for all the packet types, see [Protocol Reservation Configuration on page 383](#).

As part of the mnemonic filtering procedure, you enter the `PROTocolRsrv <tag>` action option to apply protocol reservation to designated packets. The tag name identifies those packets that receive a specified percentage of bandwidth when passing through the specified WAN port and when meeting the mask conditions set up with the Filter Service `POLicy` parameter. Tag the designated packets with the identifying name by entering a name as the `<tag>` value when you enter the `PROTocolRsrv <tag>` action option. The tag name can be any alphanumeric string no longer than 15 characters.

For bridge filtering examples using the `PROTOCOLRsrv <tag>` action option and the `-PORT PROTOCOLRsrv` parameter. For details, see [Bridge Filtering Examples on page 441](#).

Sequence

When the load-balancing algorithm is operating, packets can arrive out of sequence. You can use the `Sequence` option to sequence packets to ensure that they arrive at their destination in the order they were sent.

When operating with two or more parallel lines (including bandwidth-on-demand dial-up lines), local area transport (LAT) and NETBEUI, should be packet-sequenced using the sequence policies. If all of the traffic on the port is sequenced, bandwidth-on-demand is not used for that data. Sequenced traffic is only sent on the primary path.

If you want to sequence and send LAT packets to port 4 in the order they are received, you must use:

```
ADD -Filter POLicy LATorder Sequence LAT TO !4
```

For more information on sequencing and the `POLicy` parameter, see “Filter Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Prioritization (Priority Queueing)

The **Prioritization** option allows you to prioritize different packet types transmitted over wide area networks. You can assign priorities to packets according to their protocol type. Prioritization is a filtering component and must be configured separately. For complete information on data prioritization, see [Multiprotocol Data Prioritization on page 398](#).

For a more information about the `-Filter MASK` and `-Filter POLicy` parameters, see “Filter Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

24.5

Built-in Bridge Masks

Enterprise OS software supports several built-in predefined selection criteria, or masks. All bridge masks are associated with DataLink level as the protocol.

The following table lists the built-in DataLink masks.

If you want to display the built-in masks, you must use:

```
SHOW -Filter MASK BuiltIn
```

Table 50: Built-in Bridge Masks

Built-in Mask	Equivalent	Packet Type
BC	DataLink.DestinationAddr=BroadCast	Bcast
MC	DataLink.DestinationAddr=MultiCast	Mcast
ATALK	DataLink.Protocol=AT	AT
AARP	DataLink.Protocol=AARP	AppleTalkARP
ARP	DataLink.Protocol=ARP	ARP
DLTEST	DataLink.Protocol=DLTest	DLTest
IP	DataLink.Protocol=IP	IP
IPX	DataLink.Protocol=IPX	Novell IPX
LAT	DataLink.Protocol=LAT	LAT
NMIP	DataLink.Protocol=NetMapIP	NetMapIP

Built-in Mask	Equivalent	Packet Type
STP	DataLink.Protocol=STP	Spanning Tree

24.6

User-Defined Bridge Masks

When you use the `ADD MASK` command, you must specify a location. The location is usually expressed as a hexadecimal value representing the offset from the beginning of a packet at which a specified pattern of data is compared to the contents of a packet. The packet is selected if it matches the pattern of data at the specified location.

You also can specify a location in the mnemonic form: `<protocol>. <field>`. This format allows you to use encapsulation-independent relative offsets. You do not need to determine frame formats or specific offsets. All bridge mnemonics are associated with DataLink as `<protocol>`. Different mnemonic values are allowed for the `<field>` and `<match>` locations.

The following table lists valid locations that match the DataLink protocol. You can use these fields to specify an address, instead of specifying the offset of a particular field.

Table 51: User-Defined Bridge Masks and DataLink Locations

Field	Description	Matching Value
DstAddr	Destination Address at DataLink layer	<code><MAC_address></code> ALL <code><StationGroup></code>
SrcAddr	Source Address at DataLink layer	<code><MAC_address></code> <code><StationGroup></code>
Address	Either Destination or Source Address at DataLink layer	<code><MAC_address></code> <code><StationGroup></code>
Protocol	Packet protocol type	<code><numerical_value></code>
LENgth	Frame size, including padding	<code><numerical_value></code>
DSAP	Destination service access point	<code><numerical_value></code>
SSAP	Source service access point	<code><numerical_value></code>
LSAP	Link service access point, destination or source SAP	<code><numerical_value></code>
OUI	Organizationally unique ID	<code>%<hexadecimal_number></code>
LanID	LAN identifier in a source-routed frame	<code><numerical_value></code>
DATA=[%] <code><offset></code> [: [%] <code><length></code>]	Offset from start of DataLink data	<code>%<hexadecimal_number></code> <code><numerical_value></code>
[%] <code><offset></code> [: [%] <code><length></code>]	Offset from start of DataLink header	<code>%<hexadecimal_number></code> <code><numerical_value></code>

If you want to display a list of valid locations supported by the bridge, you must enter the `SHow -Filter MNEmonics` command. ALL is a valid match mnemonic for certain field categories. When you specify ALL, any value in the location is considered to match the criteria. Field mnemonics indicate encapsulation-independent relative offset. The software recognizes the encapsulation and locates the `<field>` at the correct offset.

24.7

Bridge Filtering Examples

This section provides examples of bridge filtering.

Bridge Filtering Example 1

To enable filtering and to stop checking policies after a policy that matches the packet is found, you must use:

```
SETDefault -Filter CONTROL = (Enabled, MatchOne)
```

Bridge Filtering Example 2

Displaying all masks:

To display all masks, you must use:

```
SHow -Filter MASK
```

Displaying built-in masks:

To display all built-in masks, you must use:

```
SHow -Filter MASK BuiltIn
```

Displaying a specific mask:

To display a specific mask, you must use:

```
SHow -Filter <maskname>
```

Bridge Filtering Example 3

Displaying all policies:

To display all policies, you must use:

```
SHow -Filter POLIcy
```

Displaying a specific policy:

To display a specific policy, you must use:

```
SHow -Filter <polIcyname>
```

Bridge Filtering Example 4


To discard all source-routed IP packets, you must use:

```
ADD -Filter POLIcy dissr_ip Discard ip, allrt
```

Bridge Filtering Example 5

This example describes how to discard all IP packets from port 1 to port 2. IP packets are selected for special action. The selection is further qualified by specifying from port 1 to port 2. The action is designated as discard. Because built-in masks are defined for IP packets (see [Built-in Bridge Masks on page 439](#)), you do not need to use the ADD MASK command.

Figure 139: Discarding IP Packets



To define the policy, you must use:

```
ADD -Filter POLicy noip Discard ip FROM !1 TO !2
```

Bridge Filtering Example 6

To check all policies, you must use:

```
SHOW -Filter POLicy
```

Bridge Filtering Example 7

To add a mask that selects packets destined to %080002123456, you must use:

```
ADD -Filter MASK to_atlas dl.dstaddr = %080002123456
```

Bridge Filtering Example 8

To add a mask that selects packets with a value greater than %45 at the first byte of data, you must use:

```
ADD -Filter MASK some_data dl.data+%0>%45
```

After the mask is added, the message that some_data was added appears on the screen.

Bridge Filtering Example 9

To bridge IP traffic among ports 1, 3, 5, and 6, you can use either command A or B below. Command A is preferred, because the built-in mask is encapsulation-independent. Command B forwards IP packets with Ethernet II encapsulation.

Command A:

```
ADD -Filter POLicy ipgroup Forward ip AMONG !1, !3, !5, !6
```

Command B:

```
ADD -Filter MASK ethernet_ip %c = %0800
```

```
ADD -Filter POLicy ipgroup Forward ethernet_ip AMONG !1, !3, !5, !6
```

Bridge Filtering Example 10

To isolate traffic between two groups of networks, you must use:

```
ADD -Filter MASK any %0 | %ff = %ff
```

```
ADD -Filter POLicy wall Discard any BETWEEN !1, !2 AND !3, !4
```

Packets with any value at offset %0 meet the condition of mask any. Any packet received on port 1 or port 2 and sent to port 3 or port 4 is discarded, but packets received on port 1 and sent to port 2 are not discarded.

Similarly, packets received on port 3 and sent to port 4, or packets that are received on port 4 and sent to port 3, are not discarded.

Bridge Filtering Example 11

The following example illustrates the use of the logical OR operator. The following commands filter all packets that contain 500 (hexadecimal) or more bytes by applying the mask 11111111 to the byte at offset 500. If any value is present at that location, the filtering condition is met.

```
ADD -Filter MASK tail %500 | %ff = %ff
```

```
ADD -Filter POLicy drop Discard tail
```

Suppose the value 10110010 is present at offset 500 hexadecimal. When the logical OR operates on this value against the mask 11111111, the result is:

```
10110010 OR 11111111 = 11111111
```

Because the result is the same as the mask, the condition is met.

If no value is present at that location, the result is always false. Packets that contain more than 500 hexadecimal bytes should be blocked.

Applying a logical OR to any value and a mask of 11111111 always has a result of 11111111; if any value is present at byte 500, the condition is met. This means that any packet that contains 500 (hexadecimal) or more bytes is filtered.

Bridge Filtering Example 12

The following example shows the use of one logical operator:

```
ADD -Filter MASK andmask %a&%80 = %80
```

```
ADD -Filter MASK ormask %a | %fe = %fe
```

```
ADD -Filter MASK notmask %a! = %8c
```

```
ADD -Filter POLicy together Discard andmask ormask notmask
```

In this example, all packets that meet the following three conditions are filtered:

Condition 1:

This condition, %A:&%80, is met if the most significant bit of byte A is 1. It applies the logical AND operator to the value found at byte A and the mask 10000000. Suppose the value at byte A is 10111000:

	10111000
AND	10000000
	10000000

Because the result, 10000000, equals the mask, 10000000, the condition is met.

Condition 2:

This condition, %A:%FE, is met if the least significant bit of byte A is 0. It applies the logical OR operator to the value found at byte A and the mask 11111110. Suppose the value at byte A is 10111000:

	10111000
OR	11111110
	11111110

Because the result, 11111110, equals the mask, 11111110, the condition is met.

Condition 3:

This condition, %A:!%8C, is met if byte A of the packet does not equal 8C. It compares the value found at byte A to the mask 10001100. Suppose the value at byte A is 10111000; because 10111000 is not equal to 10001100, this condition is met.

If a packet meets all three of these conditions, it is filtered. The packet used in this example meets all three conditions, because the value at byte A is assumed to be %B8; therefore, it is filtered.

A packet with the value 8F at byte A satisfies conditions 1 and 3, but does not meet condition 2; it is not filtered, but is forwarded to the appropriate destination.

Bridge Filtering Example 13

To add one specific address to the "accounting" station group, you must use:

```
add -Filter StationGroup accounting %080002123456
```

Bridge Filtering Example 14

To discard any traffic destined to the "accounting" station group, you must use:

```
add -Filter MASK to_accounting datalink.dstaddr = accounting
add -Filter POLicy block_account Discard to_accounting
```

Before entering these commands, you must enter the addresses of the stations belonging to the station group "accounting" by using the `ADD -Filter StationGroup` command.

Bridge Filtering Example 15

To delete one specific address from the "accounting" station group, you must use:

```
DElete -Filter StationGroup accounting %080002123456
```

Bridge Filtering Example 16

To delete the "accounting" station group, you must use:

```
DElete -Filter StationGroup accounting
```



NOTE: Before executing this command, you must delete all members of the station group accounting and delete any masks using the station group "accounting".

Bridge Filtering Example 17

To delete all members from the "accounting" station group, you must use:

```
DElete -Filter StationGroup accounting ALL
```

Bridge Filtering Example 18

To show the names of all station groups and the number of addresses in them, you must use:

```
SHow -Filter StationGroup
```

Bridge Filtering Example 19

To change the name of the "bldg_100" station group to the "bldg_200" station group, you must use:

```
CHange -Filter StationGroup bldg_100 bldg_200
```

Bridge Filtering Example 20

This example illustrates how to allow all broadcasts from port 1 to go across a WAN dial-up link on port 4 only if the link is up, and be discarded if the link is down. You could set the WAN port to DOD. You can then add a filter policy with a built-in mask "BC" by using:


```
ADD -Filter POLICY DROPBC DODDISCARD BC FROM !1 TO !4
```

Bridge Filtering Example 21

To create a mnemonic filter by using the `PROTOCOLRsrv <tag>` action option to allot 10 percent of the bandwidth to packets destined for a certain address that are passing through WAN port 3, you can use [Creating a Mnemonic Filter by Using the PROTOCOLRsrv TAG Action Option on page 445](#).

Bridge Filtering Example 22

This example shows how to use the `PROTOCOLRsrv <tag>` action option to reserve a specified percentage of bandwidth for bridged packets of specified lengths bridged outbound through a WAN port.

In this example, in a bridge configured for IP traffic, you want to reserve the following percentages of bandwidth for packets of the following lengths:

- 50% of the bandwidth for packets less than 100 bytes
- 25% of the bandwidth for packets between 100 and 400 bytes
- 20% of the bandwidth for packets greater than 400 bytes
- 5% of the bandwidth is reserved as a default for untagged traffic.

To reserve bandwidth for packets of different lengths, you can use [Reserving Bandwidth For Packets of Different Lengths on page 446](#).

24.7.1

Creating a Mnemonic Filter by Using the PROTOCOLRsrv <TAG> Action Option

You can use this procedure to create a mnemonic filter with the `PROTOCOLRsrv <tag>` action option, to allot 10% of the bandwidth to packets destined for a certain address that are passing through WAN port 3.

Procedure:

1. Add a filter mask with the name `DSTA_Mask` for a destination address of `%0800AABB1111` by entering:

```
ADD -Filter MASK DSTA_MASK DL.DA = %0800AABB1111
```
2. Add a filter policy that assigns the "dstpol" name to the policy, select the "dsta_tag" name tag for the `PROTOCOLRsrv <tag>` action option, and add the "dsta_mask" mask by entering:

```
ADD -Filter POLICY dstpol PROTOCOLRsrv DSTA_TAG DSTA_MASK
```
3. Enable the Filter Service by entering:

```
SETDefault -Filter CONTROL = Enable
```
4. Assign 10% of bandwidth to the `PROTOCOLRsrv <tag>` "dsta_tag" for port 3 by entering:

```
ADD !3 -PORT PROTOCOLRsrv DSTA_TAG 10
```
5. Set `PROTOCOLRsrv` as the `-PORT QueueCONTROL` parameter option for port 3 by entering:

```
SETDefault !3 -PORT QueueCONTROL = PROTOCOLRsrv
```

After you made these entries, any packet forwarded by the system matching the mask criteria is allotted 10% of the bandwidth in accordance with its name tag ("dsta_tag") and bandwidth allocation.

24.7.2

Reserving Bandwidth For Packets of Different Lengths

Procedure:

1. Add a user-defined mask for each packet length condition that must be met, by entering the following commands:

```
ADD -Filter MASK MYMASK1 IP.PACKETLEN <100
ADD -Filter MASK MYMASK2 IP.PACKETLEN 100-400
ADD -Filter MASK MYMASK3 IP.PACKETLEN >400
```

2. Add filter policies to use the filter masks by entering the following commands:

```
ADD -Filter POLIcy POLICY_x PROTOcolRsrv MYTAG_A MYMASK1
ADD -Filter POLIcy POLICY_y PROTOcolRsrv MYTAG_B MYMASK2
ADD -Filter POLIcy POLICY_z PROTOcolRsrv MYTAG_C MYMASK3
```

3. Select BRidging as the type of packet filtering to occur by entering:

```
SETDefault -Filter SElection = BRidging
```

4. Enable the Filter Service by entering:

```
SETDefault -Filter CONTrol = Enable
```

5. Define the percentage of bandwidth to be reserved for each of the policies entered in [step 4](#), and enter name tags that match those policies, by entering the following commands:

```
ADD !3 -PORT PROTOcolRsrv MYTAG_A 50
ADD !3 -PORT PROTOcolRsrv MYTAG_B 25
ADD !3 -PORT PROTOcolRsrv MYTAG_C 20
```

6. Specify the PROTOcolRsrv option for the -PORT Service QueueCONTrol parameter by entering:

```
SETDefault !3 -PORT QueueCONTrol = PROTOcolRsrv
```

After you made these entries, all IP packets of lengths less than 100 bytes going outbound WAN port 3 get 50% of the bandwidth. Any IP packets of a length between 100 and 400 bytes get 25% of the bandwidth, and IP packets of a length greater than 400 bytes get 20% of the bandwidth.

Five percent of the bandwidth is reserved by default for untagged traffic. If the full 100% of bandwidth is allocated by the commands for various filtering conditions, the system normalizes the amount of bandwidth allotted for each condition so that there is always a reserve of 5% for untagged traffic.

Chapter 25

User Datagram Protocol Broadcast Helper Configuration


The User Datagram Protocol (UDP) Broadcast Helper feature allows applications in the Transmission Control Protocol/Internet Protocol (TCP/IP) stack to forward broadcast packets through a gateway to another network segment. The broadcast packets are typically requests from clients for access to servers, which may contain address, configuration, or name information.

When boot servers are located through a router on another network, UDP Broadcast Helper helps BOOTP and DHCP clients locate the server(s) and retrieve address, configuration, and name information. Without the implementation of this feature, broadcast packets are not propagated outside of the same network.

Broadcast packets generally do not traverse a router; however, there are some situations in which it is useful to propagate broadcast packets to other networks.

For example, in the topology shown in the following figure, a client on network 129.1.0.0 may require access to a time server on network 129.2.0.0. Normally broadcast requests from the client on network 129.1.0.0 would not be forwarded to servers on network 129.2.0.0; however, you can configure UDP Broadcast Helper to allow the forwarding of broadcast requests to servers on network 129.2.0.0.

Figure 140: UDP Broadcast Helper Topology Example



UDP applications are identified within a packet by “well-known” port numbers. You can configure the EOS device to allow broadcast packets to well-known port 37, which is the port number mapped to built-in name TIME for the time service, through to network 129.2.0.0.

A common application for UDP Broadcast Helper is related to the Bootstrap Protocol (BOOTP) and the Dynamic Host Configuration Protocol (DHCP). UDP Broadcast Helper for BOOTP and DHCP assists clients with finding their boot servers when the boot servers are located through a router and on another network segment.

Motorola Solutions implements the BOOTP and DHCP relay agents in the EOS software, allowing existing BOOTP clients to interoperate with DHCP servers. The clients and associated servers do not need to reside on the same IP network or subnet, and changes to the client’s initialization software are unnecessary.

For further information about how UDP works with BOOTP and DHCP, see [Relaying BOOTP and DHCP Traffic on page 451](#).

BOOTP and DHCP Protocols

The BOOTP Protocol is built on the client-server model and allows a single BOOTP reply to specify many items needed for a client to boot, including the client IP address, the address of a gateway, and the address of a server.

The DHCP Protocol is an extension of the BOOTP Protocol and is also built on the client-server model. DHCP is specifically designed for servers in large network environments that have nomadic users and complex TCP/IP software configurations.

DHCP not only allows a host to automatically allocate reusable IP addresses and additional configuration parameters needed for client operations, it also allows the client/server host to configure host parameters not directly related to the IP Protocol. This feature allows the host to exchange packets with any other host on the Internet. However, DHCP does *not* register newly configured hosts with the Domain Name System and is *not* used to configure EOS devices.

The Motorola Solutions implementation of UDP Broadcast Helper feature includes the BOOTP and DHCP relay agent, which allows clients and their associated servers not residing on the same IP network or subnetwork to communicate. Without the relay agent, every subnet that has BOOTP and DHCP clients would be required to have a BOOTP and DHCP server.

Although the BOOTP and DHCP Protocols use the same UDP port numbers (67 and 68), they have some important differences:

- DHCP allows IP addresses to be “leased” for a fixed length of time.
Groups of hosts that do not need permanent IP addresses can lease an address from a limited pool of addresses. Also, a host that is only temporarily connected to the network can be assigned an IP address because the addresses can be reused when they are no longer needed by the original host.
- DHCP packet length is longer than BOOTP.
The additional packet length allows a DHCP server to provide the client with all the IP configuration parameters that it needs to operate.
- DHCP is a more complicated protocol than BOOTP.
DHCP has seven message types; BOOTP uses only two. In addition, DHCP requires complex state machines.

25.1

Configuring UDP Broadcast Helper

UDP Broadcast Helper allows you to configure up to 32 UDP ports on your EOS device using the `ADD -UDPHelp ActivePorts` command.

UDP Broadcast Helper supports several names of well-known services. The names of these services are mapped to specific UDP port numbers. (The name-to-UDP port mappings are also referred to as *built-in names*.) You can configure UDP ports using these built-in names.

The following table lists the supported service names, the UDP port numbers to which they are mapped, and the mnemonic name for each name-to-UDP port mapping. The UDP ports and built-in name mappings are reserved; you cannot change or reconfigure them.

Table 52: Supported Service Name-to-UDP Port Mappings


UDP Port Description	UDP Port Number (Decimal)	Mnemonic Name
Daytime	13	DAYTIME
Time	37	TIME
Host name server	42	IEN116
Domain name server	53	DNS

UDP Port Description	UDP Port Number (Decimal)	Mnemonic Name
TACACS – database service	65	TACACS
Bootstrap protocol server	67 ¹	BPSEVER
Trivial file transfer	69	TFTP
HOSTS2 name server	81	HOSTS2
NIC host name server	101	NIC
Simple file transfer protocol	115	SFTP

¹ BOOTP and DHCP use the same UDP port numbers: server port (67 decimal) and client port (68 decimal).

 **NOTE:** A UDP port is part of an entity address and not related to an interface (port) on the device. For this reason, in the command syntax, you do not precede the UDP port with an exclamation point (!).

Figure 141: UDP Broadcast Helper Configuration



For information on the UDPHELP Service parameters, see the “UDPHELP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

For information on the `-IP DefaultTTL` parameter, see “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Prerequisites:

Ensure that:

- You are logged on to the device with Network Manager privilege.
- Ports and paths are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Client and server LANs and WANs are set up. See [IP Routing Configuration on page 119](#).
- You examined your network configuration and determined the devices on which UDP Broadcast Helper should be enabled.
- You determined what services are available through UDP Broadcast Helper. The number of services you want to configure determines the number of UDP ports you must configure.
- For each UDP port you intend to use, you determined which networks or servers should receive related broadcast packets.
- You determined the IP addresses of the networks and servers that should receive broadcast packets.

Procedure:

1. Enable UDP Broadcast Helper by entering:

```
SETDefault -UDPHELP CONTrol = Enable
```

2. If you want to specify a UDP port by UDP port number, add the UDP port to the active ports list by entering:

```
ADD -UDPHELP ActivePorts {<UDP port> | <name>}
```

Step example:

To add UDP port 100, enter:

```
ADD -UDPHELP ActivePorts 100
```

To add a UDP port with the built-in name TIME, enter:

```
ADD -UDPHELP ActivePorts TIME
```

TIME is the name of a service that has a UDP port number mapped to it. In addition to specifying this UDP port by its built-in name, you can also specify this UDP by the port number mapped to this service. You can enter:

```
ADD -UDPHELP ActivePorts 37
```

To add a UDP port with a name you define, for example, AAA, enter:

```
ADD -UDPHELP ActivePorts AAA
```

3. If you want to specify a UDP port by a defined name, perform the following actions:

- a. Map the name to a UDP port number by entering:

```
ADD -UDPHELP Name <name_string> <UDP_port>
```

Step example:

To map the defined name AAA to UDP port number 100, enter:

```
ADD -UDPHELP Name AAA 100
```

- b. Add the UDP port to the active ports list by entering:

```
ADD -UDPHELP ActivePorts {<UDP port> | <name>}
```

4. For each UDP port you added to the active ports list, set up a list of networks and servers that should receive UDP broadcast packets.

After you add a UDP port to the active ports list, the EOS device automatically forwards broadcast packets destined for the UDP port to all interfaces. You do not need to set up a list of networks and servers that should receive UDP broadcast packets. However, Motorola Solutions strongly recommends limiting the networks and server that receive UDP broadcast packets to help prevent broadcast storms and loops.

You can use any of the following syntaxes to add up to 32 addresses to the forward address list:

```
ADD -UDPHELP ForwardAddress <UDP_port_or_name> <IP_address>
```

```
ADD -UDPHELP ForwardAddress <UDP_port_or_name> <IP_address> <subnet_mask>
```

```
ADD -UDPHELP ForwardAddress <UDP_port_or_name> <IP_address> <subnet_mask> [Ones | Zeroes]
```

```
ADD -UDPHELP ForwardAddress <UDP_port_or_name> <list_of_interfaces>
```

Step example:

To add server 1 to a list for UDP port AAA, enter:

```
ADD -UDPHELP ForwardAddress AAA 30.0.0.1
```

The device forwards broadcast packets destined for UDP port AAA to server 1 only.

To add network 2 to a list for UDP port TIME, enter:

```
ADD -UDPHELP ForwardAddress TIME 20.0.0.0
```

The device forwards broadcast packets destined for UDP port TIME to all nodes on network 2.

The device does not rebroadcast packets through Frame Relay interfaces. You must add the IP address of each server to the list of servers that must receive UDP broadcast packets.

5. To limit the reach of a broadcast packet and the potential duration of broadcast storms, specify the default number of seconds that pass before a broadcast packet is discarded by entering:

```
SETDefault -UDPHELP TTLOverride = <seconds>(1-255)
```

Upon receiving a client request packet, the device assigns the packet a time-to-live (TTL) value. The device assigns the lowest TTL value among the following possible sources:

- The TTL value of the incoming request packet minus one
- The TTL value configured by the `-UDPHELP TTLOverride` parameter
- The TTL value configured by the `-IP DefaultTTL` parameter
- If the TTL value configured by the `-UDPHELP TTLOverride` parameter is the lowest, the device forwards the packet with the TTL value configured by this parameter, which overrides the other TTL values.

25.2


Relaying BOOTP and DHCP Traffic

UDP Broadcast Helper allows you to set up BOOTP and DHCP so clients can boot from an unspecified server, which may be located through a router and on another network segment. The EOS device forwards the BOOTPREQUEST packet and DHCP messages from a booting client to a server that can respond with the client's IP address.

If your network is quickly growing or changing, you may want to use the UDP Broadcast Helper for BOOTP instead of configuring a client to boot from one particular server, in which case you would have to reconfigure the client to boot from another server if the network configuration changes.

By supporting both the BOOTP and DHCP relay agents, the EOS software allows existing BOOTP clients to interoperate with DHCP servers. BOOTP and DHCP clients and their associated servers often do not reside on the same IP network or subnetwork. If the EOS software does not provide support for a relay agent, every subnet that has BOOTP and DHCP clients is required to have a BOOTP and DHCP server.

Figure 142: UDP Broadcast Helper for BOOTP Configuration



Prerequisites:

Ensure that:

- You are logged on to the device with Network Manager privilege.
- Ports and paths are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Client and server LANs and WANs are set up. See [IP Routing Configuration on page 119](#).
- You examined your network configuration and determined the devices on which UDP Broadcast Helper for BOOTP should be enabled.
- You determined which networks or servers should receive BOOTPREQUEST packets.
- You determined the IP addresses of the networks or servers that should receive BOOTPREQUEST packets, if possible.

Procedure:

1. Enable UDP Broadcast Helper by entering:

```
SETDefault -UDPHELP CONTrol = Enable
```

2. Add a UDP port for the BOOTP or DHCP server to the active ports list by entering one of the following command:.

- ADD -UDPHELP ActivePorts bpserver
- ADD -UDPHELP ActivePorts 67

You can specify either the built-in name BPSEVER or the UDP port number 67, which is mapped to built-in name BPSEVER. Both BOOTP and DHCP use the same UDP port numbers.

3. For UDP port 67 or BPSEVER, set up a list of networks and servers that should receive the BOOTPREQUEST broadcast packets.

If your device is configured to boot from a server that is accessed through a Frame Relay interface, you must perform this step. The device does not rebroadcast BOOTPREQUEST packets over Frame Relay interfaces.

For Frame Relay networks, the EOS device duplicates the packet and forwards it to each configured or dynamically learned neighbor.

You must configure the ForwardAddress parameter to eliminate unnecessary LAN broadcast packets by using one of the following syntaxes:

```
ADD -UDPHELP ForwardAddress <UDP_port_or_name> <IP_address>
```

```
ADD -UDPHELP ForwardAddress <UDP_port_or_name> <IP_address><subnet_mask>
```

```
ADD -UDPHELP ForwardAddress <UDP_port_or_name> <IP_address><subnet_mask> [Ones | Zeroes]
```

```
ADD -UDPHELP ForwardAddress <UDP_port_or_name> <list_of_interfaces>
```

You can add up to 32 addresses to the forward address list.

If you know the specific IP address of the server (or the network IP address where the servers resides) from which the client should obtain its IP address, add the address to the list.

Step example:

If the address of the server that responds to the BOOTPREQUEST packets is 10.1.0.1, you can add this address to the list by entering:

```
ADD -UDPHELP ForwardAddress 67 10.1.0.1
```

To forward BOOTPREQUEST packets to all servers on a specific network, enter:

```
ADD -UDPHELP ForwardAddress 67 10.0.0.0
```

The EOS device stores address 10.255.255.255 in the list, meaning that all servers (hosts) on network 10 receive the BOOTPREQUEST packet.

4. Optional: Configure the EOS device to detect unauthorized BOOTP and DHCP servers by entering:


```
ADD -UDPHELP AuthDHCPServer <IP_address>
```

You must specify the addresses of authorized servers. You can add up to 32 servers to the list.

Any BOOTPREPLY or DHCP OFFER packet received with an IP source address that does not match any server's IP address on the list is discarded, a system message is entered, and an SNMP trap is sent.

For information about the trap, see the description of the `AuthDHCPServer` parameter in "UDPHELP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

25.3

Verifying the UDP Broadcast Helper Configuration

Procedure:

- To check the settings of all parameters associated with UDP Broadcast Helper and UDP Broadcast Helper for BOOTP, enter:

```
SHoW -UDPHELP CONFIguration
```

- To obtain statistics related to UDP Broadcast Helper and BOOTP, enter:

```
SHoW -SYS STATistics -UDPHELP
```

For further information, see [Statistics Displays on page 967](#).

25.4

UDP Broadcast Helper for BOOTP Customization


You can customize the UDP Broadcast Helper configuration for BOOTP by configuring the `BootpMaxHops` and `BootpThreshold` parameters in the UDPHELP Service. The `BootpMaxHops` parameter limits the number of hops that a BOOTPREQUEST packet can make on a network. The `BootpThreshold` parameter prioritizes and forwards BOOTPREQUEST packets to a server according to a predetermined plan and determines which clients are booted first.

25.4.1

Limiting the Number of Hops

By configuring the `BootpMaxHops` parameter and limiting the number of hops, you can control how far a BOOTPREQUEST packet can travel on a network. For example, if your network configuration is similar to that shown in the following figure, you can set the `BootpMaxHops` value on Router C so that clients in a given area of the network can only boot from a specific server or servers.

Figure 143: Hops Limitation for BOOTPREQUEST Packets



You can assume that a client on Network A needs to send BOOTPREQUEST packets to Server B on Network E. Because you do not know the IP address of Server B and you have not configured the `ForwardAddress` parameter on any of the routers, each router continues to forward the packet out each of its ports and flood the network with packets.

You can configure the `BootpMaxHops` parameter to control the flood of packets.

For additional information on the `BootpMaxHops` parameter, see “UDPHELP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Prerequisites:

Ensure that you configured UDP Broadcast Helper for BOOTP. See [Relaying BOOTP and DHCP Traffic on page 451](#).

Procedure:

1. On port 1 of Router C, configure the `BootpMaxHops` parameter to 2 by entering:

```
SETDefault !1 -UDPHELP BootpMaxHops = 2
```
2. On port 2 of Router C, configure the `BootpMaxHops` parameter to 1 by entering:

```
SETDefault !2 -UDPHELP BootpMaxHops = 1
```


When Router C receives BOOTREQUEST packets from the clients on Network A, it forwards the packets to the server on Network E. However, Router C receives and discards the BOOTPREQUEST packets from the clients on Network B because the `BootpMaxHops` parameter value is set to 1 on port 2. Router C discards the BOOTPREQUEST packets because the packets already traversed one gateway, which is Router B.

25.4.2

Determining the Booting Order

By configuring the `BootpThreshold` parameter in the UDPHELP Service, you can determine which clients are booted first. For example, if your network configuration is similar to that shown in the following figure, you can set the `BootpThreshold` value on each port so that clients are booted according to a predetermined plan.

Figure 144: Determining Which Clients are Booted First



You can assume that, according to your predetermined plan, you want the clients on Network A to be booted first, then the clients on Network B, and then the clients on Network D. You need to set the `BootpThreshold` parameter on ports 1, 2, and 3 to different values so that the EOS device prioritizes and forward the BOOTPREQUEST packets to the server in the proper order.

When all the clients send out BOOTPREQUEST packets (the Seconds Elapsed Field in the BOOTPREQUEST packet is initially set to 0) at the same time, the EOS device forwards the packets received on port 1 because the Seconds Elapsed Field and `BootpThreshold` value match. The device discards the packets received on port 2 and 3 because the Seconds Elapsed Field in these packets is less than the `BootpThreshold` value configured for ports 2 and 3.

The clients on Networks B and D increase the Seconds Elapsed Field value in the BOOTPREQUEST packets and resend the packets. When the Seconds Elapsed Field value is greater than or equal to the `BootpThreshold` value on port 2, the device forwards the packets from the clients on Network B to the server on Network C. The device continues to discard the BOOTPREQUEST packets from Network D until the Seconds Elapsed Field value is greater than or equal to the `BootpThreshold` value for port 3.

For additional information on the `BootpThreshold` parameter, see “UDPHELP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Prerequisites:

Ensure that:

- UDP Broadcast Helper for BOOTP is configured.
- You determined which clients you want to boot first.

Procedure:

1. Set the `BootpThreshold` value on port 1 to the lowest value of all three ports by entering:
`SETDefault !1 -UDPHELP BootpThreshold = 0`
2. Set the `BootpThreshold` value on port 2 to the next lowest value of all three ports by entering:
`SETDefault !2 -UDPHELP BootpThreshold = 40`
3. Set the `BootpThreshold` value on port 3 to a value greater than that set for ports 1 and 2 by entering:
`SETDefault !3 -UDPHELP BootpThreshold = 100`

Chapter 26

Virtual Private Networks (VPNs) Configuration

A VPN (virtual private network) is a connection that has the appearance and many of the advantages of a dedicated link, but occurs over a shared network. By using a technique called *tunneling*, data packets are transmitted across a public routed network, such as the Internet or other commercially available network, in a private “tunnel” that simulates a point-to-point connection.

VPNs provide the following advantages:

- Network traffic from many sources can travel through separate tunnels across the same infrastructure.
- Network protocols can traverse incompatible infrastructures.
- Traffic from many sources can be differentiated and directed to specific destinations.

VPNs are a cost-effective alternative for providing remote access or remote office connectivity to a central site. Typically a company is required to use dedicated leased lines, packet-switching services, and/or direct dial-up connections to enable remote users and remote offices to connect to a central site. A VPN provides a less expensive method of providing this connectivity.

The Internet service provider (ISP) is an important element in of a VPN. By providing local access for any remote user or remote office, the ISP’s network replaces the leased lines, packet-switching services, and direct dial-up connections. Instead of directly managing remote access WAN lines, a company can outsource this responsibility to an ISP, resulting in fewer WAN issues to track and potentially significant cost savings.

26.1

Tunneling


To ensure security and multiprotocol support for a VPN, you must create a tunnel to the central site. Tunneling allows you to encapsulate IP and non-IP packets, and to obtain access to the central site network through a firewall.

Tunnel Components

Basic components of a tunnel:

- A tunnel initiator (TI)
- A routed network
- An optional tunnel switch
- One or more tunnel terminators (TT)

Figure 145: Basic VPN Connection



Tunnel initiation and termination can be performed by a variety of network devices and software.

A tunnel can be started by any one of the following devices:

- An end user's laptop equipped with an analog PC modem card and VPN-enabled dial-up software.
- A VPN-enabled extranet device on an enterprise branch or home office LAN.
- A VPN-enabled access concentrator at a network service provider point of presence (POP).

A tunnel can be ended by a tunnel terminator or switch on an enterprise network or on an ISP's network.


In addition, VPNs usually contain one or more security servers, such as the RADIUS server. Along with the conventional application of firewalls and address translation, VPNs can provide data encryption, authentication, and authorization.

Tunnel Switching

You can also employ tunnel switching, to enhance the security of your network. For example, you can configure a device outside the corporate firewall as a tunnel switch with the tunnel terminator located inside the firewall. You can then configure the firewall to permit only those tunnels with the source as the tunnel switch and the destination as the tunnel terminator. Tunnel switching with this type of device-specific configuration removes the need to allow all tunnel requests through the firewall. In addition, the user community does not need to be modified when you make structural changes to the central site, because access to the central site is tunneled through one device.

In the following figure, one device is the tunnel terminator and another is the tunnel switch. In this configuration, both devices use a RADIUS server for authentication.

Figure 146: Basic Tunnel Switching Components



After the RADIUS server authenticates user names that were submitted by the tunnel switch, it instructs the tunnel switch where to build the outgoing tunnel (to the tunnel terminator). You do not need to create virtual ports on the tunnel switch; this is accomplished automatically based on the information received from the RADIUS server.

26.1.1

Authentication Mechanisms

When creating a virtual port, you can choose one of two authentication mechanisms: proprietary SysCallerID or standards-based PPP.

SysCallerID Authentication

When all sites in your VPN use Enterprise OS products, you can configure SysCallerID as the authentication mechanism. If SysCallerID is used, PPP uses that identification parameter as a “caller ID” to identify itself to its peer when establishing a PPP Link Control Protocol (LCP) link. This identification allows the central site device to map incoming calls from the remote sites. The SysCallerID parameter is an Enterprise OS proprietary PPP authentication mechanism.



NOTE: If you use virtual leased line (VLL) configurations, SysCallerID is required.

For more information about configuring a port to use SysCallerID, see [Protocol Reservation Configuration on page 383](#) and “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Standards-based PPP Authentication

When all sites in your VPN are not using Enterprise OS products, you need to use the AuthLocalUser parameter and optionally, the AuthRemoteUser parameter in the PPP service to support standards-based PAP and CHAP authentication.

The AuthLocalUser (ALU) parameter certifies the device to a PPP peer. In a dial scenario, the calling device defines the ALU parameter. The AuthRemoteUser (ARU) parameter identifies the user and password used

by the PPP peer. The information entered in this device's ARU parameter is the same as that entered in the calling party's ALU.

For more information about configuring PAP and CHAP, see [Wide Area Networking Configuration with PPP on page 341](#) and the *Enterprise OS Software Version 18.2 Reference Guide*.

26.1.2

Multiprotocol Tunneling

Multiprotocol tunnels require IP connectivity. Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP) both use IP as their carrier protocol. If the tunnel endpoints cannot establish IP connectivity between them, they are not able to establish tunnels.

VPN tunnels carry PPP links. Anything that can be accomplished with PPP links can be accomplished over PPTP and L2TP tunnels. This includes multiprotocol routing, bridging, and bandwidth-on-demand, dial-on-demand and protocol reservation.

PPTP/L2TP defines a method for transferring Point-to-Point Protocol (PPP) datagrams through a tunnel over IP. Tunneling PPP does not change PPP, but provides a vehicle by which PPP data units (PDUs) can be carried between two peers.

A PPTP/L2TP connection is defined by two parallel components: a control connection and a data pipe. For PPTP, the control connection operates over TCP and passes call control and management packets over the TCP session. The data pipe operates over IP to transfer data packets encapsulated using Generic Routing Encapsulation Protocol Version 2 (GRE V2). For L2TP, both the control connection and the data pipe operate over the UDP session.

You can configure PPTP/L2TP tunnel connections between an Enterprise OS device and a VPN enabled access concentrator. In this scenario, the Enterprise OS device is acting as a tunnel terminator, which only receives inbound calls. An Enterprise OS device can also receive inbound calls from VPN-capable RAS clients to provide remote access services.

In addition, an Enterprise OS device expands the use of PPTP/L2TP so a tunnel can be established between two peer Enterprise OS devices. In this scenario, the Enterprise OS device is able to issue outbound calls so either side can be the tunnel initiator or tunnel terminator. This is a router-to-router configuration, and both peers play the same role.

For more information about configuring PPTP and L2TP tunnels, see [Security Features Configuration on EOS Devices on page 473](#) and the *Enterprise OS Software Version 18.2 Reference Guide*.

26.1.3

Tunnel Peers

Of the two tunnel peers, the tunnel terminator must have a fixed, publicly accessible IP address. If both tunnel peers have fixed public IP addresses, then you can use the `-L2T Service AccessList` parameter at both ends to explicitly limit who can access tunnels. This can provide additional security by preventing unauthorized tunnels from connecting. If both peers have fixed public IP addresses, the tunnel can be initiated from either end.

26.1.4

Compression

You can use the Motorola Solutions proprietary implementation of the Stac LZS algorithm to compress site-to-site tunnels. Note, however, that attempting to enable compression between Win32 (Win95/Win98/WinNT) clients and RAS concurrently with encryption leads to time-outs during the PPP option negotiation. You may enable compression on the Win32 client, or enable encryption, or neither, but not both.

26.1.5

Network Address Translation

Tunneling is frequently used to connect two or more private networks across a public network. Network Address Translation (NAT) applied on the public interface of the tunnel endpoint can prevent the advertisement of the private network's IP address(es) to the public network.

If session establishment occurs only from the tunnel initiators' site(s) to a central tunnel termination site, then the remote branches may use the same LAN address range, tunnel into RAS at the central site, and perform NAT from their private addresses to the address dynamically acquired from RAS.

For more information about configuring Network Address Translation, see [Network Address Translation \(NAT\) Configuration on page 243](#) and the *Enterprise OS Software Version 18.2 Reference Guide*.

26.2

VPN Strategies

The EOS software supports two basic VPN strategies:

Remote Access Services (RAS)

You can use this configuration to connect individual remote users to a central site.

Site-to-site

You can use this configuration to connect a remote office to a central site.



NOTE: Site-to-site VPNs are also commonly referred to as router-to-router, LAN-to-LAN, and VLL (virtual leased line) VPNs.

Remote Access Services VPNs


A RAS VPN allows an enterprise to replace dedicated dial ports with Internet connectivity through Internet service providers (ISPs), thereby reducing capital costs and providing access to points of presence across the globe. RAS VPNs reduce line charges and equipment costs by ensuring that whenever employees travel, they can get onto the corporate intranet by making a local call.

The EOS software supports two types of RAS VPNs:

Remote/mobile client dial-up VPN

Enables remote or traveling employees to dial into the corporate intranet via an ISP connection.

Figure 147: Remote/Mobile Client Dial-up RAS VPN



RAS extranet VPN

Enables corporate employees and partners to access corporate resources from remote LANs as if those users were physically present on the corporate LAN. See [Figure 145: Basic VPN Connection on page 457](#).

Both RAS VPN configurations require the authentication of the tunnel-initiating client before the tunnel can be established. The authentication mechanism can be either an external RADIUS server or a database in the tunnel terminator.

If the remote clients cannot support the tunneling protocol, the ISP must have tunnel-enabled access servers. In this configuration, the tunnel setup proceeds as follows:

- The remote user dials into the ISP's access server.
- The access server recognizes (based on a user ID, for instance, or on the user's choice from a menu) that this connection should be tunneled to the central site.
- The access server establishes the tunnel with the central site.
- The remote user establishes a session directly with the central site via the tunnel, just as if the two were directly attached.


While this configuration has the advantage that no special software is required on the remote client, the remote user can only dial into properly-equipped access servers.

Site-to-Site VPNs

In a site-to-site VPN, the configuration of the VPN depends on the type of Internet access the remote office has. The remote office can use either:

- A dial-on-demand configuration that accesses the Internet only when users at the remote office require a connection to the central site.
- A virtual leased line (VLL) configuration where the ISP connection is persistent at both ends.

Figure 148: Site-to-Site VPNs



In this configuration, the remote user (the client), such as a router, or an appropriately-configured personal computer, supports the tunneling protocol. The ISP does not have to support tunneling in any way.

The remote user dials the ISP. After the connection is set up, the remote user and the central site establish the tunnel, using authentication based on a user ID and password. A digital certificate may also be used.

The remote user and the central site may also negotiate encryption. After the tunnel was established, communications proceed as usual.

26.3

VPN Software Protocols

The earliest VPN software used the Point-to-Point Tunneling Protocol (PPTP). This was a defacto standard that was challenged by another protocol, Layer 2 Forwarding (L2F). The industry called for a single standard, which caused the two protocols to be merged into a more secure and improved tunneling protocol called Layer 2 Tunneling Protocol (L2TP). L2TP was defined and is currently implemented by many networking companies.

PPTP uses GRE (General Routing Encapsulation) to transport data and TCP port 1723 for the control frames. Most other VPN software uses UDP (Unnumbered Datagram Protocol) for transport. Some VPN software companies use a combination of these protocols or implemented proprietary VPN protocols. It is important that both ends of the tunnel use the same tunneling protocol.

The following table lists the ports used by various tunneling protocols. An Enterprise OS device can support PPTP and L2TP and can be a tunnel initiator, tunnel terminator, or tunnel switch for those protocols.

Table 53: Well-Known Tunneling Ports

Protocol	Port
Point-to-Point Tunneling Protocol (PPTP)	TCP port 1723
General Routing Encapsulation (GRE)	IP protocol Type 47

26.4

Remote Access for Individual Users (RAS VPNs)

In place of setting up multiple remote access servers at the central site, VPNs allow remote users to dial a local ISP. Using a VPN for remote access is particularly useful if you have remote users at a great distance from the central site. For example, users in Europe can call a local number instead of dialing in to the central site in New York.

The following two examples illustrate RAS VPN configurations.


RAS VPN Example 1: Remote/Dialup VPN

In this example, the ISP is configured to create a tunnel from the ISP's access server to the central site.



NOTE: You can also use this method for a remote office if you do not want to configure tunneling on the device at the remote office.

Figure 149: ISP to Central Site Tunnel



The connection process typically follows this order:

- The remote user (the client) dials the ISP.
- The ISP assigns an IP address to the remote user client.
- The ISP checks its authentication server for the user, and creates a PPTP (or L2TP) tunnel to the central site based on authentication data.
- The central site checks its authentication server to verify that this user can access the network and forwards the data.


No special configuration is required on the remote user computer except the configuration required to dial into the ISP's access server.

At the central site, you must perform [Configuring the Central Site of a Remote/Dialup VPN on page 464](#).

RAS VPN Example 2

In this example, the remote workstation is configured to create a tunnel directly to the central site.

Figure 150: Remote Workstation to Central Site Tunnel



You should see the documentation for your workstation or consult your operating system vendor for instructions on how to configure your workstation as the remote PPTP/L2TP client.

The connection process typically follows this order:

- The remote client dials the ISP.

- The ISP assigns an IP address to the client.
- The remote client sends data to the IP address of the central site.
- The Windows workstation client creates a PPTP tunnel to the central site based on authentication data.
- The central site checks its authentication server to verify that this user can access the network and forwards the data.

At the central site, you must perform [Configuring the Central Site of a RAS VPN on page 464](#).

26.4.1

Configuring the Central Site of a Remote/Dialup VPN

Process:

1. Configure the L2Tunnel Service to enable the device as a tunnel terminator. See [Configuring a PP Tunnel Terminator on page 535](#) and [Configuring a Router-to-Router Tunnel Initiator/Terminator on page 537](#).
2. Configure the firewall device (if present) or the device to allow tunnel traffic to pass. See [Setting Up Internet Firewalls on page 555](#).
3. Configure the RAS service to allow authentication of the user by a server, such as a RADIUS server. See [Configuring Remote Access with Tunnel Terminator on page 467](#).

You can also configure the firewall and RAS functions on the device.

26.4.2

Configuring the Central Site of a RAS VPN

Process:


1. Configure the L2Tunnel service to enable the device as a tunnel terminator. See [Configuring a PP Tunnel Terminator on page 535](#) and [Configuring a Router-to-Router Tunnel Initiator/Terminator on page 537](#).
2. Configure the firewall device (if present) or the device to allow tunnel traffic to pass. See [Setting Up Internet Firewalls on page 555](#).
3. Configure the RAS service to allow authentication of the user by a server, such as a RADIUS server. See [Configuring Remote Access with Tunnel Terminator on page 467](#).
4. Enable PPP encryption to allow encryption keys to be used by MPPE. See [Point-to-Point Protocol Communication on page 344](#).

26.5

VPN for a Remote Office (Site-to-Site VPNs)

You can create a VPN to connect a remote office device to the central site through the ISP using tunneling protocols such as the point-to-point tunneling protocol (PPTP). The following figure shows a typical configuration. In this configuration, the tunnel is established between the remote office and the central site. The ISP provides access to the shared network but does not interact in the tunneling setup.

Figure 151: Remote Office Tunnel



The connection process typically follows this order:

- The remote office device dials the ISP.
- The ISP assigns an IP address to the remote office device.
- The remote office device sends data to the IP address of the central site.
- A PPTP/L2TP tunnel is created between the remote site and the central site, and the data is forwarded through the firewall of the central site.
- The data is decrypted by the central site.

To configure the remote office device, see [Configuring the Remote Office Route in a Site-to-Site VPN on page 465](#).

To configure the central site device in a Site-to-Site VPN, see [Configuring the Central Site Device in a Site-to-Site VPN on page 466](#).

26.5.1

Configuring the Remote Office Route in a Site-to-Site VPN

Process:

1. Perform one of the following actions:
 - Configure dial-up to the ISP. See [Configuring WAN Resources on page 362](#).
 - Configure virtual leased line configuration. See [Configuring a PPTP Virtual Leased Line on page 539](#).

The ISP assigns an IP address to the client, or you configure an IP address that is applicable to the ISP's network using the IP service. For details, see [IP Routing Configuration on page 119](#) and [Network Address Translation \(NAT\) Configuration on page 243](#).

2. Create a virtual port specifying the SysCallerID (SCID) of the central site device. See [Identifying Remote Sites with SCID and CLIP on page 370](#).
3. Add a dial number list to the virtual port, specifying the IP Address of the central site and the type PPTP. See [Configuring the Dial List on page 367](#).
4. Configure the L2Tunnel Service to enable the device as a tunnel initiator. See [Configuring a Router-to-Router Tunnel Initiator/Terminator on page 537](#).

After the remote site dials the ISP, any data that is sent to the IP address of the central site creates a PPTP tunnel between the two sites.

26.5.2

Configuring the Central Site Device in a Site-to-Site VPN

Process:

1. Configure the L2Tunnel Service to enable the device as a tunnel terminator. See [Configuring a PP Tunnel Terminator on page 535](#) and [Configuring a Router-to-Router Tunnel Initiator/Terminator on page 537](#).
2. Configure the firewall device, if present, or the device to allow tunnel traffic to pass. See [Setting Up Internet Firewalls on page 555](#).

Chapter 27

Remote Access Services Configuration


This chapter describes how to configure the Remote Access Service (RAS) to enable remote users to connect to a central site in a virtual private network (VPN).

27.1

Configuring Remote Access with Tunnel Terminator

The following figure illustrates a topology in which a device is configured for distributed remote access with a tunnel terminator.

Figure 152: Distributed Remote Access



In this configuration, the Windows client dials into an Internet Service Provider (ISP) Point-of-Presence (POP) that can act as a Point-to-Point Protocol (PPTP) or Layer 2 Tunneling Protocol (L2TP) line server (LS).

The LS at the ISP site sets up a PPTP/L2TP tunnel using the EOS device as a tunnel terminator or packet processor (PP) and forwards all PPP packets from the Windows client to the device using the PPTP/L2TP encapsulation. The PPTP/L2TP tunnel is transparent to the remote client.

A PPP session is set up directly between the remote client and the EOS device RAS server, and the client can run IP protocol to access the enterprise network. In this configuration, the client also makes an MLP connection with the ISP, and both PPP connections are bundled together at the EOS device RAS server.

 **NOTE:** The configuration of the WAN port in the EOS device is not relevant to the RAS setup.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your devices are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).

 **NOTE:** The maximum RAS user name length is 64 rather than 16 characters and the password length is 32 characters rather than 8 characters.

- For multiple logins with the same user name and password, each login with the same user name is bound to one RAS port. The only exception is Multilink (MLP), when multiple logins can be bound to the same RAS port in these cases:

- The End Point Discriminators (EPDs) are the same
- MLP is selected and the EPDs are the same



NOTE: For a connection to be considered an MLP, the MLP option must be selected and the EPDs must be specified.

Procedure:

1. Configure the router LAN ports by entering:

```
SETDefault !1 -IP NETAddr = 130.1.1.1 255.255.255.0
```

The **NETAddr** parameter assigns an IP address to the specified port and configures the directly-connected IP network or subnet.

2. Configure RAS IP by performing the following actions:

- a. Establish which IP network to use for RAS.

RAS is supported on Ethernet LANs only; therefore the IP network you specify in the following command must be an Ethernet network.



NOTE: Specify only the IP network. A network mask is not required.

Step example:

```
SETDefault -RAS IPNetwork = 130.1.1.0
```

- b. To set up the IP pool option when there is no IP pool, enter:

```
SETDefault -RAS IPAddrPool = None (default)
```

- c. To set up the IP pool option by using the local DHCP server, enter:

```
SETDefault -RAS IPAddrPool = LocalDhcpServer
```

- d. Enter the range of IP addresses to be allocated by the DHCP server for RAS:

```
ADD !<portlist> -DHCP AddressPool <IPaddr1> - <IPaddr2> [!<profileid>]
```

- e. Configure NetBIOS and DNS (for RAS) in DHCP by entering the following commands:

```
SETDefault !<profileid> -DHCP DNS
```

```
SETDefault !<profileid> -DHCP NetBios
```

3. Enable IP routing for RAS by entering:

```
SETDefault CONTROL = ( [ R0ute | NoR0ute ], [ RelaySrcRoute | NoRelaySrcRoute ],  
[ SplitLoad | NoSplitLoad ], [ Filtering | NoFiltering ], [ SECurity |  
NoSECurity ], [ FwdSubnetBcast | NoFwdSubnetBcast ], [ FwdAllSubnetBcast |  
NoFwdAllSubnetBcast ] )
```

4. If you want to configure static virtual RAS ports for an internal database, perform the following actions:

- a. Add virtual RAS ports by entering:

```
ADD !<vport> -P0rt VirtualPort RAS
```

- b. For each user, enter:

```
ADD !<vport> -PPP AuthRemoteUser ("username", "password")
```

Each user requires a unique virtual port number, username and password.

- c. Optional: Enable an authentication protocol by entering:

```
SETDefault !<vport> -PPP AuthProTocol = None|Ppp|Chap
```

- d. Enable the virtual port by entering:


```
SETDefault !<vport> -Port CONTROL = Enable
```

- e. Configure the idle timer value for each remote user by entering:

```
SETDefault !<vport> - Port DialIdleTimer = <seconds>
```

5. If you want to configure the RADIUS server (external database), perform the following actions:

- a. Set the security type to RADIUS by entering:

```
SETDefault -RAS SecurityType = radius
```

- b. Set the primary authentication and accounting server to the designated RADIUS server by entering:

```
SETDefault -RAS PrimACntSrvr = [<IP_Address>]
```

```
SETDefault -RAS PrimAuthSrvr = [<IP_Address>]
```

The RADIUS server is external to the router. When you set -RAS SecurityType to radius, at a minimum you must configure the IP addresses for the primary authentication and accounting servers. For specific configuration details for your RADIUS server, consult the manufacturer's documentation.

- c. Optional: Set the secondary authentication and accounting servers to the designated RADIUS server by entering the following commands:

```
SETDefault -RAS SecACntSrvr = [<IP_Address>]
```

```
SETDefault -RAS SecAuthSrvr = [<IP_Address>]
```

- d. Set the secret string on the router to match the RADIUS server's RADIUS Client.

- e. Ensure the UDP ports for authentication (default value 1645) and accounting (default value 1646) match the values defined in the RADIUS server.

- f. If the values in [step 5e](#) do not match, change the values in the router by entering the following commands:

```
SETDefault -RAS AuthUdpport = <UDP_port_number>
```

```
SETDefault -RAS ACntUdpport = <UDP_port_number>
```

When you use an external user database such as a RADIUS server, ports are configured dynamically as they are needed. You do not need to configure any ports.

6. Optional: Configure a port use limit by entering:

```
ADD -PO PortLimit RAS <minimum> <maximum>
```

You can limit the number of ports that are used for RAS connections. By default there is no limit other than the limit to the number of possible virtual ports in your configuration. Using the PortLimit parameter, however, you can guarantee that a minimum number of ports is available for RAS connections, and no more than your specified number of ports is used for RAS connections.

7. Optional: Add RAS traps to the list of traps sent to the SNMP Network Manager by entering:

```
ADD -SNMP TrapProfile "<TrapProfileName>" REMote
```

8. Enable RAS by entering:

```
SETDefault -RAS CONTROL = Enable
```

9. To configure the PPTP/L2TP tunnel, perform one of the following actions:

- Enable the PPTP function by entering:

```
SETDefault -L2T CONTROL = Enable Protocol = PPTP
```

- Enable the L2TP function by entering:

```
SETDefault -L2T CONTROL = Enable Protocol = L2TP
```

- Enable both PPTP and L2TP functions by entering:

```
SETDefault -L2T CONTROL = Enable Protocol = ALL
```

10. Add an access list entry for the IP address of a PPTP/L2TP tunnel and ensure that flow control is enabled by entering:

```
Add -L2Tunnel AccessList <IPaddress> [ <Network_Mask>] [Protocol = PPTP]  
[FlowControl=<Enabled / Disabled>]
```

Flow control applies to PPTP tunnels only. The IP address and the network mask are set to that of a PPTP/L2TP tunnel terminator. In this configuration, this IP address is the remote ISP's IP address. For more information about configuring PPTP/L2TP tunnels, see [L2Tunnel Connections Configuration on page 535](#).

11. If you use L2TP as the tunneling protocol, set the local user name of the router by entering:

```
SetDefault -L2Tunnel L2TPLocalUser = (" <userid>", " <password>")
```

The assigned user ID is used as the "host name," which is required by the L2TP protocol during tunnel establishment.

27.2

NCPs Configuration

When a router establishes a Link Control Protocol (LCP) connection with a RAS client, it sends out the following Network Control Protocol (NCP) Configuration request packets:

- IPCP
- OSI

These request packets are sent out regardless of whether the RAS client supports these NCPs or not. The EOS software implements a scheme where the NCPs are configurable. Only the configured NCPs are negotiated.

The NCP configuration list is in the PPP component and is system-wide. You can enable or disable an NCP by specifying the appropriate element in the list. If an NCP is disabled, the router does not send out a configuration request packet for that NCP and ignores the incoming configuration request for that NCP.

To configure NCPs, you can use the PPP service `NcpProtocolId` parameter. You can use the `SETDefault` command to modify the list and the `SHow` command to view the list. Since this command is system-wide, the command `SHowD -PPP CONF` does not apply. The default configuration is to enable the IPCP NCP and disable the others.

To set up the NCP configuration list, you can use the `SETDefault -PPP NcpProtocolId` command.

To display the list, you can use: `SHow -PPP NcpProtocolId`

27.3

RADIUS Framed Net Mask and Framed Route Attributes

With the implementation of the RADIUS framed net mask and framed route, the RAS client can be a router that connects its network to the router's network. An IP host on the router's LAN can now talk to an IP host on the RAS router's LAN.

For a RAS client to be considered as a router, both the framed IP address and the framed net mask attributes must be configured on the RADIUS server.

RADIUS Server Configuration Example

To configure the RADIUS server, you must add a static route entry to the routing table. For example, you can use:

Dest. addr = Framed IP Addr & Framed Net Mask

Gateway = !virtual Port of this RAS router

Metric = 0

Owner = STATIC

This routing entry is not saved in CCS and is deleted when the RAS connection is torn down.

For each framed route entry defined in the RADIUS server, a static route entry is added to the routing table with the following information:

Dest. addr = Framed Route's destination address & net mask

Gateway = !virtual Port of this RAS router

Metric = Framed Route's Metric

Owner = STATIC

Since all packets destined for the framed route's destination address must go through the RAS virtual port, the framed route's gateway is ignored, and is replaced by the RAS virtual port.

The framed route format must follow the RFC 2138 recommendation. For example:

140.1.2.0 255.255.255.0 140.1.2.1 2

or

140.1.2.0/24 140.1.2.1 2

This routing entry is not saved in CCS and is deleted when the RAS connection is torn down.

The RAS IP address is not added to the ARP table as in the RAS Client case since it is not a host. When you enter the `SHoW -RAS Connection` command, the RAS IP address has the format Framed IP addr/Framed Net Mask (# of netmask bits). For example:

140.1.2.160/24.

RADIUS NAS-Port-Type Attribute

The RADIUS NAS-Port-Type attribute indicates the type of NAS physical port that is authenticating the user. It can have the following values:

- 0Async
- 1Sync
- 2ISDN Sync
- 3ISDN Async V.120
- 4ISDN Async C.110

When the RAS component receives a login request from a remote client, it extracts this information from the PPTP port CB and sends it in the RADIUS Access Request to the RADIUS server. Similarly, when the remote client closes down the RAS connection, this information is added to the accounting stop packet to be sent to the RADIUS server.

27.4

Logging Messages

You can configure RAS to log system messages to the AuditLog Service, to the local console, or both. Log messages contain crucial information such as the date and time, interface, incoming and outgoing, packet header summary, and reason the message was generated. Each message contains two types of codes: facility and priority. The facility code tells syslog what subsystem the message is from and the priority code tells syslog how important the message is (ranging from Log(0), emergency, which is the highest priority to Log(7), debug, which is the lowest priority).

Procedure:

- To log messages to the AuditLog Service, enable the service by entering:
`SETDefault -AuditLog CONTrol = (COnfig, MESSAGES, SEcurity)`
- To specify whether messages are logged to the AuditLog Service, the local console, or both, enter:
`SETDefault -RAS Log = ([Syslog | NoSyslog ], [ ConSole | NoConSole ],[CoNnect | NoCoNnect ], [ AuthFail | NoAuthFail ], [ RsrcFail | NoRsrcFail ])`

27.4.1

Specifying Log Content

Procedure:

- To generate a log message when a session is disconnected, enter:
`SETDefault -RAS Log = NoCoNnect`
- To log authentication failures, enter:
`SETDefault -RAS Log = AuthFail`
- To log connections that fail due to lack of resources, enter:
`SETDefault -RAS Log = RsrcFail`

Chapter 28

Security Features Configuration on EOS Devices

This chapter describes how to configure security features on EOS devices.

28.1

Protocol Authentication with Manual Security Policies

EOS supports manual security policies which you can use to implement protocol authentication via IPSec signatures. The current EOS release supports PIM protocol authentication to safeguard the integrity of PIM traffic; prevent fraudulent routing attacks; and ensure that PIM-enabled devices receive information only from trusted sources. PIM control packets are authenticated using authentication header (AH) as the authentication mechanism and HMAC-SHA-1 as the authentication algorithm.

When PIM protocol authentication is configured:

- Incoming IP packets are processed based on the SPI value of the packet, authenticated, and forwarded back to the IP layer.
- AH headers are added to outgoing packets, and the current SA is used to authenticate the packet.

EOS generates syslog messages and traps to log authentication failures. For details, see [Supported MIBs and Traps on page 845](#) and [Syslog Messages on page 997](#).

EOS supports the following manual security policy control states:

Clear

In this state, the device or port receives both authenticated and unauthenticated PIM packets and transmits only unauthenticated PIM packets.

Transit

In this state, the device or port receives both authenticated and unauthenticated PIM packets and transmits only authenticated PIM packets.

Secure

In this state, the device or port receives and transmits only authenticated PIM packets.



NOTE: The `GblManPolState` parameter configures the control state for global (per device) security policies, and the `ManPolState` parameter configures the control state for port-based security policies.

You can configure protocol authentication by creating global manual security policies and authentication keys and binding them to security associations (SAs) using the EOS CRYPTO service. Once the keysets and SAs are created, you can activate the protocol authentication by setting the security policy control state to “Transit” and then to “Secure.”



NOTE: Motorola Solutions recommends that during configuring security policies, CRYPTO control in the PORT service should be disabled. You can enable it only after completing all security policy configurations.

28.1.1

Configuring PIM Protocol Authentication

To configure PIM protocol authentication, you must use this procedure for each EOS device in the PIM domain.



NOTE: In the current implementation, PIM protocol authentication is performed on a per-device basis using global manual security policies; the same security policy and AH header are active on all PIM ports on the device. EOS also supports protocol authentication on a per-port basis, but per-port protocol authentication is not currently used in ASTRO® 25 systems.

Procedure:

1. Define a global manual security policy for PIM protocol authentication by entering:

```
ADD -CRYPTO GblManPOLicy <policy_name> <priority> <mode> ({<protocol> <src_ipaddr/
mask> (<dst_ipaddr/mask>)}) [<encrypt_alg>] [<auth_alg>]
```

where:

<priority> is 1-999, 1=highest

<mode> is AhXport

<protocol> is PIM | ANY

<auth_alg> is MD5 | SHA1

You must specify AhXport as <mode>, PIM as <protocol>, and SHA1 as the authentication algorithm (<auth_alg>). In addition, you should specify the source and destination IP address filters as IP addresses with masks. This allows IPsec to accommodate both unicast and multicast outgoing packets.

The <policy_name> must be unique per device.

Step example:

```
ADD -CRYPTO GblManPOLicy policy1 1 AhXport PIM 0.0.0.0/0 0.0.0.0/0 SHA1
```

2. Enable the security policy by entering:

```
ADD -CRYPTO GblManPolCtrl <policy_name> [Disabled | Enabled]
```

Step example:

```
ADD -CRYPTO GblManPolCtrl policy1 Enabled
```

3. Add an authentication keyset to the device by entering:

```
ADD CRYPTO ManKeySet <keyset_name> unique name (1-15 characters)
[AuthKey "<ascii-text>" | %<hex-string>"<ascii-text>" | %<hex-string>]
```

For security purposes, you must enter the key twice.

Step example:

```
ADD -CRYPTO ManKeySet keyset1 AuthKey "key1" "key1"
```

4. Verify that the configuration is correct and matches across the network by performing the following actions:

- a. To view the keys in encrypted format, on each PIM-enabled device for which you configured PIM protocol authentication, enter:

```
SHow -CRYPTO ManKeySet [<keyset_name>]
```

- b. Compare the encrypted authentication keys listed in the "Signature" field to verify that they match.



WARNING: It is important to determine that configurations are correct and match across the network before moving the devices into transit and then secure state. A misconfiguration or mismatched keyset can cause PIM protocol authentication to fail, causing dropped packets and network down times.

5. Add a new global security association (SA) to bind the manual keying information to a security policy by entering:

```
ADD -CRYPTO GblManKeyInfo <policy_name> (<keyset_name> [SpiAh <spi_in> <spi_out>])
```



NOTE: The *<spi_in>* and *<spi_out>* values must be unique in the system. In addition, to avoid potential misconfiguration, we recommend that you set *<spi_in>* and *<spi_out>* to the same value.

If this is the only SA associated with the specified security policy, it becomes the active global transmit SA (used to authenticate outgoing packets) once the device is moved to transit state. If more than one SA is associated with the specified security policy, the SA with the lowest *<spi_out>* value is used as the active global transmit SA.

Step example:

```
ADD GblManKeyInfo policy1 keyset1 SpiAh 300 300
```

6. Change the control state of the global manual security policy to “Transit” by entering:

```
ADD -CRYPTO GblPolState <policy_name> Transit
```

Step example:

```
ADD -CRYPTO GblPolState policy1 Transit
```

The device remains in transit state until you explicitly move it to another state by using the `GblPolState` parameter.

7. Change the control state of the global manual security policy to “Secure” by entering:

```
ADD -CRYPTO GblPolState <policy_name> Secure
```

Step example:

```
ADD -CRYPTO GblPolState policy1 Secure
```

8. Verify that all PIM packets are authenticated and that the receive security parameters index (SPI) value is the same as the transmit SPI across all devices by entering:

```
SHow -CRYPTO ManPeerReport
```

A display similar to the following appears:

Peer IP	LastAhSPI	Result
10.0.0.1	300	Pass
10.0.0.2	300	Pass

The values in the `LastAhSPI` column should match the *<spi_in>* value you configured for the SA and the `Result` column should list `Pass` indicating that the authentication processing was successful.

28.1.2

Changing the Protocol Authentication Keyset (Rekeying)

Procedure:

1. Add a new authentication keyset to the device by entering:

```
ADD -CRYPTO ManKeySet <keyset_name> unique name (1-15 characters)
[AuthKey "<ascii-text>" | %<hex-string>" <ascii-text>" | %<hex-string>]
```



NOTE: For security purposes, you must enter the key twice.

Step example:

```
ADD -CRYPTO ManKeySet keyset2 AuthKey "key2" "key2"
```

2. Add a new global security association (SA) to bind the global security policy to a new keyset and a new set of SPI values by entering:

```
ADD -CRYPTO GblManKeyInfo <policy_name> (<keyset_name> [SpiAh <spi_in> <spi_out>])
```

The *<spi_in>* and *<spi_out>* values must be unique in the system. In addition, to avoid potential misconfiguration, we recommend that you set *<spi_in>* and *<spi_out>* to the same value.

Step example:

```
ADD GblManKeyInfo policy2 keyset2 SpiAh 500 500
```

3. Specify that the new SA you created in [step 2](#) is used to authenticate outgoing packets by entering:

```
ADD -CRYPTO GblManXmitSA <policy_name> <spi_out>
```

Step example:

```
ADD -CRYPTO GblManXmitSA policy2 500
```

4. Verify that all PIM packets are authenticated and that the receive security parameters index (SPI) value is the same as the transmit SPI across all devices by entering:

```
SHow -CRYPTO ManPeerReport
```

A display similar to the following appears:

Peer IP	LastAhSPI	Result
10.0.0.1	500	Pass
10.0.0.2	500	Pass

The values in the *LastAhSPI* column should match the *<spi_in>* value you configured for the SA and the *Result* column should list *Pass*, indicating that the authentication processing was successful.

5. Delete the obsolete SA by entering:

```
DELeTe -CRYPTO GblManKeyInfo <policy_name> <spi_out> [<spi_in>]
```

Step example:

```
DELeTe -CRYPTO GblManKeyInfo policy1 300
```

28.1.3

Configuring PIM Authentication Rollback

You can perform this procedure to configure the PIM protocol to roll back to non-authenticated operation without incurring downtime.

Procedure:

1. Change the control state of the global manual security policy to *Transit* by entering:

```
ADD -CRYPTO GblPolState <policy_name> Transit
```

Step example:

```
ADD -CRYPTO GblPolState policy1 Transit
```

The device now transmits authenticated packets using the current SA and receives unauthenticated packets. If a manual policy exists, the device also receives authenticated packets matching any SA. If no manual policy exists, the device drops authenticated packets.

2. Change the control state of the global manual security policy to *Clear* by entering:

```
ADD -CRYPTO GblPolState <policy_name> Clear
```

Step example:


```
ADD -CRYPTO GblPolState policy1 Clear
```

The device now receives and transmits clear packets and, if a manual policy exists, receives authenticated packets matching any SA. If no manual policy exists, the device drops authenticated packets.

3. Delete the obsolete SA by entering:

```
DELEte -CRYPTO GblManKeyInfo <policy_name> <spi_out> [ <spi_in>]
```

Step example:

```
DELEte -CRYPTO GblManKeyInfo policy1 300
```

28.1.4

Configuring X509 Certificates for the Router

Procedure:

1. Unlock the PKI database by entering:

```
SETDefault -PKI CONTrol = Unlocked
```

2. Generate the key pair by using `KeyPair` parameter.

Public-key cryptography uses the two following types of keys:

- Public key known to other parties and used in shared secret computations
- Private key that is generated on the device and never leaves its secure storage

The public part of the keypair is contained in the certificate request.

3. Provide the information to be included in certification request.

At least one of the following parameters must be provided:

- Subject name in the form of Common Name (the `SubjectName` parameter)
- External IP address of the tunnel interface (`IPADDRESS` parameter)
- DNS name (`DNSName` parameter)
- Email address (`EmailADDRESS` parameter)

4. Create the certification request (or a CSR) by using the `ADD -PKI CertReq` command.

The CSR is a message destined for the certificate authority (CA). It contains the identity of the device, defined in [step 2](#), and its public key.

You can obtain the location of the generated CSR by entering `SHOW -PKI CertReq`

5. Have the certificate request signed by the appropriate CA.

Based on information contained in the CSR, the CA forms a new certificate which is then signed with the CA's private key.

6. Import the signed certificate (the CSR response) to the router by entering:

```
ADD CERTificate <profile> <Self/TrustedCA/UnTrusted> InputFile <local_file_name>
```

The EOS device distinguishes the three following kinds of certificates:

Self – The device's own certificates. CA's responses for the device-issued CSRs.

TrustedCA – Certificates of the entities trusted without further checking.

UnTrusted – All other certificates (including any intermediate CA certificates). The validity of these certificates depends on the device's ability to build certificate chain leading to the Trusted CA certificate.

The device's own certificate must be imported as `Self`.

Each certificate installed on the EOS device has a unique name, called *profile*. It is used in the certificate operations for adding, deleting or showing status.

Before any changes to the certificate database can be made, the database must be unlocked by issuing the `SetDefault -PKI CONTrol = Unlocked` command.

Once the database is unlocked, the certificates can be added and deleted. Newly added certificates are marked as 'disabled' and are unusable for the PKI service until they are enabled. This is to enable the user to install and review certificates on the device before they influence the PKI clients. The newly installed certificates are enabled when the database is locked again (`SetDefault -PKI CONTrol = Locked`)

7. Import the remaining certificates in the certificate chain.

This includes all the intermediate CA certificates (if any) and the trusted CAs.

8. Verify completeness and validity of each certificate in the chain by entering:

```
SHoW -PKI CERTIFICATE LIst
```

9. Ensure that all certificates are displayed as *Valid*.

28.1.4.1

Configuring X509 Certificates (Example)

Procedure:

1. Unlock the PKI database by entering:

```
SETDefault -PKI CONTrol = Unlocked
```

2. Generate new ECDSA-384 KeyPair by entering:

```
ADD -PKI KeyPair MyKeyPair ECDSA 384
```

3. Set the parameters that are included in the CSR by entering the following commands:

```
SETDefault -PKI SubjectName = cn=John Doe,o=ACME Inc.,ou=Headquarters,c=US
```

```
SETDefault -PKI EmailADDRESS = none
```

```
SETDefault -PKI IPADDRESS = none
```

```
SETDefault -PKI DNSName= none
```

4. Generate the CSR by entering:

```
ADD -PKI CertReq MyCr
```

5. Display the generated CSR by entering:

```
SHoW -PKI CertReq MyCr
```

6. Verify that the displayed information is valid and transfer the CSR from the indicated location to the CA.

7. Obtain the certificate based on the generated CSR.

The process of signing the CSR is outside of this manual's scope.

8. Obtain certificate of the CA and, if applicable, intermediate CAs, and transfer them to the router.

9. Import the appropriate certificates to form a full certification path.

The CA and router's own certificate must be installed. The commands vary depending on the certificate:

```
ADD -PKI CERTIFICATE OwnCert Self InputFile a:/primary/selfcert.pem
```

```
ADD -PKI CERTIFICATE IcaCert UnTrusted InputFile a:/primary/icacert.pem
```

```
ADD -PKI CERTificate CACert TrustedCA InputFile a:/primary/cacert.pem
```

10. Display the certificates and verify that the information is correct.

When importing the CA certificate, utmost care needs to be taken to verify the validity of the CA certificate fingerprints:

- Display the fingerprint of the certificate by issuing `SHOW - PKI CERTificate TrustedCa`
- Obtain the fingerprint using a different communication channel than the one used for obtaining the certificate (for example, a phone).
- Compare the certificates. If they do not match, delete the offending certificate at once.

11. Lock the PKI database to make the newly added certificates usable for the PKI service:

```
SETDefault -PKI CONTrol = Locked
```

12. Verify the validity of each imported certificate:

```
SHow -PKI CERTificate LIst
```

28.2

IP Traffic Protection with Dynamic Security Policies

Dynamic security policies provide protection for sensitive IP traffic traversing unsecured networks, such as the Internet, with a greater level of security than manual key policies. When you create a dynamic security policy, you specify:

- The type of IP traffic to be protected
- How the identified IP traffic is protected
- Whether pre-shared keys or PKI certificates are used as authentication mechanism when establishing IKE security associations.



NOTE: It is recommended that while you are configuring security policies you disable CRYPTO control in the PORT service and enable it only after you completed all security policy configurations.

EOS supports IKEv1 security associations for IPv4 and IKEv2 security associations for IPv6.

IKEv2 dynamic security policies for IPv6 are supported only on the GGM 8000 gateway and the MNR S6000 router. The MNR S2500 does not support dynamic security policies for IPv6.

28.2.1

Configuring a Dynamic-Key Security Policy for IPv4

Procedure:

1. Define the traffic that needs to be protected by configuring the `SelectorList` parameter.

CRYPTO selector lists are used to determine which traffic is secured by a given dynamic policy. The selector list specifies one or more types of traffic to include (or exclude) and is linked to the dynamic policy by its name. You must enter the selector list before you add the dynamic policy.



NOTE: When you use the FRF.17, you must set the `SelectorList` parameter to `ANY 0.0.0.0 0.0.0.0/0` to include all non-control data.

2. Define the type of protection by using the `TransformList` parameter.

This defines how the data traffic is protected.

Dynamic policies allow a great variety of security transforms to be used to protect IP traffic. These transforms are specified in CRYPTO transform lists, which are named lists of protocol-transform

combinations. Like selector lists, transform lists are included by name, and you must enter them before you add the dynamic policy.

3. Define how the IKE/ISAKMP negotiation is protected by using the `IKEProfile` parameter.

The `IKEProfile` parameter defines a group of settings for CRYPTO to use when establishing an IKE security association. The settings include authentication method, encryption algorithm, hash algorithm and, optionally, the lifetime and Diffie-Hellman group to use in negotiations.

By providing the authentication method parameter, you decide whether to use PreSharedKey or X509 certificates to secure the IKE exchange.



NOTE:

If X509 certificates are in use, authentication method must match the type of the key pair generated when configuring the X509 certificates for the router. For more information, see [Configuring X509 Certificates for the Router on page 477](#).

4. Optional: Perform one of the following actions:

- If you chose PreSharedKey in [step 3](#), define the `PreSharedKey` parameter for authentication.
- If you did not choose the PreShared Key in [step 3](#), configure the X509 certificates for the router. See [Configuring X509 Certificates for the Router on page 477](#).

The `PreSharedKey` parameter defines the pre-shared keys used when establishing IKE security associations by using the pre-shared key authentication method. The key is associated with the peer or peers by using the `<peer_ID>` value specified with the `PreSharedKey` parameter. You can enter key values as quoted ASCII text, or as a series of hexadecimal digits preceded by %.

You can configure large networks easily by using the same key values across many EOS devices. By specifying the peer ID as an IP address with a subnet mask, all the peers falling within the subnet can share a single key. The Phase 1 ID 0.0.0.0/0 matches any IP address to facilitate a global shared key.



IMPORTANT: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the pre-shared key. For details about bringing up a device in FIPS mode, see the FIPS Security Policy.

5. Bind the information together by using the `DynamicPolicy` parameter:

```
ADD !<portlist> DynamicPOLicy <policy_name> <priority> [<version>] <mode>
<selctrlist_name> <xfrmlist_name> [<pfs>] [<lifetime>] [<preconnect>]
[<replay_win_size>]
```

where:

`<policy_name>` is a unique name (1-15 characters)

`<priority>` is 1-9999 (1 is the highest)

`<version>` is IKEv1 | IKEv2

`<mode>` is Tunnel | Xport | FRF17

`<selctrlist_name>` is the name of SelectorList to match

`<xfrmlist_name>` is the name of TransformList to use

`<lifetime>` is

GlobalLifeTime | ([5-1440 min | 1-504 hr | 1-21 dy] | [1000-4000000 kb])

`<preconnect>` is PreConnect [Peer IP]

`<replay_win_size>` is Winsize [0 | 64 | 128 | 256 | 512 | 1024]

The `DynamicPOLicy` parameter adds dynamic-key security policies to one or more ports. Dynamic policies specify whether to use IPSec tunnel, IPSec transport, or FRF17 mode, which selector list to use to match IP traffic, and which transform list to use when encrypting or authenticating packets.



NOTE: The anti-replay feature applies only to IPSec (Tunnel or Xport) security policies. If you do not specify a `WinSize` value by using the `ADD -CRYPTO DynamicPOLicy` command, the default anti-replay window size (1024 bits) is used.

For the FRF.17 security policies, EOS automatically sets the `WinSize` value to 0 to disable the anti-replay feature. If you attempt to add an FRF.17 policy with a `WinSize` value, the `ADD -CRYPTO DynamicPOLicy` command fails.

When a dynamic policy is created, it is given a unique name. This name is used to identify the policy in subsequent commands. The policy is also assigned a unique priority from 1 to 9999 to determine the preference between policies.



NOTE: Traffic that matches more than one policy is always secured by the policy with the lowest priority. Since dynamic policies may exist on several ports, their priority values must be unique across all of the ports on the system.

When you use pre-connect, the `<portlist>` is not supported. You must enter the dynamic policies separately for each port.

6. Enable security policy checking on the port by entering:

```
SETDefault !<portlist> -CRYPTO CONTro1 = Enabled
```



NOTE: If you enable security policy checking (set the `CRYPTO` service `CONTro1` parameter to Enabled) on an FRF.17 link, all (non-control) data is encrypted.

Step example:

To enable security policy on port 1, enter:

```
SETDefault !1 -CRYPTO CONTro1 = Enabled
```

28.2.2

Configuring a Dynamic Security Policy for IPv6

Procedure:

1. Define the traffic that needs to be protected by configuring the `SelectorLIst` parameter.

For more information, see "CRYPTO Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

CRYPTO selector lists are used to determine which traffic is secured by a given dynamic policy. The selector list specifies one or more types of traffic to include (or exclude) and is linked to the dynamic policy by its name. You must enter the selector list before you add the dynamic policy.

2. Define the type of protection by using the `TransformLIst` parameter.

This defines how the data traffic is protected.

Dynamic policies allow a great variety of security transforms to be used to protect IP traffic. These transforms are specified in CRYPTO transform lists, which are named lists of protocol-transform combinations. Like selector lists, transform lists are included by name, and you must enter them before you add the dynamic policy.



NOTE: The DES encryption algorithm is not supported for IKEv2. When you create transform lists for IKEv2 security policies, do not specify encryption protocol-transforms (by using the `<enc_transform>` value of the `TransformLIst` parameter) that use DES encryption.

3. Define how the IKE/ISAKMP negotiation is protected by using the `IKEProfile` parameter.

For more information, see “CRYPTO Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

The `IKEProfile` parameter defines a group of settings for CRYPTO to use when establishing an IKE security association. The settings include authentication method, encryption algorithm, hash algorithm and, optionally, the lifetime and Diffie-Hellman group to use in negotiations.



NOTE: The `<version>` value of the `IKEProfile` parameter must be set to `IKEv2`, and the `<encrypt_alg>` value of the `IKEProfile` parameter cannot be set to `DES`. The `DES` encryption algorithm is not supported for `IKEv2` profiles.

4. Define the `PreSharedKey` parameter for authentication.

For more information, see “CRYPTO Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

The `PreSharedKey` parameter defines the pre-shared keys used when establishing IKE security associations by using the pre-shared key authentication method. The key is associated with the peer or peers using the `<peer_ID>` value specified by using the `PreSharedKey` parameter. You can enter key values as quoted ASCII text, or as a series of hexadecimal digits preceded by %.

You must set the `<version>` value of the `PreSharedKey` parameter to `IKEv2`.

You can configure large networks easily by using the same key values across many EOS devices. By specifying the peer ID as an IPv6 address with a prefix, all the peers falling within the subnet can share a single key. The `Phase 1 ID ::` matches any IP address to facilitate a global shared key.



IMPORTANT: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the pre-shared key. For details about bringing up a device in FIPS mode, see the FIPS Security Policy.

5. Bind the information together by entering:

```
ADD !<portlist> DynamicPOLicy <policy_name> <priority> [<version>] <mode>
<selctrlist_name> <xfrmlist_name> [<pfs>] [<lifetime>] [<preconnect>]
[<replay_win_size>]
```

where:

- `<policy_name>` is a unique name (1-15 characters)
- `<priority>` is set from 1-9999 (1 is the highest priority)
- `<version>` is `IKEv1` | `IKEv2`
- `<mode>` is `Tunnel` | `Xport` | `FRF17`
- `<selctrlist_name>` is the name of `SelectorList` to match
- `<xfrmlist_name>` is the name of `TransormList` to use
- `<lifetime>` is
`GlobalLifeTime` | ([5-1440 min | 1-504 hr | 1-21 dy] | [1000-4000000 kb])
- `<preconnect>` is `PreConnect` [Peer IP]
- `<replay_win_size>` is `Winsize` [0 | 64 | 128 | 256 | 512 | 1024]

The `DynamicPOLicy` parameter adds dynamic-key security policies to one or more ports. Dynamic policies specify whether to use IPSec tunnel, IPSec transport, or FRF17 mode, which selector list to use to match IP traffic, and which transform list to use when encrypting or authenticating packets.



NOTE: The `IKEv2` protocol does not support `FRF.17` security policies or transforms that use the `DES` encryption algorithm.

The anti-replay feature applies only to IPSec (`Tunnel` or `Xport`) security policies. If you do not specify the `WinSize` value by using the `ADD -CRYPTO DynamicPOLicy` command, the default anti-replay window size (1024 bits) is used.

When a dynamic policy is created, it is given a unique name. This name is used to identify the policy in subsequent commands. The policy is also assigned a unique priority from 1 to 9999 to determine the preference between policies.



NOTE: Traffic that matches more than one policy is always secured by the policy with the lowest priority. Since dynamic policies may exist on several ports, their priority values must be unique across all of the ports on the system.

When you use pre-connect, the `<portlist>` is not supported. You must enter the dynamic policies separately for each port.

6. Enable security policy checking on the port by entering:

```
SETDefault !<portlist> -CRYPTO CONTrol = Enabled
```

Step example:

To enable security policy checking on port 1, enter:

```
SETDefault !1 -CRYPTO CONTrol = Enabled
```

28.2.3

Customized Security Associations

Two optional parameters are provided to customize the security associations created by dynamic policies:

DynpolPFS

Provides higher security by renegotiating a shared secret between CRYPTO peers each time a new key is needed. Since generating a shared secret demands intense numerical calculations (known as Diffie-Hellman calculations), using this option may cause reduced performance during renegotiation.

DynpolLifeTime

Determines the amount of time a security association can exist before it must renegotiate a new security association. You can specify the `DynpolTime` in units of minutes (m), hours (h), or days (d). By default, policies use the value specified in the `GlobalLifeTime` parameter.

28.2.4

Dynamic Key Configuration Examples

The following examples illustrate configurations of different topologies employing dynamic security policies:

- IPsec tunnel mode (GGSN)
To configure the GGSN and CBR, see [Configuring the GGSN and CBR on page 484](#).
- IPsec tunnel mode (GGSN) - configuration with X509 certificates
To set up a Dynamic Key Configuration with X509 certificates, see [Setting Up the Dynamic Key Configuration with X509 Certificates on page 486](#).
- FRF17 mode (CWR)
To configure core routers in a fully-meshed cooperative WAN routing (CWR) topology, using FRF17 mode with a pre-shared key and IKE pre-connect negotiation, see [Setting Up a CWR Topology for Encryption on page 488](#).
- IPsec tunnel mode (fully meshed)
To configure routers in a fully meshed topology with a dynamic key, see [Configuring Routers in Fully Meshed Topology with a Dynamic Key on page 491](#).
- IPsec tunnel mode over IPv6 infrastructure (IPv4in IPv6 tunnel with IKEv2)
To create an IPv6 IPsec tunnel with an IKEv2 Dynamic Key and configure routers, see [Creating an IPv6 IPsec Tunnel with an IKEv2 Dynamic Key on page 496](#)
- IP Sec transport mode


To create a hub and spoke topology and configure routers, see [Creating a Hub and Spoke Topology on page 498](#).

28.2.4.1

Configuring the GGSN and CBR

The following figure depicts a topology between a GGSN and a CBR.

Figure 153: GGSN Topology: GGSN and CBR



This configuration has the following properties:

- IPsec tunnel mode, with an IPIP tunnel between the GGSN and the CBR.
- Dynamic Keys using IKE.
- Phase 1 IKE Profile that uses the pre-shared keys, 256-bit AES, SHA and Group 5 PFS.
- Phase 2 TransformList that uses the 256-bit AES encryption and SHA authentication.



NOTE: If you configure by using configuration files, the `ADD DynamicPOLicy`, `ADD PreSharedKey`, and `SETDefault -CRYPTO CONTROL = Enabled` commands should be in the `XGSN.cfg` file.

Procedure:

Configuring GGSN

1. Configure IPIP point-to-point virtual ports for the topology by entering the following commands:

```
ADD !v1 -Port VirtualPort IPIP 10.99.1.254 10.99.1.1
ADD !v2 -Port VirtualPort IPIP 10.99.2.2 10.99.2.21
ADD !v10 -Port VirtualPort IPIP 10.99.10.10 10.99.10.29
```

2. Configure the security policy by performing the following actions:

- a. Specify a global lifetime of eight hours to be used with all dynamic policies by entering:

```
SETDefault -CRYPTO GlobalLifeTime = 8 hr
```

- b. Add a selector list to specify the traffic to which the policies apply by entering:

```
ADD -CRYPTO SelectorList sl 1 Include ANY 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values `0.0.0.0/0` are used.

- c. Add a transform list that specifies the Phase 2 security association by entering:

```
ADD -CRYPTO TransformList tl 1 ESP-AES/256 ESP-SHA
```

This is the description of the protocols and transforms proposed when the security associations are negotiated.

- d. Define an IKE profile that describes the Phase 1 security association by entering:

```
ADD -CRYPTO IKEProfile 3 PreSharedKey AES/256 SHA Group5 1 dy
```


This is used by IKE to secure its own negotiation and is not used to secure the data traffic

- e. Add dynamic-key security policies to the virtual ports configured in [step 1](#), by entering the following commands:

```
ADD !v1 -CRYPTO DynamicPOLicy dp2 2 t sl t1 5 min pc 10.99.1.1 WinSize 512
```

```
ADD !v2 -CRYPTO DynamicPOLicy dp3 3 t sl t1 5 min pc 10.99.2.21 WinSize 512
```

```
ADD !v10 -CRYPTO DynamicPOLicy dp1 1 t sl t1 5 min pc 10.99.10.29 WinSize 512
```

The **PreConnect** option is used to specify that the policies are pre-connect policies. The **WinSize** option is used to configure the anti-replay window size.

- f. Define pre-shared keys to be used when establishing IKE security associations with the preshared key authentication method by entering the following commands:

```
ADD -CRYPTO PreSharedKey 10.99.1.1 "12345"
```

```
ADD -CRYPTO PreSharedKey 10.99.2.21 "12345"
```

```
ADD -CRYPTO PreSharedKey 10.99.10.29 "12345"
```



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the **FipsPreShrdKey** parameter to add the preshared key. The syntax of the **FipsPreShrdKey** parameter requires you to enter the key twice for added security.

- g. Enable CRYPTO control on the CRYPTO ports by entering the following commands:

```
SETDefault !v10 -CRYPTO CONTROL= e
```

```
SETDefault !v1 -CRYPTO CONTROL= e
```

```
SETDefault !v2 -CRYPTO CONTROL= e
```

- h. Check the CRYPTO configuration by entering:

```
SHow -CRYPTO CONFIguration
```

Configuring the CBR

3. Configure IPIP point-to-point virtual ports by entering the following commands:

```
ADD !v1 -Port VirtualPort IPIP 10.99.1.1 10.99.1.254
```

```
ADD !v2 -Port VirtualPort IPIP 10.99.2.21 10.99.2.2
```

```
ADD !v10 -Port VirtualPort IPIP 10.99.10.29 10.99.10.10
```

4. Configure the security policy by performing the following actions:

- a. Specify a global lifetime of eight hours to be used with all dynamic policies by entering:

```
SETDefault -CRYPTO GlobalLifeTime = 8 hr
```

- b. Add a selector list to specify the traffic to which the policies apply by entering:

```
ADD -CRYPTO SelectorList sl 1 Include ANY 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values **0.0.0.0/0** are used.

- c. Add a transform list that specifies the Phase 2 security association by entering:

```
ADD -CRYPTO TransformList tl 1 ESP-AES/256 ESP-SHA
```

This is the description of the protocols and transforms proposed when the security associations are negotiated.

- d. Define an IKE profile that describes the Phase 1 security association by entering:

```
ADD -CRYPTO IKEProfile 3 PreSharedKey AES/256 SHA Group5 1 dy
```

This is used by IKE to secure its own negotiation and is not used to secure the data traffic.

- e. Add dynamic-key security policies to the virtual ports configured in [step 3](#), by entering the following commands:

```
ADD !v1 -CRYPTO DynamicPOLicy dp2 2 t sl t1 5 min pc 10.99.1.254 WinSize 512
ADD !v2 -CRYPTO DynamicPOLicy dp3 3 t sl t1 5 min pc 10.99.2.2 WinSize 512
ADD !v10 -CRYPTO DynamicPOLicy dp1 1 t sl t1 5 min pc 10.99.10.10 WinSize 512
```

The `PreConnect` option is used to specify that the policies are pre-connect policies. The `WinSize` option is used to configure the anti-replay window size.

- f. Define pre-shared keys to be used when establishing IKE security associations using the preshared key authentication method by entering the following commands:

```
ADD -CRYPTO PreSharedKey 10.99.1.254 "12345"
ADD -CRYPTO PreSharedKey 10.99.2.2 "12345"
ADD -CRYPTO PreSharedKey 10.99.10.10 "12345"
```



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- g. Enable CRYPTO control on the CRYPTO ports by entering the following commands:

```
SETDefault !v1 -CRYPTO CONTrol= e
SETDefault !v2 -CRYPTO CONTrol= e
SETDefault !v10 -CRYPTO CONTrol= e
```


- h. Check the CRYPTO configuration by entering:

```
SHow -CRYPTO CONFiguration
```

28.2.4.2

Setting Up the Dynamic Key Configuration with X509 Certificates

Figure 154: GGSN Topology: GGSN and CBR (X509 Certificates)



This configuration has the following properties:

- IPsec tunnel mode, with an IPIP tunnel between the GGSN and the CBR
- Dynamic Keys using IKE
- Phase 1 IKE Profile that uses the X509 certificates with ECDSA-384 keypair

- Phase 2 TransformList that uses the 256-bit AES encryption and SHA authentication



NOTE: If you configure by using configuration files, the `ADD DynamicPOLicy`, `ADD PreSharedKey`, and `SETDefault -CRYPTO CONTROL = Enabled` commands should be in the `XGSN.cfg` file.

Procedure:

1. Configure the X509 certificates. See [Configuring X509 Certificates \(Example\) on page 478](#).

Ensure that all certificates are valid.

2. Configure IPIP point-to-point virtual ports for the topology by entering the following commands:

```
ADD !v1 -Port VirtualPort IPIP 10.99.1.254 10.99.1.1
```

```
ADD !v2 -Port VirtualPort IPIP 10.99.2.2 10.99.2.21
```

```
ADD !v10 -Port VirtualPort IPIP 10.99.10.10 10.99.10.29
```

3. Configure the security policy by performing the following actions:

- a. Specify a global lifetime of eight hours to be used with all dynamic policies by entering:

```
SETDefault -CRYPTO GlobalLifeTime = 8 hr
```

- b. Add a selector list to specify the traffic to which the policies apply by entering:

```
ADD -CRYPTO SelectorList sl 1 Include ANY 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values `0.0.0.0/` are used.

- c. Add a transform list that specifies the Phase 2 security association by entering:

```
ADD -CRYPTO TransformList tl 1 ESP-AES/256 ESP-SHA
```

This is the description of the protocols and transforms proposed when the security associations are negotiated.

- d. Define an IKE profile that describes the Phase 1 security association by entering:

```
ADD -CRYPTO IKEProfile 3 ECDSA-384 AES/256 SHA Group5 1 dy
```

This is used by IKE to secure its own negotiation and is not used to secure the data traffic.

- e. Add dynamic-key security policies to the virtual ports configured in [step 2](#) by entering the following commands:

```
ADD !v1 -CRYPTO DynamicPOLicy dp2 2 t sl t1 5 min pc 10.99.1.1 WinSize 512
```

```
ADD !v2 -CRYPTO DynamicPOLicy dp3 3 t sl t1 5 min pc 10.99.2.21 WinSize 512
```

```
ADD !v10 -CRYPTO DynamicPOLicy dp1 1 t sl t1 5 min pc 10.99.10.29 WinSize 512
```

The `PreConnect` option is used to specify that the policies are pre-connect policies. The `WinSize` option is used to configure the anti-replay window size.

- f. Define pre-shared keys to be used when establishing IKE security associations using the preshared key authentication method by entering the following commands:

```
ADD -CRYPTO PreSharedKey 10.99.1.1 "12345"
```

```
ADD -CRYPTO PreSharedKey 10.99.2.21 "12345"
```

```
ADD -CRYPTO PreSharedKey 10.99.10.29 "12345"
```



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- g. Enable CRYPTO control on the CRYPTO ports by entering the following commands:

```
SETDefault !v10 -CRYPTO CONTROL= e  
SETDefault !v1 -CRYPTO CONTROL= e  
SETDefault !v2 -CRYPTO CONTROL= e
```

h. Check the CRYPTO configuration by entering:


```
SHOW -CRYPTO CONFIGURATION
```

28.2.4.3

Setting Up a CWR Topology for Encryption

The following figure illustrates a fully-meshed cooperative WAN routing (CWR) topology that uses a FRF17 mode with pre-shared key and IKE pre-connect negotiation.

Figure 155: CWR Topology (FRF17)



This topology demonstrates the following configuration properties:

- FRF17 mode.
- Dynamic Keys that use IKE.
- Phase 1 IKE Profile that uses pre-shared keys, 3DES, SHA and Group 5 PFS.

- Phase 2 TransformList that uses FRF.17 for encryption (256-bit AES) and authentication (MD5).

Procedure:

Configuring Core Routers (Both devices should be configured identically, except where noted in the procedure.)

1. Configure virtual ports. See [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
2. Configure the system IP addresses by performing the following actions:

- a. On the core router 1, enter:

```
SETDefault -IP SystemIP = 10.1.253.9
```

- b. On the core router 2, enter:

```
SETDefault -IP SystemIP = 10.1.253.17
```



IMPORTANT: If you change the system IP address on an EOS device, the SNMPv3 credentials must be reestablished with that device. If the SNMPv3 users are configured on the device prior to the system IP address change, you must issue the `ResetV3` command to reset the SNMPv3 data, then reconfigure the SNMPv3 users with the appropriate privilege levels. For more information about how to configure the SNMPv3 users, see [Predefined SNMPv3 USM Users Management on EOS Devices on page 670](#) and [General SNMPv3 USM Users Management on EOS Devices on page 679](#).

3. Add the static routes to the site routers by entering the following on each router:

```
ADD -IP ROute 10.101.2.1 255.255.255.255 20.1.3.1 2
```

```
ADD -IP ROute 10.101.2.2 255.255.255.255 20.1.1.1 2
```

4. Configure the security policy on each router by performing the following actions:

- a. Specify a global lifetime of eight hours to be used with all dynamic policies by entering:

```
SETDefault -CRYPTO GlobalLifeTime = 8 hr
```

- b. Add a selector list to specify the traffic to which the policies apply by entering:

```
ADD -CRYPTO SelectorLIst s1 1 Include ANY 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values `0.0.0.0/` are used.

- c. Add a transform list that specifies the Phase 2 security association by entering:

```
ADD -CRYPTO TransformLIst t1 1 FRF-AES/256 FRF-SHA
```

This is the description of the protocols and transforms proposed when the security associations are negotiated.

- d. Define an IKE profile that describes the Phase 1 security association by entering:

```
ADD -CRYPTO IKEProfile 3 PreSharedKey AES/256 SHA Group5 1 dy
```

This is used by IKE to secure its own negotiation and is not used to secure the data traffic.

- e. Add dynamic-key security policies to the virtual ports configured in [step 1](#) by entering the following commands:

```
ADD !v101 -CRYPTO DynamicPOLicy dp1 1 FRF17 s1 t1 pc 10.101.2.1
```

```
ADD !v103 -CRYPTO DynamicPOLicy dp2 2 FRF17 s1 t1 pc 10.101.2.2
```

The `PreConnect` option is used to specify that the policies are pre-connect policies. The `WinSize` option is used to configure the anti-replay window size.

- f. On each router, define pre-shared keys to be used when establishing IKE security associations by using the preshared key authentication method and entering:

```
ADD -CRYPTO PreSharedKey 10.101.2.1 "12345"
```

```
ADD -CRYPTO PreSharedKey 10.101.2.2 "56789"
```



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- g. On each router, enable CRYPTO control on the CRYPTO ports by entering:

```
SETDefault !v101 -CRYPTO CONTrol= e
```

```
SETDefault !v103 -CRYPTO CONTrol= e
```

- h. Check the CRYPTO configuration by entering:

```
SHOW -CRYPTO CONFIguration
```

Configuring Site Routers (Both devices should be configured identically, except where noted in the procedure.)

5. Configure virtual ports. See [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).

6. Configure the system IP addresses by performing the following actions:

- a. On the site router 1, enter:

```
SETDefault -IP SystemIP = 10.101.2.1
```

- b. On the site router 2, enter:

```
SETDefault -IP SystemIP = 10.101.2.2
```



IMPORTANT: If you change the system IP address on an EOS device, the SNMPv3 credentials must be reestablished with that device. If the SNMPv3 users are configured on the device prior to the system IP address change, you must issue the `ResetV3` command to reset the SNMPv3 data, then reconfigure the SNMPv3 users with the appropriate privilege levels. For more information about how to configure the SNMPv3 users, see [Predefined SNMPv3 USM Users Management on EOS Devices on page 670](#) and [General SNMPv3 USM Users Management on EOS Devices on page 679](#).

7. Add the static routes to the core routers by performing the following actions:

- a. On site router 1, enter the following commands:

```
ADD -IP ROute 10.1.253.9 255.255.255.255 20.1.3.2 2
```

```
ADD -IP ROute 10.1.253.17 255.255.255.255 20.1.3.2 2
```

- b. On site router 2, enter the following commands:

```
ADD -IP ROute 10.1.253.9 255.255.255.255 20.1.1.2 2
```

```
ADD -IP ROute 10.1.253.17 255.255.255.255 20.1.1.2 2
```

8. On each router, configure the security policy by performing the following actions:

- a. Specify a global lifetime of eight hours to be used with all dynamic policies by entering:

```
SETDefault -CRYPTO GlobalLifeTime = 8 hr
```

- b. Add a selector list to specify the traffic to which the policies apply by entering:

```
ADD -CRYPTO SelectorList sl 1 Include ANY 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values `0.0.0.0/0` are used.

- c. Add a transform list that specifies the Phase 2 security association by entering:

```
ADD -CRYPTO TransformList tl 1 FRF-AES/256 FRF-SHA
```

This is the description of the protocols and transforms proposed when the security associations are negotiated.

- d. Define an IKE profile that describes the Phase 1 security association by entering:

```
ADD -CRYPTO IKEProfile 3 PreSharedKey AES/256 SHA Group5 1 dy
```

This is used by IKE to secure its own negotiation and is not used to secure the data traffic.

- e. On site router 1, add dynamic-key security policies to the virtual ports configured in [step 5](#) by entering the following commands:

```
ADD !v101 -CRYPTO DynamicPOLicy dp1 1 FRF17 sl t1 pc 10.1.253.9
```

```
ADD !v101 -CRYPTO DynamicPOLicy dp2 2 FRF17 sl t1 pc 10.1.253.17
```

The `PreConnect` option is used to specify that the policies are pre-connect policies.

- f. On site router 2, add dynamic-key security policies to the virtual ports configured in [step 5](#) by entering the following commands:

```
ADD !v103 -CRYPTO DynamicPOLicy dp1 1 FRF17 sl t1 pc 10.1.253.9
```

```
ADD !v103 -CRYPTO DynamicPOLicy dp2 2 FRF17 sl t1 pc 10.1.253.17
```

The `PreConnect` option is used to specify that the policies are pre-connect policies.

- g. On site router 1, define pre-shared keys to be used when establishing IKE security associations by using the preshared key authentication method and entering the following commands:

```
ADD -CRYPTO PreSharedKey 10.1.253.9 "12345"
```

```
ADD -CRYPTO PreSharedKey 10.1.253.17 "12345"
```



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- h. On site router 2, define pre-shared keys to be used when establishing IKE security associations by using the preshared key authentication method and entering the following commands:

```
ADD -CRYPTO PreSharedKey 10.1.253.9 "56789"
```

```
ADD -CRYPTO PreSharedKey 10.1.253.17 "56789"
```



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- i. On site router 1, enable CRYPTO control on the CRYPTO ports by entering:

```
SETDefault !v101 -CRYPTO CONTrol= e
```

- j. On site router 2, enable CRYPTO control on the CRYPTO ports by entering:

```
SETDefault !v103 -CRYPTO CONTrol= e
```

- k. Check the CRYPTO configuration by entering:

```
SHow -CRYPTO CONFiguration
```

28.2.4.4


Configuring Routers in Fully Meshed Topology with a Dynamic Key

The following figure illustrates an example of a fully meshed topology between three routers that use the following:

- IPSec Tunnel mode for the tunnels.

- Dynamic Keys using IKE.
- Pre-shared keys, DES, MD5 for Phase 1 IKE Profile.
- ESP for encryption (192-bit AES) and authentication (MD5) for Phase 2 TransformList.
- RIP as the routing protocols over the tunnels.

Figure 156: Fully Meshed Topology Between Three Routers (Dynamic Key)



Procedure:

Configuring Router 1

1. Add an IPIP point-to-multipoint tunnel virtual port by entering:
`ADD !v1 -Port Virtual Port IPIP P2MP`
2. Assign an IP address to the local LAN interface by entering:
`SETDefault !1 -IP NETAddress = 11.0.0.1`
3. Assign an IP network address to the Internet interface by entering:
`SETDefault !2 -IP NETAddress = 101.1.1.1`
4. Assign an IP network address to the IPIP P2MP tunnel interface by entering:
`SETDefault !v1 -IP NETAddress = 50.0.0.1`
5. Specify the mappings of the peer Tunnel IP address to the peer Internet interface IP addresses by performing the following actions:
 - a. For router 2, enter:
`ADD -IP ADDRESS 50.0.0.2 ipip 102.1.1.1`
 - b. For router 3, enter:
`ADD -IP ADDRESS 50.0.0.3 ipip 103.1.1.1`
6. Add a default route to the Internet (assuming !2 is a PPP port) by entering:
`ADD -IP ROute 0.0.0.0 !2 1`
7. Enable IP routing by entering:
`SETDefault -IP CONTROL = ROute`
8. Configure the security information by performing the following actions:

- a. Add a selector list to specify the traffic to which the policies apply by entering:

```
ADD -CRYPTO SelectorList sl10 10 include any 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values 0.0.0.0/ are used.

- b. Add a transform list that specifies the Phase 2 SA by entering:

```
ADD -CRYPTO TransformList tl10 10 ESP-AES/192 ESP-MD5
```

This is the description of the security for the actual data packets over the tunnel.

- c. Define a common pre-shared key shared by all devices that need to communicate with each other by entering:

```
ADD -CRYPTO PreSharedKey 0.0.0.0/0 "secretkey"
```

The 0.0.0.0/0 mask is used to select all devices.



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- d. Define an IKE profile that describes the Phase 1 SA by entering:

```
ADD -CRYPTO IKEProfile 10 PreSharedKey des md5
```

This is used by IKE to secure its own negotiation, and is not used to secure the data traffic.

- e. Bind all the information together by entering:

```
ADD !v1 -CRYPTO DynamicPOLicy pol_ea10 10 Tunnel sl10 tl10
```

- f. Enable CRYPTO control on the tunnel port by entering:

```
SETDefault !v1 -CRYPTO CONTROL = e
```

- g. Check the CRYPTO configuration by entering:

```
SHow -CRYPTO CONFIGuration
```

9. Enable RIP Talk and Listen on the tunnel port by entering:

```
SETDefault !v1 -RIP CONTROL= (ta, li)
```

Configuring Router 2

10. Add an IPIP point-to-multipoint tunnel virtual port by entering:

```
ADD !v1 -Port VirtualPort IPIP P2MP
```

11. Assign an IP address to the local LAN interface by entering:

```
SETDefault !1 -IP NETaddress = 12.0.0.1
```

12. Assign an IP network address to the Internet interface by entering:

```
SETDefault !2 -IP NETaddress = 102.1.1.1
```

13. Assign an IP network address to the IPIP P2MP tunnel interface by entering:

```
SETDefault !v1 -IP NETaddress = 50.0.0.2
```

14. Specify the mappings of the peer Tunnel IP address to the peer Internet interface IP addresses by performing the following actions:

- a. For router 2, enter:

```
ADD -IP ADDRESS 50.0.0.1 IPIP 101.1.1.1
```

- b. For router 3, enter:

```
ADD -IP ADDRESS 50.0.0.3 IPIP 103.1.1.1
```

15. Add a default route to the Internet (assuming !2 is a PPP port) by entering:

```
ADD -IP ROUTE 0.0.0.0 !2 1
```

16. Enable IP routing by entering:

```
SETDefault -IP CONTROL = ROUTE
```

17. Configure the security information by performing the following actions:

- a. Add a selector list to specify the traffic to which the policies apply by entering:**

```
ADD -CRYPTO SelectorList sl10 10 include any 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values 0.0.0.0/ are used.

- b. Add a transform list that specifies the Phase 2 SA by entering:**

```
ADD -CRYPTO TransformList tl10 10 ESP-AES/192 ESP-MD5
```

This is the description of the security for the actual data packets over the tunnel.

- c. Define a common pre-shared key shared by all devices that need to communicate with each other by entering:**

```
ADD -CRYPTO PreSharedKey 0.0.0.0/0 "secretkey"
```

The 0.0.0.0/0 mask is used to select all devices.



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the pre-shared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- d. Define an IKE profile that describes the Phase 1 SA by entering:**

```
ADD -CRYPTO IKEProfile 10 PreSharedKey des md5
```

This is used by IKE to secure its own negotiation and is not used to secure the data traffic.

- e. Bind all the information together by entering:**

```
ADD !v1 -CRYPTO DynamicPOLicy pol_ea10 10 Tunnel sl10 tl10
```

- f. Enable CRYPTO control on the tunnel port by entering:**

```
SETDefault !v1 -CRYPTO CONTROL = e
```

- g. Check the CRYPTO configuration by entering:**

```
SHOW -CRYPTO CONFIGuration
```

18. Enable RIP Talk and Listen on the tunnel port by entering:

```
SETDefault !v1 -RIP CONTROL= (ta, li)
```

Configuring Router 3

19. Add an IPIP point-to-multipoint tunnel virtual port by entering:

```
ADD !v1 -PORT VirtualPort IPIP P2MP
```

20. Assign an IP address to the local LAN interface by entering:

```
SETDefault !1 -IP NETaddress = 13.0.0.1
```

21. Assign an IP network address to the Internet interface by entering:

```
SETDefault !2 -IP NETaddress = 103.1.1.1
```

22. Assign an IP network address to the IPIP P2MP tunnel interface by entering:

```
SETDefault !v1 -IP NETaddress = 50.0.0.3
```

- 23. Specify the mappings of the peer Tunnel IP address to the peer Internet interface IP addresses by performing the following actions:**

- a. For router 2, enter:**

```
ADD -IP ADDRESS 50.0.0.1 IPIP 101.1.1.1
```

- b. For router 3, enter:**

```
ADD -IP ADDRESS 50.0.0.2 IPIP 102.1.1.1
```

- 24. Add a default route to the Internet (assuming !2 is a PPP port) by entering:**

```
ADD -IP ROUTE 0.0.0.0 !2 1
```

- 25. Enable IP routing by entering:**

```
SETDefault -IP CONTROL = ROUTE
```

- 26. Configure the security information by performing the following actions:**

- a. Add a selector list to specify the traffic to which the policies apply by entering:**

```
ADD -CRYPTO SelectorList sl10 10 include any 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values 0.0.0.0/ are used.

- b. Add a transform list that specifies the Phase 2 SA by entering:**

```
ADD -CRYPTO TransformList tl10 10 ESP-AES/192 ESP-MD5
```

This is the description of the security for the actual data packets over the tunnel.

- c. Define a common pre-shared key shared by all devices that need to communicate with each other by entering:**

```
ADD -CRYPTO PreSharedKey 0.0.0.0/0 "secretkey"
```

The 0.0.0.0/0 mask is used to select all devices.



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- d. Define an IKE profile that describes the Phase 1 SA by entering:**

```
ADD -CRYPTO IKEProfile 10 PreSharedKey des md5
```

This is used by IKE to secure its own negotiation and is not used to secure the data traffic.

- e. Bind all the information together by entering:**

```
ADD !v1 -CRYPTO DynamicPOLicy pol_ea10 10 Tunnel sl10 tl10
```

- f. Enable CRYPTO control on the tunnel port by entering:**

```
SETDefault !v1 -CRYPTO CONTROL = e
```

- g. Check the CRYPTO configuration by entering:**

```
SHOW -CRYPTO conf
```

- 27. Enable RIP Talk and Listen on the tunnel port by entering:**

```
SETDefault !v1 -RIP CONTROL= (ta, li)
```


28.2.4.5

Creating an IPv6 IPsec Tunnel with an IKEv2 Dynamic Key

The following figure illustrates an example with a single tunnel between two routers. This configuration has the following properties:

- IPsec IPv6 Tunnel mode for the tunnel.
- Dynamic Keys using IKEv2.
- Pre-shared keys, AES/256, SHA1 for Phase 1 IKE Profile.
- ESP for encryption (256-bit AES) and authentication (SHA1) for Phase 2 TransformList.

Figure 157: IPv4 in IPv6 Tunnel with IKEv2 (Dynamic Key)



Procedure:

Configuring Router 1

1. Enable IP routing by entering:
`SETDefault -IP CONTROL = R0ute`
2. Enable IPv6 routing on port 3 by entering:
`SETDefault !3 -IPv6 CONTROL = R0ute`
3. Assign an IP address to the local LAN interface by entering:
`SETDefault !1 -IP NETAddr = 10.1.251.51 255.255.255.0`
4. Assign an IPv6 network address to the Internet interface by entering:
`ADD !3 -IPv6 NETAddr fd00:0:03001::1`
5. Add an IPIIPv6 point-to-point tunnel virtual port by entering:
`ADD !v101 -Port VirtualPort IPIIPv6 fd00:0:03001::1 fd00:0:0:3002::101`

6. Assign an IP network address to the IPIPV6 tunnel interface by entering:

```
SETDefault !v101 -IP NETAddr = 172.24.24.5 255.255.255.252
```

7. Add an IPv6 route to Router 2 by entering:

```
ADD !3 -IPv6 ROUTe fd00:0:0:3002::101/128 fd00:0:0:3001::254 0
```

8. Add a route to the port 1 (!1) LAN of Router 2 by entering:

```
ADD -IP ROUTe 10.1.1.0 255.255.255.0 172.24.24.6 0
```

9. Configure the security information by performing the following actions:

- a. Add a selector list to specify the traffic to which the policies apply by entering:

```
ADD -CRYPTO SelectorList sl10 10 include any 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values 0.0.0.0/ are used.

- b. Add a transform list that specifies the Phase 2 SA by entering:

```
ADD -CRYPTO TransformList tl10 10 ESP-AES/256 ESP-SHA
```

This is the description of the security for the actual data packets over the tunnel.

- c. Define a common pre-shared key shared by all devices that need to communicate with each other by entering:

```
ADD -CRYPTO PreSharedKey IKEv2 :: "secretkey"
```

The IPv6 address :: is used to select all devices.



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- d. Define an IKE profile that describes the Phase 1 SA by entering:

```
ADD -CRYPTO IKEProfile IKEv2 10 PreSharedKey AES/256 SHA
```

This is used by IKE to secure its own negotiation, and is not used to secure the data traffic.

- e. Bind all the information together by entering:

```
ADD !v101 -CRYPTO DynamicPOLicy dp1 10 IKEv2 Tunnel sl10 tl10
```

- f. Enable CRYPTO control on the tunnel port by entering:

```
SETDefault !v101 -CRYPTO CONTROL = e
```

- g. Check the CRYPTO configuration by entering:

```
SHow -CRYPTO CONFIguration
```

Configuring Router 2

10. Enable IP routing by entering:

```
SETDefault -IP CONTROL = ROUTe
```

11. Enable IPv6 routing on port 3 by entering:

```
SETDefault !3 -IPv6 CONTRo1 = ROUTe
```

12. Assign an IP address to the local LAN interface by entering:

```
SETDefault !v1 -IP NETAddr = 172.24.24.6 255.255.255.252
```

13. Assign an IPv6 network address to the Internet interface by entering:

```
ADD !3 -IPv6 NETAddr fd00:0:0:3002::101
```

14. Add an IPIPv6 point-to-point tunnel virtual port by entering:

```
ADD !v1 -Port VirtualPort IPIPv6 fd00:0:0:3002::101 fd00:0:0:3001::1
```

15. Assign an IP network address to the IPIPv6 tunnel interface by entering:

```
SETDefault !v1 -IP NETAddr = 172.24.24.6 255.255.255.252
```

16. Add an IPv6 route to Router 2 by entering:

```
ADD !3 -IPv6 ROUTe fd00:0:0:3001::1/128 fd00:0:0:3002::254 0
```

17. Add a route to the port 1 (!1) LAN of Router 2 by entering:

```
ADD -IP ROUTe 0.0.0.0 0.0.0.0 172.24.24.5 0
```

18. Configure the security information by performing the following actions:**a. Add a selector list to specify the traffic to which the policies apply by entering:**

```
ADD -CRYPTO SelectorList sl10 10 include any 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values 0.0.0.0/ are used.

b. Add a transform list that specifies the Phase 2 SA by entering:

```
ADD -CRYPTO TransformList tl10 10 ESP-AES/256 ESP-SHA
```

This is the description of the security for the actual data packets over the tunnel.

c. Define a common pre-shared key shared by all devices that need to communicate with each other by entering:

```
ADD -CRYPTO PreSharedKey IKEv2 :: "secretkey"
```

The IPv6 address :: is used to select all devices.



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

d. Define an IKE profile that describes the Phase 1 SA by entering:

```
ADD -CRYPTO IKEProfile IKEv2 10 PreSharedKey AES/256 SHA
```

This is used by IKE to secure its own negotiation, and is not used to secure the data traffic.

e. Bind all the information together by entering:

```
ADD !v1 -CRYPTO DynamicPOLicy dp1 10 IKEv2 Tunnel sl10 tl10
```

f. Enable CRYPTO control on the tunnel port by entering:

```
SETDefault !v1 -CRYPTO CONTrol = e
```

g. Check the CRYPTO configuration by entering:

```
SHOW -CRYPTO CONFiguration
```


28.2.4.6**Creating a Hub and Spoke Topology**

The following figure illustrates an example of a hub and spoke topology between three routers. This configuration has the following properties:

- L2TP or PPTP for tunnels.
- IPSec Transport mode.
- Dynamic Keys using IKE.

- Phase 1 IKE Profile using pre-shared keys, DES, MD5.
- Phase 2 TransformList using ESP for encryption (128-bit AES).
- OSPF as the routing protocols over the tunnels.

Figure 158: Hub and Spoke Topology Between Three Routers



You can use this procedure to configure Routers 1, 2, and 3 in a hub and spoke topology. All three devices should be configured identically, except where noted in the following procedure.

Procedure:

1. Configure the PTPP or L2TP tunnels. See [L2Tunnel Connections Configuration on page 535](#).
2. Configure the routing policies by performing the following actions:
 - a. Add a default route to the Internet (assuming !2 is a PPP port) by entering:


```
ADD -IP R0ute 0.0.0.0 !2 1
```
 - b. Enable IP routing by entering:


```
SETDefault -IP CONTrol = R0ute
```
3. Configure the security policy by performing the following actions:
 - a. Add a selector list to specify the traffic to which the policies apply by entering:


```
ADD -CRYPTO SelectorList sl10 10 include any 0.0.0.0/0 0.0.0.0/0
```

In this case traffic to all the IP addresses in the network is to be encrypted, so the values 0.0.0.0/ are used.
 - b. Add a transform list that specifies the Phase 2 SA by entering:


```
ADD -CRYPTO TransformList tl10 10 ESP-AES/128 ESP-MD5
```

This is the description of the security for the actual data packets over the tunnel.
 - c. Define a common pre-shared key shared by all devices that need to communicate with each other by entering:


```
ADD -CRYPTO PreSharedKey 0.0.0.0/0 "secretkey1234567"
```

The 0.0.0.0/0 mask is used to select all devices.



NOTE: If you bring up the device in compliance with Federal Information Processing Standards (FIPS) 140-2, you must use the `FipsPreShrdKey` parameter to add the preshared key. The syntax of the `FipsPreShrdKey` parameter requires you to enter the key twice for added security.

- d. Define an IKE profile that describes the Phase 1 SA by entering:

```
ADD -CRYPTO IKEProfile 10 PreSharedKey des md5
```

This is used by IKE to secure its own negotiation and is not used to secure the data traffic.

- e. Bind all the information together by entering:

```
ADD !2 -CRYPTO DynamicPOLicy pol_ea10 10 Xport s110 t110
```



NOTE: For PPTP/L2TP that uses CRYPTO transport mode, you must configure the `DynamicPOLicy` on the physical port, not the virtual port.

- f. Enable CRYPTO control on the CRYPTO port by entering:

```
SETDefault !2 -CRYPTO CONTrol= e
```

- g. Check the CRYPTO configuration by entering:

```
SHow -CRYPTO CONFIguration
```

4. Enable OSPF on the virtual ports by performing the following actions:

- a. For Router 1, enter the following commands:

```
SETDefault !v1 -Ospf CONTrol = e
```

```
SETDefault !v2 -Ospf CONTrol = e
```

- b. For Router 2, enter:

```
SETDefault !v1 -Ospf CONTrol = e
```

- c. For Router 3, enter:

```
SETDefault !v1 -Ospf CONTrol = e
```



NOTE: This assumes that port !2 is not running the OSPF and direct policy is not configured.

28.3

MD5 Authentication Configuration for OSPF and BGP

EOS now implements authentication of the OSPF and BGP routing protocols using the Message Digest 5 (MD5) algorithm. MD5 authentication for OSPF and BGP prevents fraudulent route updates from being received and ensures that OSPF and BGP devices receive information only from trusted sources. EOS uses MD5 signatures to protect OSPF and BGP route table integrity. The message digest is created using a key and a message, but the key itself is not sent.

You can create a set of MD5 keys by adding them on a per-port (OSPF) or per-peer (BGP) basis with OSPF service `SecretKey` parameter, or the BGP service `PeerAuthKey` parameter. For details, see the descriptions of these parameters “OSPF Service Parameters” and “BGP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. Keys do not expire and are available for use until they are deleted. EOS supports up to 32 MD5 keys for both OSPF and BGP.

When you add a key, you specify a key ID that is used to identify the specific key used to provide the MD5 digest for authentication when a message is received by the device’s OSPF neighbor or BGP peer.

- When an OSPF device receives a message, it compares the received MD5 digest with a calculated digest using the key ID from the received packet to access the appropriate key for MD5 validation.

- When a BGP device receives a message, it compares the received MD5 digest with a calculated digest using the currently active key. BGP compares the active key, followed by all other configured keys for that peer, until a match is found (or not).



NOTE: EOS implements TCP level authentication to secure BGP sessions as specified in RFC 2385. The MD5 digest is carried in a TCP segment and acts as a signature for that segment, incorporating information known only to the connection endpoints.

You can specify the active transmit key with the OSPF or BGP service `ActiveKey` parameter. For details, see the description of this parameter “OSPF Service Parameters” and “BGP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. If you do not specify an active key, the device sets the key with the lowest key ID as the active transmit key when the OSPF neighbor or BGP peer transitions from clear state to transit or secure state, and generates a syslog message.

Any key in the configured key may be used to receive at any time. The MD5 keys and corresponding keyset IDs must match across OSPF neighbors and BGP peers. The receiver uses the key ID value to determine which key to use for authentication.

28.3.1

MD5 Authentication Control States

An OSPF port or BGP peer can operate in one of three MD5 authentication control states:

Clear

In this state, the port/peer:

- Transmits unsigned/clear packets.
- Receives both signed and unsigned/clear packets.
- Does not exercise MD5 authentication on either transmitted or received packets.

Transit

This state enables initial provisioning, migration, and rollback of the MD5 authentication feature without associated downtime. In the transit state, the port/peer:

- Transmits signed packets and exercises MD5 authentication on these packets.
- Receives both signed and unsigned/clear packets and does not exercise MD5 authentication on these packets.

The transit state timer, configured with the OSPF or BGP service `AuthCont` parameter, specifies how long the device is in transit state. When the timer expires, the device transitions to secure state and begins receiving only signed messages.

Secure

In this state, the port/peer:

- Transmits signed packets and exercises MD5 authentication on these packets.
- Receives signed packets and exercises MD5 authentication on these packets.

You can configure the state of the OSPF port or BGP peer using the OSPF or BGP service `AuthCont` parameter. For details, see the description of this parameter in the “OSPF Service Parameters” and “BGP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. To view the current state of the OSPF port or BGP peer, you can use the following commands:

```
SHoW ! <port> -OSPF AuthCont
```

```
SHoW ! <Peer_ID> -BGP AuthCont
```

You can add keys regardless of the state of the OSPF or BGP device (clear, transit, or secure). However, you cannot delete the active key if the port/peer is in secure or transit state. When the device transitions from clear state to secure or transit state and no active key is specified, the device uses the key with the lowest key ID and generates a syslog message.

State Persistence Across Reboots

When an EOS device reboots, it reads the MD5 authentication control state and key information from flash. The following table describes the possible states on reboot based on whether or not the MD5 keys are found and what state is stored in flash.

Table 54: MD5 Authentication States after Reboot

State Stored in Flash	Keys Found?	Result on Reboot
Clear	Yes	Device starts in clear state.
Clear	No	Device starts in clear state.
Transit or Secure	No	Device starts in clear state and syslog message is generated indicating that the device rebooted in transit or secure state and no keys were found.
Transit	Yes	Device starts in transit state, transit state timer is started, and syslog message is generated indicating that the device is reentering transit state after reboot.
Secure	Yes	Device starts in secure state.
None	No	Device starts in clear state.

28.3.2

Configuring MD5 Authentication for OSPF

For details about the parameters used in this procedure, see “OSPF Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



IMPORTANT: All the configured parameters for the keys should match across all OSPF devices on the subnet.

Procedure:

1. Add keys by entering:

```
ADD !<port> -OSPF SecretKey <keyID> <"string"> <"string">
```

You can add up to 32 MD5 keys, using key IDs in the range 0-255.

Step example:

```
ADD !1 -OSPF SecretKey 9 "key1" "key1"
```

2. Specify the active transmit key by entering:

```
SETDefault !<port> -OSPF ActiveKey = <keyID>
```

If you do not specify an active transmit key, the device uses the key with the lowest key ID as the transmit key.

Step example:

```
SETDefault !1 -OSPF ActiveKey = 9
```

3. Verify that the configuration is correct and matches across the network by performing the following actions:

- a. On each OSPF neighbor for which you configured MD5 authentication, view the keys in encrypted format by entering:

```
SHow !<port> -OSPF SecretKeyEnc
```

- b. Compare the keys and verify that they match.
- c. On each OSPF neighbor for which you configured MD5 authentication, view neighbor status information by entering:

```
SHow !<port> -OSPF NeighborStatus
```

- d. Compare the values listed in the KeyID column to verify that the neighbors are receiving on the same key.



WARNING: It is important to determine that configurations are correct and match across the network before configured transit timer expires and moves the neighbors to secure state. A misconfiguration or mismatched keyset can cause OSPF to fail and bring the network down.

4. Put the device in transit state and set the transit state timer by entering:

```
SETDefault !<port> -OSPF AuthCont = MD5Transit <number> <TransitMinutes / TransitHours / TransitDays>
```



WARNING: If you upgrade or migrate a system, ensure that you configure transit state timers long enough to accommodate the upgrade/migration. If the transit state timer(s) expire before the upgrade/migration is complete, OSPF fails and the network goes down.

Step example:

```
SETDefault !1 -OSPF AuthCont = MD5Transit 6 TransitHours
```

The device transitions to secure state when the transit state timer expires. If your upgrade/migration finishes much earlier than anticipated, you can set the AuthCont parameter to MD5Secure to transition to secure mode earlier.

28.3.3

Configuring MD5 Authentication for BGP

For details about the parameters used in this procedure, see “BGP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



WARNING: All the configured parameters for the keys should match across the BGP peers.

Procedure:

1. Add keys by entering:

```
ADD !<Peer_ID> -BGP PeerAuthKey <keyID> <"string"> <"string">
```

You can add up to 32 MD5 keys, using key IDs in the range 0-255.

Step example:

```
ADD !10.1.251.12 -BGP PeerAuthKey 9 "key1" "key1"
```

2. Specify the active transmit key by entering:

```
SETDefault !<Peer_ID> -BGP ActiveKey = <keyID>
```

Step example:

```
SETDefault !10.1.251.12 -BGP ActiveKey = 9
```

If you do not specify an active transmit key, the device uses the key with the lowest key ID as the transmit key and generate a syslog message.

3. Verify that the configuration is correct and matches across the network by performing the following actions:
 - a. On each BGP peer for which you configured MD5 authentication, to view the keys in encrypted format, enter:

```
SHoW !<Peer_ID> -BGP PeerAuthKeyEnc
```

- b. Compare the keys and verify that they match.
- c. On each BGP peer for which you configured MD5 authentication, view peer status information by entering:

```
SHoW !<Peer_ID> -BGP PEER
```

- d. Compare the values listed in the KeyID column to verify that the peers are receiving on the same key.



WARNING: It is important to determine that configurations are correct and match across the network before configured transit timer expires and moves the peers to secure state. A misconfiguration or mismatched keyset can cause BGP to fail and bring the network down.

4. Put the device in transit state and set the transit state timer by entering:

```
SETDefault !<Peer_ID> -BGP AuthCont = MD5Transit <number> <TransitMinutes / TransitHours / TransitDays>
```



WARNING: If you upgrade or migrate a system, make sure that you configure transit state timers long enough to accommodate the upgrade/migration. If the transit state timer(s) expire before the upgrade/migration is complete, BGP fails and the network goes down.

Step example:

```
SETDefault !10.1.251.12 -BGP AuthCont = MD5Transit 6 TransitHours
```

The device transitions to secure state when the transit state timer expires. If your upgrade/ migration finishes much earlier than anticipated, you can set the AuthCont parameter to MD5Secure to transition to secure mode earlier.

28.3.4

MD5 Authentication Statistics

You can view MD5 authentication statistics by issuing the following commands:

- For OSPF:

```
SHoW STATistics -OSPF
```

Authentication Error – Number of packets received that fail authentication.

- For BGP:

```
SHoW -TCP CONNEctions long
```

This command returns the following MD5 authentication statistics:

- Sent – The number of segments sent with the MD5 header option set.
- RecvdGood – The number of segments received with a matching MD5 digest.
- Mismatched – The number of segments received with a mismatched MD5 digest.
- Missing – The number of segments received for which the MD5 header option was expected but was not set.
- Ignored – The number of segments received with the MD5 header option set, but the receiver was not configured for the MD5 option.

28.3.5

MD5 Authentication Traps

EOS devices send the following traps pertaining to MD5 authentication for OSPF and BGP:

- If a device configured for OSPF MD5 authentication receives a packet that it cannot authenticate, it sends an `mnrOspfMD5KeyMismatch` trap. For details, see the description of this trap in “MIBs and Traps Supported on EOS Devices” in the *Enterprise OS Software Version 18.2 User Guide*.



NOTE: In the case where an OSPF neighbor is transmitting in secure mode and changes to clear mode, the receiving device issues an `mnrOspfMD5KeyMismatch` trap. Since the clear packet does not contain a keyID, this trap lists “-1” as the keyID.

By default, the `mnrOspfMD5KeyMismatch` trap has an SNMP throttle time of 1800 seconds (30 minutes). You can use the SNMP service `TrapThrottle` parameter to override this default value. For details, see the description of the `TrapThrottle` parameter in “SNMP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

- If a device configured for BGP MD5 authentication goes down as a result of an MD5 authentication failure, the device sends an `MNR_BGP_PEER_DOWN` trap. For details, see the description of this trap in “MIBs and Traps Supported on EOS Devices” in the *Enterprise OS Software Version 18.2 User Guide*.

28.4

Secure Shell (SSH) Support

EOS implements the Secure Shell (SSH) client/server architecture to protect network communication. SSH provides secure encrypted communications between two untrusted hosts over an insecure network. Secure copy (SCP) copies files between hosts on a network, using SSH for data transfer. It provides the same security as SSH.

EOS devices are used within the SSH architecture in the following ways:

- As SCP file transfer servers, supporting one secure file transfer operation at a time. SCP file transfer connection setup is noninteractive (requiring no password entry) and uses public key authentication.
- As SSH clients/servers, supporting four concurrent sessions.
 - If the device is functioning as an SSH client, SSH connection setup is interactive: the user must enter a password to establish the connection.
 - If the device is functioning as an SSH server, SSH sessions can be interactive or noninteractive, depending on whether the SSH client is another EOS device or not.
 - If the SSH client is another EOS device, SSH sessions are interactive: the user must enter a password to establish the connection.
 - If the SSH client is not another EOS device, SSH connection setup can be either interactive (requiring password entry) or noninteractive (not requiring password entry).



NOTE: For SSH sessions in which the SSH server is a RADIUS server and the SSH client is an EOS device configured with SystemIP, you must verify the following:

- The RADIUS server has a route back to the device SystemIP.
- The SystemIP address has been added as a client on the RADIUS server.

28.4.1

SSH Public Key Management

EOS uses public key authentication when a device is running SSH in either Client or Server mode. The public keys are managed in such a way that authentication can be performed securely and persistently. Two types of public keys are supported

SSH user public keys

EOS devices functioning as SSH servers use these keys to authenticate users attempting to establish a connection from an SSH client other than another EOS device. The SSH user public keys are stored on the device that is functioning as the SSH server. When a user from a non-EOS SSH client tries to

authenticate on the EOS device SSH server, it sends its public key in the SSH request message. The EOS device matches the stored public key with the one received to authenticate the user.

SSH server (known hosts) keys

EOS devices functioning as SSH clients use these keys to authenticate SSH servers. The SSH server (known host) keys are stored on the device that is functioning as the SSH client and used to authenticate the server when an SSH connection is made.

Generating Public Keys

You can generate the public keys on the device that you are using as the SSH server. This may or may not be an EOS device.

Generating Public Keys on an EOS Device

To generate public keys on an EOS device functioning as an SSH server, you can use:

```
GenSshKey [<algorithm> <bits>]
```

where *<algorithm>* is the algorithm used to generate the key (RSA or DSA) and *<bits>* is the length of the key, in bits (1024 or 2048).

For details, see “GenSshKey” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: Beginning with the ASTRO® 25 7.14 system release, EOS does not support 2048-bit DSA keys. If you want to generate a key using the DSA algorithm, you must set the key length to 1024 bits (GenSshKey DSA 1024).

Generating Public Keys on a Non-EOS SSH Server/Client

To generate public keys on a non-EOS SSH server or client, you can use either the UNIX OpenSSH Client or the Windows PuTTY Client.

- In UNIX, you can generate the key using the `ssh-keygen` command and store it in the user’s account at `~/.ssh/id_rsa` and `~/.ssh/id_dsa`. For details, see the OpenSSH Client documentation. The format for these files is `algorithm key username@client (IP Address)`.
- In Windows you can generate the key using the PuTTYgen program. The public key is displayed or stored on a file. For details, see the PuTTY Client documentation.



NOTE: You should edit the COMMENT to contain the user name and client IP address.

Loading Public Keys

You can load public keys onto an EOS device functioning as an SSH client or as an SSH server with a non-EOS SSH client using the CRYPTO service `ADD PublicKey` command. See “PublicKey” in the *Enterprise OS Software Version 18.2 Reference Guide* for details.

Verifying Public Keys

The following EOS commands are useful in verifying the presence of the public keys required to make an SSH connection:

- `SHoW -CRYPTO PublicKey`

You can issue this command on an EOS device functioning as an SSH client to display the SSH known host public keys for SSH servers that are stored on the device. In order to make an SSH connection between the SSH client device and the SSH server, the SSH known host public key for the server must be present on the SSH client device. If the SSH server is another EOS device, you can view the SSH known host public key for the SSH server by issuing the `ShowSshKey` command from the SSH server device, as described below.

You can issue this command on an EOS device functioning as an SSH server to display the SSH user public keys for users on non-EOS SSH clients which are stored on the device. In order to make an SSH

connection between the user on the non-EOS SSH client and the EOS SSH server, the SSH user public key for the non-EOS SSH client must be present on the SSH server device.

- **ShowSshKey**

You can issue this command on an EOS device functioning as an SSH server to display the SSH known host public key for that device. The SSH known host public key for the SSH server must be stored on the SSH client device in order to make a EOS device-to-EOS device SSH connection.

28.4.2

Starting an SSH Session (EOS SSH Client and EOS SSH Server)

If you use an EOS device as the SSH client and another EOS device as the SSH server, you can use this procedure to start an SecureShell (SSH) session.

Procedure:

1. Issue the `SHoW -SYS NetAccessTimer` command and ensure that the console and telnet timeout values are set to 15 minutes or higher.

For details about how to set the timeout values, see the description of the `NetAccessTimer` parameter in the “System Service Parameters” chapter of the *Enterprise OS Software Version 18.2 Reference Guide*.

2. On the EOS device that you use as the SSH server, generate the SSH server (known host) public key by entering:

```
GenSshKey {[ <algorithm>] [ <bits>]}
```

where:

<algorithm> specifies the algorithm to be used to generate the key (RSA or DSA)

<bits> specifies the key length (1024 or 2048)

For further details, see “GenSshKey” in the *Enterprise OS Software Version 18.2 Reference Guide*.

EOS does not support 2048-bit DSA keys. If you want to generate a key using the DSA algorithm, you must set the key length to 1024 bits (`GenSshKey DSA 1024`).

If both an RSA and a DSA key exist on the device, the RSA key takes precedence.



IMPORTANT: If the console or telnet session times out or is interrupted by the `Ctrl+C` command during SSH key generation, the key generation computation continues, but the generated keys are not stored. If you attempt to initiate new key generation during this computation process, the system returns an error message similar to the following:

Another instance of SSH key generation already in progress. Try again later.

Step example: `GenSshKey RSA 2048`

3. On the EOS device that you use as the SSH client, load the public key generated on the SSH server device in [step 2](#) into device memory by entering:

```
ADD -CRYPTO PublicKey <IPAddr> <RSA / DSA> "<key>" "<fingerprint>"
```

where:

<IPAddr> is the IP address of the SSH server device

<RSA / DSA> is the algorithm used to generate the key in [step 2](#)

"<key>" is the known host public key

"<fingerprint>" is the fingerprint

For further details, see “PublicKey” in the *Enterprise OS Software Version 18.2 Reference Guide*.

You can issue the `ShowSshKey` command on the SSH server device to display the `<key>` and `<fingerprint>`, then copy and paste this information into the `ADD -CRYPTO PublicKey` command on the SSH client device.

Step example:

- On the SSH server device (with IP address 10.79.138.232):

```
ssk
server ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEA3qH0qBCUF07dcFVa6CcJsn+FR4HbeRdY9TY76DMcPgDyx
PbVDqu+bLXjfz9i0CVSCzWU8jj9Ndj5Jcp4yXmvYoTBG+JzDLVmi/5Pt9ep8BwYDEfYKH64
x5vx5p0sRXn579ytzX+J1N44XUoF0p4sBtvFtwCdereusupHJiMfU=
The Server rsa key fingerprint is:
fa:64:75:80:f5:59:cd:9a:b0:99:7b:54:f3:21:0a:25
```

- On the SSH client device:

```
ADD -CRYPTO PublicKey 10.79.138.232 RSA
"AAAAB3NzaC1yc2EAAAABIwAAAIEA3qH0qBCUF07dcFVa6CcJsn+FR4HbeRdY9TY76DMcPgDy
xPbVDqu+bLXjfz9i0CVSCzWU8jj9Ndj5Jcp4yXmvYoTBG+JzDLVmi/5Pt9ep8BwYDEfYKH6
4x5vx5p0sRXn579ytzX+J1N44XUoF0p4sBtvFtwCdereusupHJiMfU="
"fa:64:75:80:f5:59:cd:9a:b0:99:7b:54:f3:21:0a:25"
```

4. On the SSH client device, verify that the SSH server known host public key was added by entering:

```
Show -CRYPTO PublicKey
```

Step example: `Show -CRYPTO PublicKey`

```
-----SSH Public Key Information-----
Server (Known Hosts) Public Keys :
ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEA3qH0qBCUF07dcFVa6CcJsn+FR4HbeRdY9TY76DMcPg
DyxPbVDqu+bLXjfz9i0CVSCzWU8jj9Ndj5Jcp4yXmvYoTBG+JzDLVmi/
5Pt9ep8BwYDEfYKH64x5vx5p0sRXn579ytzX+J1N44XUoF0p4sBtvFtwCdereusupHJiMf
U= 10.79.138.232
```

5. On both devices, enable SSH by entering:

```
SETDefault -SYS NetAccess = Ssh
```

6. On the SSH client device, initiate the SSH session by entering:

```
SSH <user>@<IP_address>
```

where:

`<user>` is the user name.

`<IP_address>` is the IP address of the SSH server device.

Step example: `SSH root@10.79.138.232`

7. When prompted, enter the password for the specified `<user>` to complete the SSH connection.

28.4.3

Starting and SSH Session (EOS SSH Client and Non-EOS SSH Server)

If you use an EOS device as the SSH client and a non-EOS device as the SSH server, you can use this procedure to start an SecureShell (SSH) session.

Procedure:

1. On the EOS device functioning as an SSH server, generate the SSH server (known host) public key by performing one of the following actions:

- In UNIX, use the UNIX OpenSSH Client to generate the key by using the `ssh-keygen` command and store it in the user's account at `~/.ssh/id_rsa` and `~/.ssh/id_dsa`. See the OpenSSH Client documentation for details.

The format for these files is: `algorithm key username@client (IP Address)`. For details, see the OpenSSH Client documentation.

- In Windows, generate the key by using the PuTTYgen program.

The public key is displayed or stored on a file. For details, see the PuTTY Client documentation.



NOTE: You should edit the COMMENT to contain the user name and client IP address.

2. On the EOS device, load the public key that you generated on the SSH server in [step 1](#) into device memory by entering:

```
ADD -CRYPTO PublicKey <IPAddr> <RSA / DSA> "<key>" "<fingerprint>"
```

where:

`<IPAddr>` is the IP address of the SSH server

`<RSA / DSA>` is the algorithm used to generate the key in [step 1](#)

`<key>` is the known host public key

`<fingerprint>` is the fingerprint

For further details, see "PublicKey" in the *Enterprise OS Software Version 18.2 Reference Guide*.

3. Enable SSH on the EOS device by entering:

```
SETDefault -SYS NetAccess = Ssh
```

4. On the EOS device, initiate the SSH session by entering:

```
SSH <user>@<IP address / Target name>
```

where:

`<user>` is the username

`<IP address / Target name>` is the IP address or hostname of the SSH server

5. When prompted to complete the SSH connection, enter the password for the specified `<user>`.
6. If the SSH connection fails, verify that the SSH server (known host) public key for the SSH server is stored on the SSH client device.

For details, see [SSH Public Key Management on page 505](#).

28.4.4

Starting an SSH Session (Non-EOS SSH Client and EOS SSH Server)

If you use a non-EOS device as the SSH client and an EOS device as the SSH server, you can use this procedure to start an SecureShell (SSH) session.

Procedure:

1. On the EOS device functioning as an SSH server, generate the SSH server (known host) public key by performing one of the following actions:
 - In UNIX, use the UNIX OpenSSH Client to generate the key by using the `ssh-keygen` command and store it in the user's account at `~/.ssh/id_rsa` and `~/.ssh/id_dsa`. See the OpenSSH Client documentation for details.

The format for these files is: `algorithm key username@client (IP Address)`. For details, see the OpenSSH Client documentation.

- In Windows, generate the key by using the PuTTYgen program.
The public key is displayed or stored on a file. For details, see the PuTTY Client documentation.



NOTE: You should edit the COMMENT to contain the user name and client IP address.

2. On the EOS device, load the public key that you generated on the SSH client in [step 1](#) into device memory by entering:

```
ADD -CRYPTO PublicKey <username@IPAddr> <RSA / DSA> "<key>" "<fingerprint>"
```

where:

- <username> is the user name
- <IPAddr> is the IP address of the SSH server
- <RSA / DSA> is the algorithm used to generate the key in [step 1](#)
- <key> is the known host public key
- <fingerprint> is the fingerprint

For further details, see “PublicKey” in the “CRYPTO Service Parameters” chapter of the *Enterprise OS Software Version 18.2 Reference Guide*.

3. Enable SSH on the EOS device by entering:

```
SETDefault -SYS NetAccess = Ssh
```
4. Initiate the SSH session from the non-EOS SSH client by using the syntax appropriate for your client type (Unix, Windows, or Linux).
5. If the SSH connection fails, verify that the SSH user public key for the non-EOS SSH client is stored on the SSH server device.

For details, see [SSH Public Key Management on page 505](#).

28.4.5

Stopping SSH Sessions

Interrupting an SSH session depends on whether the SSH client is an EOS device, a UNIX client, or a Windows client.

Procedure:

- To stop an SSH session from an EOS device, enter one of the following commands:
 - `li`
 - `logout`
 - `exit`
- To stop an SSH session from a UNIX OpenSSH client, enter one of the following commands:
 - `~.`
 - `li`
 - `logout`
 - `exit`
- To stop an SSH session from a Windows PuTTY client, enter one of the following commands:
 - `~`
 - `li`
 - `logout`

- exit

28.4.6

SCP Sessions

Since EOS devices are implemented as SCP servers, SCP sessions are not initiated from EOS devices, and there is no EOS command to initiate an SCP session. You can initiate the SCP session from your SCP client, using the syntax appropriate for your client type (Unix, Windows, or Linux).

28.5

Password Protection

With the implementation of the password protection feature, security parameters are stored in encrypted form, and you can configure those parameters in a secure manner by entering encrypted strings rather than clear text strings.

The protection scheme:

- Provides EOS parameters to configure secured parameters.
Parameters are provided in the AC, BGP, GGSN, LDAP, and OSPF services to support password protection.
 - AC service
The `SEcretEnc` parameter allows you to enter the RADIUS secret value in encrypted form. For details, see “AC Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.
 - GGSN service
The `AAASEcretEnc` parameter allows you to enter the GGSN RADIUS secret in encrypted form. For details, see “GGSN Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.
 - BGP Service
The `PeerAuthKey` parameter allows you to enter a BGP peer authorization key. You cannot view the key in clear format. Instead, you must use the `PeerAuthKEyEnc` parameter to view the key in encrypted format. For details, see “BGP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.
 - LDAP Service
The `BindDNEnc` parameter requires you to enter the password used to authenticate the bind request in encrypted form. In addition, you can use the optional `Encrypted` keyword with the `BindDN` parameter to enter the password in encrypted form. For details, see “LDAP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.
 - OSPF Service
The `SecretKey` parameter allows you to enter an OSPF secret key. You cannot view the key in clear format. Instead, you must use the `SecretKEyEnc` parameter to view the key in encrypted format. For details, see “OSPF Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.
- Displays passwords either in encrypted form or as a string of asterisks
Depending on the security policy of the network, you can configure security parameters with either clear text entry or using the new encrypted options. However no passwords are displayed in clear text; only encrypted code is displayed. Security parameters are also stored in a secure manner. The memory region used by the security parameters is protected and no local copy of that region is allowed in the software.
- Enables offline configuration of security parameters
For each service that supports password protection, a corresponding parameter was added for each security parameter to allow you to configure the parameter by entering an encrypted string rather than a clear text string. You can use the new secured parameters in the device configuration (`boot.cfg`) file.
- Supports security parameter zeroization of the RADIUS secret.

For details about zeroization, see [Zeroization of Critical Security Parameters on page 512](#).

28.6

Encrypted Values for Password-Protected Parameters

You can obtain the encrypted strings that you enter when configuring password-protected parameters in the following ways:

- You can use a software tool on a trusted machine to generate an encrypted string from a clear text input.
- You can use clear text configuration on one device, then issue the `SHoW` command for the corresponding encrypted parameter to obtain the encrypted value.

For example:

- You can configure the OSPF secret key parameter in clear text on one EOS device by using the `ADD -OSPF SecretKey` command.
- You can issue the `SHoW -OSPF SecretKEyEnc` command to display the secret key as an encrypted string.
- You can add the OSPF secret key in encrypted format on another EOS device (or edit it into the `boot.cfg` file) by using the `ADD -OSPF SecretKEyEnc` command.

28.7

Zeroization of Critical Security Parameters

In the EOS software, key encryption key (KEK) is used to encrypt and decrypt critical security parameters (CSPs) such as keys, secrets, and passwords to prevent unauthorized access to these parameters.

- The `KEKZeroize` command zeroizes or destroys all the KEK-related information stored on the device, including the KEK and the CSPs for its registered IPSec client and SSH.
- The `ZER0ize` command zeroizes all CSPs on the device, including those that are zeroized by the `KEKZeroize` command, as well as SNMPv3 passphrases and OSPF/BGP/PIM authentication keys.



WARNING: Exercise extreme caution when using the `ZER0ize` command as it destroys all keys and secrets on the device, causing all links that depend on those keys to go down. The `ZER0ize` command is intended for use in maintenance situations or when a security breach occurred, not during normal operation.

28.8

Device Encryption Troubleshooting

This section provides troubleshooting information for EOS device encryption.

28.8.1

CRYPTO MIB Trap Overview

EOS devices support the following traps from the CRYPTO MIB:

- `CryptoMibSACreated` – This routine trap indicates that a phase 2 security association was established. The device sends this trap whenever a manual or dynamic rekey occurs or encryption is enabled by using the following command:

```
SETDefault -CRYPTO CONTrol = Enabled
```

For dynamic keys, the `CryptoMibSACreated` trap is sent after phase 2 negotiation completes and the first security association of a bundle is created. For manual keys (statically configured), this trap is sent when the security association is configured.

Since security associations are created on a per-permanent virtual circuit (PVC) basis, it is typical to see a large number of `CryptoMibSACreated` traps.

- `CryptoMibSADeleted` – This routine trap indicates that a phase 2 security association was deleted. The device sends this trap whenever a manual or dynamic rekey occurs or encryption is disabled via the following command:

```
SETDefault -CRYPTO CONTrol = Disabled
```

For dynamic keys, this trap is sent after the last security association of a bundle is deleted. For manual keys (statically configured), this trap is sent when the security association is deleted.

The *TerminateReason* variable indicates the reason that the security association was deleted. For information about what action (if any) to take based on the value of the *TerminateReason* variable, see [Responses to Crypto MIB Traps on page 513](#).

Since security associations are created on a per-permanent virtual circuit (PVC) basis, it is typical to see a large number of `CryptoMibSADeleted` traps.

- `CryptoMibSAException` – This trap indicates that a security association exception was encountered. The *Exception* variable indicates the reason for the exception. For details about what action (if any) to take based on the value of the *Exception* variable, see [Responses to Crypto MIB Traps on page 513](#).

28.8.2

Responses to Crypto MIB Traps

The following table provides guidance about what action (if any) to take if you see one of the CRYPTO MIB traps.

Table 55: Crypto MIB Trap Troubleshooting

Trap Name	When does this trap occur?	What should I do if I see this trap?
SACreated	When a manual or dynamic rekey occurs or the CRYPTO service CONTrol parameter is set to Enabled.	Nothing. CryptoMibSACreated traps occur during rekeying. The number of these traps that the device sends depends on the number of secure links and the key lifetime. A lack of these traps indicates that rekeying is not occurring and may warrant further investigation.
SADeleted	When a manual or dynamic rekey occurs or the CRYPTO service CONTrol parameter is set to Disabled.	In general, nothing. CryptoMibSADeleted traps occur during rekeying. The number of these traps that the device sends depends on the number of secure links and the key lifetime. A lack of these traps indicates that rekeying is not occurring and may warrant further investigation. A TerminateReason variable value of NoResource could indicate a problem. Seeing a single CryptoMibSADeleted trap with a TerminateReason variable

Trap Name	When does this trap occur?	What should I do if I see this trap?
SAException	In various situations, as indicated by the <code>Exception</code> variable: (1) <code>other</code> – An exception occurred for a reason other than those described below. (2) <code>ahHashFailure</code> – N/A (3) <code>ahDuplicateSequence</code> – N/A (4) <code>ahOutOfRange</code> – N/A (5) <code>ahUnknownAlgorithm</code> – N/A (6) <code>espCryptoFailure</code> – This exception indicates that an error occurred in encryption/decryption. For example a packet may be corrupted, or the header padding or alignment may be incorrect. This is an internal problem that in most cases indicates a software defect or a capacity issue.	value of <code>NoResource</code> is generally not a cause for concern. However if you see several such traps, you should evaluate the context in which these traps are sent. If you see the traps in conjunction with another problem, for example, you may want to contact the Centralized Managed Support Operations (CMSO).
	(7) <code>espHashFailure</code> – This exception is seen on the receiving end of the security association. It indicates that the packet was decrypted but the hash comparison failed. It can be manifested as intermittent rekey failures on one side of the security association, which result in packet drops and are reported in the CRYPTO statistics as “Recv Err Unknown SPI” or “IKE up/FRF.17 EN,NT OP”.	The appropriate action depends on the value of the <code>Exception</code> variable: Seeing a single <code>CryptoMibSAException</code> trap with an <code>Exception</code> variable value of <code>espCryptoFailure</code> is generally not a cause for concern. However if you see several such traps, you should evaluate the context in which these traps are sent. If you see the traps in conjunction with another problem, for example, you may want to contact the Centralized Managed Support Operations (CMSO).
	(8) <code>espDuplicateSequence</code> – This exception pertains only to IPsec. It indicates that a duplicate sequence number was received. The duplicate se-	Seeing a single <code>CryptoMibSAException</code> trap with an <code>Exception</code> variable value of <code>espHashFailure</code> is generally not a cause for concern. However if you see several such traps, you should evaluate the context in which these traps are sent. If you see the trap in conjunction with CRYPTO statistics “Recv Err Unknown SPI” or “IKE up/FRF.17 EN,NT OP”, the problem may be due to incorrect configuration of the encryption key lifetime parameter. Contact the Centralized Managed Support Operations (CMSO) for a procedure to correct the misconfiguration and update the configuration file. If you see the traps in conjunction with another problem, you may want to contact the Centralized Managed Support Operations (CMSO). Seeing a single <code>CryptoMibSAException</code> trap with an <code>Exception</code> variable value of <code>espDuplicateSequence</code> is generally not a cause for concern. How-

Trap Name	When does this trap occur?	What should I do if I see this trap?
	quence is generally caused by a condition such as looping. In general this is a network rather than an encryption issue.	ever if you see several such traps, you should evaluate the context in which these traps are sent. If you see the traps in conjunction with another problem, for example, you may want to contact the Centralized Managed Support Operations (CMSO).
	(9) <code>espOutOfRange</code> – This exception pertains only to IPSec and is not seen in currently-deployed systems. It indicates that the IPSec sequence number is out of range.	Seeing a single <code>CryptoMibSAException</code> trap with an <code>Exception</code> variable value of <code>espOutOfRange</code> is generally not a cause for concern. However if you see several such traps, you should evaluate the context in which these traps are sent. If you see the traps in conjunction with another problem, for example, you may want to contact the Centralized Managed Support Operations (CMSO).
	(10) <code>espUnknownAlgorithm</code> – This exception indicates that an unknown encryption algorithm was encountered. It is caused by misconfiguration or failed negotiation.	Check the configuration (<code>boot.cfg</code>) file to ensure that the encryption algorithm is configured correctly.
	(11) <code>policyInvalidSpi</code> – This exception indicates an IKE negotiation failure, which could be caused by a software defect or software failure. Under normal circumstances, you should not see this exception because the SPI value is automatically set to 0 in the configuration files.	Seeing a single <code>CryptoMibSAException</code> trap with an <code>Exception</code> variable value of <code>policyInvalidSpi</code> is generally not a cause for concern since IKE automatically retries negotiation. However if the security association does not come back up, then the link itself is down. In this case, check the configuration (<code>boot.cfg</code>) file and try rebooting the device.
	(12) <code>policySAExpired</code> – This exception indicates an IKE negotiation failure. The security association timed out, which could be the result of a re-key failure.	Seeing a single <code>CryptoMibSAException</code> trap with an <code>Exception</code> variable value of <code>policySAExpired</code> is generally not a cause for concern since IKE automatically retries negotiation. However if the security association does not come back up, then the link itself is down. In this case, check the configuration (<code>boot.cfg</code>) file and try rebooting the device.
	(13) <code>policyClearText</code> – This rare exception indicates an IKE negotiation failure, probably resulting from a software defect. The device failed to recognize the header indicating that the packet is encrypted.	Seeing a single <code>CryptoMibSAException</code> trap with an <code>Exception</code> variable value of <code>policyClearText</code> is generally not a cause for concern since IKE automatically retries negotiation. However if the security association does not come back up, then the link itself is down. In

Trap Name	When does this trap occur?	What should I do if I see this trap?
		this case, check the configuration (boot.cfg) file and try rebooting the device.
	(14) isakmpFailure – This exception indicates an IKE negotiation failure resulting from a pre-shared key mismatch or from one of the devices on either side of the security association being unreachable.	Seeing a single CryptoMibSAException trap with an Exception variable value of isakmpFailure is generally not a cause for concern since IKE automatically retries negotiation. However if the security association does not come back up, then the link itself is down. In this case, check the configuration (boot.cfg) file; check the pre-shared key; and try rebooting the device.

28.8.3

FRF.17, IPSec, and IKE Error Logging

EOS supports FRF.17, IPSec, and IKE error logging. These messages are logged to the local audit log and to any configured remote syslog servers.

To view the audit log, you must use:

```
SHoW -AuditLog LocalLog
```

To view the syslog, you must execute the appropriate commands from the syslog server.

Logging Level Configuration

EOS supports the logging level defined by the CRYPTO service LogLevel parameter:

- 0 – No CRYPTO message logging.



IMPORTANT: The default setting for the LogLevel parameter is 0 (no CRYPTO message logging). In order for EOS to log CRYPTO audit log messages, you must set the LogLevel parameter to a non-default value (1, 2, 3, 4, or 5).

- 1 – Log only authentication failures.
- 2 – Log other negotiation failures as well as authentication failures.
- 3 – Log Phase 2 session creation and deletion messages as well as authentication failures and other negotiation failures.
- 4 – Log IKE (Phase 1) session creation and deletion messages as well as Phase 2 session creation and deletion messages, authentication failures, and other negotiation failures.
- 5 – Log session rekeying messages as well as IKE (Phase 1) session creation and deletion messages, Phase 2 session creation and deletion messages, authentication failures, and other negotiation failures.

FRF.17 Failures

EOS logs FRF.17 failures in the following format:

- Date/time
- SPI value
- Source and destination IP addresses
- Port

- Reason for discard

If some of the above information cannot be determined; for example, when a received packet contains an unknown SPI, a short form of the error log is written out containing the following information:

- Date/time
- Port
- Reason for discard

Table 56: FRF.17 Syslog Messages and Reasons for Packet Discard

Syslog Message	Reason for FRF.17 Discard
auth fail	FRF.17 authorization failure
crypto fail	FRF.17 encryption failure
unknown SPI	Unknown SPI
unknown Peer	Unknown peer
invalid pkt	Invalid packet
FRF17 en nt op	FRF.17 enabled but not operational
FRF17 dis	FRF.17 disabled
no IKE Info	No IKE information available
clear text pkt	Clear text error
misc fail	Other miscellaneous FRF.17 error

IPSec Error Logging

EOS logs IPSec failures in the following format:

- Date/time
- SPI value
- Source and destination IP addresses
- Sequence number
- Reason for discard

Table 57: IPSec Syslog Messages and Reasons for Packet Discard

Syslog Message	Reason for IPSec Discard
corrupt hdr	Corrupted ESP or AH header
auth fail	Authentication failure
crypto fail	Encryption failure
unknown SPI	Invalid SPI in incoming packet
invalid IP	Invalid IP in incoming packet
dupl pkt	Duplicate packet (possible replay - repeated sequence number)
policy disc	Policy is discarded
bad packet	Bad packet format

Syslog Message	Reason for IPSec Discard
bad padlen	Bad packet payload length
queue overflow	Queue overflowed
seq num cycled	Sequence number cycled
clear text err	Error in clear text
no SAs	No security associations available
enqueue err	Queue overflow
misc fail	Other miscellaneous IPSec error
icv mismatch	Authentication failure due to mismatch in ICV (integrity check value). The ICV is a variable-length field computed over the ESP header, Payload, and ESP trailer fields. An ICV mismatch means that the ICV value computed by the processing node (router, host) over the above fields of the packet does not agree with the value that is contained in the ICV field of the packet originally computed by the source.

In addition, EOS logs the following informational IPSec syslog messages:

- sa up (security association up)
- sa down (security association down)
- enter transit state
- exit transit state
- spi not equal to Transmit SA (SPI doesn't match transmit security association)

These messages have the same format as syslog messages, but they are informational warnings only and do not result in packet discard.

IKE Error Logging

EOS logs IKE failures in the following format:

- Date/time
- Source and destination IP addresses
- Cookie pair
- Name of payload responsible for the rejection, if applicable
- Reason for discard, as indicated by the syslog message

EOS now supports the IKE audit log / syslog messages listed in the following table.

Table 58: IKE Syslog Messages and Reasons for Packet Discard

Syslog Message	Reason for IKE Discard
malformed payload	Malformed payload received
invalid payload	Invalid payload received
doi not supp	Received unsupported digital object identifier
sit not supp	Received unsupported SIT
invalid spi	Invalid SPI in incoming packet

Syslog Message	Reason for IKE Discard
invalid prot id	Received invalid protocol ID
invalid trans id	Received invalid transform ID
attr mismatch	Phase 2 attribute mismatch
invalid cert enc	Received invalid certificate
missing payload	Payload missing in incoming packet
create payload err	Security association payload creation failed
auth fail	Authentication failure
encryption fail	Encryption failure
decryption fail	Decryption failure
diffie hellman err	Diffie-Hellman algorithm error
mem all err	Memory allocation error
software err	Phase 1 IKE internal error
enc key err	Encryption key error
selector mismatch	Selector mismatch
unknown peer	Unknown peer
misc fail	Other miscellaneous IKE error
PH1 SA Lookup Fail	Phase 1 security association lookup failure

28.9

CRYPTO Service Overview

The CRYPTO service provides three main areas of security: *authentication*, which validates the communicating parties; *integrity*, which makes sure the data was not altered; and *confidentiality*, which ensures the data cannot be intercepted and viewed. The CRYPTO service works with the existing Internet infrastructure using encapsulation. It secures a packet of data by encrypting it before sending it over the Internet. On the receiving end, a CRYPTO-compliant device decrypts the data.

EOS supports encryption of data over Ethernet and Frame Relay links using the following protocols under the CRYPTO service:

IPSec

Provides security at the network layer. Because it is integrated into IP itself, IPSec adds security to any link, regardless of what application you use.

IP Security Protocol (IPSec) tunnel mode

Tunnel mode security associations protect IP traffic forwarded by the EOS device on IPSec tunnel ports. EOS supports IPSec tunnel mode security associations configured on ports with either IPv4 or IPv6 addresses. The IKEv2 protocol is supported for security associations configured on IPv6 ports. When an MNR S6000 router or GGM 8000 gateway receives a packet that must be transmitted securely, the security policy database (SPD) triggers the IKEv1/IKEv2 protocol. If the tunnel endpoint is an IPv6 address, the IKEv2 protocol is used; if the tunnel endpoint is an IPv4 address, the IKEv1 protocol is used.

IPSec transport mode

Transport mode security associations protect traffic that is viewed on an end system from an IPSec perspective. For example, you can use transport mode security with PPTP/L2TP/IPIP tunnels or to serve network management traffic such as Telnet or SNMP.

FRF.17 with IKE negotiation

Provides security at the transport layer and adds security to Frame Relay links only. IKE negotiates the shared session keys used to encrypt or decrypt data messages with the peer and passes these keys to the FRF.17 protocol. IKE-negotiated session keys allow the FRF.17 protocol to further secure its data channel.

IKE negotiation is initiated in one of two ways:

Data-packet-triggered

When IP/CRYPTO receives a packet to be transmitted on a port with security enabled, it starts IKE negotiation and queues or drops the data packet. When IKE negotiation is complete, the queued data packets are transmitted.

Pre-connect

Peers for which a Frame Relay port is specified using the `ADD -CRYPTO PreSharedKey` command are identified as pre-connect peers. The EOS device attempts IKE negotiation as soon as it has a route to the IKE peer instead of waiting for a data packet to trigger IKE negotiation. Pre-connect IKE negotiation is required in failover architectures such as CWR.

28.9.1

Encapsulation Security Payload

Encapsulation security payload (ESP) is used to provide data confidentiality via encryption. For outbound traffic, ESP encrypts the IP payload and inserts an ESP header between the IP header and the payload. For inbound traffic, ESP decrypts the IP payload and removes the ESP header. You can apply ESP alone or with authentication headers.

The EOS software supports DES, 3DES, and AES encryption algorithms. DES is the US Data Encryption Standard (DES). 3DES has three stages, as indicated by the name. These stages include an encryption stage, a decrypting stage, and another encryption stage. 3DES keys must be at least 16 bytes long. The 3DES key is constructed using the first and third 8 bytes for the encrypt phase, and the second 8 bytes for the decrypt phase. AES encryption is supported with three key lengths: 128-bit, 192-bit, and 256-bit.

Key lengths are enforced when they are entered. Warning messages inform you when the entered key does not meet the requirements. Entered keys longer than the supported maximum length for the chosen encryption algorithm are truncated as necessary.



NOTE: Encryption **cannot** be exported without a legal export license.

28.10

Key Encryption Key Support

EOS supports Key Encryption Key (KEK), a mechanism that prevents certain long-lifetime and authentication keys (such as IKE pre-shared keys and SSH private keys) from being available on EOS devices in plain text. KEK is used by services such as CRYPTO and SNMPv3 to prevent access to keys both in memory and in flash. The KEK encrypts keys entered in plain text, stores the encrypted versions in a security file, and decrypts the keys when they are needed.

KEK is created and maintained in the EOS software by using the following two commands:

- `KEKGenerate` – Generates a new random KEK. You can use this command in the event that no KEK was configured on the device.
- `KEKZeroize` – Zeroizes or destroys all the KEK-related information stored on the device, including the KEK and the keys encrypted using KEK.

For details about the KEK commands, see “Commands” in the *Enterprise OS Software Version 18.2 Reference Guide*.

KEK Logs

When the `KEKGenerate` or `KEKZeroize` command is run, a message is returned to the console and a message is logged to `AuditLog`.

- When the `KEKGenerate` command is run, a message similar to the following is returned on the console:

```
Mon Jan 15 11:16:55 2007 New KEK was Generated
```

and the following message is logged to `AuditLog`:

```
650: New Key Encryption Key Generated
```

- When the `KEKZeroize` command is run, a message similar to the following is returned on the console:

```
Mon Jan 15 11:17:50 2007 KEK was zeroized
```

and the following message is logged to `AuditLog`:

```
651: New Key Encryption Key was Zeroized
```

Chapter 29

IP Security Options

This chapter describes how to configure your router to implement IP security options to protect datagrams at specified classification levels under the protection rules of specific authorities. These security features, which comply with RFC 1108, are necessary for any network implementing IP security options; for example, Department of Defense networks.

Your system can use Internet Protocol (IP) security options in an internetwork to:

- Transmit the common security labels from source to destination.
- Ensure that the route taken by the datagram is protected to the level required by all protection authorities indicated in the datagram.

This chapter also describes the use of source IP spoofing as a common type of security violation and provides the Internet Computer Emergency Response Team (CERT) recommendations for preventing this type of network attack.

The following terms are used in relation to IP security:

Basic security options

Identifies the U.S. classification level at which the datagram is to be protected and the authorities whose protection rules apply to each datagram.

Classification level

Specifies the U.S. classification level (top secret, secret, confidential, and unclassified) at which the datagram must be protected.

Classless Interdomain Routing (CIDR)

A method of using IP addresses without regard to traditional address classes to help solve the problem of the lack of class B network numbers.

Extended security options

Permits additional security labeling information beyond what is presented in the basic security option to meet the needs of additional registered authorities.

Label

Refers to the classification level and protection authority characteristics of a datagram.

Protection authority

Identifies the agency that specifies the protection rules for transmission and processing of information contained in the datagram. Examples of protection authorities include the following:

GENSER: the point of contact for this authority is the Designated Approving Authority per Department of Defense (DOD) 5200.29.

SIOP-ESI: The point of contact for this authority is the Department of Defense, Organization of the Joint Chiefs of Staff.

SCI: The point of contact for this authority is the Director of Central Intelligence.

NSA: The point of contact for this authority is the National Security Agency.

DOE: The point of contact for this authority is the Department of Energy.

Source IP spoofing

A common type of security violation in which an intruder accesses a protected domain by using the IP address of a trusted machine.

29.1

Configuring IP Security Parameters for End Systems

This procedure can be used to configure IP security parameters for end system configurations (IP routing is not used). The parameters related to IP security that you may need to configure depend on the type of networking devices involved and the amount of security required.



NOTE: This is an example of how to use the IP security option commands and is not a standard configuration procedure. Depending on the type of network security you require, your configuration procedure may differ from the one provided.

This procedure describes how to configure your system to transmit and receive datagrams with a TopSECRet classification level, how to accept datagrams with any combination of GENSER and SIOP-ESI protection authorities, and how to attach a TopSECRet GENSER label to datagrams originated by the system. The IP address was assigned to the system by using the SETDefault !0 -IP NETAddr command.

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- Your system has the following initial settings for the IP security option parameters:
 - SecLEVel is set to UNCLass for all ports.
 - SecAuthIn is set to GENSER for all ports.
 - SecCONTRol is set to NoEXTended for all ports.
 - SecLabelSys is set to UNCLass GENSER for all ports.
 - SecLabelValues is set to RFC1108.

If you use your system as an end system (the system receives packets and sends packets through the same interface; packets are not routed through the router to another interface), see the configuration example provided in this procedure.

If you use your system as an IP router, see the configuration example provided in [Configuring IP Security Options for Routers on page 524](#).

Procedure:

1. Configure all system ports to transmit and receive TopSECRet datagrams by entering:

```
SETDefault !0 -IP SecLEVel = TopSECRet
```
2. Specify the protection authorities that can be present in datagrams received on all ports by performing the following actions:
 - a. Change the default setting by entering:

```
DELeTe !0 -IP SecAuthIn GENSER
```
 - b. Set the Security Authorization SIOP-ESI protection by entering:

```
ADD !0 -IP SecAuthIn GENSER SIOP ANY
```
3. Configure the classification level and protection authority label for datagrams originated by the system by entering:

```
SETDefault !0 -IP SecLabelSys = TopSECRet GENSER
```
4. Enable the system to perform security processing of packets received from a file server by entering:

```
SETDefault -IP SecFileServer = Yes
```

The default for the `SecFileServer` parameter is `No`. For the system to communicate with the file server when IP security options are enabled, you must set the `SecFileServer` parameter to `Yes`. For more information, see the description of the `SecFileServer` parameter in the “IP Service Parameters” chapter in the *Enterprise OS Software Version 18.2 Reference Guide*.

5. Enable security options on the system by entering:

```
SETDefault -IP CONTROL = SECURITY
```

For information on how to check your configuration, see [Verifying IP Security Options on page 528](#).


29.2

Configuring IP Security Options for Routers

You can use this procedure to configure IP security parameters for router configurations. The parameters related to IP security that you may need to configure depend on the type of networking devices involved and the amount of security required.

 **NOTE:** This procedure is an example of how to use the IP security option commands and is not a standard configuration procedure. Depending on the type of network security you require, your configuration procedure may differ from the one provided.

Figure 159: IP Security Options Configuration Example



The illustrated configuration allows the following communications:

- PCs with the supercomputer
- Workstations with other workstations as well as the supercomputer
- Mainframes with the supercomputer

The devices perform the following operations:

- Mainframes generate the label “TSEC SIOPI SCI.”
- Workstations on port 2 and 4 generate the label “SEC GENSER” and require the basic option to be first in the IP header.
- PCs can neither generate nor receive labels. A default label of “CONF GENSER” is generated for them. This label is stripped before a datagram is sent on port 3.

- Supercomputer assigns labels based on host addressing. It needs to generate datagrams with labels “SEC GENSER” when communicating with workstations, “CONF GENSER” for PCs, and “TSEC SIOP SCI” for mainframes.

Port 1 connects to a supercomputer, which assigns labels based on host addressing. You want to configure the system to allow the supercomputer to communicate with the workstations on ports 2 and 4, the PCs, and the mainframes. You need to configure port 1 to transmit to and receive from the supercomputer datagrams with a TopSECRet, SECRet, or CONFidential classification level and with protection authority flags SIOP-ESI and SCI both set, or just GENSER set.

Port 2 connects to workstations, which require the basic security option to be the first option in the IP header. You want to configure the system to allow the workstations to communicate with the supercomputer and the workstations on port 4. You must configure port 2 to transmit to and receive from the workstations datagrams with a SECRet classification level and with the GENSER protection authority flag. You must also configure the port so that the basic security option is the first option in the IP header of datagrams transmitted on this port.

Port 3 connects to PCs. You want to configure the system to allow the PCs to communicate only with the supercomputer. You must configure port 3 to transmit and receive datagrams with a CONFidential classification level and with GENSER protection authority flag. Because the PCs can neither generate nor accept a security label, the system must attach a default label (CONFidential GENSER) to datagrams received on port 3 and destined for the supercomputer, and strip the label from datagrams destined for the PCs.

Port 4 connects to mainframes and workstations. You want to configure the system to allow the mainframes to communicate only with the supercomputer, and the workstations to communicate both with the supercomputer and the workstations on port 2. You must configure port 4 to transmit and receive datagrams with a TopSECRet or SECRet classification level, and with SIOP-ESI and SCI protection authorities or GENSER set.

Prerequisites:

Ensure that your system has the following initial settings for the IP security option parameters:

- SecLEVel is set to UNCLass for all ports.
- SecAuthIn and SecAuthOut are set to GENSER for all ports.
- SecLabelDefault is set to NONE for all ports.
- SecCONTrol is set to NoEXTended, NoBasicFirst, NoLabelAdd, and NoLabelStrip for all ports.
- SecLabelSys is set to UNCLass GENSER for all ports.
- SecLabelValues is set to RFC 1108.

Procedure:

Configuring port 1

1. Specify the range of security levels of the datagrams that can be transmitted and received on port 1 by entering:

```
SETDefault !1 -IP SecLEVel = CONFidential TopSECRet
```

The system can receive or transmit datagrams from port 1 with classification levels of CONFidential, SECRet, or TopSECRet.

2. Specify the protection authorities that can be present in datagrams received on port 1 by entering:

```
ADD !1 -IP SecAuthIn SIOP SCI
```

The system can receive datagrams from the network on port 1 with SIOP-ESI and SCI set, or just GENSER set. GENSER appears in the SecAuthIn table by default.

3. Specify the protection authorities that can be present in datagrams transmitted on port 1 by entering:

```
ADD !1 -IP SecAuthOut SIOP SCI
```

The system can transmit datagrams to the network on port 1 with SIOP-ESI and SCI set, or just GENSER set.

4. Configure a single classification level and protection authority label for datagrams originated by the system and transmitted on port 1 by entering:

```
SETDefault !1 -IP SecLabelSys = CONFidential GENSER
```

Any datagram generated by the system, including Internet Control Message Protocol (ICMP) messages, have this label when transmitted on port 1.

5. Configure the system so that a label is attached to datagrams before transmission over port 1 by entering:

```
SETDefault !1 -IP SecCONTROL = LabelAdd
```

You configure the SecCONTROL parameter because PCs cannot generate labels. The system must be configured to attach a label to a datagram destined for the supercomputer. The label that is attached to the datagram before transmission to the supercomputer is based on the value of the SecLabelDefault parameter, which you set on port 3 of the system.

Configuring port 2

6. Specify the security level of datagrams that can be transmitted and received on port 2 by entering:

```
SETDefault !2 -IP SecLEVel = SECRet
```

The system can receive or transmit datagrams from port 2 with classification level of SECRet.

7. Configure port 2 so that the basic security option is the first option in the IP header of datagrams transmitted on port 2 by entering:

```
SETDefault !2 -IP SecCONTROL = BasicFirst
```

For datagrams transmitted on port 2, the workstations require that the basic security option is the first option in the IP header.

8. Configure a single classification level and protection authority label for datagrams originated by the system and transmitted on port 2 by entering:

```
SETDefault !2 -IP SecLabelSys = SECRet GENSER
```

Any datagrams generated by the system, including ICMP messages, have this label when transmitted over port 2.

Configuring port 3

9. Specify the security level of datagrams that can be transmitted and received on port 3 by entering:

```
SETDefault !3 -IP SecLEVel = CONFidential
```

The system can receive or transmit datagrams by this port with classification level of CONFidential.

10. Configure port 3 to strip the security label from datagrams transmitted to the PCs by entering:

```
SETDefault !3 -IP SecCONTROL = LabelStrip
```

Because PCs cannot receive labels in datagrams, the system must strip the label before transmission on port 3.

11. Configure port 3 to attach a default label to datagrams received from PCs by entering:

```
SETDefault !3 -IP SecLabelDefault = CONFidential GENSER
```

PCs cannot generate or transmit labels; therefore, the system must attach a default label of CONFidential and GENSER. The datagram can then be properly routed to port 1 and the supercomputer.

12. Configure port 3 so that datagrams originated by the system and transmitted on port 3 do not have labels by entering:

```
SETDefault !3 -IP SecLabelSys = NONE
```

Because PCs cannot receive labels, you must set the `SecLabelSys` parameter to `NONE`.

Configuring port 4

13. Specify the security level of datagrams that can be transmitted and received on port 4 by entering:

```
SETDefault !4 -IP SecLeVeL = SEcRet TopSEcRet
```

The system can receive or transmit datagrams on this port with classification levels of `SEcRet` and `TopSEcRet`.

14. Specify the protection authorities that can be present in datagrams received on port 4 by entering:

```
ADD !4 -IP SecAuthIn SIOP SCI
```

The system can receive datagrams from the network on port 4 with `SIOP-ESI` and `SCI` set, or just `GENSER`. `GENSER` appears in the `SecAuthIn` table by default.

15. Specify the protection authorities that can be present in datagrams transmitted on port 4 by entering:

```
ADD !4 -IP SecAuthOut SIOP SCI
```

To transmit datagrams to the mainframes and workstations connected to port 4, datagrams must have the `SIOP-ESI` and `SCI` protection authority flags set in the security label of the datagram. Using this protection authority, the system can receive datagrams from the supercomputer and route them to the mainframes and workstations on port 4. With the `GENSER` authority (the default), the system can transmit datagrams between the supercomputer and workstations on ports 2 and 4.

16. Configure a single classification level and protection authority label for datagrams originated by the system and transmitted on port 4 by entering:

```
SETDefault !4 -IP SecLabelSys = SEcRet GENSER
```

Any datagrams generated by the system, including ICMP messages, have this label when transmitted on port 4.

Enabling IP security processing

17. Enable security options on the system by entering:

```
SETDefault -IP CONTroL = SECurity
```

After enabling security options, change the default of the `SecFileServer` parameter to `Yes` to ensure proper communication with the file server. For more information, see the `SecFileServer` parameter in “IP Service Parameters” chapter in the *Enterprise OS Software Version 18.2 Reference Guide*.

18. Display the configuration settings by entering:

```
SHoW -IP CONFIguration
```

For information on how to check your configuration, see [Verifying IP Security Options on page 528](#).

29.2.1

Configuring Extended Security Option Labels

For environments requiring extra security measures, you can add extended security labels to IP packets leaving specific ports. With this option, you can add a string to outgoing IP packets so that only specific hosts can accept the packets. You can use the extended security label options with the normal IP security options described earlier in this chapter, or you can use them independently.

The extended security label option is not required for the majority of configurations. If you use this option, be careful to configure it correctly to obtain the desired effect.

For information about the parameters, see the “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Procedure:

1. If you have basic IP security control enabled, make sure the `SecCONTRol` parameter is set to `EXTended` and `NoBasicFirst` for port 2 by entering:

```
SETDefault !2 -IP SecCONTRol = (EXTended, NoBasicFirst)
```

If you do not specify these values, the extended labels are discarded.

2. Specify the extended security label to be added to packets using:

```
SETDefault -IP SecLabelXtra = "<string>"
```

Specify the string as values of individual bytes given as a decimal number, with each byte separated by a slash (/). The total number of bytes must be a multiple of four, and the string must end with a slash. No syntax checking is performed; therefore, you must specify the string correctly.

3. Configure port 2 to add the label specified by the `SecLabelXtra` parameter to all packets sent over this port by entering:

```
SETDefault !2 -IP SecCONTRol = LabelXtraAdd
```

4. Configure port 1 to strip the extended label for all packets sent over this port by entering:

```
SETDefault !1 -IP SecCONTRol = LabelExtStrip
```

29.3

Verifying IP Security Options

You can use this procedure to check your configuration before using the IP security settings on your system

Procedure:

1. Verify the IP security configuration by entering:

```
SHow -IP CONFIguration
```

The system displays all IP configuration settings. If any of the IP security settings are incorrect, you can reconfigure them.

2. Check the IP security configuration by entering:

```
SecCheck
```

Because no port number is specified, the system checks all port configurations. If you specify a port number with this command, the system checks the configuration for that port.

This command is a diagnostic tool and does not check for all possible configurations. If the system finds misconfigurations, warning messages are displayed.

29.4

ICMP Error Messages

The following ICMP error messages may be generated as a result of security processing:

- ICMP Parameter Problem Missing Option, Type 12 Code 1, Pointer =130
- No security options in packet, but security options required on the port.
- ICMP Parameter Problem Missing Option, Type 12 Code 1, Pointer = malformed option
- Malformed security option, for example, an invalid label.

- ICMP Destination Unreachable Communication Administratively Prohibited, Type 3
- Code 9 for the network, Code 10 for the host. Security label out of range.

These ICMP error messages may be generated because of an incorrect configuration or may indicate an unauthorized break into a secure network. The network manager can decipher ICMP message codes using a network analyzer.

To display a count of ICMP error messages, you must use:

```
SHoW -SYS STATistics -IP
```

For a sample display and explanation of the entries, see [Statistics Displays on page 967](#).

29.5


Security Attacks Prevention on IP Routers

Source IP spoofing is a common type of security violation. The Internet CERT summarized the danger of how IP spoofing is used in network attacks. This section describes how to configure Enterprise OS software to provide security against this type of attack.

IP Spoofing

To generate an attack, the intruder creates packets with spoofed source IP addresses. In this type of attack, the intruder transmits packets from outside the “protected” domain that claim to be from a trusted machine inside the “protected” domain (for example, the packet contains the source IP address of a trusted machine). If the router is not configured to filter incoming packets whose source address is in the local domain, it forwards the traffic, and the targeted system may become compromised. A router generally forwards this traffic because it only examines the destination IP address—not the source IP address—when it makes its forwarding decision. The following figure illustrates the operation of a spoofed source IP address attack.

Figure 160: Spoofed Source IP Address Attack



Attacks are aimed at applications that use authentication based on source IP addresses. If successful, an attack leads to unauthorized user and possibly root access on the targeted system. It is important to note that the described attack is possible even if no reply packets can reach the attacker. Also, disabling source routing at the router does not provide protection from this type of attack.

Examples of configurations that are potentially vulnerable to attack include:

- Routers to external networks that support multiple internal interfaces
- Routers with two interfaces that support subnetting on the internal network
- Proxy firewalls where the proxy applications use the source IP address for authentication

Hijacking Tool

After intruders achieved root access on a system, they use a tool to dynamically modify the UNIX kernel. This modification allows them to hijack existing terminal and logon connections from any user on the system. In taking over existing connections, intruders can bypass one-time passwords and other strong authentication schemes by tapping the connection after the authentication is complete.

For example, a legitimate user may connect to a remote site through a logon or terminal session. An intruder can hijack the connection after the user completed authentication to the remote location. The site would now be compromised. Currently, the hijacking tool is used primarily on SunOS 4.1.x systems. However, system features that make this attack possible are not unique to SunOS.


Attack Prevention

To prevent this type of attack, the CERT Coordination Center recommends that network security personnel take these measures:

1. Install a filtering router that restricts the input to the external interface (known as an input filter) by not allowing a packet through if it has a source address from the internal network.
2. Filter outgoing packets that have a source address different from the internal network to prevent an attack originating from the local site.

The following figure illustrates the CERT recommendations.


Figure 161: CERT Recommended Filters



CERT recommends an alternative solution if a router does not support filtering on the in-bound side. The spoofed IP packets may be filtered by installing a second router between the original external interface (a) and the outside connection (b). This router can then be configured to block all packets that have a source address in the internal network on the outgoing interface (c) connected to the original router.

The following figure illustrates the alternative CERT recommendation.

Figure 162: Alternative CERT Configuration



29.5.1


Secure Configuration Solutions

The following examples illustrate how you can configure Enterprise OS software to support the CERT Advisory recommendations. Each of these examples assumes that the value of the `-IP FilterDefAction` parameter is configured to `Forward`. However, none of these examples prevent a source IP spoofing attack originating from the local site.

Noncontiguous IP Networks

The following figure illustrates a two-router solution where the internal network is configured with noncontiguous IP network numbers. The filters are installed on the border router, which can only have two interfaces. In a two-port router, an output filter on one port is equivalent to an input filter on the other port.

Figure 163: Noncontiguous IP Networks




You must add the following filters to the border router to prevent an external attack:

```
ADD -IP FilterAddrs 10.0.0.0/0.255.255.255 > 10.0.0.0/0.255.255.255 Discard
ADD -IP FilterAddrs 20.0.0.0/0.255.255.255 > 20.0.0.0/0.255.255.255 Discard
ADD -IP FilterAddrs 30.0.0.0/0.255.255.255 > 30.0.0.0/0.255.255.255 Discard
ADD -IP FilterAddrs 10.0.0.0/0.255.255.255 <> 20.0.0.0/0.255.255.255 Discard
ADD -IP FilterAddrs 10.0.0.0/0.255.255.255 <> 30.0.0.0/0.255.255.255 Discard
ADD -IP FilterAddrs 20.0.0.0/0.255.255.255 <> 30.0.0.0/0.255.255.255 Discard
```

This configuration prevents the external attack and allows the internal router to route traffic between networks 10.0.0.0, 20.0.0.0, and 30.0.0.0. This configuration also works for the cascade topology shown in the following figure.


Figure 164: Noncontiguous IP Networks (Alternative Topology)



Subnets on the Internal Network

The following figure illustrates a two-router solution when the internal network is configured with multiple subnets of the Class B network address, 130.5.0.0.

Figure 165: Subnets on the Internal Network



You must add the following filter to the border router to prevent an external attack:


```
ADD -IP FilterAddrs 130.5.0.0/0.0.255.255 > 130.5.0.0/0.0.255.255 Discard
```

This configuration prevents the external attack and allows the internal router to route traffic between all subnetworks of 130.5.0.0. In this example, a single filter can protect multiple subnets.

Multiple Contiguous IP Networks

The following figure illustrates a two-router solution where the internal network is configured with contiguous IP network numbers. Assume the service provider provided the subscriber with the Classless Interdomain Routing (CIDR) Protocol block 200.5.0.0/255.255.0.0.

Figure 166: Multiple Contiguous IP Networks



You must add the following filter to the border router to prevent an external attack:

```
ADD -IP FilterAddr 200.5.0.0/0.0.255.255 > 200.5.0.0/0.0.255.255 Discard
```

This configuration prevents the external attack and allows the internal router to route traffic between supernets of 200.5.0.0/255.255.0.0. In this example, a single filter can protect multiple contiguous IP networks numbers assigned as a CIDR block.

Alternative Two-Router Configurations


Various Motorola Solution routers can be configured for security. In some cases, however, routers from two different vendors may be optimal because a bug or back door that allows entry by a hacker in one vendor's code may not exist in the other vendor's code.

In many cases, the network topology can have the following characteristics:

- An external link to the Internet, which is a simple serial link to the network provider's router.
- The inside network consists of a few noncontiguous networks or subnets of a single network number.

The following figure illustrates this common configuration. The external router is configured with the required filters. The external router is also configured with a default route pointing to the Internet. The service provider installs static routes in their router that point to the customer's network. For this configuration, it is not necessary to run a routing protocol over the external link. If the network connectivity is more complex and you are connected using a multipoint technology such as Frame Relay, you can run the Border Gateway Protocol version 4 (BGP-4) to provide the required connectivity.


Figure 167: Two-Router Configuration



Firewall Configurations

Many firewall configurations require the use of two routers. The following figure illustrates a typical Internet firewall using two routers.

Figure 168: Internet Firewall



In this example, the routers create a packet filtering firewall while the bastion gateway functions as an application gateway firewall. In addition to using routers, creating a secure Internet firewall requires packet filtering and applications gateways. For information about filtering, see [Mnemonic Filtering Configuration on page 433](#).

Chapter 30

L2Tunnel Connections Configuration

This chapter describes how to configure a router as a tunnel terminator packet processor; how to configure a router as a tunnel initiator/terminator in a router to router configuration; and how to configure virtual leased lines with the Point-to-Point Tunneling Protocol (PPTP) and Layer 2 Tunneling Protocol (L2TP).

PPTP/L2TP defines a method for transferring Point-to-Point Protocol (PPP) datagrams through a tunnel over IP. Tunneling PPP does not change PPP but provides a vehicle by which PPP data units (PDUs) can be carried between two peers: a line server (LS) and a packet processor (PP). The LS-PP pair defines the endpoints of a PPTP/L2TP connection.

A PPTP/L2TP connection is defined by two parallel components: a control connection and a data pipe. Both operate between the same LS-PP pair. For PPTP, the control connection operates over TCP and passes call control and management packets over the TCP session. The data pipe operates over IP to transfer data packets encapsulated using Generic Routing Encapsulation Protocol Version 2 (GRE V2). For L2TP, both the control connection and the data pipe operate over the UDP session.

When a router adopts the PPTP/L2TP protocol, it functions as a packet processor. You can configure PPTP/L2TP tunnel connections between a router and a server (PPTP only) or a hub (acting as a line server). In this scenario, the router acts as a tunnel terminator, which only receives inbound calls. A router can also receive inbound calls from VPN-capable RAS clients to provide remote access services.


In addition, S6000 routers expand the use of PPTP/L2TP so a tunnel can be established between two peer routers. In this scenario, the router is able to issue outbound calls so either side can be the tunnel initiator or tunnel terminator. This is a router-to-router configuration, and both peers play the same role.

30.1

Configuring a PP Tunnel Terminator

When you use a router as a packet processor, you employ a hub as a line server, as shown in the following figure, and configure the PPPT/L2P connections between the router and the hub.

Figure 169: PPTP/L2TP Tunnel Connections Between Router and Hub



By default, the LS with any IP address are able to connect to the router using PPTP/L2TP. Flow control for all PPTP/L2TP sessions is disabled by default. If you need to enable flow control or if you want to restrict which line servers can have PPTP/L2TP connections with the router, you can configure an access list. After

you configured an access list, the router accepts PPTP/L2TP connections only from an LS whose IP address was specified. The flow control for each PPTP connection is configurable.



NOTE: The configuration for L2TP is identical to that for PPTP, except that L2TP has its own flow control mechanism embedded in the L2TP protocol. No flow control configuration is required for L2TP tunnels. You also need to define L2TPLocalUser and L2TPRemoteUser on both routers.

Procedure:

1. Enable the L2Tunnel service by entering one of the following commands:

- SETDefault -L2Tunnel CONTrol = Enabled Protocol = PPTP
- SETDefault -L2Tunnel CONTrol = Enabled Protocol = L2TP
- SETDefault -L2Tunnel CONTrol = Enabled Protocol = ALL

2. Configure the access list if you want to restrict the incoming peer tunnel initiators or if you want to enable the PPTP flow control by entering:

```
ADD -L2Tunnel AccessList <IP_Address> [<Subnet_Mask>]  
[FlowControl=Enabled | Disabled] [Protocol = PPTP | L2TP | ALL]
```

You can use this command if you configure the IP address or IP address range of acceptable tunnel initiators. By default, FlowControl is disabled.

Step example:

To enable a line server with an IP address of 129.213.48.6 to have a PPTP connection with the router, enter:

```
ADD -L2Tunnel AccessList 129.213.48.6 255.255.255.255 Protocol = PPTP
```

You can also create an access list for a group of IP addresses representing an entire subnet. To create an access list for an entire subnet 129.213.48.0 for both PPTP and L2TP, enter:

```
ADD -L2Tunnel AccessList 129.213.48.0 255.255.255.0 Protocol = ALL
```

3. Repeat [step 2](#) for each IP address or IP address range you want to add to the access list.
4. To display the access list that is configured, enter:

```
SHow -L2Tunnel AccessList
```

5. To remove a previously-configured access list entry, enter:

```
DElete -L2Tunnel AccessList <IP_Address>
```

6. If not yet configured, configure either SysCallerID or AuthRemoteUser of the virtual port you chose to be bound to the incoming virtual path by performing one of the following actions:

- Enter:

```
ADD !<port> -PORT VirtualPort SCID "<SysCallerID>"
```
- Enter the following commands:

```
ADD !<port> -PORT VirtualPort PPP  
ADD !<port> -PPP AuthRemoteUser ("<userid>", "<password>")
```

Step example:

If you want to create virtual port !V1 to connect to the LS box LS1, you should configure something similar to:

```
ADD !V1 -PORT VirtualPort SCID "LS1"
```

or

```
ADD !V1 -PORT VirtualPort PPP
```

```
ADD !V1 -PPP AuthRemoteUser ("LS1", "LS1PW")
```

7. If you used a SysCallerID in [step 6](#), configure SysCallerID of the router by entering:

```
SETDefault -SYS SysCallerID=" <string>"
```

Step example:

If the SysCallerID for this router is "ST", enter:

```
SETDefault -SYS SysCallerID="ST"
```

8. If you used AuthRemoteUser in [step 6](#), configure local user names and passwords for the port that is going to connect to the peer by entering:

```
SETDefault !<port> -PPP AuthLocalUser ("userid", "password")
```

For additional parameters that you can use optionally to configure the L2TP connection, see "L2T Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

If the user name and password for PPP negotiation are "ST" and "STPW", enter:

```
SETDefault !V1 -PPP AuthLocalUser=("ST","STPW")
```

9. To display statistics for all pptp connections that are currently active, enter:

```
SHow -L2Tunnel pptpSTATS
```

10. To display information about the state of each PPTP connection, enter:

```
SHow -L2Tunnel pptpSTATUS
```

11. To display information about the state and statistics of L2TP connections, enter the following commands:

```
SHow -L2Tunnel L2TPTunnels
```

```
SHow -L2Tunnel L2TPStats
```

12. If L2TP is used as the tunneling protocol, set the L2TP local and remote user name of the router by entering the following commands:

```
SETDefault -L2T L2TPLocalUser = ("userid", "Password")
```

```
ADD -L2T L2TPRemoteUser = ("userid", "Password")
```

The assigned L2TPLocalUser name is used as the "Host Name" during the tunnel establishment required by the L2TP protocol.

30.2

Configuring a Router-to-Router Tunnel Initiator/Terminator

When two routers are used in a router-to-router configuration, either router can issue dial commands. When a router issues a dial command, an outgoing call request message is sent to the peer. The peer responds by sending back an outgoing call reply message. A session is established within a PPTP/L2TP tunnel. (This setup mechanism is unlike the LS-PP scenario in which incoming call messages are exchanged.) Follow this procedure to configure a router as a tunnel initiator/terminator.

Procedure:

1. Enable the L2Tunnel service by entering one of the following commands:
 - SETDefault -L2Tunnel CONTROL = Enabled Protocol = PPTP
 - SETDefault -L2Tunnel CONTROL = Enabled Protocol = L2TP
 - SETDefault -L2Tunnel CONTROL = Enabled Protocol = ALL

2. Configure the access list for restricting the incoming peer tunnel initiators or enabling the PPTP flow control by entering:

```
ADD -L2Tunnel AccessList <IP_Address> [<Subnet_Mask>]  
[FlowControl=Enabled | Disabled] [Protocol = PPTP | L2TP | ALL]
```

You can use this command if you configure the IP address or IP address range of acceptable tunnel initiators. By default, FlowControl is disabled.

Step example:

To enable a line server with an IP address of 129.213.48.6 to have a PPTP connection with the router, enter:

```
ADD -L2Tunnel AccessList 129.213.48.6 255.255.255.255 Protocol =PPTP
```

You can also create an access list for a group of IP addresses representing an entire subnet.

To create an access list for an entire subnet 129.213.48.0 for both PPTP and L2TP, enter:

```
ADD -L2Tunnel AccessList 129.213.48.0 255.255.255.0 Protocol = ALL
```

3. Repeat [step 2](#) for each IP address or IP address range you want to add to the access list.
4. To display the access list that is configured, enter:

```
Show -L2Tunnel AccessList
```
5. To remove a previously-configured access list entry, enter:

```
DELeTe -L2Tunnel AccessList <IP_Address>
```
6. If not yet configured, configure either SysCallerID or AuthRemoteUser of the virtual port you chose to be bound to the incoming virtual path by performing one of the following actions:

- Enter:

```
ADD !<port> -PORT VirtualPort SCID "<SysCallerID>"
```
- Enter the following commands:

```
ADD !<port> -PORT VirtualPort PPP  
ADD !<port> -PPP AuthRemoteUser ("<userid>", "<password>")
```

Step example:

For example, if you want to create virtual port !V1 to connect to the LS box LS1, you should configure something similar to:

```
ADD !V1 -PORT VirtualPort SCID"LS1"
```

or

```
ADD !V1 -PORT VirtualPort PPP
```

```
ADD !V1 -PPP AuthRemoteUser ("LS1", "LS1PW")
```

7. If you used a SysCallerID in [step 6](#), configure SysCallerID of the router by entering:

```
SETDefault -SYS SysCallerID="<string>"
```

Step example:

If the SysCallerID for this router is "ST", enter:

```
SETDefault -SYS SysCallerID="ST"
```

8. If you used AuthRemoteUser in [step 6](#), configure local user names and passwords for the port that is going to connect to the peer by entering:

```
SETDefault !<port> -PPP AuthLocalUser ("userid", "password")
```

For additional parameters that you can use optionally to configure the L2TP connection, see "L2T Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

If the user name and password for PPP negotiation are "ST" and "STPW", enter:

```
SETDefault !V1 -PPP AuthLocalUser=("ST","STPW")
```

9. Configure a DialNoList for each dial out virtual port by entering one of the following commands:

- ADD !<port> -PORT DialNoList "<@IP_Address>" Type=PPTP
- ADD !<port> -PORT DialNoList "<@IP_Address>" Type=L2TP

The DialNoList contains the IP address of the physical interface of the peer router.

Step example:

If the IP address of the physical interface of the peer router or is 129.213.48.6 and it is intended to use PPTP tunneling and to connect it through virtual port !V1, enter:

```
ADD !V1 -PORT DialNoList "@129.213.48.6" Type=PPTP
```

10. To establish the PPTP connection, for each router enter:

```
Dial !<port>
```

If the physical interface is a dial-up line, you must configure your communication resources to use your dial-up lines before issuing the Dial command. For information about configuring communication resources for dial-up use, see [Protocol Reservation Configuration on page 383](#).

Step example:

To dial from virtual port !V1, enter:

```
Dial !<port>
```


30.3

Configuring a PPTP Virtual Leased Line

When you configure a router as a tunnel initiator/terminator (router-to-router scenario), the Dial command establishes the tunnel. In this scenario, tunnel usage may be disrupted due to heavy traffic or poor line quality. In reliability sensitive environments, you may choose to configure the PPTP/L2TP tunnel to operate as a leased line.

If the tunnel goes down, you can preconfigure the environment so that a backup line can be dialed, as shown in the following figure. All traffic is then sent through the backup line. While the backup line is in use, the router tries to reestablish the PPTP/L2TP tunnel. When the tunnel is reestablished, all the traffic is switched from the backup line to the tunnel. The backup line is then automatically torn down.

Figure 170: Virtual Leased Line and Backup Connection



The PPTP/L2TP virtual leased line is brought up automatically after it was configured, but you must configure parameters on both routers in order to bring up the line.

Procedure:

1. If not enabled, enable the L2Tunnel service and choose the desired protocol by entering:

```
SETDefault -L2Tunnel CONTROL = Enabled
```

2. Add a virtual leased line by entering:

```
ADD -L2Tunnel VLeasedLine <IP_Address>
```

where *<IP_Address>* is the IP address of the physical port of the peer.

On the peer router, there can be more than one physical port configured with different IP addresses. You should choose the IP address of the physical port through which the tunnel/TCP session is established.

In some environments, you may not know the IP addresses of the remote peers in advance—if the remote peer is located across the internet and its IP address is assigned dynamically by an Internet Service Provider (ISP), for example. In such environments, the IP address of the peer's physical port cannot be preconfigured in step 2. In these cases, you can configure a VLeasedLine IP address entry of 0.0.0.0, meaning that any IP address from a remote peer is accepted. There should be one 0.0.0.0 IP address entry for each incoming remote peer with an unknown IP address.

If you need to establish a virtual leased line over PPTP to the peer with physical port IP address 129.213.48.6, enter:

```
ADD -L2Tunnel VLeasedLine 129.213.48.6
```

3. To configure the backup line for virtual port under which PPTP is running, enter:

```
SETDefault !<port> -PORT DialCONTROL = DisasterRcvry
```


Step example:

If virtual port !V1 is intended to configured a backup line, enter:

```
SETDefault !V1 -PORT DialCONTRol = DisasterRcvry
```

4. Add the dial number for the backup line by entering:

```
ADD !<port> DialNoList "<phone-no>" [Type = Modem | Bri | Sw56]
```

Step example:

If the backup line is going from port V1 through modem to the peer's physical port with dialed phone number 9241234, enter:

```
ADD !V1 DialNoList "9241234" Type=Modem
```

5. If not yet configured, configure either SysCallerID or AuthRemoteUser of the virtual port that you choose to be bound to the incoming virtual path by performing one of the following actions:

- Enter:

```
ADD !<port> -PORT VirtualPort SCID "<SysCallerID>"
```

- Enter the following commands:

```
ADD !<port> - PORT VirtualPort PPP
```

```
ADD !<port> -PPP AuthRemoteUser ("<userid>", "<password>")
```



NOTE: At least one endpoint must use SysCallID to authenticate; the other endpoint can use either SysCallID or AuthRemoteUser.

Step example:

If you want to create virtual port !V1 to connect to the LS box LS1, you should configure something similar to one of the following:

```
ADD !V1 -PORT VirtualPort SCID"LS1"
```

or

```
ADD !V1 -PORT VirtualPort PPP
```

```
ADD !V1 -PPP AuthRemoteUser ("LS1", "LS1PW")
```

6. If you used a SysCallerID in [step 5](#), configure SysCallerID of the router by entering:

```
SETDefault -SYS SysCallerID="<string>"
```

Step example:

If the SysCallerID for this router is "ST", enter:

```
SETDefault -SYS SysCallerID="ST"
```

7. If you used AuthRemoteUser in [step 5](#), configure local user names and passwords for the port that is going to connect to the peer by entering:

```
SETDefault !<port> -PPP AuthLocalUser = ([ "<userid>", | None], "<password>")
```

Step example:

If the user name and password for PPP negotiation are "ST" and "STPW", enter:

```
SETDefault !V1 -PPP AuthLocalUser =("ST", "STPW")
```

8. Configure the parameters of the path from which the backup line is dialed.
9. If the physical interface is a dial-up line, configure your communication resources to use this line before issuing the DIal command.

For information about configuring communication resources for dial-up use, see [Protocol Reservation Configuration on page 383](#).



NOTE: You should not configure an access list in this configuration unless you know the range of the IP addresses that the ISP assigns to the remote peer.

10. If L2TP is used as the tunneling protocol, set the L2TP local user name of the router by entering:

```
SETDefault -L2T L2TPLocalUser = "Name"
```


The assigned L2TPLocalUser name is used as the "Host Name" during tunnel establishment which is required by the L2TP protocol.

30.4

Configuring Tunnel Switching

To set up tunnel switching, private LAN tunnel users must first establish the tunnels (inbound tunnels) with the router. As shown in the following figure, after the router authenticates the users, a second (outbound) tunnel is established between Router A and Router B. From then on, all the data packets from the remote users are switched to Router B.

Figure 171: Configuring Tunnel Switching



Procedure:

1. To configuring PPTP Tunnel Switching (Router A), perform the following actions:

- a. Enable the L2TP Service by entering:

```
SETDefault -l2t cont=e p=all
```

- b. Set up IP for physical port !1 on Router A by entering the following commands:

```
SETDefault !1 -ip net=126.1.1.1 255.255.0.0
```

```
SETDefault -ip cont=ro
```

- c. Set up TunnelSwitch on port !v1 on the router by entering:

```
ADD !v1 -po VirtualPort TunnelSwitch
```

- d. Add the DialNumberList to the TunnelSwitch port.

The server address of DialNumberList must be the IP address of the tunnel terminator.

Step example:

Assuming the IP address of port !1 on Router B is 126.1.2.2, enter:

```
ADD !v1 -po DialNumberList"@126.1.2.2" type=pptp
```

- e. Enable L2TP tunnel switching by entering:

```
ADD !v1 -po DialNumberList"@126.1.2.2" type=L2TP
```

- f. Add a remote user into the AuthRemoteUser of the TunnelSwitch port by entering the following commands:

```
ADD !v1 -ppp aru "demo" "demopw"
```

```
SETDefault !v1 -po cont=e
```

2. To configure tunnel switching by using the RADIUS Server, perform the following actions:

- a. Enable L2T service by entering:

```
SETDefault -l2t cont=e p=all
```

- b. Set up IP for physical port !1 of Router A by entering the following commands:

```
SETDefault !1 -ip net=126.1.1.1 255.255.0.0
```

```
SETDefault -ip cont=ro
```

- c. Set up RAS service on Router A using the IP address 126.1.1.2 for the RADIUS server by entering the following commands:

```
SETDefault -ras PrimAuthSrvr =126.1.1.2
```

```
SETDefault -ras PrimACcntSrvr =126.1.1.2
```

```
SETDefault -ras Secret ="your secret"
```

```
SETDefault -ras SecurityType=RADIUS
```

```
SETDefault -ras cont=e
```

- d. Follow the instructions that accompanied your RADIUS server to set up your RADIUS server.

- e. Define the following tunneling attributes for the users performing tunnel switching:

Tunnel_Type: PPTP or L2TP

Tunnel_Medium_Type: IP

Tunnel_Server_Endpoint : 126.1.2.2 (IP address of the tunnel terminator)

Tunnel_Preference: 0

If the value of Tunnel_Server_Endpoint is not a local IP address of Router A, Router A dynamically allocates a TunnelSwitch port.

30.5

Tunnel Security

L2TP implements two levels of tunnel security: tunnel authentication and tunnel encryption.

L2TP tunnel authentication

Also known as tunnel challenge is similar to the CHAP in PPP and serves the same purpose. However tunnel authentication is operated at the L2TP protocol level (the device level) while CHAP in PPP is operated at the PPP protocol level (the end user level). The L2TPLocalUser and L2TPRemoteUser parameters are equivalent to AuthLocalUser and AuthRemoteUser respectively in the PPP service.

L2TP tunnel encryption

Can be specified at the data packet level, control packet level or both.

- If the L2TPSecLevel parameter is set to None, no tunnel encryption is performed.
- If the L2TPSecLevel parameter is set to Data, only data packets are encrypted.

- If the `L2TPSecLevel` parameter is set to `Control`, only control packets are encrypted.
- If the `L2TPSecLevel` parameter is set to `Both`, both data and control packets are encrypted.

Chapter 31

Tunnel Route Shortcut (TRSC) Endpoints Configuration

This chapter describes how to configure a point-to-multi-point (P2MP) tunnel to discover a Tunnel Route Shortcut (TRSC) endpoint within a virtual private network (VPN). It also describes how to configure the Next Hop Routing Protocol (NHRP) to dynamically establish P2MP tunnels.

Data Forwarding in a Tunnel

IP-Over-IP tunnels allow you to build VPNs across distant geographical internets. One of the most common topologies for a VPN is a central site network that uses VPN tunnels to connect to remote sites; for example, a corporation that has a central office and remotely located branches.

Most often, the VPN is set up so there is a statically-configured tunnel between the central office network and each remote site. Using this configuration, the data must travel an extra hop—through the central office network—in order for two remote sites to communicate. As the VPN grows, this configuration can add many additional hops between the source and destination endpoints.

One way to make data forwarding more efficient in this topology is to make the entire VPN fully meshed. In this case, each pair of routers must be explicitly defined. Obviously, creating a fully meshed VPN is a time-consuming task that is not easily scalable.

Using TRSC and NHRP, however, you can configure the remote sites to automatically discover each other. Once discovered, these sites can establish a dynamic virtual tunnel between themselves. NHRP requires very little configuration and allows the routers to dynamically learn the best next hop node. By using the dynamic virtual tunnel, TRSC solves the problem of multiple hops and inefficient routed pathways.

31.1

NHRP

The TRSC Service uses NHRP to allow a tunnel endpoint to resolve the best next hop endpoint so that a dynamic tunnel can be established in a P2MP VPN. Before it can establish a TRSC, each endpoint router must terminate at least one statically-configured tunnel (the default path), and all participating nodes must be connected. Each router in the VPN works as both an NHRP client and an NHRP server.

When routing a packet, the routing protocol first looks up the next hop for a destination. If the routing protocol determines the next hop gateway is via an IPIP port (the VPN), the router attempts to retrieve the NHRP shortcut. If a shortcut is not available, the router forwards the packet to the default (static) tunnel path towards the destination.

In addition to forwarding the packet, the router kicks off an NHRP Resolution Request for the destination using the VPN path. The router issues an NHRP Resolution Request when the following are true:

- There is a data packet to be routed.
- The target destination is located on the VPN network, via an IPIP port route entry.
- NHRP shortcut information for the route is not available in the route table.

When a router receives a valid NHRP Resolution Reply, the router saves the NHRP information in its `ClientCache` table. It also adds the learned NHRP information to its local routing and address tables, using NHRP as the owner type.

When the router receives an NHRP Resolution Request, it does one of three things:

- Drops the packet.

The router first attempts to look up a route for the target destination. If a route does not exist, the router drops the packet without further action.

- Forwards the request.
If the route exists, and it is toward an IPIP port, and the TRSC Service is enabled on that outbound IPIP port, the router attaches its own address to the packet and forwards it.
- Replies to the request using an NHRP Request Reply.
The router determines that it is the next hop server (NHS) because the destination network is located via a non-IPIP port or a port that does not have TRSC enabled. The router (NHS) replies to the Resolution Request using an NHRP Resolution Reply along the reverse path from which the Resolution Request was received.
The NHS stores the source *<network_layer_address>* and *<media_layer_address>* pair in its address table. The NHS also makes an entry in its *ServerCache* table, and tracks the *ValidTime* and *HoldTime*.


31.2

Configuring NHRP

Next Hop Routing Protocol (NHRP) allows router endpoints in a P2MP VPN to discover and resolve the best next hop tunnel end-node, thereby establishing a dynamic virtual tunnel on demand, in a partially meshed network.

The following figure is an example of a central site network that is connected in a star configuration to two remote offices.

Figure 172: Central Site VPN with Two Remote Sites



Router 2 is the center site in the VPN. To establish a connected VPN without using NHRP, Router 2 requires two neighbor address configurations, one each for neighbor (Router 1 and Router 3).

Using NHRP, the routing table and address table of each router is enhanced with NHRP-learned shortcut information to help make the routing path more efficient. Because the NHRP shortcut is stored in the routing table, whenever the route is outgoing to an IPIP port, the router attempts to use its NHRP shortcut instead of a statically-configured path.

Port 1 of Router 1 (169.70.1.2), port 2 of Router 2 (139.70.1.2), and port 8 of Router 3 (149.70.1.2) are physical paths that connect to the global Internet. Layered on top of these paths, an IP-over-IP VPN (using the address 60.10.10.xx) must be created that, in turn, connects the three LAN segments, Lan_1, Lan_2, and Lan_3.

T1 and T2 are static IP-over-IP tunnels, manually configured by the system administrator. Without TRSC/ NHRP, data traffic flowing from PC1 to PC3 must traverse Router 2, because this router is the only connected tunnel node. After TRSC and NHRP are enabled on the routers (1, 2, and 3), traffic from PC1 to PC3 triggers the NHRP discovery mechanism and eventually establishes the dynamic virtual tunnel, T', between Router 1 and Router 3.

Procedure:

Configuring Router 1

1. Assign a NETaddress to the physical path that connects to the global Internet.

Step example:

```
SETDefault !1 -IP NETAddr=169.70.1.2
```

2. Enable IP routing by entering:

```
SETDefault -IP CONTrol=ROute
```

3. Create a P2MP virtual port, V1.

Step example:

```
ADD !V1 -Port VP IPIP P2MP
```

4. Assign a VPN NETaddress to V1.

Step example:

```
SETDefault !V1 -IP NETAddr=60.10.10.1
```

5. Enable OSPF operation on port V1, so that Router 1 can exchange routing information with other connected VPN nodes.

Step example:

```
SETDefault !V1 -OSPF CONTrol=Enable
```

6. Configure the neighboring tunnel endpoint address for VPN port V1.

Step example:

```
ADD -IP ADDRess 60.10.10.2 IPIP 139.70.1.2
```

7. Configure static and default routes on the VPN port and the physical port so that V1 is connected to the VPN and Router 2 is connected to the global Internet.

Step example:

```
ADD -IP ROute 0.0.0.0 169.70.1.1 3
```

```
ADD -IP ROute 60.10.10.3 60.10.10.2 5
```

8. Configure the private segment, Lan_1, then assign an IP NETaddress and enable OSPF on port 3, so the routing information in the private segment is advertised into the VPN, and routing information from the VPN can be learned.

Step example:

```
SETDefault !3 -IP NETAddr=10.0.0.2
```

```
SETDefault !3 -OSPF CONTrol=Enable
```

9. Enable TRSC to discover the shortcut to reach other tunnel endpoints.

Step example:

```
SETDefault !V1 -TRSC CONTrol=Enable
```

Configuring Router 2

10. Assign a NETaddress to the physical path that connects to the global Internet.

Step example:

```
SETDefault !2 -IP NETAddr=139.70.1.2
```

11. Enable IP routing by entering:

```
SETDefault -IP CONTrol=R0ute
```

12. Create a P2MP virtual port, V1.

Step example:

```
ADD !V1 -PO VP IPIP P2MP
```

13. Assign a VPN NETaddress to V1.

Step example:

```
SETDefault !V1 -IP NETAddr=60.10.10.2
```

14. Enable OSPF operation on port V1, so that Router 2 can exchange routing information with other connected VPN nodes.

Step example:

```
SETDefault !V1 -OSPF CONTrol=Enable
```

15. Configure the neighboring tunnel endpoint address for VPN port V1.

Step example:

```
ADD -IP ADDRess 60.10.10.1 IPIP 169.70.1.2
```

```
ADD -IP ADDRess 60.10.10.3 IPIP 149.70.1.2
```

16. Configure static and default routes on the VPN port and the physical port so that V1 is connected to the VPN and router Router 3 is connected to the global Internet.

Step example:

```
ADD -IP ROUTe 0.0.0.0 139.70.1.1 3
```

17. Configure the private segment, Lan_2. Assign an IP NETaddress and enable OSPF on port 1, so the routing information in the private segment is advertised into the VPN, and routing information from the VPN can be learned.

Step example:

```
SETDefault !1 -IP NETAddr=20.0.0.2
```

```
SETDefault !1 -OSPF CONTrol=Enable
```

18. Enable TRSC to discover the shortcut to reach other tunnel endpoints.

Step example:

```
SETDefault !V1 -TRSC CONTrol=Enable
```

Configuring Router 3

19. Assign a NETaddress to the physical path that connects to the global Internet.

Step example:

```
SETDefault !8 -IP NETAddr=149.70.1.2
```

20. Enable IP routing by entering:

```
SETDefault -IP CONTrol=R0ute
```

21. Create a P2MP virtual port, V1.

Step example:

```
ADD !V1 -PO VP IPIP P2MP
```

22. Assign a VPN NETaddress to V1.

Step example:

```
SETDefault !V1 -IP NETAddr=60.10.10.3
```


23. Enable OSPF operation on port V1, so that Router 1 can exchange routing information with other connected VPN nodes.

Step example:

```
SETDefault !V1 -OSPF CONTrol=Enable
```

24. Configure the neighboring tunnel endpoint address for VPN port V1.

Step example:

```
ADD -IP Address 60.10.10.2 IPIP 139.70.1.2
```

25. Configure static and default routes on the VPN port and the physical port so that V1 is connected to the VPN and Router 1 is connected to the global Internet.

Step example:

```
ADD -IP ROUTe 0.0.0.0 149.70.1.1 3
```

```
ADD -IP ROUTe 60.10.10.1 60.10.10.2 5
```

26. Configure the private segment, Lan_3, then assign an IP NETaddress and enable OSPF on port 4, so the routing information in the private segment is advertised into the VPN, and routing information from the VPN can be learned.

Step example:

```
SETDefault !4 -IP NETaddr=30.0.0.2
```

```
SETDefault !4 -OSPF CONTrol=Enable
```

27. Enable TRSC to discover the shortcut to reach other tunnel endpoints.

Step example:

```
SETDefault !V1 -TRSC CONTrol=Enable
```

31.2.1

Address and Routing Table Displays

The address table of configured Router 1 is as follows:

Address	Port	Media	Address	Owner
60.10.10.1	V1	Local		IPIP Static
60.10.10.2	V1	External	139.70.1.2	IPIP Static
60.10.10.3	V1	External	149.70.1.2	IPIP NHRP

The address table of Router 2 is as follows:

Address	Port	Media	Address	Owner
60.10.10.2	V1	Local		IPIP Static
60.10.10.1	V1	External	169.70.1.2	IPIP Static
60.10.10.3	V1	External	149.70.1.2	IPIP Static

The address table of Router 3 is as follows:

Address	Port	Media	Address	Owner
60.10.10.3	V1	Local		IPIP Static
60.10.10.2	V1	External	139.70.1.2	IPIP Static
60.10.10.1	V1	External	169.70.1.2	IPIP NHRP

The route entries on Router 1 with regard to PC1, PC2, and PC3 are as follows:

Destination	Mask	Gateway	Source
10.0.0.0	255.0.0.0	10.0.0.2	Connected
20.0.0.0	255.0.0.0	20.0.0.2	OSPF
30.0.0.0	255.0.0.0	60.10.10.3	OSPF

The routing entry on Router 2 with regard to PC1, PC2, and PC3 is as follows:

Destination	Mask	Gateway	Source

10.0.0.0	255.0.0.0	60.10.10.1	OSPF
20.0.0.0	255.0.0.0	20.0.0.2	Connected
30.0.0.0	255.0.0.0	60.10.10.3	OSPF

The routing entries on Router 3 with regard to PC1, PC2, and PC3 are as follows:

Destination	Mask	Gateway	Source
10.0.0.0	255.0.0.0	60.10.10.3	RIP
20.0.0.0	255.0.0.0	60.10.10.2	RIP
30.0.0.0	255.0.0.0	60.10.10.3	Connected

Chapter 32

Internet Firewalls Configuration

This chapter describes how to configure an Internet firewall, provides a conceptual overview of firewalls, and gives guidelines for operating and managing firewalls successfully.

A firewall allows users inside a private network to have outbound access, while restricting outside users from inbound access. The types of incoming and outgoing traffic can be identified as follows:

- Inside-originated request to an outside service
- Outside reply to the inside-originated request
- Outside-originated request to an inside response
- Inside reply to the outside request

Firewalls are typically constructed on bastion hosts and a multiprotocol router. The bastion hosts, usually UNIX, are configured to prevent them from being compromised by outsiders and to provide detailed logging of system activity for security monitoring. The host may serve as an externally accessible server for FTP, email, or the Internet.

Another common firewall component is a packet-filtering router. The packet-filtering router has extensive filtering capabilities to limit the type and direction of traffic that passes through it. The router usually is not the object of an attack, but can serve as a barrier to other, more desirable targets, or as the basis of a denial-of-service attack.

The following terms are used to describe the firewall feature:

Firewall

An EOS device and/or workstation with multiple network interfaces that controls and limits specific protocols, types of traffic within each protocol, types of services, and direction of the flow of information.

Secure logging

A method that takes an audit trail of system activity received from a bastion host and places it in a secure location.

IP source address spoofing

Spoofing uses “forged” source addresses to make an outside packet appear to come from the inside network so that the firewall allows it to have access to the private network. Spoofing works on the principle that, by default, EOS devices perform route lookups only on the destination IP address in each packet, paying no attention to the incoming interface of the source IP address.

Packet-Filtering Routers

The packet-filtering router can make a permit or deny decision for each packet it receives. The router examines each datagram to determine if it matches one of its packet filtering rules. The filtering rules are based on the packet header information that is made available to the IP forwarding process. This information consists of the following items:

- IP source address
- IP destination address
- Incoming interface of the packet
- Outgoing interface of the packet
- Encapsulated protocol (TCP, UDP, ICMP, or IP tunnel)
- TCP/UDP source port
- TCP connection Establishment packets

- TCP/UDP destination port
- ICMP message type

If a match is found and the rule permits the packet, the packet is forwarded according to the information in the routing table. If a match is found and the rule denies the packet, the packet is discarded, unless it is redirected by either `NextPort` or `NextHop`. If there is no matching rule, a user-configurable “default action” parameter determines whether the packet is forwarded or discarded.

The majority of Internet firewall systems use only a packet-filtering router. Other than the time spent designing the filters and configuring the router, little or no cost is required to implement packet filtering because the feature is included as part of standard router software releases. Because Internet access is usually provided over a WAN interface, there is little impact on router performance if traffic loads are moderate, and few filters are defined. A properly designed firewall using a packet-filtering router can be transparent to end users and applications, so it does not require specialized user training or require that specific software be installed on each host.

32.1

Firewall Filter Types

You can configure your firewall with the following types of filters:

- Service-independent filters
- Predefined (service-dependent) filters
- FTP filters
- TCP stateful filter
- Generic filters

32.1.1

Service-Independent Filters

Some types of Internet attacks are popular, but difficult or impossible to specify using only generic packet-header information. These attacks are generally service-independent and are difficult to specify because filtering rules require additional information that can only be learned by looking in the routing table. For these types of attacks, a separate control (`-FireWall CONTrol`) was created for each of the known attack types. The `-FireWall CONTrol` parameter works at the IP layer and allows you to control the following items:

- Filtering
- Source IP address spoofing
- TCP/IP tiny fragment attacks
- Packets that contain IP options such as source-route, record-route, and time-stamp
- IP-over-IP tunnels
- ICMP messages (protection against denial of service attacks)

For more information on the `-FireWall CONTrol` parameter, see “FireWall Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

32.1.2

Predefined (Service-Dependent) Filters

Some of the important services commonly used over Internet are: Telnet, SMTP, NNTP, FTP, HTTP, Gopher, and DNS. The parameters in the Firewall Service are designed for filtering that service; one parameter

controls incoming connections and one parameter controls outgoing connections. You can configure each parameter separately on a per-interface (per-port) basis.

 **NOTE:** The filtering operations performed by these service-dependent filters can also be performed using generic filters, which are described in [Generic Filters on page 554](#).

Service-dependent filters, and their related parameters, provide the following benefits:

- The parameters are designed for connection flows instead of packet flows. Connection flows deal with issues such as outbound FTP sessions or inbound Telnet sessions. Packet flows (the basic concept of generic filter rules) deal with inbound TCP packets or outbound UDP packets. A connection flow (such as inbound Telnet) usually involves bidirectional packet flows (such as both inbound TCP packets and outbound TCP packets).
- These parameters, such as `FTPOut`, automatically take care of both outbound and inbound TCP packets.
- These parameters allow a user who is not an expert on packet formats, or who may not be aware of the protocol details and port number schemes, to easily configure a packet-filtering router.
- Some services, such as DNS, are quite complex and they are difficult to specify using the generic filter rules.

Be aware that providing packet-filtering rules for a particular service does not mean the service is secured. Each service has its own weaknesses and security holes that are beyond the ability of a packet-filtering router to control. In general, permitting outbound connections is very safe. Outbound connections are connections initiated from internal networks to the Internet; they do not permit Internet-initiated connections into internal networks. Permitting inbound connections is much more risky. Consider using application-level proxy services that further enhance your firewall.

TCP-based services such as Telnet, FTP, SMTP, NNTP, DNS, Gopher, and Archie are much safer than non-TCP-based services. Avoid using non TCP-based services on your Internet connections. Because of the CPU requirements of `DenySrcSpoofing`, we recommend that you turn this parameter on only where absolutely required. Interfaces receiving external traffic should have `DenySrcSpoofing` turned on to repel source spoofing attacks. Internal interfaces may not require this type of checking.

32.1.3

FTP Filters

FTP is a TCP-based service; it is unusual because it uses two or more simultaneous TCP connections. The first TCP connection is initiated from client to server. This connection, usually called the *command channel*, carries commands and replies. The second TCP connection, usually called the *data channel*, is dedicated to transferring data. The second TCP connection is made in two ways: regular FTP and passive FTP.

Regular FTP

Occurs when a client issues a `PORT` command on the command channel to the server and the server opens the data TCP session. The port number of the client's desired data TCP socket is embedded in the `PORT` command.

Passive FTP

Occurs when a client issues a `PASV` command and, if the server responds positively, the client initiates the data TCP session. The TCP port number is embedded in the command and reply.

The `FTPIn` parameter and the `FTPOut` parameter understand both forms of FTP sessions. The FTP filters permit server-to-client data TCP connections when they detect the `PORT` command. They also permit a client-to-server data TCP connection when they detect a `PASV` command. The TCP port number is extracted from the `PORT` (or `PASV`) command so that only a specific data connection is allowed; no persistent holes in the firewall occur.

An FTP session may involve several data TCP connections, therefore, the `FTPIn` parameter and `FTPOut` parameter constantly monitor the active command channels for `PORT` and `PASV` commands and readjust their permissions window accordingly.

When the command channel is closed, all associated data channels are also closed.

32.1.4

TCP Stateful Filters

The EOS TCP stateful firewall feature mitigates the risk from the following types of attacks on TCP stacks:

- Rejecting or injecting malicious data on an existing TCP session.
In a typical TCP sequence, the initial sequence numbers of both ends of the connection are exchanged during the initialization period. After initialization, all data segments are referenced with these sequence numbers. The TCP stateful firewall checks the sequence numbers of all TCP packets to make sure that they are correct and drops any acknowledgements for data which was not transmitted.
- Denial of service attacks in which the server is flooded with connection requests.
A denial of service attack occurs when the attacker floods the listener with several SYN packets. The listener typically creates state and allocates memory resources for each SYN request. Since the listener has a finite amount of resources it can allocate to TCP connections, the large number of SYN request packets from the attacker causes the listener to run out of resources, and it stops accepting new connections. By defining a range of IP addresses and TCP port numbers (source and destination) on which stateful inspection is performed, you can limit the number of open connections to any open listener.

32.1.5

Generic Filters

The purpose of generic firewall filters is to permit or deny traffic for which predefined filters were not provided. Motorola Solutions's simple yet powerful filter language allows you to write your own specialized filters, each of which may be comprised of a number of rules within the generic filters.

- Rules can be specified on a per-interface basis.
- Rules can be applied to incoming traffic, outgoing traffic, or both.
- Rules are based on easy-to-understand names and values, instead of hex numbers, offsets, or bit-masks.
- Rules can provide comprehensive logging for both permitted or denied packets.
- Rules can permit or deny packets based on any combination of source address, destination address, protocols, source TCP/UDP port, destination TCP/UDP port, ICMP message types, and TCP "Establish" keyword to differentiate the direction of TCP connections from the value of the SYN bit.

The syntax for writing your own generic firewall filter is as follows:

```
SETDefault !<port> -FW InFilter = <string>
```

```
SETDefault !<port> -FWOutFilter = <string>
```

where <string> is the name of a file containing a list of filter rules. This file:

- Observes standard DOS filename syntax.
- Contains all of the generic In or Out rules for the port.
- Is not limited in terms of the number of rules it can contain.
- Can be edited off-line and transmitted via TFTP into the EOS device's ./FILTERS directory.

To define the filter rules, you must use:

```
ADD -FW FILTERS <file_name> (<rules> <;rules> . . . )
```

where the syntax for the filter rules (<rules>) is:

```
Permit | Deny [Log[0-7]]  
[From <IPaddr>/<mask (0-32 for IPv4; 0-128 for IPv6)>| <AddressList>]
```

```
[To <IPaddr>/<mask (0-32 for IPv4; 0-128 for IPv6)>| <AddressList>]
[Secure | NoSecure | Both]
[NextHop <IPaddr> [<IPaddr> [<IPaddr>]]] [NextPort !<Port> [!<Port> [!<Port>]]]
[TCP |UDP | ICMP | <protocol_value> | [!]<UserDefService>] [<options>]
```

If TCP is specified, the syntax for *<options>* is:

```
[Src <compare> <port / port1-port2>]
[Dst <compare> <port / port1-port2>]
[Estab | StateFul]if UDP [Src <compare> <port>]
[Dst <compare> <port>]
```

If UDP is specified, the syntax for *<options>* is:

```
[Src <compare> <port / port1-port2>]
[Dst <compare> <port / port1-port2>]
```

If ICMP is specified, the syntax for *<options>* is:

```
<ICMP type [ <ICMP_code>]>
```

For details see the description of the `FILTERs` parameter in “FireWall Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

32.2

Setting Up Internet Firewalls


To configure an EOS device to perform firewall functions, you must set parameters in the FireWall Service.

The following figure shows two levels of firewall protection set up using two routers (A and B).

The first firewall is Router A, which connects the Internet to a server subnet. The server subnet is where most Internet servers, such as the mail server and the WWW server, are located. The firewall on Router A is enabled on port 3.

The second firewall is Router B, which connects the server subnet to the internal corporate network. The firewall on Router B is enabled on port 7. Traffic can flow freely between the other interfaces on Router B because the firewall is not enabled on those interfaces. This permits Router B to not only perform firewall functions but also perform high-speed routing for internal networks.

Figure 173: Firewall Setup Example



Although Router B is a secondary firewall, it is really the primary defense for internal networks since the internal servers are directly reachable from the Internet and they have a much higher chance of being compromised. Access from these internal servers to internal networks should be limited, and the servers should be configured using secured applications.

When firewalls are implemented, your operations, administration, and maintenance (OAM) procedures must keep working. Examples of OAM procedures include Telnet, Internet Control Message Protocol (ICMP or PING), File Transfer Protocol (FTP), Trivial File Transfer Protocol (TFTP), and Simple Network Management Protocol (SNMP). For information on how to configure a firewall for these services, see `TelnetIn` and `TelnetOut`, `FTPIn` and `FTPOut`, `TFTP`, `ICMP`, and `SNMP` parameters in “FireWall Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

In this procedure, 123.0.1.0/24 is the IP address of the management subnet, and 123.0.1.12/32 is the IP address of the SNMP management station. The syntax “*a.b.c.d/x*” denotes the address *a.b.c.d* with the top *x* bits significant for comparison.

 **NOTE:** The management station need not be directly attached to one of the firewall devices. The management subnet could be anywhere on the corporate network. For the purpose of this procedure, we assume only that there is such a network with a common IP prefix/mask.

All of the filters used in the following steps are designed to control traffic to and from the EOS devices (routers). At this stage, no attempt is made to control traffic through the routers; that procedure is covered in [Blocking Unwanted Traffic on page 560](#).

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths are configured as described in [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).

Process:

1. To define the basic stance of your firewall and decide whether log messages are or are not recorded, perform the following actions:

- a. On Router A, enter the following commands:

```
SETDefault !<port> -FireWall DefActionIn = ([Permit | Deny], [Log | NoLog])  
SETDefault !<port> -FireWall DefActionOut = ([Permit | Deny], [Log | NoLog])
```
- b. On Router B, enter the following commands:

```
SETDefault !<port> -FireWall DefActionIn = ([Permit | Deny], [Log | NoLog])  
SETDefault !<port> -FireWall DefActionIn = ([Permit | Deny], [Log | NoLog])
```

The Deny stance means that after all of the filters were applied to a packet, and no actions were taken, the packet must be dropped. You can explicitly deny specific types of traffic within your rules, which would stop traffic that you find dangerous or unnecessary before the system has to check that traffic against all of the other rules.

For more information, see the DefActionIn, DefActionOut, and Log parameters in “FireWall Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.
2. If OSPF is the routing protocol in use, to allow OSPF packets to come in and go out on all LAN interfaces, perform the following actions:
 - a. On Router A, enter:

```
ADD !3 -FireWall OSPF Permit
```
 - b. On Router B, enter:

```
ADD !7 -FireWall OSPF Permit
```
3. If RIP is the routing protocol in use, to allow RIP packets to come in and go out on all LAN interfaces, perform the following actions:
 - a. On Router A, enter:

```
ADD !3 -FireWall RIP Permit
```
 - b. On Router B, enter:

```
ADD !7 -FireWall RIP Permit
```
4. If BGP-4 or EGP is used, write a generic filter to allow any of these protocols to work properly. See [Generic Filters on page 554](#).
5. To configure Telnet on both routers, enter:

```
ADD !7 -FireWall TELnetOut Permit From 123.0.1.0/24 to 123.45.67.2
```

This command allows users on the management LAN to Telnet into Router A on Router B. It also guarantees that Router B does not block traffic from management stations to Router A. No blocking actions are taken on Telnet traffic to Router B itself, because the firewall is not enabled on port 8.
6. To configure TFTP on both routers (accept TFTP packets from Router A to the management station, on Router B), enter:

```
ADD !7 -FireWall TFTP Permit From 123.45.67.2 To 123.0.1.0/24
```

EOS devices sometimes use TFTP to perform file transfers or network booting. The EOS device always functions as a TFTP client, while the management station functions as a TFTP server.

This command allows both TFTP request packets (123.45.67.2 -> 123.0.1.0/24) and TFTP response packets (123.0.1.0/24 -> 123.45.67.2). Router B does not have a problem accessing the management server because the firewall is not enabled on port 8.
7. To configure ICMP on both routers, allow Router A to send and receive ICMP (Ping) messages, on Router B, enter:

```
ADD !7 -FireWall ICMP Permit
```
8. To allow Router A to be SNMP manageable, on Router B, enter the following commands:

```
ADD !7 -FireWall SNMP Permit From 123.45.67.1 To 123.0.1.12/32
```

```
ADD !7 -FireWall SNMP Permit From 123.0.1.12/32 To 123.45.67.1
```

SNMP (another UDP-based protocol) has two sides. On the receive side, there are management requests. On the transmit side, there are two kinds of packets: responses to management requests and traps. The following two filters allow Router A to send SNMP traffic to and receive traffic from the SNMP management station. Motorola Solutions is explicitly denying SNMP that originates elsewhere on the Internet.

9. To use FTP to move files between any EOS device and an FTP server, on Router B, enter:

```
ADD !7 -FireWall FTPIn Permit From 123.45.67.1 To 123.0.1.0/24
```

Router A needs to be able to FTP to and from the management subnet.

10. To verify the configuration on both routers, see [Verifying the Firewall Configuration on page 558](#).
11. To check the overall status, see [Checking the Overall Status on page 559](#).
12. To block unwanted traffic on Router A and Router B, see [Blocking Unwanted Traffic on page 560](#).
13. To define a service, see [User-Defined Services on page 561](#).
14. To use IP address grouping, see [IP Address Grouping on page 563](#).
15. To use RealPlayer predefined service, see [The RealPlayer Predefined Service on page 563](#).
16. To enable security, see [Enabling Security on page 566](#).

32.2.1

Verifying the Firewall Configuration

After you defined the OAM filters, you can use this procedure to turn on IP Firewall and verify the configuration of both routers.

Procedure:

1. Turn on the firewall on Router A by entering:

```
SETDefault !3 -FireWall CONTrol = Filter
```
2. Turn on the firewall on Router B by entering:

```
SETDefault !7 -FireWall CONTrol = Filter
```
3. Verify current functionality on both routers by performing the following actions on Router A:
 - a. Ensure that both routers have all of their routes in the IP routing table by entering:

```
SHow -IP AllRoutes
```
 - b. If you use OSPF, ensure that both routers have all of their OSPF adjacencies by entering:

```
SHow -OSPF NeighborStatus
```
 - c. Verify that Telnet works between the routers and the management subnet.
 - d. Verify that TFTP works between the routers and the management subnet by entering:

```
COpy 123.0.1.76:test a:test
```
 - e. Verify that PING works between the routers and the management subnet by entering:

```
PING 123.0.1.17
```
 - f. Verify that SNMP management of the routers works:
You must verify that the routers still “green” in the management application.
You must take down an interface on the router by removing a cable, and verify that you see a trap on the SNMP management station.

- g. Verify that you can use FTP from the routers by setting up the firmware and performing PUT and GET tests by entering the following commands:

PUT <filename>

GET <filename>

4. Repeat [step 3](#) on Router B.

32.2.2

Checking the Overall Status

Procedure:

To check the overall status of your configuration, display the settings of the parameters in the FireWall Service by entering:

SHoW !<port> -FireWall CONFiguration

For the following predefined services, only the ports and their corresponding rules are displayed: RealPlayer, TraceRouteIn, TraceRouteOut, DNSClientSvsIn, DNSClientSvsOut. For all other predefined services, permit, deny, log and active time fields are also displayed. A display similar to the following appears:

```
=====Firewall Configuration for Port !2=====
```

CONTRol (1st priority)-----Permit-Deny-Log			
Filter	-	-	-
IgnoreSrcSpoofing	-	-	-
DenyTinyFragment	-	0	0
DenyAnyFragment	-	0	0
DenySrcRoute	-	0	0
DenyRecordRoute	-	0	0
DenyTinyStamp	-	0	0
DenyIPTunnel	-	0	0
SuppressICMP (0 ICMP msgs issued)	-	0	0
Services (2nd priority)-----Active@-----idle--Permit-Deny-Log			
FTPOut	Permit Log	Mar25 12:21	19min
Filters (3rd Priority)-----Bytes-Active@-----idle--Permit-Deny-Log			
InFilter good		34	Mar25 12:21
OutFilter		0	-
DefAction (4th priority)-----Permit-Deny-Log			
Deny	0	0	0

32.2.3

Blocking Unwanted Traffic

You can use this procedure to block unwanted traffic on Router A and Router B.



NOTE: TCP-based services such as TELnet, FTP, SMTP, NNTP, DNS, Gopher, and Archie are much safer than non-TCP-based services. Avoid using non TCP-based services on your Internet connections. Turn on `DenySrcSpoofing` on interfaces receiving external traffic to repel source spoofing attacks. Internal interfaces may not require such checking.

The filters in this procedure apply to traffic through the EOS device, not to and from traffic as described earlier.

Procedure:

1. Configure external protection for Router A by performing the following actions:

- a. Block hacker tricks by entering on Router A:

```
SETDefault !3 -Firewall CONTrol = DenySrcSpoofing
```

These are nonintrusive filters that adversely affect traffic only from those individuals who try to break into your site. All normal traffic proceeds as usual.

- b. Block hacker tricks by entering on Router B:

```
SETDefault !1 -FireWall CONTrol = DenySrcSpoofing
```

- c. To allow remote secondary name servers to talk to your external name server, configure Router A by entering:

```
ADD !1 -Firewall DNSSvrSvr Permit From 123.45.67.90/32 To <IPaddr>
```

```
ADD !1 -Firewall DNSSvrSvr Permit From 123.45.67.90/32 To <IPaddr>
```

These commands allow TCP-based domain name service (DNS) server-to-server traffic to be sent between 123.45.67.90 (the IP address of the name server in this example) and several remote name servers (for example, two off-site secondary name servers). Multiple name servers can be accommodated by adding more of these commands.

2. Allow mail to be sent from the Internet to an external mail host on both routers by performing the following actions:

- a. To configure Router A, enter the following commands:

```
ADD !3 -FireWall SMTPOut Permit From 123.45.67.90/32
```

```
ADD !3 -FireWall SMTPIn Permit To 123.45.67.90/32
```

- b. To configure Router B, enter the following commands:

```
ADD !7 -Firewall SMTPIn Permit From 123.45.67.90/32 To 123.0.0.0/8
```

```
ADD !7 -Firewall SMTPOut Permit From 123.0.0.0/8
```

SMTP is a particularly vulnerable service on many UNIX workstations because users run it everywhere, and common vendor-supplied versions have many security holes in their implementations. Because of this vulnerability (and the difficulty of keeping all of the internal machines up-to-date with the latest version of "sendmail"), Motorola Solutions only allows hosts on the Internet to establish SMTP connections to the external mail host.

Router B must be configured to allow SMTP connections from anywhere on the corporate network (123.0.0.0/8 in this example) to anywhere on the Internet. You can also force all internal mail to go through the external mail host, but that is not usually a requirement. Allow SMTP to come in from the corporate backbone. On the other LAN interfaces, with "special treatment" for the perimeter network, only accept SMTP traffic in from the main mail host, and only to machines within the corporate address (123.0.0.0/8). On all other interfaces, SMTP is allowed to go out as long as it is coming from within the corporate address.

3. Configure HTTP and WWW connections on both routers (allow remote HTTP connections to and from the WWW server only; allow internal WWW browsers to go out) by performing the following actions:
 - a. To configure Router A, enter the following commands:

```
ADD !3 -Firewall HTTPIn Permit To 123.45.67.91/32
ADD !3 -Firewall HTTPOut Permit
```

On Router A, allow HTTP traffic from the Internet to the external WWW server only. Block traffic to the rest of the corporate network.
The first filter allows Internet to access your WWW server; the second filter allows your internal users to originate HTTP traffic to anywhere on the Internet.
 - b. To configure Router B, enter:

```
ADD !7 -Firewall HTTPOut Permit
```

This command explicitly permits HTTP out through Router B. If you do not take this step, the `DefaultAction` parameter blocks all HTTP traffic in and out on all ports.
4. Optional: Allow remote news feeds to get to the external news server by performing the following actions:
 - a. Allow the internal Network News Transfer Protocol (NNTP) server to connect over Router A LAN interface in the outgoing direction by entering:

```
ADD !3 -Firewall NNTPOut Permit From 123.45.67.89/32
```
 - b. Allow the external NNTP servers to make connections through the Router A Internet link, but only to the external NNTP server, by entering the following commands:

```
ADD !3 -Firewall NNTPIn Permit From <IPaddr> To 123.45.67.89/32
```

where `<IPaddr>` is the IP address of the first external news feeder.

```
ADD !3 -Firewall NNTPIn Permit From <IPaddr> To 123.45.67.89/32
```

where `<IPaddr>` is the IP address of the second news feeder.
If you have 10 external news feeds, then you need 10 filter rules like the first two.
 - c. Allow NNTP traffic to cross Router B by entering the following commands:

```
ADD !7 -Firewall NNTPIn Permit
ADD !7 -Firewall NNTPOut Permit
```

If you do not explicitly permit NNTP using these rules, the `DefaultAction` parameter settings deny it.

32.2.4

User-Defined Services

Several predefined services are provided, but you can also create your own service definition. For example, you could allow HTTP on several unofficial but well-known TCP ports: 8080, 8000, 8001, and 8888. After you create this service definition, you can see the service name.

Defining a Service

The ability to define a service allows you to reference a service name when you define filter rules.

To define a service, you must use:

```
ADD -FireWall UserDefService <service_name> <protocol> [Src <compare> <port / port1-port2>] [Dst <compare> <port / port1-port2>]
```

The `<port / port1-port2>` option lets you specify a single port or a range of ports.

For example, you can define a service by entering:


```
ADD -FireWall UserDefService my_telnet TCP Src > 1024 Dst = 23
```

You can then use the service name in place of the *<protocol>* option when defining rules. After a service is defined, you can interchange the Src and the Dst ports value by using !serviceName.

For example, using !my_telnet, then means TCP Src = 23 Dst > 1024.

The example in the following figure illustrates how to configure a service and then use the service name when you define packet filtering rules.

Figure 174: User-Defined Service Example



Assume that there is no predefined filter for Telnet and a client wants to Telnet to the server.

To configure the Telnet service and define packet filtering rules, you can add a service name by using:

```
ADD -FireWall uds my_telnet TCP Src 2000-2100 Dst = 23
```

If you specify my_telnet as the name when defining packet filtering rules, it means that the user protocol is TCP, destination port is 23, source port can be in the range 2000-2100, and the source is allowed to initiate a connection.

Configuring a Firewall with a User Defined Service Name

After you define a service name, you can use service rules to create a firewall. To configure a firewall on port 4 to allow a Telnet session between the client and the server (initiated from the client) you must use the following commands:

On the output filter, you must enter:

```
permit my_telnet from 129.213.0.0/16 to 170.1.1.1
```

On the input filter, you must enter:

```
permit !my_telnet from 170.1.1.1/16 to 129.213.0.0/16
```

Differences Between a Predefined Service Filter and a User-Defined Service Based Filter

There is a difference between a user-defined service filter and predefined service filter in terms of the number of rules to be applied on a port. For example, if TelnetOut was allowed, you must use the following predefined service filter:

```
ADD !4 -FireWall TELnetOut Permit From 129.213.0.0/16 to 170.1.1.1
```

You must use this filter as the output filter rule, and it internally creates another rule in the input direction to permit the traffic from 170.1.1.1 to 129.213.0.0/16. Using the predefined filter eliminates the need to create an additional rule.

With a user-defined service, on the other hand, you would need to create two rules in the input and output directions. However, the user-defined service is useful when creating filters for protocols for which no predefined filters exist.

32.2.5

IP Address Grouping

IP address grouping allows you to identify a group of IP addresses (an address list) which need the same type of access by a single name. You can then use this name when you define filters.

To add IP addresses to an address list, you must use:

```
ADD -FireWall AddressList <AddressListName><IPAddress/mask>
```



NOTE: The *<IP Address>* can be either an IPv4 or an IPv6 address, and the *<mask>* can be 0-32 bytes for IPv4 addresses or 0-128 bytes for IPv6 addresses.

For example, to add an IP address to the address list named `my_list`, you must use:

```
ADD -FireWall AddressList my_list 10.1.1.10
```

You can then create rules by entering:

```
permit from my_list to 10.1.1.1
```

and

```
permit from 10.1.1.1 to my_list
```

Assuming that `my_list` has the three IP addresses 20.1.1.1, 30.1.1.1, 40.1.1.1, the following three rules are created internally:

```
permit from 10.1.1.1 to 20.1.1.1
```

```
permit from 10.1.1.1 to 30.1.1.1
```

```
permit from 10.1.1.1 to 40.1.1.1
```

The last rule corresponds to “permit from 10.1.1.1 to `my_list`”

If you define a rule by entering:

```
permit from my_list to my_list
```

The following rules are created internally:

```
permit from 20.1.1.1 to 30.1.1.1
```

```
permit from 20.1.1.1 to 40.1.1.1
```

```
permit from 30.1.1.1 to 20.1.1.1
```

```
permit from 30.1.1.1 to 40.1.1.1
```

```
permit from 40.1.1.1 to 20.1.1.1
```

32.2.6

The RealPlayer Predefined Service

The RealPlayer application includes the RealAudio and RealVideo applications. This section describes the operation of RealAudio. Operation of RealVideo is same as the RealAudio.

To support RealAudio and RealVideo through the firewall, you should enable the traffic on the following range of ports:


- TCP port 7070 for incoming and outgoing traffic
- Negotiated UDP port during TCP session setup, for incoming traffic from the server

The TCP port is used by the client (RealPlayer) to initiate a conversation with an external RealServer, to authenticate the player to the server, and to pass control messages during playback including pausing or stopping the audio stream.

The range of UDP ports, on the other hand, carry the incoming stream. These ports begin to carry traffic only after the client and server performed the authentication routine, and should be enabled only for incoming traffic. The UDP port on the client to which this traffic should be sent is made known to the server during TCP packet exchange.

The following figure illustrates the RealAudio protocol sequence between the client and server.

Figure 175: RealAudio Protocol Sequence



The protocol sequence steps are as follows:

1. The client (RealPlayer) sends an initial TCP request to the server (RealServer) with source port 7070 and destination port 7070.
2. The RealServer sends a TCP response back to the RealPlayer with source port 7070 and destination port 7070.
3. The connection is established.
4. If the RealPlayer decides to get the audio streams delivered at its UDP port, the RealServer sends the UDP stream from any of the ports greater than 1023 to a client port, which was identified by the client during the initial exchange of TCP packets.
5. If the RealPlayer decides to get the audio streams delivered at its TCP port, the RealServer sends the TCP streams from its port 7070 to the client port 7070.

Firewall Configuration

You can configure the FireWall Service to permit users inside a private network to access the RealServer. This configuration does not allow RealServer to initiate a connection to the RealPlayer. Since the FireWall Service is not configured on port 1, there are no restrictions on the traffic that is allowed to pass through the port. However, since the FireWall Service is configured on port 4 and the default action is deny, you should add the following output filter rule on port 4:

```
Add !4 -fw RealPlayer permit from 129.213.0.0/16 to 170.1.1.1/16 ClientUDPPort 9000
```

When you add this rule, it creates the following filters internally to support the RealAudio data flow:

- An output filter to permit TCP connection from client to server


```
Permit from 129.213.0.0/16 to 170.1.1.1 TCP Src=7070  
Dst =7070
```

- An input filter to permit TCP connection from RealServer to RealPlayer

```
Permit from 170.1.1.1 To 129.213.0.0/16 TCP Src=7070  
Dst =7070 Estab
```

- An input filter to permit UDP streams from server to client

```
Permit from 170.1.1.1 To 129.213.0.0/16 UDP Dst = 9000
```

Additional Predefined Services

TraceRoute Support

To allow the TraceRoute outbound where the user on the internal network is running TraceRoute to the external destination through the FireWall Service, the following packets are allowed through the firewall:

- Constructed UDP packets outbound.
- Relevant ICMP packets ("TTL exceeded" and "port unreachable") back inbound.

The EOS device generates a TraceRoute packet by creating a UDP packet using 33434, 33435, 33436 as the destination port numbers. The source port is created randomly.

DNS Client To Server Support

To allow the DNS query outbound where the user on the internal network sends DNS query to the external destination through the FireWall, the TCP/UDP packets at the following ports are allowed through the FireWall:

- TCP/UDP packets from the client to the server with Dst = 53
- TCP/UDP packets from the server to the client with Src = 53 Dst = *<client's_port>* specified in the request

Secure HTTP (S-HTTP/HTTPS) Support

The HTTPS protocol is HTTP on top of Secure Socket Layer (SSL). SSL was designed by Netscape and specifies a protocol for providing data security layered between application protocols (such as HTTP or NNTP) and TCP/IP. This security protocol provides data encryption, server authentication, message integrity, and optional client authentication for a TCP/IP connection.

SSL was designed to provide security between client and server and to avoid any kind of "3-way man-in-the-middle" attack. Therefore SSL cannot be proxied through traditional application-level firewalls, because SSL considers a proxy server to be a middleman. The simplest alternative to this problem is to use a packet-filtering firewall. The SHTTP server listens at the TCP port 443; therefore, in order to allow the SHTTP traffic to pass through the firewall, the traffic destined to and from port 443 should be allowed unrestricted through it.

BGP-4 Support

BGP is a TCP-based protocol with no client-server relationship (it has a peer-to-peer relationship). Each of the BGP peers listens on the TCP port 179. When a BGP peer wants to establish a TCP connection, it uses port 179 as the destination port; the source port can be anything.

Finger Support

The Finger Service provides information about users such as a person's real name and a user name login time. The user may want to limit incoming finger requests. Outgoing finger requests are mildly problematic. The attacks are possible through the data in the finger response. The attack can be in the form of server sending immense data to the client or sending certain control characters to reprogram client's keys or sending a command that mails the password file to the server. You should be aware of these consequences which may happen due to allowing finger in or out.

WHOIS Support

WHOIS is commonly used to obtain public information about hosts, networks, and domains. Sites generally do not provide their own WHOIS server; they merely access the WHOIS servers at the NICs. There are no known security problems with WHOIS clients. Security problems that occurred are data driven.

SOCKS Support

SOCKS is a networking proxy protocol that enables hosts on one side of the SOCKS server to gain full access to hosts on the other side of the SOCKS server without requiring direct IP reachability. SOCKS redirects connection requests from hosts on opposite sides of a SOCKS server. The SOCKS server authenticates and authorizes the requests, establishes a proxy connection, and relays data. The SOCKS server listens on the standard TCP port 1080, however SOCKS can be configured to listen to some other port as well.

32.2.7

Enabling Security

Procedure:

- To allow a predefined service access through the firewall, enter:

```
Add !<port> -FireWall <pre-defined_service_name> permit | deny [Log [0-7]] [From <IPaddr/mask>][To <IPaddr/mask>] [NextHop <IPaddr>][Secure | NoSecure | Both]
```

- To allow only packets with ESP headers through the firewall, enter:

```
Add !1 -FireWall TELnetIn permit from 10.1.1.1 to 20.2.2.2 Secure
```

This command allows only secure TelnetIn packets in through the firewall.

- To allow both Telnet in and Telnet out packets through the firewall, enter:

```
Add !1 -FireWall TELnetOut permit from 10.1.1.1 to 20.2.2.2 Secure
```



NOTE: You must specify each predefined service parameter individually.

32.3

Filters Management

Each filter has a name assigned to it. The syntax of a filter name is the same as a DOS file name; it can be up to eight characters followed by up to a three-character extension. File names are case-insensitive. When a new filter is created with the same name as an existing filter, the new filter replaces the old one, in memory and on the disk.

A filter must have at least one rule defined within it. Each rule must begin with one of two keywords: Permit or Deny. There is no limit to how many rules can be defined in a filter; the software continues to accept new rules as long as there is memory or disk space available for them.

For detailed information on rule syntax and corresponding values, see the description of the `FILTERs` parameter in “FireWall Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

32.3.1

Filter Definition with the ADD Filter Command

To define a filter, you must use:

```
ADD Filter <filter_name> [<rule_1> <rule_2> <more_rules> ...]
```

A filter begins with the left parenthesis, and terminates with the right parenthesis. In between the parenthesis, any number of rules are allowed. Empty rules are ignored.

Each rule must be terminated by a new-line character or multiple rules may be specified on one line when each rule is terminated by a semicolon (;). Rule syntax checking is performed when rules are entered. Any syntax error is identified immediately and triggers the following actions:

- Displays error descriptions about the kind of syntax problem
- Discards the current line of input
- Displays help information
- Prompts for continued input

A rule that begins with the pound (#) sign indicates that it is a comment, and the software ignores the whole line.

TCP Stateful Firewall Filter

A TCP stateful firewall filter defines a range of IP addresses and TCP ports (source and destination) over which a stateful packet inspection is performed on TCP packets. For further details about stateful firewall filters, see [TCP Stateful Filters on page 554](#).



IMPORTANT: The stateful filter does not filter the traffic; therefore, normal filter rules to permit or deny traffic still apply. However you can apply stateful filtering for TCP-based service filters.



NOTE: The stateful firewall filter must see both directions of traffic. For example, if the EOS device is the transit for the traffic and you have a connection coming in on virtual port !V12 and going out on virtual port !V14, you must turn on the firewall InFilter for both port !V12 and port !V14. The OutFilter is required only for sessions to the EOS device; for example, for SSH connections to the device.

Configuring a User-Defined TCP Stateful Firewall Filter

To configure a user-defined TCP stateful firewall filter, you must use:

```
ADD -FireWall FILters <filter_name> (<rules> [; <rules>]...)
```

and define rules that enable the TCP stateful firewall checks.

For example:

```
ADD -FireWall FILters myRule
rule tcp_in: Permit TCP log From 20.20.20.55 To 20.20.20.96 StateFul
rule tcp_out: Permit TCP log From 20.20.20.96 To 20.20.20.55 StateFul
```

This command specifies that a stateful packet inspection is performed on TCP packets from IP address 20.20.20.55 to IP address 20.20.20.96 and from IP address 20.20.20.96 to IP address 20.20.20.55. The rule should be applied on all ports that see TCP packets for this connection, using one or both of the following commands:

```
SETDefault !<port> -FireWall InFilter = <filter_name>
SETDefault !<port> -FireWall OutFilter = <filter_name>
```

For example:

```
SETDefault !V12 -FireWall InFilter = "tcp_in"
SETDefault !V12 -FireWall InFilter = "tcp_out"
```

You can limit the maximum number of TCP connections to a server by using:

```
SETDefault -FireWall TCPParams <number_of_connections>
```

This is a global value and applies to all connections. When the configured number of connections is reached, all subsequent connections are rejected.

To view the state of the TCP connections, you must enter:

```
SHow -FireWall TCPState
```

A display similar to the following appears:

```
Server(connections)ClientSPort DPortStateTime Left
20.20.20.55 (1) 20.20.20.96537623ESTABLISHED431992(secs)
```

In addition, when the FireWall service FwSecControl parameter is set to ICSALog, syslog messages are generated when a packet is denied by default configuration of generic filters. For further details, see the description of the FwSecControl parameter in “FireWall Service Parameters” of the *Enterprise OS Software Version 18.2 Reference Guide*.

Configuring a Predefined Service Filter for TCP Stateful Firewall Filtering

To configure a predefined service filter for TCP stateful firewall filtering, use the keyword StateFull in the incoming and outgoing service filter definitions. Stateful firewall filtering is supported for the following TCP-based services:

- BGP (BGP; there is no directional sense associated with the BGP filter)
- DNS server-to-server (DNSSvrSvr)
- Finger (FingerIn, FingerOut)
- FTP (FTPIIn, FTPOut)
- Gopher (GopherIn, GopherOut)
- HTTP (HTTPIn, HTTPOut)
- NNTP (NNTPIn, NNTPOut)
- POP (POPIIn, POPOut)
- SHTTP (SHTTPIn, SHTTPOut)
- SMTP (SMTPIn, SMTPOut)
- SOCKS (SOCKS)
- Telnet (TELnetIn, TELnetOut)
- WAIS (WAISIn, WAISOut)
- Whols (WhoIsOut)

For the Telnet service, you must use:

```
ADD !<port> -FireWall TELnetIn Permit From <IPaddr>/<mask> To <IPaddr>/<mask> StateFull
ADD !<port> -FireWall TELnetOut Permit From <IPaddr>/<mask> To <IPaddr>/<mask> StateFull
```

For example:

```
ADD !1 -FireWall TELnetIn Permit From 20.20.20.55 To 20.20.20.96 StateFull
ADD !1 -FireWall TELnetOut Permit From 20.20.20.55 To 20.20.20.96 StateFull
```

This command specifies that a stateful packet inspection is performed on port 1 on incoming Telnet connections from IP address 20.20.20.55 to IP address 20.20.20.96 and outgoing Telnet connections from IP address 20.20.20.96 to IP address 20.20.20.55.

For details about the predefined service filter parameters, see “FireWall Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

32.3.2

Filter Creation with an Off-line Editor

You can use any PC or workstation text editor to create and edit your filter files. The workstation is also a good place to back up the filters file. Because filters can be complex, you must do the following:

- Examine and carefully edit the filters on your workstation.
- Add comments to them using the pound (#) sign.
- Use TFTP to transfer the final results to the FILTERS directory on the EOS device.

After the file transfers are complete, you must either reboot the device or issue the `REStart` command to restart the firewall with the newer set of filters.

You must use the `TEst` command to test filters with test packets generated by the EOS device. The system then reports whether the packet was permitted or denied.

The `REStart` command examines the filter file, detects any syntax errors, and provides the line number, the offending keywords, and other applicable help information. If there is a syntax error in the filter file, none of the defined filters take effect. As a result, you are responsible for making appropriate corrections to the filter file off-line and re-entering the `REStart` command.

A filter file must contain only filter rules conforming to the syntax specified in the description of the `FILters` parameter in the “FireWall Service Parameters” of the *Enterprise OS Software Version 18.2 Reference Guide*. Blank lines and comment lines starting with the pound sign (#) are ignored.

32.3.3

Filters Display and Deletion

Displaying Filters

To display all the filters that are currently defined, you must use:

```
SHow -Firewall Filter
```

The display shows the names, sizes, and creation dates for all of the filters that are currently stored on disk.

To display the contents of a particular filter, you must use:

```
SHow -Firewall Filter <filter_name>
```

Deleting Filters

Filters must be deleted individually from the system.

To delete a filter, you must use:

```
DElete -Firewall Filter <filter_name>
```

Deleting a filter only removes it from local storage, but not from the memory. To remove it from memory, you must use the `REStart` command.

32.3.4

Filters Assignment to Interfaces

Each interface (port) can have two filters associated with it: one filter that applies to traffic received on that port, and one that applies to traffic to be transmitted on that port.

The command syntax for `InFilter` and `OutFilter` is:

```
SETDefault !<port> -Firewall InFilter = <filter_name>
```

```
SETDefault !<port> -Firewall OutFilter = <filter_name>
```

For more information on the `InFilter` and `OutFilter` parameters, see “FireWall Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

All incoming packets, including broadcast, unicast, and multicast IP packets, received on the `<port>` are subject to filtering operations as specified in the `InFilter` parameter. Packets going to the EOS device, as well as packets to be forwarded by the EOS device, are all subject to filtering, including all the incoming routing protocol packets such as OSPF, RIP, or others.

All outgoing packets, including broadcast, unicast, and multicast IP packets, to be transmitted over `<port>` are subject to the filtering operation as specified in the `OutFilter` parameter. Packets originating from the

EOS device (for example, routing protocol packets), as well as forwarded packets (from another source), are all subject to the same filtering operations.

To remove a filter from the interface, you must use:

```
SETDefault !<port> -Firewall InFilter = "" (empty name)
```

If an assigned <filter_name> is not found in local storage, no filter operation is activated.

32.3.5

Filters Activation and Deactivation

To enable or disable all firewall filtering on an interface, you must use:

```
SETDefault !<port> -FireWall CONTrol = Filter | NoFilter
```

You can enable each interface independently of other interfaces. Enabling filtering on an interface enables ALL of the applicable firewall filters including:

- Service-independent filters, such as SourceSpoofing and IPTunnel.
- Generic filters, such as those defined in the InFilter and OutFilter parameters.
- Default actions as defined in the DefAction parameter and predefined filters.

Only packets forwarded by IP forwarding routine, as well as those destined to or originated from local system, are subject to the filtering action; bridged packets are not subject to this service. Bridging filters are the appropriate place to define rules for filtering bridged packets.

32.3.6

Firewall Filters Versus IP Filters

There are several differences between firewall filters and traditional IP filters used on the EOS devices:

- Traditional IP filters are output filters only. You cannot configure them to perform input filtering, an important first level of defense. Not only is input filtering necessary in blocking certain attacks (such as IP source address spoofing), it also protects the device itself.
- You cannot specify separate filters on a per-interface basis using traditional IP filters.
- Traditional IP filters contain non user-friendly hex numbers, offsets, and masks. This syntax is complex and confusing to design and maintain.
- Traditional IP filters cannot effectively deal with packets containing IP options.
- Traditional IP filters reorder the filtering rules to their own preference, causing confusion in some situations, and the defeat of the filter design in other situations.

Because of the special role traditional IP filters play, they remain unchanged and operate in parallel with the firewall. Some specific services – such as packet priority, protocol reservation, and dial-on-demand discard – are available only from the traditional IP filter. For those services, you must continue to use existing IP filters.

32.3.7

Execution Order for Traffic with Enabled Filters

When IP and Firewall filters are enabled, the sequence of execution for through traffic is as follows:

- receive packet → input firewall filter → “forwarding decision” → traditional IP filter → output firewall filter → transmit packet

For traffic coming into and terminating at the EOS device, the sequence is:

- receive packet → input firewall filter → “internal process” (no IP filter is applied)

For traffic originating from within the EOS device, the sequence is:

- “internal process” → output firewall filter → transmit packet
Traditional IP filters only apply to “through” traffic. Traffic *to* the EOS device is unfilterable.

32.4

System Logs Setup

You can configure IP Firewall to log system messages to the AuditLog Service, to the local console, or both. Log messages contain information about the system messages such as the date and time, interface, direction of the packets (ingoing or outgoing), packet header summary, and reason.

Each message contains two types of codes: facility and priority. The facility code tells syslog what subsystem the message is from and the priority code tells syslog how important the message is (ranging from Log(0), emergency, which is the highest priority to Log(7), debug, which is the lowest priority). Logging can be done on permitted or denied packets.

You can set up your firewall to log system messages in one of three ways:

- If you want log messages sent to your local console port, you must use:

```
SETDefault -FireWall Log = Console
```

- If you want log messages sent to the AuditLog service, you must use:

```
SETDefault -FireWall Log = Syslog
```

The AuditLog Service controls the delivery of messages to the syslog server(s). Most UNIX workstations come with syslog support. Consult your workstation manual for more details. If you choose this logging method, you need to enable the AuditLog Service and configure the IP address of your syslog server by using:

```
SETDefault -AuditLog CONTROL = (CONFIG, MESSAGES, SECURITY)
```

```
SETDefault -AuditLog LogServerAddr = <IP_address>
```

- If you want to send log messages to your local console port and to your syslog server, you must enter:

```
SETDefault -FireWall Log = (Syslog, Console)
```

If you want to limit the number of messages that are logged to a specified number per second, you can set the FireWall service LogThrottle command accordingly. The FwSecControl parameter also allows you to specify whether or not ICSA syslog messages are logged. For details, see the descriptions of the FwSecControl and LogThrottle parameters in “FireWall Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Specifying Log Content

You can enable specific information to be logged. For example, to see all of the denied source-spoofing packets, you must enter:

```
SETDefault -FireWall Log = SrcSpoofing
```

If you want to log all incoming FTP connections, add a Log option to your -FireWall FTPIn command.

For example:

```
ADD !2 -FireWall FTPIn Permit Log
```

You can log the summary of the packet or detailed contents (the first 64 bytes) of the packet by using:

```
SETDefault -FireWall Log = SUMMARY | DETAIL
```

For more information on selectively logging messages, see the Log parameter in the “FireWall Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Log Description

The log generated by the firewall has the following format:

```
<Date, time> Firewall <Tx / Rx> | !<interface> <Source_IP_addr> <Source Port> ->  
<Destination_IP_addr> <Destination_port> <Protocol> <action> <Filter_type>
```

The fields in the log contain the following information:

Table 59: Description of the Fields in a Log

Value	Description
<Date, time>	Specifies the actual date and time when the firewall received or transmitted a packet.
<Tx / Rx>	Specifies whether the packet was transmitted or received through the firewall. Tx indicates transmitted, Rx indicates received.
!<interface>	Indicates the interface number on the EOS device on which the packet filtering rules were applied for the packet. In the receiving direction, the interface number is the interface on which the packet was received. In the transmitting direction, it represents the interface on which the packet was transmitted.
<Source_IP_addr>	Indicates the source IP address of the packet. This can be an IPv4 address or an IPv6 address.
<Source_port>	Indicates the source port for the TCP/UDP packet. For the well known applications, the name of the application is specified instead of actual TCP/UDP port numbers. For example, if there is an incoming Telnet packet received by the firewall, the log says “Telnet” instead of 23 (which is the actual TCP port for the Telnet application). However, with the increasing use of random UDP ports by applications (multimedia applications, Archie services, DNS services, and so on). it may not be possible for the firewall to determine the type of traffic running through the port. For example, if you are running RealAudio traffic – which uses destination TCP port 7070 (for the server) and any random TCP source port above 1023 – through the firewall, and if that port happens to be 1525 (which is the well known port for Archie), the log indicates that the Src port is Archie and the Dst port as RealAudio. When such confusion occurs, you should analyze the previous log history to find out the kind of traffic going through your firewall.
<Destination_IP_addr>	Indicates the destination IP address for the packet. This can be an IPv4 address or an IPv6 address.
<Destination_port>	Indicates the destination port for the TCP/UDP packet. For the well know applications, the name of the application is specified instead of the TCP/UDP port number. For example, if there is an outgoing Telnet packet transmitted through the firewall the log says “Telnet” instead of 23 which is the actual TCP port for telnet. However, with the increasing use of random UDP ports by applications (multimedia applications, Archie services, DNS services, and so on), it may not be possible for the firewall to determine the type of traffic running through it. For example, if a you are running Real Audio traffic – which uses TCP port 7070 and any random UDP port above 1023 – through the firewall , and if you specify port 1525 as the UDP port on which to receive the Re-

Value	Description
	alAudio stream, the log indicates the destination port as Archie and the Src port as RealAudio. When such confusion occurs, you should analyze the previous log history to find out the king of traffic going through the firewall.
<i><Protocol></i>	Indicates the protocol carried within the IP packet: TCP, UDP, or ICMP. AH/ESP protocols are specified by name. Any other protocols are specified by numeric value. For ICMP packets, this field is followed by the ICMP packet type.
<i><action></i>	Indicates whether the firewall permitted or denied the packet. Permit means the packet was allowed to pass through the firewall, Deny means that the packet was dropped by the firewall.
<i><Filter_type></i>	Indicates the type of firewall filter which was applied to the packet when making the permit/deny decision. For example, if you configured a TelnetIn service filter and the incoming IP packet complies with this rule, then the log indicates TelnetIn. Other keywords include DefActionIn, DefActionOut, InFilter, and OutFilter. However, for the DNSClientSvrIn, DNSClientSvrOut, RealPlayer, TraceRouteIn, and TraceRouteOut services, the <i><filter_type></i> is always either InFilter or OutFilter.

Firewall Log Examples

Example 1:

The firewall received a packet on port 4 from 10.0.0.2 which was addressed to 10.0.0.255 and it was a RIP packet. The protocol was UDP. The firewall permitted this packet and this action was taken as a result of DefActionIn.

```
[161] S6000 # May 4 10:16 FireWall Rx !4 10.0.0.2 (RIP)->10.0.0.255 (RIP) UDP Permit
DefActionIn
```

Example 2:

The firewall received a packet on virtual port 1 with the source address of 129.213.205.110 and with a DNS source port 53. The destination address was 129.213.119.38 and the destination port was 32877. The packet was permitted through the firewall as a result of InFilter rules.

```
May 5 13:54 FireWall Rx !V1 12921.3.205.110 (DNS)-> 129.213.119.38 (32877) UDP Permit
InFilter
```

Chapter 33

Access Control Service Configuration with RADIUS

This chapter describes how to configure the Access Control (AC) Service for RADIUS Access Control.

An EOS device itself can perform Telnet, Console or SSH access authentication, or you can use a RADIUS server.

- The EOS device performs local authentication by searching a local database for users and their passwords.
- RADIUS authentication provides another means of authentication, utilizing an external RADIUS server with its own database of users, passwords, and access levels. This centralizes and simplifies the task of administrating access to large number of devices on a network.

With RADIUS configured, user account information is sent to a RADIUS server when a user attempts to access the EOS device. Based on the server response, the device takes the following actions:

- Declines access if the RADIUS account is invalid.
- Provides user or network manager level access depending on the user's access level.
- If no response is received, tries a secondary RADIUS server, if configured.
- If no response is received from the RADIUS Server(s) and a resolution order of "RADIUS Local" was configured, the user is authenticated locally.

33.1

Configuring RADIUS Authentication



NOTE: If the primary RADIUS authentication server becomes unreachable, logins are authenticated via the secondary RADIUS authentication server, specified with the `SecAuthSrvr` parameter, and that server becomes the current RADIUS authentication server. However all subsequent logins still attempt to authenticate via the primary RADIUS authentication server, and if a login attempt is successful, authentication switches back to the primary RADIUS authentication server, and that server becomes the current RADIUS authentication server again.

Procedure:

1. Specify the IP address of Primary Authentication server by entering:

```
SETD -AC PrimAuthSrvr = <IP_address>
```

Step example:

```
SETD -AC PrimAuthSrvr = 192.147.72.84
```

2. Specify Resolution Order by entering:

```
SETD -AC RESolutionOrder = [NONE | Local | Radius | Radius Local]
```

Step example:

```
SETD -AC RESolutionOrder = Radius
```

The `RESolutionOrder` parameter sets the order for authentication:

- If you set `RESolutionOrder` to `Radius Local`, the EOS device attempts to authenticate using the RADIUS server. If there is no response from the RADIUS server, authentication is accomplished locally.

- If you set `RESolutionOrder` to `Radius`, the EOS device authenticates using only the RADIUS server, and local access control is disabled for all users except `root`.



NOTE: For the `root` user, there is no difference between the `Radius Local` and `Radius` settings. Regardless of how the `RESolutionOrder` parameter is set, the EOS device first attempts to authenticate using the RADIUS server; if there is no response, authentication is accomplished locally. This prevents the `root` user from being denied access when `RESolutionOrder` is set to `Radius` and the RADIUS server is down.

33.2

RADIUS Parameters in the AC Service

All RADIUS parameters available in the AC service menu are also available in the Remote Access service (RAS) menu. When a common parameter is set in one service, it is also set in the other service. The parameters are `ACntUdpport`, `AuthUdpport`, `PrimACntSrvr`, `PrimAuthSrvr`, `RetransTimer`, `SecACntSrvr` and `SecAuthSrvr`.



NOTE: At this time, you cannot set RADIUS Accounting parameters in the AC service. These parameters are only applicable in the RAS service. RADIUS Authentication only (not RADIUS Accounting) is supported in the AC service.

For further details, see “AC Service Parameters” and “RAS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Chapter 34

Lightweight Directory Access Protocol Configuration

This chapter describes how to configure the Lightweight Directory Access Protocol (LDAP) service to support the CCGW and Firewall clients. LDAP defines a standard protocol for accessing diverse directory services. It allows you to store device configurations or security policies on the server and download these configurations when required.

The following terms are used to describe the LDAP feature:

Attributes

Holds descriptive information about an entry. Each attribute consists of an attribute type and one or more attribute values. The attribute type identifies the class of information given by that attribute (for example, telephone number). The attribute value is the particular instance appearing that entry (for example, 326-7500). Attributes generally have short, mnemonic names, for example "cn" is the abbreviation for commonName.

Object class

Defines the type of data that an entry can contain (for example, does the entry represent a person, an organization, a country, and so forth). Each object class defines a set of required and optional attributes. This attribute list defines the kind of data that you can store on the entry.

Schema

The type of information that can be stored in the directory is defined by the directory server's schema. The schema is defined by the following elements:

- Object classes
- Attributes

Schema Extension

Customizes new object classes and attributes to store specialized information. The schema supplied with the Standard directory server is limited and possibly cannot store the information you may want to store.

Relative Distinguished Name (RDN)

The name of the entry itself, before the entry ancestors are appended to the string to form the full distinguished name.

Distinguished Name

The string representation for the entry names in a directory server directory. The distinguished name of an entry is the string concatenation of the RDNs from the entry to the root of the tree.

34.1

Configuring LDAP for CCGW Clients

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- IP routing was configured. See [IP Routing Configuration on page 119](#).

Procedure:

1. Specify the IP address and TCP port of the LDAP server and specify CCGW as a service for which this server IP address must be used when sending out a request by entering:

```
ADD -LDAP ServerAddress <IP_address> [TCP Port] "ccgw"
```

where:

<IP_address> is the IP address of the LDAP server

TCP Port is the TCP port at the LDAP server where the connection must be made

2. Specify the DN value for the LDAP bind request and specify CCGW as the service configuration that should be downloaded from the LDAP server by entering:

```
ADD -LDAP BindDN "<DN string>" "<password>" "ccgw" [Encrypted]
```

where:

<DN string> is the DN value, in ASCII text format

<password> is a string of 0-16 characters

The DN value is used to authenticate an area in the tree where the ADD/DELETE/Search operations can be performed. You must use the optional `Encrypted` keyword if you want to enter the password in encrypted form.

3. Specify the UDP port to be used by the Notifier to send notifications to the CCGW by entering:

```
SETDefault -CCGW NMUDPPort = <UDP Port #>
```

where <UDP Port #> is the UDP port number.

34.2

LDAP Working with CCGW Clients

CCGW commands are executed in the device `boot.cfg` file upon bootup. After the `boot.cfg` file was executed, EOS initializes CCGW, which sets up communication with the network management system's LDAP server. EOS then listens to the UDP notification port and configures the CCGW based on the network management system's LDAP configuration.

A CCGW configuration file (CCGWDB) is created on the EOS flash. This file contains the latest CCGW configuration from the network management LDAP server. It is updated when changes are received from the network management LDAP server after the CCGW was rebooted and is used for bootup persistence. If the CCGW cannot communicate with the LDAP server, it reads the configuration from the EOS flash.

If the CCGW receives a configuration from the LDAP server, and the required hardware does not exist on the box, the CCGW:

- Accepts the configuration.
- Sends a state change trap (`ccgwChannelStateChangeTrap`) with channel state DOWN, and fault cause 6 (User Configured Channel).

For details about the hardware required to support CCGW functionality, see [CCGW Hardware Requirements on page 727](#).



IMPORTANT: If a change is made to the CCGWDBVersion, DBSchemaVer, DSR, SiteID, SiteType, or ZoneID parameter from the LDAP network management server (to correct a misconfiguration, for example), you must perform one of the following steps after the change is made:

- Remove the CCGWDB file from flash, then reboot the CCGW.
- Resynchronize the CCGWDB file with the LDAP server by using the following commands to disable and re-enable the CCGW service:

```
SETDefault -CCGW CONTROL = Disable
```

```
SETDefault -CCGW CONTROL = Enable
```

The CCGWDB file is refreshed from the LDAP server. When a CCGW is replaced, the LDAP server is the source for refreshing the CCGWDB file.



CAUTION: When upgrading or migrating a CCGW, you should remove the CCGWDB file from flash before upgrading the software on the CCGW.

34.3

Configuring LDAP for Firewall Clients

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths are set up. See [Basic Ports and Paths Configuration on page 50](#) and [Virtual Ports and Multiple Instance Lists Configuration on page 77](#).
- IP routing was configured. See [IP Routing Configuration on page 119](#).

Procedure:

1. Configure the LDAP server address by entering:

```
ADD -LDAP ServerAddress <Name / IP> [TCP port] "<service name>"
```

You can assign the IP address or the domain name of the server. If you assign the domain name, you must configure the name service using the IP service parameters. You must also set the value of the TCP port to the value configured on the LDAP server where it listens to the LDAP requests. If you do not assign the TCP port, a default value of 389 is used.

Step example:

To assign an IP address of 129.213.200.10 for the LDAP ServerAddress to fetch the Firewall configuration and use the default TCP port value of 389, enter:

```
ADD -LDAP ServerAddress 129.213.200.10 "firewall"
```

2. Configure the RootDN value by entering:

```
ADD -LDAP RootDN "<DN_string>" "<service_name>"
```

Step example:

To assign a DN value of "o=Moca, c=us", enter:

```
ADD -LDAP RootDN "o=Moca, c=us" "firewall"
```

3. Optional: Configure the BindDN value by entering:

```
ADD -LDAP BindDN "<DN_string>" "<password>" "<service_name>" [Encrypted]
```

Step example:

To assign a DN value of "o=Moca, c=us" and password as "Moca", enter:

```
ADD -LDAP BindDN "o=Moca, c=us" "Moca" "firewall"
```

You must use the optional `Encrypted` keyword if you want to enter the password in encrypted form.

If you do not assign a BindDN value, a NULL value is used.

4. Create a UserProfile by entering:

```
ADD -LDAP UserProfile "<user_name>" "<service_name>" <SysInit/LogOn> <User/System>
```

Step example:

To download the firewall policies for the user xxx at the time of logon, enter:

```
ADD -LDAP UserProfile "xxx" "firewall" LogOn User
```

To download the firewall system configuration for the user system_xxx, enter:

```
ADD -LDAP UserProfile "system_xxx" "firewall" SysInit System
```

5. Add the FireWall Service in the LDAP list.

Step example:

```
Add -LDAP LocalOrLdap "firewall"
```

6. If you configure LDAP for the first time on the device, from the command line, enter:

```
REStart
```

Running this command downloads all the firewall policies for the users configured in the `UserProfile` for which the value of `SysInit/LogOn` is set to `SysInit`.

34.4

LDAP Working with Firewall Clients

In a typical LDAP model for an enterprise, you can create the policies offline on an NMS/workstation and store them on the LDAP server. The LDAP server can either be owned by an enterprise or it may reside outside an enterprise.

The network management station creates the security policies for its enterprise users and stores them on the directory server. You can then download the policies on the device controlling the various desktops. There are two possible ways to download these policies on the device:

- At the device initialization time.

This approach is more suitable when the policies are group-based (same policies for a group of people). In this case, the port/system based policies should be downloaded on the device at the time of initialization.

- When the remote user dials into the network.

This policy is more suitable when users have different security policies. When the remote user dials-in, the EOS device does the authentication. The authentication information can either be stored locally or on the remote RADIUS server. When the user authentication is successful, LDAP server is contacted to get the firewall policies for the user. User authentication is not considered complete until firewall policies are obtained.

The LDAP policy manager provides a transparent interface for the Firewall component to interact with the LDAP client to store its policies on the LDAP server. This section describes how to configure an EOS device to use LDAP to download the Firewall configuration from the LDAP server.

34.4.1

Attributes and Object Classes

There is no standard way to load the extended schema definitions on an LDAP server. Different server vendors may have different mechanisms to do this. This section describes how to load the EOS firewall schemas on the Netscape directory server.

Instructions for how to configure the Netscape directory server are beyond the scope of this document; however, this section does provide procedures for adding the attributes and object classes as well as for populating the LDAP Server database.

The schema of a directory defines the set of objects that can be created in that directory and the set of attributes that can be used to describe those objects. Extension of a directory schema is required when the standard schema supplied with the directory server does not meet your need. You can define a schema by defining new attributes, by adding attributes to an object type, or by defining a new object type.

To store the firewall configuration on the LDAP server, the following attributes and object classes are defined:

router-cmd, serviceType, policyName

The attribute definitions are as follows:

```
router-cmd ATTRIBUTE
  WITH ATTRIBUTE-SYNTAX
    octetStringSyntax
:: = {attributeType 1}
servicetype ATTRIBUTE
  WITH ATTRIBUTE-SYNTAX
    octetStringSyntax
:: = {attributeType 2}
policyname ATTRIBUTE
  WITH ATTRIBUTE-SYNTAX
    octetStringSyntax
:: = {attributeType 3}
The object class created are:
policy, firewallpolicy
The definitions of the object classes are as follows:
policy OBJECT-CLASS
  SUBCLASS OF top
  MUST CONTAIN {
    policyName}
  MAY CONTAIN {
    description,
  }
:: = {objectClass 1}
firewallpolicy OBJECT-CLASS
  SUBCLASS OF policy
  MUST CONTAIN {
    router-cmd,
    serviceType}
:: = {objectClass 2}
```

You should use this schema to access a directory via LDAP to search or retrieve directory information for the FireWall service.

34.4.2

Adding Attributes

Procedure:

1. Go to the Netscape Directory Server screen and click the **Schema**.
2. In the **Attribute Name** window perform the following actions:
 - a. Type in: policyname

- b. Click **Add New Attribute**.
3. In the **Attribute Name** window perform the following actions:
 - a. Type in: `router-cmd`
 - b. Click **Add New Attribute**.
4. In the **Attribute Name** window perform the following actions:
 - a. Type in: `servicetype`
 - b. Click **Add New Attribute**.

34.4.3

Creating an Object Class

Procedure:

1. Go to the Netscape Directory Server screen and click **Schema**.
2. Click **Create Object class**.
3. Create a new object class called `policy` by performing the following actions:
 - a. In the **ObjectClass Name** window, type in: `policy`.
 - b. Select **person** as the parent object class.
 - c. Add the `polycname` attribute.
 - d. Click **Create New ObjectClass**.
4. Create a new object class called `firewallpolicy` by performing the following actions:
 - a. In the **ObjectClass Name** window, type in: `firewallpolicy`.
 - b. Select **policy** as the parent object class.
 - c. Add the `router-cmd` and `servicetype` attributes.
 - d. Click **Create New ObjectClass**.

34.4.4

Database Populating on the LDAP Server

Before you can download the firewall configuration from the LDAP server, you must add the configuration into the LDAP server database.

The first step towards populating the database is to select the directory information tree. The LDAP directory services organizes directory entries using a Directory Information Tree (DIT).

The design that the network/system administrator uses for the DIT is determined by the types of information that must be stored in the directory; for example, the physical nature of the enterprise and the applications that are used with the directory services.

As is the case with many directory services and file systems, the DIT is usually viewed as an inverted tree. The root of the tree is at the top and individual directory entries are at the lowest points on the tree. The root of the tree is represented by a special entry whose distinguished name (DN) is called the directory suffix. An important part of the DIT tree design is selecting this DN.

The best practice for choosing a directory suffix is to attempt to align it with a DNS name or internet domain name associated with an enterprise. This is the best approach regardless of the nature of the enterprise.

For example, if you wanted to store the Firewall policy for each of its users in Motorola Solutions, you could define the base tree or the root DN and store the policies on a per-user basis. However you can define the tree in any form; for example, you could store the policies on the basis of devices, locations, or IP addresses.

After you selected the appropriate directory tree, you can add the database using the LDAP Data Interchange Format (LDIF) defined in the Netscape directory server.

34.4.5

Database Creation in LDAP Data Interchange Format

LDAP Data Interchange Format (LDIF) consists of one or more entries separated by a blank line. Each LDIF entry consists of an optional entry ID, a required distinguished name, one or more object classes, and multiple attribute definitions.

The basic form of an LDIF entry is:

```
[ <id>]
dn: <distinguished name>
objectClass: <object_class>
objectClass: <object_class>
...
<attribute_type>: <attribute_value>
<attribute_type>: <attribute_value>
```

Only the DN and at least one object class definition are required. In addition, any attributes required by the object classes that are defined for an entry must also be defined on the entry. All other attributes and object classes are optional.

LDIF is a standard way of adding entries in the Netscape directory server. For more details, see the appropriate Netscape directory server manuals.

The following three entries describe the LDIF fields to add the firewall rule “add !1 -fw POPIn permit from 10.1.1.1” for the user called user10 in the tree whose root DN is set to “o=Moca, c=us”. The firewall rule is applied to a fixed port, hence this rule should be downloaded at the system initialization time.

```
#Add user name user10
dn: cn=user10, o=Moca, c=us
cn: user10
sn: user10
userpassword:user10
objectclass:top
objectclass:person
#Add a new policy called fwpolicy
dn: policyname=fwpolicy, cn=user10, o=Moca, c=us
cn:user10
sn:user10
userpassword:user10
policyname: fwpolicy
objectclass:top
objectclass:policy
# Add a firewall rule "add !1 -fw POPIn permit from 10.1.1.1"
dn: router-cmd=add !1 -fw POPIn permit from 10.1.1.1,
policyname=fwpolicy,
cn=user10, o=Moca, c=us
cn:user10
sn:user10
policyname:fwpolicy
router-cmd:add !1 -fw POPIn permit from 10.1.1.1
servicetype:firewall
objectclass:top
objectclass:firewallpolicy
```

The following four entries describe the LDIF fields to add the firewall rule “add !user_xxx -fw rp permit from 30.1.1.1” and “add !user_xxx -fw TraceRouteIn permit from 30.1.1.1” for the user called user_xxx in the tree whose root DN is set to “o=Moca, c=us”. The firewall rule is applied to a dynamic virtual port. This rule should be downloaded at the user logon time.

```
#Add user name user_xxx
dn: cn=user_xxx, o=Moca, c=us
cn: user_xxx
sn: user_xxx
```

```

userpassword:user_xxx
objectclass:top
objectclass:person
#Add a new policy called fwpolicy
dn: policyname=fwpolicy, cn=user_xxx, o=Moca, c=us
cn:user_xxx
sn:user_xxx
userpassword:user_xxx
policyname: fwpolicy
objectclass:top
objectclass:policy
#Add a firewall rule 'add !user_xxx -fw rp permit from
30.1.1.1
dn:router-cmd=add !user_xxx -fw rp permit from 30.1.1.1,
policyname=fwpolicy,
cn=user_xxx, o=Moca, c=us
cn:user_xxx
sn:user_xxx
policyname:fwpolicy
router-cmd:add !user_xxx -fw rp permit from 30.1.1.1
servicetype:firewall
objectclass:top
objectclass:firewallpolicy
#Add a firewall rule 'add !user_xxx -fw TraceRouteIn permit
from 30.1.1.1'
dn:router-cmd=add !user_xxx -fw TraceRouteIn permit from
30.1.1.1, policyname=fwpolicy, cn=user_xxx, o=Moca, c=us
cn:user_xxx
sn:user_xxx
policyname:fwpolicy
router-cmd:add !user_xxx -fw TraceRouteIn permit from
30.1.1.1
servicetype:firewall
objectclass:top
objectclass:firewallpolicy

```

In general, to create a directory using LDIF on the Netscape directory server, follow these guidelines:

- Create an ASCII file containing the entries that need to be added in LDIF format. Each entry should be separated from the next by an empty line.
- Begin each directory in the database with the topmost (root) entry. The root point of the directory must represent a suffix set for the directory server. The suffix actually specifies the distinguished name used for the local database. Incoming queries must have a suffix matching this value. For our example, this is o=Moca, c=us.
- When creating a directory, make sure that an entry representing a branch point is created before you attempt to create new entries under that branch. In our example, you should create the following branch points, in order before storing the Firewall policies:

```

dn: cn=user10, o=Moca, c=us
<list_of_attributes_and_object_classes>
dn: policyname=fwpolicy, cn=user10,
o=Moca c=us
<list_of_attributes_and_object_classes>

```

Chapter 35

GGSN Management Configuration

This chapter describes how to configure Gateway GPRS Support Node (GGSN) functionality on EOS devices. GGSN functionality enables data capability, thereby allowing a mobile subscriber to access the customer enterprise network (CEN) to which it belongs.

GGSN functionality is supported on both MNR S6000 routers and GGM 8000 gateways.

In addition to basic GGSN functionality, EOS supports the following optional GGSN features:

Dynamic System Resilience (DSR)

Implements APN to multiple IPIP binding. When the link for one of the tunnels fails, the GGSN switches over to the other IPIP tunnel, thereby preserving connectivity between the GGSN and the CEN. For more information, see [APN Configuration to Multiple IPIP Binding on page 614](#).

High Availability GGSN

Implements GGSN hardware redundancy, by using the VRRP protocol, and a database synchronization protocol that provides high availability across failures by enabling the active GGSN to copy the context information to the standby GGSN on a real-time basis. For more information, see [Configuring Devices as HA GGSN Peers on page 614](#).

Conventional Data Redundancy (HA and DSR)

Provides high availability conventional data and dynamic system resilience (HA and DSR). It ensures reliable data delivery by preventing single points of hardware failure north of an RF site from causing loss of data functionality within a conventional data system. The HA and DSR feature provides conventional data failover recovery mechanisms for both single zone and multiple zone data systems, and eliminates both the GGSN and the PDG as single points of failure in a conventional data system. For more information, see [Conventional Data Redundancy \(HA and DSR\) Configuration Example on page 636](#).

GGSN Charging

Allows the GGSN to send billing information about GTP usage to a charging gateway. For more information, see [Configuring GGSN Charging on page 659](#).

Dynamic DNS (DDNS) Update

Allows you to retrieve the IP address of a particular mobile subscriber from the DNS server on the CEN. For more information, see [GGSN Configuration for Dynamic DNS on page 660](#).

GGSN Hardware Redundancy

This feature is implemented on Dimetra systems only. It enables the MNR GGSN to be deployed with a hot standby — if the active GGSN fails, the standby GGSN takes over. For more information, see [Configuring GGSN Hardware Redundancy for Dimetra Systems on page 661](#).

Overload Protection

Monitors device CPU utilization, memory usage, queue drops, and informs registered applications (OPM clients; in this case, GGSN) when any of these parameters meets or exceeds a configured threshold. For more information, see [Overload Protection Configuration on page 752](#) and “GGSN Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Motorola Solutions GGSN functionality requires the following:

- MNR S6000 router or GGM 8000 gateway with an IP path to each Packet Data Gateway (PDG) and Customer Border Router (CBR).
The PDG consists of the Packet Data Router (PDR) and the Radio Network Gateway (RNG). It performs the functions of a Serving GPRS Support Node (SGSN) and also handles mobility.
- Enterprise OS (EOS) software GW or GS package.
The EOS software PW and PS packages do not support GGSN functionality.

- EOS software configuration (GGSN service) with the following:
 - Virtual ports to the CEN domain configured.
 - Access point names (APNs) configured to map the virtual ports to the CENs.
 - APN profiles created to configure IP address allocation.
 - GGSN control enabled.

The GGSN software configuration is contained in the GGSN configuration file (`XGSN.cfg`). For more information, see [GGSN Configuration File Management on page 601](#).

35.1

Basic GGSN Functionality

GGSN functionality enables a Global System for Mobile communication (GSM) network to support packet data capability. The GGSN tunnels packet data through private networks to mobile subscribers, thereby allowing the mobile subscribers to access the customer enterprise networks (CENs) to which they belong.



NOTE: A mobile subscriber is typically comprised of a mobile computer attached to a mobile radio via a serial or USB connection. The mobile computer is unaware of its mobility; the mobile radio performs all mobility tasks on behalf of the mobile computer.

35.1.1

Data Flow

When a mobile computer needs to access its CEN, the following occurs:

- The mobile computer starts a PPP session with the mobile radio. Various protocols are employed between the mobile radio and the PDG.
- The PDG combines the information carried by these protocols with the information provided in the Home Location Register (HLR) within the User Configuration Server (UCS) and creates a GPRS Tunneling Protocol (GTP) tunnel to the GGSN. This tunnel is identified by a unique tunnel ID (TID) and is used to provide a data path between the mobile subscriber and its CEN. Each tunnel is created for the exclusive use of one mobile subscriber using one packet data protocol (PDP) and is also called a PDP context. The interface between the PDG and the GGSN is called the Gn interface.
- If the mobile subscriber does not have a proposed IP address, preassigned by the UCS before the PDP context create reaches the GGSN, the GGSN allocates an address through an “internal” DHCP mechanism or requests an IP address from an external DHCP or RADIUS server. For details, see [Configuring APNs on the GGSN on page 602](#).
- The response to the PDG carries the IP address assigned to the mobile subscriber. The PDG in turn conveys the IP address to the mobile subscriber via a sequence of protocols until finally the mobile computer learns the address via PPP.



NOTE: To prevent duplicate active IP addresses, the GGSN checks existing PDP contexts of the same APN to make sure the address it is about to allocate (or the proposed address) is not already in use.

- The GGSN delivers data packets received from the PDG to the CEN through an IP-IP tunnel to the Customer Border Router (CBR). You set up IP-IP tunnels by creating IP-IP type virtual ports and specifying access point names (APNs) to associate each port with a CEN. Each IP tunnel termination is a virtual port on the GGSN and requires its own IP address. A GGSN may support several APNs. Each APN belongs to a single CEN, and each mobile subscriber is assigned to an APN to identify the CEN to which it belongs. The interface between the GGSN and the CEN is called the Gi interface.
- The GGSN receives data packets from the CEN via the IP-IP tunnel (Gi interface) and forwards the packets to the PDG via the appropriate GTP tunnel (Gn interface) based on the destination IP address


contained in the packet. The incoming virtual port identifies the APN, and the destination IP address in the data packet identifies the PDP context (GTP tunnel) to which the packet should be forwarded.



NOTE: GTP can support several protocols (IP, PPP, and IPv6). However this release of the Motorola Solutions GGSN supports only the IP protocol on the Gn interface and the IP-IP protocol on the Gi interface.

The following figure summarizes the data flow and DHCP/RADIUS address allocation in GGSN functionality.

Figure 176: GGSN Functionality – Data Flow and IP Address Allocation



35.1.2

RADIUS Re-Authentication

If the GGSN receives a context create request after a context already exists and is in the UP state, the GGSN performs a RADIUS re-authentication. If the authentication fails, the GGSN fails the context. During the re-authentication process, the GGSN continues to forward the payload traffic.

35.1.3

RADIUS Authentication Server Switchover Algorithm

EOS supports the following RADIUS authentication server switchover algorithm:

- Initially the GGSN RADIUS client sends all of its authentication and IP address requests to the primary authentication server.
- If the GGSN RADIUS client does not receive a response to a request within the amount of time specified by the `AAAReXmitTime` parameter, it repeats the request to the primary server for the number of times specified by the `AAAReXmitCnt` parameter. If all these requests fail, the GGSN RADIUS client sends the request once to the secondary authentication server.
- If the secondary server does not respond within the amount of time specified by the `AAAReXmitTime` parameter, the authentication fails.
- After the first such failure of the primary authentication server, the GGSN RADIUS client still attempts to authenticate or obtain IP addresses first with the primary server, attempting to authenticate or obtain IP addresses via the secondary server only after the number of retries specified with the `AAAReXmitCnt` parameter failed.
- If the GGSN RADIUS client tries and fails to authenticate via the primary authentication server for the number of times specified by the `AAASwitchCnt` parameter in a row (for the same or different mobile subscribers), the next authentication request initiates a switchover to the secondary authentication server. The GGSN RADIUS client then authenticates exclusively via the secondary server until the primary server comes back up again and the switchover algorithm forces the GGSN RADIUS client to switch back to the primary server.



NOTE: If you want the GGSN RADIUS client to authenticate via the primary server again when it comes back up, you should unplug the secondary server until at least the number of authentication requests specified by the `AAASwitchCnt` parameter failed. The failure scenario described above occurs, and the GGSN RADIUS client switches back to the primary authentication server.

To determine which authentication server that the GGSN RADIUS client is using currently, you must use the following command:

```
SHoW !<APN> -GGSN CENSTATistics
```

where `<APN>` is the name of the APN. The APN statistics include a `RADIUS` entry, which lists the IP address of the current RADIUS authentication server.

35.1.4

DHCP Server Switchover Algorithm

EOS supports the following DHCP server switchover algorithm:

- Initially the GGSN unicasts a DHCP discover packet to obtain an IP address for a mobile subscriber from the primary DHCP server, and the specified primary DHCP server is designated as the current server.
- If the primary server does not respond within the amount of time specified by the `DHCPReXmitTime` parameter, the GGSN repeats the request to the primary server for the number of times specified by the `DHCPReXmitCnt` parameter. If all these requests fail, the GGSN sends the request once to the secondary DHCP server.

- If the secondary server does not respond within the amount of time specified by the `DHCPReXmitTime` parameter, the IP address request fails.
- After the first such failure of the primary DHCP server, the GGSN still attempts to obtain an IP address first from the primary server, attempting to obtain IP addresses from the secondary server only after the number of retries specified with the `DHCPReXmitCnt` parameter failed.
- If the GGSN tries and fails to obtain an IP address from the primary DHCP server for the number of times specified by the `DHCPSwitchCnt` parameter in a row (for the same or different mobile subscribers), the next mobile subscriber request for an IP address initiates a switchover to the secondary DHCP server. The GGSN then designates the secondary DHCP server as the current server and unicasts subsequent DHCP discover packets initially to the secondary server. This switchover occurs only if the number of IP address requests to the primary server specified by the `DHCPSwitchCnt` parameter in a row failed. For example, if `DHCPSwitchCnt` is set to 5, The sixth request initiates a switchover only if all of the first five requests failed.

If you take the primary DHCP server down – for a software upgrade, for example – and it stays down long enough for the number of IP address requests specified by the `DHCPSwitchCnt` parameter in a row to fail, the GGSN designates the secondary DHCP server as the *current server*. If you want the primary server to be the current server again when it comes back up, you should unplug the secondary server (which is now the *current server*) until at least the number IP address requests specified by the `DHCPSwitchCnt` parameter failed. The failure scenario described above occurs, and the GGSN designates the original primary DHCP server as the current server once again.

To determine which DHCP server that the GGSN is using currently, you must use the following command:

```
SHoW !<APN> -GGSN CENSTATistics
```

where `<APN>` is the name of the APN. The APN statistics include a DHCP entry, which lists the IP address of the current DHCP server.

35.2

APN to Multiple IPIP Binding

EOS devices support multiple IPIP tunnels per APN for redundancy. This feature supports the dynamic system resilience (DSR), HA GGSN, and conventional data redundancy (HA and DSR) features.



NOTE: Multiple IPIP tunnels are supported only for the purpose of redundancy, and only one tunnel is active at a time.

To bind an APN to multiple IPIP tunnels, you must specify a bidirectional forwarding detection (BFD) gateway IP address (the IP address of the border router) and a priority value for each tunnel when you create the APN using the GGSN service `AccessPointName` parameter.

BFD maintains the link status for each tunnel and informs the GGSN and GGSN software when a link comes up or goes down. When the GGSN receives a status change notification, it transparently uses the active IPIP tunnel with the highest priority value to connect to the CEN. If that tunnel fails, the GGSN switches over to the other IP-IP tunnel until the higher priority link is reestablished. The GGSN sends a trap to the configured network management station both when the GGSN connection to the CEN was established and when the GGSN connection to the CEN goes down. For details about the traps, see [Supported MIBs and Traps on page 845](#).

You can bind either up to four IP-IP tunnels with BFD or one IP-IP tunnel without BFD per APN. The following table describes the rules for assigning IP-IP tunnels to an APN based on what type of tunnel configuration (if any) already exists on the APN. As shown in the table, an error message is returned if you attempt any of the following IP-IP tunnel assignments:

- Assigning an IPIP tunnel without BFD when an IP-IP tunnel with BFD is already bound to the APN.
- Assigning an IPIP tunnel with BFD when an IP-IP tunnel without BFD is already bound to the APN.
- Assigning an IPIP tunnel with BFD when four IP-IP tunnels with BFD are already bound to the APN.

Table 60: APN to Multiple IPIP Tunnel Assignment Rules

New IP-IP Tunnel Assignments	IP-IP Tunnels Currently Defined for the APN					
	<i>None</i>	<i>1 IPIP Tunnel without BFD</i>	<i>1 IPIP Tunnel with BFD</i>	<i>2 IPIP Tunnels with BFD</i>	<i>3 IPIP Tunnels with BFD</i>	<i>4 IPIP Tunnels with BFD</i>
None	None	1 IPIP Tunnel without BFD	1 IPIP Tunnel with BFD	2 IPIP Tunnels with BFD	3 IPIP Tunnels with BFD	4 IPIP Tunnels with BFD
1 IPIP Tunnel without BFD	1 IPIP Tunnel without BFD	1 IPIP Tunnel without BFD ¹	Error	Error	Error	Error
1 IPIP Tunnel with BFD	1 IPIP Tunnel with BFD	Error	2 IPIP Tunnels with BFD	3 IPIP Tunnels with BFD	4 IPIP Tunnels with BFD	Error

¹ The new IPIP tunnel without BFD replaces the existing IPIP tunnel without BFD.

35.3

HA GGSN Feature

The HA GGSN feature provides GGSN hardware redundancy via the VRRP protocol. It also provides a new database synchronization protocol with high availability across failures. The active GGSN can copy the context information to the standby GGSN.

In addition, EOS supports traps reporting the redundancy status of the GGSN; the status of the connection between the active GGSN and the border router (BR) in the CEN; and the status of the context database synchronization. For details about the traps, see [Supported MIBs and Traps on page 845](#).

GGSN Hardware Redundancy

GGSN hardware redundancy enables the MNR GGSN to be deployed with a hot standby – when the active GGSN fails, a standby GGSN takes over. The redundant GGSN pair appears as a single device to outbound and inbound neighbors (the packet data gateway (PDG) and border router (BR)) but as two devices to network management and the optional GGSN charging service. The fundamental building block of GGSN hardware redundancy is the VRRP protocol.

- HA GGSN peers are configured as backup VRRP routers with different priorities to ensure that there is a predictable VRRP master when both GGSNs boot up at the same time. The active and standby GGSNs share a LAN on both the inbound and outbound interfaces. A GGSN must be the VRRP master on both LAN interfaces before it can become the active GGSN of the redundant GGSN pair.
 - On the outbound LAN interfaces, IP addresses known as the "service IPs" are shared between the two GGSNs. The "service IPs" are IP addresses on the subnets connecting the GGSN to the CEN and are specified, along with VRIDs, when the GGSNs are configured as backup VRRP routers for those subnets. The GTP protocol from the PDG communicates solely with the "service IPs", and the PDGs send all echoes and context activations/data to these addresses. The GGSN that is the VRRP master owns the "service IPs". If the VRRP master goes down, the VRRP standby takes over these IP addresses.
 - On the inbound LAN interface, IP-IP tunnels are configured on each APN (to the BRs). The IP-IP virtual ports over which the IP-IP tunnels run are associated with VRID@port number sessions corresponding to the VRIDs specified with the backup VRRP configuration. This identifies the IP-IP virtual ports as a VRRP-controlled virtual ports, with GGSN as the owner, and IP-IP notifies GGSN of

the status of the IP-IP tunnels. The tunnels are UP on the GGSN that is the VRRP master and DOWN on the GGSN that is the VRRP standby. The GGSN receiving the context activations/data from the PDG is also the GGSN that forwards/receives the data from the BR. Only the VRRP master originates the packets to the PDG.

- VRRP switchovers are managed by using event-based macros
 - When IP-IP notifies GGSN that the IP-IP tunnels corresponding to the VRRP-controlled virtual ports are UP, a macro (`port_up`, for example) is invoked, and VRRP control is enabled on the interfaces on which VRRP backup functionality was configured.
 - When IP-IP notifies GGSN that the IP-IP tunnels corresponding to the VRRP-controlled virtual ports are DOWN, a macro (`port_down`, for example) is invoked, and VRRP control is disabled on the interfaces on which VRRP backup functionality was configured.

GGSN Context Database Synchronization

EOS implements a context database synchronization protocol to support the HA GGSN feature. This is a separate protocol that runs over UDP between HA GGSN peers. The packets are sourced and received by the system IP address (SysIP) of the HA GGSN peers. In the event of a failure and switchover, the database synchronization protocol replicates the active context database from the active GGSN to the standby GGSN with no loss of contexts during the switchover.

When two GGSNs are established as HA GGSN peers (a valid peer SysIP address is configured on each GGSN using the GGSN service `HAStandBy` parameter and the GGSN service is enabled), EOS begins context database synchronization between the two peers, with the active/standby roles of the GGSNs determined by VRRP mastership.

The context database synchronization protocol employs six types of messages:

Heartbeat messages

The HA GGSN peers send these messages to each other. The heartbeat messages sent from the active GGSN to the standby GGSN list the packet data gateway (PDG) with which it is communicating and maintain a recovery number per PDG link. The PDG recovery number remains unchanged as long as the PDG is reachable by the GGSN. It changes during premature switchover when the standby GGSN becomes active before it is fully synchronized with the old active GGSN. The heartbeat messages also include a GGSN recovery number that corresponds to the sender and changes only after the sender reboots. Receiving a heartbeat message with a changed GGSN recovery number indicates to the receiving GGSN that its peer rebooted and the context database needs to be resynchronized.

Single context create/delete messages

Every time the active GGSN receives a context create message from the PDG, it replicates it on the standby GGSN and waits for an acknowledge (Ack) message. If the Ack is not received within a preconfigured number of seconds, the standby is considered down. When the active GGSN receives a context delete message, it immediately deletes the context and sends a context delete message to the standby GGSN.

Bulk context create messages

These messages are similar to the single context create messages, except there are multiple contexts in each message. Bulk context updates are performed in the following scenarios:

- The standby GGSN comes up and the active GGSN has active contexts.
- The standby GGSN comes up from a down state after missing heartbeat messages.

Bulk context create messages bundle multiple contexts from different PDGs, which a maximum packet size of 1200 bytes. The standby GGSN acknowledges the bulk context create messages with bulk Ack messages.

Context create acknowledgment success/failure messages

The standby GGSN sends a context create acknowledgment success message when it successfully received and processed a context create message from the active GGSN. The standby sends a context

create acknowledgment failure message when the context create message it receives from the active GGSN contains an error.

Context flush messages

These messages are sent by both the active and the standby GGSNs under the following conditions:

- When both HA GGSN peers come up in a down state, receive heartbeat messages, and transition to two-way communication, the active GGSN sends a context flush message to the standby GGSN.
- When the standby GGSN reboots and changes from down state to two-way state, the active GGSN sends a context flush message to the standby GGSN.
- When recovering from a split brain scenario, both HA GGSN peers send a context flush messages to clean out each other's databases when the state changes from down to two way.

Context flush acknowledge messages

These messages are sent by an HA GGSN peer after it flushed its contexts to acknowledge that the contexts were flushed.

35.4

Conventional Data Redundancy (HA and DSR) Feature

The conventional data redundancy (HA and DSR) feature ensures reliable data delivery by preventing single points of hardware failure north of an RF site from causing loss of data functionality within a conventional data system. The conventional data redundancy (HA and DSR) feature provides conventional data failover recovery mechanisms for both single zone and multiple zone data systems and eliminates both the GGSN and the PDG as single points of failure in a conventional data system.

In a data system, subscriber calls come in to the GGSN through a packet data gateway (PDG). The calls are authenticated, and an IP connection is established to the customer enterprise network (CEN). The conventional data redundancy (HA and DSR) feature implements a mobile IP protocol that enables HA conventional data to function properly when a subscriber roams to a different zone and loses communication with the originating PDG, or when a subscriber switches master sites due to a DSR failover. In both these scenarios, traffic between the subscriber and the CEN is re-routed via tunnels.

35.4.1

Additional Transport Network Redundancy

To support the conventional data redundancy (HA and DSR) feature, the transport network now includes the following redundant elements:

- Network connectivity for up to two conventional Packet Data Gateways (PDG) per zone core.
- Up to 2 GGSNs per zone core
 - 2 HA GGSNs that support one zone or the system in an HA, non-DSR system.
 - 2 HA GGSNs that act as the primary core GGSNs for one zone and as the backup core GGSNs for another zone in an HA and DSR system.
- Additional APN tunnels – that connect a border router and each GGSN zone network management (ZNM) interface by using a firewall, DMZ switch and optional peripheral network router.
 - 2 tunnels per zone in an HA, non-DSR system.
 - 4 tunnels per zone in an HA and DSR system.
- Up to two CN1 tunnels per zone core between each GGSN data interface and a border router, through the network transport in another zone core. These tunnels exit one zone core and enter the another through the exit routers. In the second zone core, the tunnels are routed through one of the CN1 paths – ultimately to a border router. DSR is required for these inter-zone tunnels, and a total of four such tunnels are supported for a DSR zone (2 originating in each zone core).

- Concurrent conventional/trunked/HPD HA data and conventional/trunked/HPD DSR data traffic.



NOTE: Starting from the 2022.HS and 2022.1 system releases the HPD feature is not supported.

- Use of bi-directional forwarding detection (BFD) to detect failure of the active CNI tunnel and switch to an alternate tunnel automatically.
- An additional VLAN in the CEN to segment the border router mobile IP traffic from the other applications in the CEN.

35.4.2

Mobile Agent Functionality

The conventional data redundancy (HA and DSR) feature introduces the concept of mobility agents. Home agent functionality is implemented on the border router. Foreign agent functionality is implemented on the GGSN. Mobility agent functionality enables EOS devices to track subscriber movements and route data back to subscribers.

Mobile IP functionality is triggered as a service by other EOS components that function as clients. When the GGSN receives a context creation trigger from the PDG (serving subscribers), the GGSN triggers a `MOBILEIP` registration request.

Mobile agent functionality allows Mobile IP configurations to ensure that valid paths from the subscriber to the CEN are maintained on HA and DSR failovers. A conventional subscriber is represented by a home agent (border router) and a foreign agent (GGSN). Mobile IP messaging ensures that the path between a subscriber's home agent and foreign agent is valid and corrected for failure scenarios.

Mobile IP Foreign Agent

Mobile IP foreign agent functionality is implemented on the GGSN. Previously, the GGSN created contexts for subscribers and handled traffic flow between the CEN and the subscribers. Now, the GGSN has the added role of establishing `MOBILEIP` registration between mobility agents. This functionality is enabled with the new `MOBILEIP` service UI and is triggered by the GGSN acting as a client to `MOBILEIP`. The foreign agent functions as a gateway for context creation requests sent from the PDG. It is responsible for relaying the registration request to the home agent, setting up the tunnel to the home agent, and forwarding packets to the subscriber. With `MOBILEIP` implemented, the GGSN can only create a context for a subscriber if the `MOBILEIP` registration succeeds. Subscriber contexts and IP addresses are maintained over a DSR failover.

In an HA configuration, the two GGSNs within a zone core use one virtual IP address on the foreign agent VLAN to communicate with the border router. The foreign agent address of the GGSN maps to this virtual address, and the VRRP configuration on the foreign agent VLAN ensures that the CEN request from the border router is sent to the HA GGSN acting as the master VRRP.

Mobile IP Home Agent

Mobile IP home agent functionality is implemented on the border router. As the mobile IP home agent, the border router:

- Acts as a proxy for the subscriber contexts, acting as home agents to the subscribers' foreign agents (preferred GGSNs). The home agent maintains a table of entries for the foreign agents and the subscriber home addresses. The border router gains knowledge of the foreign agents by using the mobile node registrations sent by the foreign agents from each zone core.
- Maintains a list of mobile nodes and the corresponding foreign agent for each node. When the border router receives a new mobile node registration, it adds an entry to its list. An existing mobile node entry is updated with the most recently-received foreign agent address. The mobile node registration request includes the following key information:
 - Source address (set to the tunnel address on the GGSN)

- Destination address (set to the tunnel address on the border router)
- Home address (subscriber IP address)
- Home agent – tunnel address on the border router
- COA (care of address) – tunnel address on the GGSN
- Identifies foreign agents from mobile node entries and routes the CEN application requests to the GGSN foreign agent. Both border routers in the CEN are on a new home agent VLAN that is governed by VRRP. The home agent VLAN isolates the mobile IP traffic from the other CEN traffic and is connected to the other CEN application VLAN via the CEN switch. The home agent address on both border routers points to a floating VRRP IP address on the home agent VLAN. The GGSNs are configured with the VIP of the home agent, and the mobile node entries are synchronized between the two border routers. In the event of a CNI path failure, the border router attempts to route the request down to the GGSN using the next prioritized CNI path.
- Acts as a DHCP server that processes DHCP requests from GGSN-enabled devices.

Tunneling

Mobile IP adds an additional encapsulation to the data packets transported between the border router and the GGSN. BFD-controlled tunnels are still used for both IPv4 and IPv6 traffic, and encryption between the border router and the firewall for both IPv4 and IPv6 traffic is unchanged from previous releases. However, a new mobile IP encapsulation is added to mobile IP control messages and data messages. Messages from the GGSN to remote services such as DHCP and RADIUS do not include the mobile IP header encapsulation.

35.4.3

Support for Multi-Hop BFD Sessions Between FA and HA VRRP Addresses

To support conventional data redundancy (HA and DSR) functionality, EOS implements the multi-hop BFD sessions between the VRRP addresses of the foreign agent and the home agent. This implementation supports many-to-many connections where one home agent pair has connections to more than one foreign agent pair and one foreign agent pair has connections to more than one home agent pair. To achieve this, EOS supports more than one BFD session per port.

Foreign agent and home agent pairs are configured as the backup VRRP routers with different priorities to ensure that there is a predictable master VRRP when both devices boot up at the same time. In addition, EOS implements a support recovery number on the home agents. This number indicates to the foreign agent that the home agent is no longer synchronized with the foreign agent. This condition may occur in the following scenarios:

- The home agent reboots and cannot establish communication with a peer.
- The standby home agent becomes the active home agent when a bulk context database update is in progress and is not completed.

When the MOBILEIP service detects that either of these conditions occurs, it calls the BFD protocol to execute a restart mechanism on all configured foreign agents.

35.4.4

Subscriber IP Address Allocation Through a DHCP Server on the Border Router

DCHP server functionality, for the purpose of allocating subscriber IP addresses, was moved from the GGSN to the border router. The border router maintains a pool of DHCP addresses for all GGSNs in the system.

From the perspective of the GGSN, the DHCP server on the border router is no different than a remote DHCP server.

As a result, the GGSN supports only the two following types of subscriber IP address assignments:

Static

Subscribers have statically assigned IP addresses and the IP address proposal is made in the context create message.

Remote

The IP address assigner is a DHCP server located either in the CEN or on the border router.

The GGSN supports two DHCP servers, one that acts as the primary server, and the other that acts as the secondary server. The GGSN also supports the switchover mechanism supported in previous releases.



NOTE: The GGSN no longer supports the subscriber IP address allocation through a RADIUS server.

In addition, DHCP address pools can be specified either on a per-port or port-independent basis. For more information, see "DHCP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

35.4.5

Replication Protocol

The conventional data redundancy (HA and DSR) feature implements a replication protocol that runs as a service, on UDP port 49517, to clients that register with it. EOS replicates application databases to another EOS device, a replicator peer, that serves as a standby in case the active device becomes inoperative.

The replication protocol covers the following:

MOBILEIP client

- Replication of home agent mobility databases
- Replication of foreign agent visitor tables

DHCP client

Replication of DHCP allocations

NAT client

Replication of TCP connections, to enable sharing of NAT sessions between the border routers.

The basic operation of the replication protocol is as follows:



NOTE: The replication protocol supports one-way replication, the active peer replicates the database to the standby peer, and a two-way replication, where both peers replicate their databases to the other peer. The type of replication that is implemented depends on the registered client. The MOBILEIP and DHCP clients implement one-way replication. The NAT client implements a two-way replication.

- A client application passes a replication IP address to the replication protocol. This address is used for communication and replication data transport.
- The replication protocol registers the IP address with VRRP.
- The VRRP protocol determines the master/backup status and passes this status back to the client application. The VRRP status (master/backup) determines the replication peer status (active/standby). The replication protocol performs a three-way state exchange control protocol (handshake) between the replicator peers to establish a connection between the peers and transition to a synchronized state (CONN-SYNC). When VRRP informs the replication protocol that it is in a backup VRRP state, the replication protocol starts the three-way handshake to establish connection with the active replicator peer. The three states are:

INIT

Both replicator peers enter this state when they receive the VRRP status indicator and start the handshake.

CONN-SYNC

Both replicator peers transition to this state when they establish connection and begin bulk data synchronization if there is data pending.

FULL

Both replicator peer enters this state when the bulk data synchronization is complete.



NOTE: As soon as the replicator completes the 3-way handshake with its replication peer, it begins periodic transmission of heartbeat packets. The presence of these packets indicates to replicator that its peer is available for replication. Absence of these packets causes the replicator to stop replication.

When the client application sends the replication data, it includes a predefined application ID to distinguish the application from other applications using the same replication service.

Application database updates can be performed as single updates or bulk updates.

- In the case of a single update, when an active replicator peer, that is a device for which the VRRP status is displayed as *master*, is in the *FULL* state, and an application adds an entry to its database, the active replicator attempts to replicate the data to its peer.
- In the case of a bulk update, when a standby replicator peer, a device for which the VRRP status is displayed as *backup*, learns of its status through the VRRP, the replication protocol initiates a connection with its active, master, peer.
 - When the connection is established, the standby replicator flushes its application database and waits for the data synchronization from the active replicator.
 - The active replicator sends any application data in its database in a bulk form, where one packet may contain multiple entries of the application's database, to the standby replicator.
 - When all entries are transported, the active replicator sends an indicator of bulk completion with a final data packet and at that point both replicator peers transition to the *FULL* state.

The EOS *REPL*ication service supports a *ReplRetry* parameter that allows you to configure a replication retry count and retransmission interval. For more information, see the “REPLicator Service Parameters” chapter in the *Enterprise OS Software Version 18.2 Reference Guide*.

- When the replicator does not receive a response from its replication peer within the configured retransmission interval, it retransmits the packets.
- If the configured retransmission interval elapses again, and the replicator still does not receive a response from its replication peer, it retransmits the packets again, up to the retry count configured with the *ReplRetry* parameter.
- The connection is determined to be down if the replicator does not receive a response from its replication peer after the maximum number of retries. At this point, the replicator does not replicate any additional data on the peer until the connection is established again.

NAT Redundancy

When NAT registers as a client of the replication protocol, TCP connections are replicated between the border routers (home agents) in a conventional data redundancy (HA and DSR) configuration. Replication of TCP connections enables the border routers to share the NAT sessions table and avoid dropping packets in TCP traffic streams since both border routers may be active at a given time.

When the border routers boot up, one is the master VRRP and one is the backup VRRP; however, both are considered active since from the perspective of the NAT client both can receive traffic at any given moment.

- At bootup, each border router replicates its TCP connection table and sends it to its replicator peer (the other border router).
- When the backup replicator peer learns of its status via VRRP, it flushes its TCP connection table and, the active replicator peer synchronizes its TCP connection table with the backup replicator peer.

TCP connection table updates are performed as single updates or bulk updates.

35.4.6

Conventional Channel Gateway Support for Conventional Data DSR

The Conventional Channel Gateway (CCGW) supports a complete DSR solution for conventional systems. Conventional data DSR support and a conventional voice DSR. For more information, see [CCGW Support for Conventional Data DSR on page 744](#).

35.4.7

Support for Intersystem Data-Connected CEN


The intersystem data-connected CEN feature enables a mobile subscriber from any CEN to roam to another ASTRO system and maintain connectivity back to its original CEN, and the associated packet data applications. The EOS OSPF service supports configurations, pertaining to Type 5 and Type 7 link state advertisements (LSAs) that enable ASTRO systems to share the same CEN LAN in a way that the traffic intended for one system is not routed to the other system under any failure and recovery scenario. For more information about the OSPF parameters that support this feature, see the descriptions of the `DirectPolicy`, `ExteriorPolicy`, `InteriorPolicy`, `NSSADefMetric`, `StaticPolicy`, and `NSSACont` parameters in “OSPF Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

35.5

GGSN Charging

As illustrated in the following figure, GGSN charging runs over the Ga interface, which can be located in the DMZ or the RNI. The charging gateways collect billing records from the GGSN and store them until they are retrieved by a billing service.

Figure 177: GGSN Charging System Architecture



A billing protocol called GTP Prime (GTP') is terminated in the GGSN and in the Charging Gateway. EOS support GTP' version 0, with 20-byte GTP' headers. Through the GTP' protocol, the GGSN reports the resource consumption on the external network by each of the mobile subscribers.

Every PDP context is assigned a unique charging identity number for billing purposes during context establishment, and a GGSN call detail record (G-CDR) is opened for each activated PDP context. As soon as a PDP context is established, the GGSN creates a traffic data volume (TDV) and starts collecting the charging information. Each G-CDR may contain multiple TDVs.

A TDV is closed when any of the following events occurs:

- Tariff period changes – You can configure up to 12 tariff periods. At each configured tariff time, the TDV is closed and another TDV is opened to log further charging information.
- Maximum record size reached – You can configure a maximum size for the G-CDR's aggregated data volume over the TDV (or multiple TDVs). When this size is reached, both the current TDV and the G-CDR are closed.
- Periodic DRT transfers – You can configure the interval at which the GGSN sends data record transfer (DRT) requests to report the charging information for each of the PDP contexts. When the configured period ends, the GGSN reports the collected charging information to the Charging Gateway and closes the current TDV and the G-CDR.
- PDP context terminated – The PDP context is terminated if any of the following occurs:
 - The mobile subscriber hangs up.
 - The connection to the PDR is lost.
 - GGSN (systemwide or per APN) or the charging function is disabled, or an individual APN is deleted.

When the PDP context is terminated, the GGSN closes the current TDV and the G-CDR and reports the collected charging information to the Charging Gateway.

 **NOTE:** The GGSN reports data volume byte counts according to the actual data packets received. If packet larger than the `volume_limit` is received, the GGSN does not close the record at the `volume_limit` but rather reports the volume in one G-CDR.

If the primary charging becomes unreachable, the GGSN switches to the secondary Charging Gateway gateway (if one is configured) and continues reporting charging information to the secondary Charging Gateway as long as it is reachable. After a switchover, the GGSN does not attempt to poll the primary Charging Gateway, nor does it switch back automatically to the primary Charging Gateway. However the GGSN switches back from the secondary to the primary Charging Gateway if any of the following events occurs:

- GTP' echo timeout – no echo response is received from the secondary Charging Gateway after six retries.
- A node alive request message is received from the primary Charging Gateway.
- A redirection request message is received from the secondary Charging Gateway with the recommended node set to the primary Charging Gateway or with no recommended node.
- The IP address of the secondary Charging Gateway is set to 0.0.0.0.
- GGSN charging control (`CHGCONTROL`) is disabled and then enabled with the primary Charging Gateway up.

GGSN Charging Gateway Switchover Algorithm

The `CHGEchoRetryCnt` and `CHGEchoRetryTm` parameters support the following switchover algorithm:

- The GGSN sends echo requests to the primary Charging Gateway at an interval specified by the `CHGEchoTime` parameter for keepalive purposes.
- If the GGSN does not receive an echo response from the Charging Gateway, it retransmits the message at an interval specified by the `CHGEchoRetryTm` parameter.
- If the GGSN does not receive an echo response from the Charging Gateway for the number of times specified by the `CHGEchoRetryCnt` parameter consecutively, it switches over to the secondary Charging Gateway and retransmits all unacknowledged DRT requests to the secondary Charging Gateway (if one is configured).

 **NOTE:** If you take the primary Charging Gateway down – for a software upgrade, for example – the charging echo retransmission is interrupted, as described above, and force a switchover to the secondary Charging Gateway. The GGSN then sends all subsequent DRT requests to the secondary Charging Gateway (if one is configured). After the primary Charging Gateway is back online, it sends a node alive request to the GGSN, and the GGSN immediately forwards subsequent DRT requests to the primary Charging Gateway.

35.6

GGSN Dynamic DNS Update Feature

The GGSN Dynamic DNS Update (DDNS) feature allows you to retrieve the IP address of a particular mobile subscriber from the DNS server on the CEN.

The data flow is as follows:

- During the context activation, the GGSN receives a PDP context create message from the PDR.
- If the optional RADIUS authentication process passes, the GGSN brings up the context.
- The GGSN stores the IP address of the mobile subscriber, which it obtains in one of the following ways:
 - Proposed in the context create message
 - Allocated from the GGSN internal IP address pool

- Allocated from the external DHCP service
- Allocated from the external RADIUS service
- The GGSN sends out UDP-based dynamic DNS update messages to the DNS server, located in the CEN, to register the following information about the mobile subscriber:
 - Fully Qualified Domain Name (FQDN), in the following format:
`<MSISDN> <DDName>`
where:
`<MSISDN>` is the mobile subscriber international ISDN number
`<DDName>` is the configured domain name for DDNS updates (the configured `AccessPointName` if no `<DDName>` was specified).
 - Assigned mobile IP address.



NOTE: In order to allow the GGSN to send dynamic DNS update messages to the DNS server, you must configure the IP address of the DNS server on the GGSN. For details, see [APN Configuration to Multiple IPIP Binding on page 614](#).

The dynamic DNS update messages sent by the GGSN include both a resource record, which specifies the IP address of the mobile subscriber for DNS forward lookup, and a PTR resource record, which specifies the FQDN of the mobile subscriber for DNS reverse lookup.

- After sending out a dynamic DNS update request, the GGSN starts a retransmission timer. If it does not receive a dynamic DNS response, the GGSN retransmits the update request three more times, at intervals of 5 seconds, and does not deactivate the PDP context even if it exhausted all the retransmissions.



NOTE: If the GGSN receives a dynamic DNS update response in the middle of a retransmission, it stops the retransmission and frees up the retransmission timer.

35.7

GGSN Hardware Redundancy



NOTE: GGSN hardware redundancy refers to the redundant GGSN feature implemented on Dimetra systems only.

The GGSN hardware redundancy feature enables the MNR GGSN to be deployed with a hot standby – when the active GGSN fails, the standby GGSN takes over. The redundant GGSN pair appears as a single device to outbound and inbound neighbors (the packet data router (PDR) and border router (BR)), and as two devices to network management and the optional GGSN charging service. When the active GGSN in a redundant GGSN pair fails, the PDRs and BRs do not notice the change of device, although the mobile subscriber must reactivate to restore the packet data context on the new active GGSN.

The redundant GGSN feature incorporates the following enhancements to the EOS code:

Synchronized VRRP sessions

Multiple VRRP (VRID@port number) sessions are synchronized so that the sessions are either all master or all backup.

Synchronized VRRP sessions linked to an IP-IP tunnel

When you add an IP-IP virtual port, you can specify a list of VRID@port number sessions that are synchronized and control the IP-IP tunnel. IP-IP receives this list and registers it as a synchronized VRRP session with VRRP.

IP-IP tunnel status notification to GGSN

When a virtual port is configured with GGSN as the owner (VRRP-controlled), IP-IP notifies GGSN of its status.

- If all ports associated with the VRRP-controlled IP-IP tunnel are UP, GGSN sends a trap indicating that the GGSN is active.

- If all ports associated with the VRRP-controlled tunnel are DOWN, GGSN sends a trap indicating that the GGSN is in standby mode.

The fundamental building block of the redundant GGSN feature is the VRRP protocol. The active and standby GGSNs share a LAN on both the inbound and outbound interfaces. A GGSN must be the VRRP master on both LAN interfaces before it can become the active GGSN of the redundant GGSN pair.

- On the outbound LAN interface, an IP address known as the “service IP” is shared between the two GGSNs. This “service IP” is bound to the currently active GGSN. The GTP protocol from the PDR communicates solely with the “service IP.”
- On the inbound LAN interface, there are IP-IP tunnels configured to each APN. These tunnels are UP on the active GGSN and down on the standby GGSN.

35.8

GGSN Configuration File Management

When GGSN is enabled on a device, the device supports two configuration files:

- `boot.cfg` – routing, port, and path configurations
The `boot.cfg` file contains the static parts of the device configuration.
- `XGSN.cfg` – GGSN, virtual port, and APN configurations
The `XGSN.cfg` file is created by a network management application. It contains the GGSN configuration parameters and the APN configuration commands for the APNs configured on the device. These are the dynamic parts of the device configuration.

If there is no `boot.cfg` file on the GGSN, the `XGSN.cfg` file does not execute. The `boot.cfg` file executes first, followed by the `XGSN.cfg` file.



NOTE: The `XGSN.cfg` file is introduced to provide GGSN configuration manageability. If you boot a device running a software package that supports GGSN (the GW and GS packages), you should use the `XGSN.cfg` file for GGSN configuration commands. You can include GGSN configuration commands in the `boot.cfg`, but a `boot.cfg` file with GGSN configuration commands causes logging errors when used on a device with a software package that does not support GGSN.

35.9

Configuring Basic GGSN Functionality on EOS Devices



NOTE: This configuration comprises the `xsgn.cfg` file of the device.
If you configure a device for conventional data redundancy (HA and DSR) or for HA GGSN, the GGSN configuration (`xsgn.cfg` commands) should be the same on both peers.

For more information about the GGSN performance statistics, see [GGSN Service Statistics on page 971](#).

For detailed descriptions and syntax for all GGSN service parameters, see “GGSN Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Process:

1. Configure the basic GGSN parameters on the device. See [Configuring Basic GGSN Parameters on page 602](#).
2. Add and configure access points. See [Configuring APNs on the GGSN on page 602](#).

Each access point is identified by an access point name (APN) which represents a CEN. For more information, see [Configuring APNs on the GGSN on page 602](#).

The GGSN does not support subscriber IP address allocation by using a RADIUS server.

To configure a GGSN for DSR, HA GGSN, or conventional data redundancy (HA and DSR) support, you bind an APN to multiple IPIP tunnels when you create the APN. In addition, for a conventional data redundancy (HA and DSR) or HA GGSN configuration, you must use the `VRID <vrid@port-no>`, `<vrid@port-no>` syntax to configure the IPIP tunnels as VRRP-controlled tunnels.

To configure a GGSN for dynamic DNS (DDNS) functionality, you must specify the IP address of the DNS server when you configure each APN.

3. Optional: Configure the GGSN charging. See [Configuring GGSN Charging on page 659](#).
4. Optional: Configure overload protection. See [Configuring EOS Devices for Overload Protection on page 752](#).
5. Enable the GGSN functionality on the device by entering:

```
SETDefault -GGSN CONTrol = Enable
```

35.10

Configuring Basic GGSN Parameters

Procedure:

1. To view the currently-configured GGSN parameter settings, enter:

```
SHow [!<APN>] -GGSN CONFIGuration
```

where `<APN>` is the access point name, as defined by the `AccessPointName` parameter.

2. Specify the maximum number of PDP contexts that the GGSN can support by entering:

```
SETDefault -GGSN MaximumPDPContexts = <value>
```

where `<value>` is the maximum number of PDP contexts supported on the GGSN (1-65535). The default `<value>` is 20000.

3. Set the TOS (type of service) value for the signal packets on the Gn interface (the interface with the PDG (Packet Data Gateway) over which GTP runs) by entering:

```
SETDefault -GGSN SignalTos = <tos>
```

where `<tos>` is the TOS-byte value (0-255). The default `<tos>` is 0.

The TOS value is a one-byte value, with the TOS priority value expressed in the first three bits. For more information about the `SignalTos` parameter, TOS priority values, and a table that maps the `SignalTos` values to the corresponding hexadecimal, see "GGSN Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

4. Set the TOS (type of service) value for the data packets on the Gn interface (the interface with the PDG (Packet Data Gateway) over which GTP runs) by entering:

```
SETDefault -GGSN DataTos = <tos>
```

where `<tos>` is the TOS-byte value (0-255). The default `<tos>` is 0.

The TOS value is a one-byte value, with the TOS priority value expressed in the first three bits. For more information about the `DataTos` parameter, TOS priority values, and a table that maps the `DataTos` values to the corresponding hexadecimal, see "GGSN Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

35.11

Configuring APNs on the GGSN

You can use this procedure to configure APN to use:

Static subscriber IP address allocation

All requesting subscribers have statically-assigned IP addresses, pre-assigned by the User Configuration Server (UCS) before the PDP context create reaches the GGSN.

Dynamic subscriber IP address allocation via a local server (IPAssigner set to LocalDhcpServer)

Subscriber IP addresses are assigned by the GGSN, acting as an "internal" DHCP server that allocates IP addresses dynamically from a specified pool. Optionally, authentication is performed by a RADIUS authentication server, located within the CEN's address space.

Dynamic subscriber IP address allocation via a DHCP server (IPAssigner set to RemoteDhcpServer)

Subscriber addresses are assigned dynamically by a DHCP server, located within the CEN's address space. Optionally, authentication is performed by a RADIUS authentication server, located within the CEN's address space.



NOTE: The GGSN does not supports subscriber IP address allocation by using a RADIUS server.

The following limitations were introduced to prevent major GGSN configuration modifications from being introduced when active contexts exist for an APN:

- If you want to change the setting of the `IPAssigner` parameter, you must:
 1. Disable APN control (by setting the `AccessPointCtl` parameter to `Disable`).

NOTE: Disabling APN control shuts down any active sessions. For this reason, we recommend that you make major configuration changes, such as changing the `IPAssigner`, during a scheduled downtime.
 2. Change the `IPAssigner` as desired.
 3. Enable APN control (by setting the `AccessPointCtl` parameter to `Enable`).
- If the `AAACONTrol` parameter is set to `Enabled`, the `IPAssigner` parameter can be set to any value. If the `AAACONTrol` is set to `Disabled`, the `IPAssigner` can be set to any value other than `RADIUS`. Therefore, if you change the setting of the `AAACONTrol` parameter from `Enabled` to `Disabled`, you must also make sure that the `IPAssigner` parameter is not set to `RADIUS` in order for the new `AAACONTrol` setting to function.

To configure a GGSN for DSR, HA GGSN, or conventional data redundancy (HA and DSR) support, you bind an APN to multiple IPIP tunnels when you create the APN. To do this, you specify a bidirectional forwarding detection (BFD) gateway IP address (the IP address of the border router) and a priority value for each tunnel when you create the APN using the GGSN service `AccessPointName` parameter.



NOTE: The APN to multiple IPIP binding is included in the `xgsn.cfg` file of the device.

Multiple IPIP tunnels are supported only for the purpose of redundancy, and only one tunnel is active at a time. For this release, EOS supports up to four IPIP tunnels per APN.

Procedure:

Perform one of the following actions:

- To configure an APN to use the static subscriber IP address allocation, perform [Configuring APNs to Use Dynamic Subscriber IP Address Allocation with Local Server on page 607](#).
- To configure an APN to use the dynamic subscriber IP address allocation by using a local server, perform [Configuring APNs to Use Dynamic Subscriber IP Address Allocation with Local Server on page 607](#).
- To configure an APN to use the dynamic subscriber IP address allocation by using a DHCP server, perform [Configuring APNs to Use Dynamic Subscriber IP Address Allocation with a DHCP Server on page 610](#).

35.11.1

Configuring APNs to Use Static Subscriber IP Address Allocation

You can use this procedure to configure an access point name (APN) on a GGSN to use a static IP address allocation.



NOTE: You should add the commands used in this procedure to the `XGSN.cfg` file of the GGSN, so that the configuration is maintained when the GGSN reboots.

Procedure:

1. Create virtual port(s) over which the IPIP tunnel(s) connecting GGSN and the CEN run by entering:

```
ADD !v<port> -Port VirtualPort IPIP <local_ipaddr> <remote_ipaddr> [VRID  
<vrid@port-no>, <vrid@port-no>]
```

where:

`<local_ipaddr>` is the physical address on the GGSN to be used to transmit to the CEN
`<remote_ipaddr>` is the endpoint address of the tunnel peer. The `<remote_ipaddr>` must be a valid address and should be unique. The IPIP tunnel peer should have the same configuration, with `<local_ipaddr>` and `<remote_ipaddr>` swapped. The optional VRID keyword indicates that the IPIP tunnel is a VRRP-controlled tunnel.
`<vrid@port-no>` specifies the virtual router identifier (VRID) and the GGSN port number with which the VRID is associated. The optional VRID keyword indicates that the IPIP tunnel is a VRRP-controlled tunnel.

If you configure an APN for a DSR, HA GGSN, or HA and DSR configuration, you must create a virtual port for each of the IPIP tunnels connecting the GGSN to the CEN. In addition, for an HA GGSN or HA and DSR configuration, you must use the

VRID `<vrid@port-no>`, `<vrid@port-no>`
syntax to configure the IPIP tunnels as VRRP-controlled tunnels.

You can use the `Port service Name` parameter to assign names to the virtual ports. For more information, see “Port Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

2. Assign an IP address to each of the virtual ports you created in [step 1](#) by entering:

```
SETDefault !v<port> -IP NETaddr <IP_address> <network_mask>
```

where:

`<IP_address>` is the IP address you want to assign to the virtual port to be associated with this CEN,
`<network_mask>` is the subnet mask.

The VPN tunnel peer must have a similar configuration and it must be on the same subnet.



NOTE: If you configure an APN for a DSR, HA GGSN, or HA and DSR configuration, with up to four IPIP tunnels bound to each APN, you must assign an IP address to each virtual port.

3. Specify the APN by which you want to identify the CEN by entering:

```
SETDefault !v<portid> -GGSN AccessPointName = "<APN>" [<BFD_target><Priority>]
```

where:

`<portid>` is the virtual port associated with the APN
`<APN>` is the access point name

The `<APN>` is comprised of just the APN Network Identifier and can be up to 63 characters long. The APN consists of one or more labels, each of which can contain only the alphabetic characters (A-Z and a-z), digits (0-9), the period (.), and the dash (-). The dash is not allowed as the first character

in the APN. In addition, the *<APN>* cannot start with the strings *rac*, *lac*, or *sgsn*, and it cannot end in *.gprs*. The case of the alphabetic characters is not significant. If you are specifying an APN with multiple labels, separate the labels within the APN with periods ("*Label1.Label2.Label3*"). For example: "*CountyFire.net*". You must enclose the name in quotation marks, and the name must be unique among APNs.

4. Perform one of the following actions:

- If you want to bind a single IPIP tunnel to the APN, enter:
`SETDefault !v<portid> -GGSN AccessPointName = "<APN>"`
- If you want to bind up to four IPIP tunnels to the APN, to support a DSR, HA GGSN, or HA and DSR configuration, enter the following command for each IPIP tunnel:
`SETDefault !v<portid> -GGSN AccessPointName = "<APN>" [<BFD_target><Priority>]`
 where:
 <BFD_target> is the border router (BFD gateway) IP address,
 <Priority> is the session priority for the IPIP tunnel.

Step example:

To bind a single IPIP tunnel configured on virtual port 1 to the CEN identified by APN "*CountyFire.net*", enter:

```
SETDefault !v1 -GGSN AccessPointName = "CountyFire.net"
```

To bind the IPIP tunnels configured on virtual ports 1 and 2 to the CEN identified by APN "*CountyFire.net*", enter the following commands:

```
SETDefault !v1 -GGSN AccessPointName = "CountyFire.net" 10.74.2.1 5
```

```
SETDefault !v2 -GGSN AccessPointName = "CountyFire.net" 10.245.2.1 9
```

The IPIP tunnel running over virtual port 1 connects the GGSN to the border router with IP address 10.74.2.1, and the IPIP tunnel running over virtual port 2 connects the GGSN to the border router with IP address 10.245.2.1. The GGSN uses the IPIP tunnel with the highest priority value to connect to the CEN. If that tunnel fails, the GGSN switches over to the other IPIP tunnel until the higher priority link is re-established.

5. Enable the APN by entering:

```
SETDefault !<APN> -GGSN AccessPointCtl = Enable
```

6. If you want to be able to retrieve the IP address of a particular mobile subscriber from the DNS server in the CEN, perform the following actions:

- a. Specify the IP address of the DNS server by entering:

```
SETDefault !<APN> -GGSN DDNServer = <IP_address>
```

where *<IP_address>* is the IP address of the DNS server on the CEN.

- b. **Optional:** Specify the domain name to be used for the dynamic DNS updates in the CEN by entering:

```
SETDefault !<APN> -GGSN DDName = "<String>"
```

where:

<APN> is the access point name,

<String> is the domain name that you want to use for the dynamic DNS updates.

You must enclose the domain name in quotation marks. If you do not specify a *DDName*, the GGSN uses the configured *AccessPointName* for the CEN (see [step 3](#)) as the domain name.

7. Specify the local server to be the IP assigner by entering:

```
SETDefault !<APN> -GGSN IPAssigner = LocalDhcpServer
```

8. Ensure that no address pools are configured (the default setting) by entering:

```
SHoW !<APN> -GGSN AddressPool
```



NOTE: Do not include the SHoW !<APN> -GGSN AddressPool command in the XGSN.cfg file

9. If any address pools are configured, delete them by entering:

```
DELeTe !<APN> -GGSN AddressPool All
```

10. If you want to use RADIUS authentication, perform the following actions:

- a. Specify the IP address of the primary RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAAPrimAUTH = <IP_address>
```

where <IP_address> is the IP address of the primary RADIUS authentication server located in the CEN.

- b. Optional: Specify the IP address of the secondary RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAASecAUTH = <IP_address>
```

where <IP_address> is the IP address of the secondary RADIUS authentication server located in the CEN.

- c. Configure the secret text string to be exchanged between the GGSN (RADIUS client) and the RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAASecret = <"string">
```

where <"string"> is a unique ASCII string (1 to 32 characters) used by the RADIUS server to authenticate the user.



NOTE: The string of six asterisks (*****) is not allowed as a RADIUS secret.

- d. Optional: Specify a non-default UDP port number for the RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAAAUTHPort = <UDP_Port_#>
```

where <UDP_Port_#> is the UDP port number for the RADIUS authentication server.

If you do not specify a port number, the default UDP port number for the RADIUS authentication server (1812) is used.

- e. Specify the maximum number of times the GGSN RADIUS client retransmits RADIUS messages by entering:

```
SETDefault !<APN> -GGSN AAAReXmitCnt = <value>
```

The default retransmission count is 3.

- f. Specify the time interval, in seconds between RADIUS message retransmissions by entering:

```
SETDefault !<APN> -GGSN AAAReXmitTime = <seconds>
```

The default retransmission interval is 2 seconds.

- g. Specify the maximum number of context creation failures allowed on the primary (or secondary) authentication server before the GGSN RADIUS client switches over permanently to the secondary (or primary) authentication server by entering:

```
SETDefault !<APN> -GGSN AAASwitchCnt = <value>
```

The default switchover count is 1.

- h. Enable RADIUS authentication for the APN by entering:

```
SETDefault !<APN> -GGSN AAAControl = Enable
```

35.11.2

Configuring APNs to Use Dynamic Subscriber IP Address Allocation with Local Server

You can use this procedure to configure an APN on a GGSN to use a dynamic IP address allocation through a local server (the GGSN acts as an "internal" DHCP server).



NOTE: It is recommended to add the commands used in this procedure to the `XGSN.cfg` file of the GGSN, so that the configuration is maintained when the GGSN reboots.

Procedure:

1. Create virtual port(s) over which the IPIP tunnel(s) connecting GGSN and the CEN run by entering:

```
ADD !v<port> -Port VirtualPort IPIP <local_ipaddr> <remote_ipaddr> [VRID
<vrid@port-no>, <vrid@port-no>]
```

where:

<local_ipaddr> is the physical address on the GGSN to be used to transmit to the CEN

<remote_ipaddr> is the endpoint address of the tunnel peer. The <remote_ipaddr> must be a valid address and should be unique. The IPIP tunnel peer should have the same configuration, with <local_ipaddr> and <remote_ipaddr> swapped. The optional VRID keyword indicates that the IPIP tunnel is a VRRP-controlled tunnel.

<vrid@port-no> specifies the virtual router identifier (VRID) and the GGSN port number with which the VRID is associated.



NOTE: If you configure an APN for a DSR, HA GGSN, or HA and DSR configuration, you must create a virtual port for each of the IPIP tunnels connecting the GGSN to the CEN. In addition, for an HA GGSN or HA and DSR configuration, you must use the VRID <vrid@port-no>, <vrid@port-no> syntax to configure the IPIP tunnels as VRRP-controlled tunnels.

You can use the `PORT service NAME` parameter to assign names to the virtual ports. For more information, see "PORT Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

2. Assign an IP address to each of the virtual ports you created in [step 1](#) by entering:

```
SETDefault !v<port> -IP NETaddr <IP_address> <network_mask>
```

where:

<IP_address> is the IP address you want to assign to the virtual port to be associated with this CEN,

<network_mask> is the subnet mask.

The VPN tunnel peer must have a similar configuration and it must be on the same subnet.



NOTE: If you configure an APN for a DSR, HA GGSN, or HA and DSR configuration, with up to four IPIP tunnels bound to each APN, you must assign an IP address to each virtual port.

3. Specify the access point name (APN) by which you want to identify the CEN by entering:

```
SETDefault !v<portid> -GGSN AccessPointName = "<APN>" [<BFD_target><Priority>]
```

where:

<portid> is the virtual port associated with the APN

<APN> is the access point name

The <APN> is comprised of just the APN Network Identifier and can be up to 63 characters long. The APN consists of one or more labels, each of which can contain only the alphabetic characters (A-Z and a-z), digits (0-9), the period (.), and the dash (-). The dash is not allowed as the first character in the APN. In addition, the <APN> cannot start with the strings `rac`, `lac`, or `sgsn`, and it cannot end in `.gprs`. The case of the alphabetic characters is not significant. If you are specifying an APN

with multiple labels, separate the labels within the APN with periods ("Label1.Label2.Label3"). For example: "CountyFire.net". You must enclose the name in quotation marks, and the name must be unique among APNs.

4. Perform one of the following actions:

- If you want to bind a single IPIP tunnel to the APN, enter:
`SETDefault !v<portid> -GGSN AccessPointName = "<APN>"`
- If you want to bind up to four IPIP tunnels to the APN, to support a DSR, HA GGSN, or HA and DSR configuration, enter the following command for each IPIP tunnel:
`SETDefault !v<portid> -GGSN AccessPointName = "<APN>" [<BFD_target><Priority>]`
where:
 - <BFD_target> is the border router (BFD gateway) IP address,
 - <Priority> is the session priority for the IPIP tunnel.

Step example:

To bind a single IPIP tunnel configured on virtual port 1 to the CEN identified by APN "CountyFire.net", you must enter:

```
SETDefault !v1 -GGSN AccessPointName = "CountyFire.net"
```

To bind the IPIP tunnels configured on virtual ports 1 and 2 to the CEN identified by APN "CountyFire.net", you must enter the following commands:

```
SETDefault !v1 -GGSN AccessPointName = "CountyFire.net" 10.74.2.1 5
```

```
SETDefault !v2 -GGSN AccessPointName = "CountyFire.net" 10.245.2.1 9
```

The IPIP tunnel running over virtual port 1 connects the GGSN to the border router with IP address 10.74.2.1, and the IPIP tunnel running over virtual port 2 connects the GGSN to the border router with IP address 10.245.2.1. The GGSN uses the IPIP tunnel with the highest priority value to connect to the CEN. If that tunnel fails, the GGSN switches over to the other IPIP tunnel until the higher priority link is re-established.

5. Enable the APN by entering:

```
SETDefault !<APN> -GGSN AccessPointCtl = Enable
```

6. If you want to be able to retrieve the IP address of a particular mobile subscriber from the DNS server in the CEN, perform the following actions:

a. Specify the IP address of the DNS server by entering:

```
SETDefault !<APN> -GGSN DDNServer = <IP_address>
```

where <IP_address> is the IP address of the DNS server on the CEN.

b. Optional: Specify the domain name to be used for the dynamic DNS updates in the CEN by entering:

```
SETDefault !<APN> -GGSN DDName = "<String>"
```

where:

<APN> is the access point name,

<String> is the domain name that you want to use for the dynamic DNS updates.

You must enclose the domain name in quotation marks. If you do not specify a DDName, the GGSN uses the configured AccessPointName for the CEN (see [step 3](#)) as the domain name.

7. Configure IP address allocation through the GGSN acting as an "internal" DHCP server by performing the following actions:

a. Specif the local server to be the IP assigner by entering:

```
SETDefault !<APN> -GGSN IPAssigner = LocalDhcpServer
```

- b. Ensure that no address pools are configured by entering:

```
SHow !<APN> -GGSN AddressPool
```



NOTE: Do not include the SHow !<APN> -GGSN AddressPool command in the XGSN.cfg file

- c. If unwanted address pools are configured, delete them by entering:

```
DELeTe !<APN> -GGSN AddressPool All
```

- d. Configure one or more address pools from which the GGSN, that acts as an internal DHCP server, can dynamically assign IP addresses to requesting subscribers for this APN by entering:

```
ADD !<APN> -GGSN AddressPool <IPaddr1> [-<IPaddr2>]
```

where:

<IPaddr1> is the start of the range of IP addresses that are to be used for local IP address assignment,

<IPaddr2> is the end of the range of the range of IP addresses that are to be used for local IP address assignment.

If you do not specify an <IPaddr2>, the address pool contains a single IP address (<IPaddr1>).

8. If you want certain subscribers to have reserved IP addresses, create a list of IP address reservations that matches subscriber IP addresses to hard-coded subscriber IMSI numbers by entering:

```
ADD !<APN> -GGSN StaticAddr <IMSI> <IPaddress>
```

where:

<IMSI> is the International Mobile Subscriber Identification (IMSI) number,

<IPaddress> is the IP address you want to reserve for that subscriber.

The IMSI number must be a 15-digit number, and each IMSI number can be associated with only one unique IP address. No two IMSIs can share the same IP address.

If you use a mixture of dynamic and static IP address allocation, you must configure address reservations for the subscribers with static IP address assignments so the statically-assigned addresses are not assigned dynamically to other subscribers. Static IP addresses should not overlap addresses in the dynamic address pool.



IMPORTANT: Because the DHCP server expects a 16-digit number in the client hardware address field, the GGSN adds an extra zero to the IMSI value passed by the subscriber. Thus if you use the DHCP server manager to configure a subscriber IP address reservation on the DHCP server (located in the CEN), you must add a zero to the end of the IMSI number in order to match the value passed by the GGSN.

9. If you want to use RADIUS authentication, perform the following actions:

- a. Specify the IP address of the primary RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAAPrimAUTH = <IP_address>
```

where <IP_address> is the IP address of the primary RADIUS authentication server located in the CEN.

- b. Optional: Specify the IP address of the secondary RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAASecAUTH = <IP_address>
```

where <IP_address> is the IP address of the secondary RADIUS authentication server located in the CEN.

- c. Configure the secret text string to be exchanged between the GGSN (RADIUS client) and the RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAASecret = <"string">
```

where *<"string">* is a unique ASCII string (1 to 32 characters) used by the RADIUS server to authenticate the user.



NOTE: The string of six asterisks (*****) is not allowed as a RADIUS secret.

- d. Optional: Specify a non-default UDP port number for the RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAAAUTHPort = <UDP_Port_#>
```

where *<UDP_Port_#>* is the UDP port number for the RADIUS authentication server.

If you do not specify a port number, the default UDP port number for the RADIUS authentication server (1812) is used.

- e. Specify the maximum number of times the GGSN RADIUS client retransmits RADIUS messages by entering:

```
SETDefault !<APN> -GGSN AAAReXmitCnt = <value>
```

The default retransmission count is 3.

- f. Specify the time interval, in seconds between RADIUS message retransmissions by entering:

```
SETDefault !<APN> -GGSN AAAReXmitTime = <seconds>
```

The default retransmission interval is 2 seconds.

- g. Specify the maximum number of context creation failures allowed on the primary (or secondary) authentication server before the GGSN RADIUS client switches over permanently to the secondary (or primary) authentication server by entering:

```
SETDefault !<APN> -GGSN AAASwitchCnt = <value>
```

The default switchover count is 1.

- h. Enable RADIUS authentication for the APN by entering:

```
SETDefault !<APN> -GGSN AAACONtrol = Enable
```

35.11.3

Configuring APNs to Use Dynamic Subscriber IP Address Allocation with a DHCP Server

You can use this procedure to configure an APN to use a dynamic IP address allocation through a DHCP server.



NOTE: It is recommended to add the commands used in this procedure to the `XGSN.cfg` file of the GGSN, so that the configuration is maintained when the GGSN reboots.

Procedure:

1. Create virtual port(s) over which the IPIP tunnel(s) connecting GGSN and the CEN run by entering:

```
ADD !v<port> -Port VirtualPort IPIP <local_ipaddr> <remote_ipaddr> [VRID  
<vrid@port-no>, <vrid@port-no>]
```

where:

<local_ipaddr> is the physical address on the GGSN to be used to transmit to the CEN

<remote_ipaddr> is the endpoint address of the tunnel peer. The *<remote_ipaddr>* must be a valid address and should be unique. The IPIP tunnel peer should have the same configuration, with *<local_ipaddr>* and *<remote_ipaddr>* swapped. The optional VRID keyword indicates that the IPIP tunnel is a VRRP-controlled tunnel.

`<vrid@port-no>` specifies the virtual router identifier (VRID) and the GGSN port number with which the VRID is associated.



NOTE: If you configure an APN for a DSR, HA GGSN, or HA and DSR configuration, you must create a virtual port for each of the IPIP tunnels connecting the GGSN to the CEN. In addition, for an HA GGSN or HA and DSR configuration, you must use the VRID `<vrid@port-no>`, `<vrid@port-no>` syntax to configure the IPIP tunnels as VRRP-controlled tunnels.

You can use the `PORT service NAME` parameter to assign names to the virtual ports. For more information, see “Port Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

2. Assign an IP address to each of the virtual ports you created in [step 1](#) by entering:

```
SETDefault !v<port> -IP NETaddr <IP_address> <network_mask>
```

where:

`<IP_address>` is the IP address you want to assign to the virtual port to be associated with this CEN,

`<network_mask>` is the subnet mask.

The VPN tunnel peer must have a similar configuration and it must be on the same subnet.



NOTE: If you configure an APN for a DSR, HA GGSN, or HA and DSR configuration, with up to four IPIP tunnels bound to each APN, you must assign an IP address to each virtual port.

3. Specify the access point name (APN) by which you want to identify the CEN by entering:

```
SETDefault !v<portid> -GGSN AccessPointName = "<APN>" [<BFD_target><Priority>]
```

where:

`<portid>` is the virtual port associated with the APN

`<APN>` is the access point name

The `<APN>` is comprised of just the APN Network Identifier and can be up to 63 characters long. The APN consists of one or more labels, each of which can contain only the alphabetic characters (A-Z and a-z), digits (0-9), the period (.), and the dash (-). The dash is not allowed as the first character in the APN. In addition, the `<APN>` cannot start with the strings `rac`, `lac`, or `sgsn`, and it cannot end in `.gprs`. The case of the alphabetic characters is not significant. If you are specifying an APN with multiple labels, separate the labels within the APN with periods (`"Label1.Label2.Label3"`). For example: `"CountyFire.net"`. You must enclose the name in quotation marks, and the name must be unique among APNs.

4. Perform one of the following actions:

- If you want to bind a single IPIP tunnel to the APN, enter:

```
SETDefault !v<portid> -GGSN AccessPointName = "<APN>"
```

- If you want to bind up to four IPIP tunnels to the APN, to support a DSR, HA GGSN, or HA and DSR configuration, enter the following for each IPIP tunnel:

```
SETDefault !v<portid> -GGSN AccessPointName = "<APN>" [<BFD_target><Priority>]
```

where:

`<BFD_target>` is the border router (BFD gateway) IP address,

`<Priority>` is the session priority for the IPIP tunnel.

Step example:

To bind a single IPIP tunnel configured on virtual port 1 to the CEN identified by APN

`"CountyFire.net"`, enter:

```
SETDefault !v1 -GGSN AccessPointName = "CountyFire.net"
```

To bind the IPIP tunnels configured on virtual ports 1 and 2 to the CEN identified by APN

`"CountyFire.net"`, enter the following commands:


```
SETDefault !v1 -GGSN AccessPointName = "CountyFire.net" 10.74.2.1 5  
SETDefault !v2 -GGSN AccessPointName = "CountyFire.net" 10.245.2.1 9
```

The IPIP tunnel running over virtual port 1 connects the GGSN to the border router with IP address 10.74.2.1, and the IPIP tunnel running over virtual port 2 connects the GGSN to the border router with IP address 10.245.2.1. The GGSN uses the IPIP tunnel with the highest priority value to connect to the CEN. If that tunnel fails, the GGSN switches over to the other IPIP tunnel until the higher priority link is re-established.

5. Enable the APN by entering:

```
SETDefault !<APN> -GGSN AccessPointCtl = Enable
```

6. If you want to be able to retrieve the IP address of a particular mobile subscriber from the DNS server in the CEN, perform the following actions:

a. Specify the IP address of the DNS server by entering:

```
SETDefault !<APN> -GGSN DDNServer = <IP_address>
```

where <IP_address> is the IP address of the DNS server on the CEN.

b. Optional: Specify the domain name to be used for the dynamic DNS updates in the CEN by entering:

```
SETDefault !<APN> -GGSN DDName = "<String>"
```

where:

<APN> is the access point name,

<String> is the domain name that you want to use for the dynamic DNS updates.

You must enclose the domain name in quotation marks. If you do not specify a DDName, the GGSN uses the configured AccessPointName for the CEN (see [step 3](#)) as the domain name.

7. Configure IP address allocation with a DHCP server by performing the following actions:

a. Specify the DHCP server to be the IP assigner by entering:

```
SETDefault !<APN> -GGSN IPAssigner = RemoteDhcpServer
```

b. Specify the IP address of the primary DHCP server by entering:

```
SETDefault !<APN> -GGSN DHCPPrimSrvr = \<IP_address> [<Gateway>]
```

where:

<IP_address> is the IP address of the primary DHCP server located in the CEN,

<Gateway> is an IP address on the same subnet as the address pool range from which the subscriber IP addresses is allocated.

c. Optional: Specify the IP address of the Secondary (backup) DHCP server by entering:

```
SETDefault !<APN> -GGSN DHCPSecSrvr = <IP_address> [<Gateway>]
```

where:

<IP_address> is the IP address of the secondary DHCP server located in the CEN

<Gateway> is as IP address on the same subnet as the address pool range from which the subscriber IP addresses is allocated.

d. Specify the maximum number of times the GGSN retransmits mobile subscriber IP address requests by entering:

```
SETDefault !<APN> -GGSN DHCPReXmitCnt = <value>
```

The default retransmission count is 3.

e. Specify the time interval, in seconds between DHCP message retransmissions by entering:

```
SETDefault !<APN> -GGSN DHCPReXmitTime = <seconds>
```


The default retransmission interval is 2 seconds.

- f. Specify the maximum number of mobile subscriber IP address request failures allowed on the primary (or secondary) authentication server before the GGSN RADIUS client switches over permanently to the secondary (or primary) authentication server by entering:

```
SETDefault !<APN> -GGSN DHCPSwitchCnt = <value>
```

If use a mixture of dynamic and static IP address allocation, you should configure address reservations for the subscribers with static IP address assignments so the statically-assigned addresses are not assigned dynamically to other subscribers. Create these reservations on the DHCP server. Static IP addresses should not overlap addresses in the dynamic address pool.



CAUTION: Do **not** include the GiAddr IP address in the address pool from which mobile subscriber IP addresses are allocated on the DHCP server. The Gidder is the gateway IP address to the DHCP server.

8. If you want to use RADIUS authentication, perform the following actions:

- a. Specify the IP address of the primary RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAAPrimAUTH = <IP_address>
```

where <IP_address> is the IP address of the primary RADIUS authentication server located in the CEN.

- b. Optional: Specify the IP address of the secondary RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAASecAUTH = <IP_address>
```

where <IP_address> is the IP address of the secondary RADIUS authentication server located in the CEN.

- c. Configure the secret text string to be exchanged between the GGSN (RADIUS client) and the RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAASecret = <"string">
```

where <"string"> is a unique ASCII string (1 to 32 characters) used by the RADIUS server to authenticate the user.



NOTE: The string of six asterisks (*****) is not allowed as a RADIUS secret.

- d. Optional: Specify a non-default UDP port number for the RADIUS authentication server by entering:

```
SETDefault !<APN> -GGSN AAAAUTHPort = <UDP_Port_#>
```

where <UDP_Port_#> is the UDP port number for the RADIUS authentication server.

If you do not specify a port number, the default UDP port number for the RADIUS authentication server (1812) is used.

- e. Specify the maximum number of times the GGSN RADIUS client retransmits RADIUS messages by entering:

```
SETDefault !<APN> -GGSN AAARemitCnt = <value>
```

The default retransmission count is 3.

- f. Specify the time interval, in seconds between RADIUS message retransmissions by entering:

```
SETDefault !<APN> -GGSN AAARemitTime = <seconds>
```

The default retransmission interval is 2 seconds.

- g. Specify the maximum number of context creation failures allowed on the primary (or secondary) authentication server before the GGSN RADIUS client switches over permanently to the secondary (or primary) authentication server by entering:

```
SETDefault !<APN> -GGSN AAASwitchCnt = <value>
```

The default switchover count is 1.

- h. Enable RADIUS authentication for the APN by entering:

```
SETDefault !<APN> -GGSN AAAControl = Enable
```

35.12

APN Configuration to Multiple IPIP Binding

To configure a GGSN for DSR, HA GGSN, or conventional data redundancy (HA and DSR) support, you bind an APN to multiple IPIP tunnels when you create the APN. To do this, you specify a bidirectional forwarding detection (BFD) gateway IP address (the IP address of the border router) and a priority value for each tunnel when you create the APN using the GGSN service `AccessPointName` parameter.

The APN to multiple IPIP binding is included in the `xgsn.cfg` file of the device.

Multiple IPIP tunnels are supported only for the purpose of redundancy, and only one tunnel is active at a time. For this release, EOS supports up to four IPIP tunnels per APN.

35.13

Configuring Devices as HA GGSN Peers

The HA GGSN peer configuration is included in the `boot.cfg` file for each HA GGSN peer.

Procedure:

1. Define macros to enable and disable VRRP control on the GGSN Ethernet ports associated with the specified virtual router identifiers (VRIDs) by entering:

```
define <enable_macro_name>=(  
SETDefault !<port> -VRRP CONTROL = Enable <VRID_1>  
SETDefault !<port> -VRRP CONTROL = Enable <VRID_2>  
)  
define <disable_macro_name>=(  
SETDefault !<port> -VRRP CONTROL = Disable <VRID_1>  
SETDefault !<port> -VRRP CONTROL = Disable <VRID_2>  
)
```

where:

<enable_macro_name> is the name you want to assign to the macro that enables VRRP control when the port comes up (for example `port_up`)

<disable_macro_name> is the name you want to assign to the macro that disables VRRP control when the port goes down (for example `port_down`)

<VRID_1> is the VRID associated with one subnet to the CEN

<VRID_2> is the VRID associated with the second subnet to the CEN

You configure these VRIDs when you configure VRRP on the HA GGSN peers. For details, see [step 4](#).

The VRIDs you specify must match the VRIDs you assign to the ports when you configure VRRP backup functionality (see [step 4a](#)).

2. As part of the IP address setup, configure a system IP address (SysIP) on each of the HA GGSN peers by entering:

```
SETDefault -IP SystemIP = <IP_address>
```

where *<IP_address>* is the SysIP.

3. Enable the loopback system IP feature for the NTP, FTP, PING, SNMP, Telnet, TFTP, and Syslog services so that the SysIP is used by sessions initiated by the GGSN by entering the following command on both HA GGSN peers:

```
SETDefault -SYS UseSystemIP = (NTP, FTP, PING, SNMP, Telnet, TFTP, SYSLOG)
```

4. Configure each HA GGSN peer as a backup VRRP router on each of the CEN subnets by performing the following actions:

- a. Configure VRRP backup router functionality on the GGSN Ethernet ports by entering:

```
ADD -VRRP !<port> BackUp AsOwner <vrid> <IP_address>
```

where:

<port> is the port on which you are configuring VRRP

AsOwner is a keyword indicating that when the port takes over the IP address it behaves as if it owns the IP address

<vrid> is the VRID to be associated with the port

<IP_address> is an IP address on the CEN subnet

- b. Set the advertisement interval for VRRP packets by entering:

```
SETDefault !<port> -VRRP AdvertiseInt = <adv_time> <vrid>
```

where:

<port> is the port on which you are configuring VRRP

<adv_time> is the advertisement interval, in milliseconds (1 - 100,000)

<vrid> is the VRID associated with the port



NOTE: You must set the advertisement interval to the same value on both HA GGSN peers.

- c. Set the VRRP hold time (the time interval before the backup VRRP declares the master VRRP down) by entering:

```
SETDefault !<port> -VRRP HoldTime = <hold_time> <vrid>
```

where:

<port> is the port on which you are configuring VRRP

<hold_time> is the hold time, in milliseconds (300 - 3000000)

<vrid> is the VRID associated with the port



NOTE: The hold time should be at least three times the advertisement interval.

- d. Set the VRRP priority by entering:

```
SETDefault !<port> -VRRP PRIOrity = <priority_number> <vrid>
```

where:

<port> is the port on which you are configuring VRRP

<priority_number> is the priority value (1-254)

<vrid> is the VRID associated with the port



NOTE: You must set the VRRP priority to the same value on each of the GGSN Ethernet ports but to different values on the HA GGSN peers.

- e. Enable VRRP on the GGSN Ethernet ports by entering:

```
SETDefault !<port> -VRRP CONTRol = Enable <vrid>
```

where:

<port> is the port on which you are configuring VRRP

<vrid> is the VRID associated with the port

5. For each GGSN Ethernet port on which VRRP is configured, set up the scheduler to invoke the `port_down` macro when the port goes down and to invoke the `port_up` macro when the port comes up by entering the following commands:

```
SETDefault -SCH EbmeCONTRol = Enable
```

```
ADD !<port> -SCH EbmeEvent PortDown <DebounceTimer> do port_down
```

```
ADD !<port> -SCH EbmeEvent PortUp <DebounceTimer> do port_up
```

where:

<port> is the GGSN Ethernet port number

<DebounceTimer> is the debounce timer (the amount of time (50-3600000 milliseconds) to wait to wait after receiving the port status change notification and before executing the macro

6. Specify the HA peer of the GGSN by entering:

```
ADD -GGSN HASTandBy <IP_address> [<Peer_SysName>] [HelloTime <seconds> <count>]
```

where *<IP_address>* is the system IP address of the HA GGSN peer.

EOS software supports only one context database synchronization peer per GGSN and only IPv4 address format.

<Peer_SysName> is an optional name used to identify the HA GGSN peer in SNMP traps. If you do not specify a *<SysName>*, the IP address of the GGSN context database synchronization peer is used in SNMP traps.

`HelloTime <seconds> <count>` is the heartbeat timer for context database synchronization. The specified *<seconds>* and *<count>* are the length of time (in seconds) and number of hellos before the GGSN declares the context database synchronization peer down and restarts the synchronization relationship. If you do not specify a `HelloTime`, the default values (3 seconds and 5 hellos) are used, meaning that the GGSN times out the HA peer in 15 seconds (3 seconds per hello x 5 hellos).




WARNING: In order to avoid unexpected synchronization interruption, you must configure the same `HelloTime` values on both HA GGSN peers.

35.13.1

Configuring GGSNs in the HA GGSN Sample Topology

The following figure illustrates a sample HA GGSN topology with two GGSNs (GGSN 1 and GGSN 2) connecting to two redundant BRs in two subnets (Subnet 1 and Subnet 2).

Figure 178: Sample HA GGSN Topology



On GGSN 1:
Ethernet port 1, with IP address 10.1.255.21, is configured for backup VRRP functionality on Subnet 1 and associated with VRID 83.
Ethernet port 2, with IP address 10.1.223.128, is configured for backup VRRP functionality on Subnet 2 and associated with VRID 73.
GGSN 2, with SysIP 10.1.225.241, is configured as the HA GGSN peer.

On GGSN 2:
Ethernet port 1, with IP address 10.1.255.22, is configured for backup VRRP functionality on Subnet 1 and associated with VRID 83.
Ethernet port 2, with IP address 10.1.223.129, is configured for backup VRRP functionality on Subnet 2 and associated with VRID 73.
GGSN 1, with SysIP 10.1.225.242, is configured as the HA GGSN peer.

On both HA GGSN Peers

The IPIP tunnel running over virtual port 101 connects the GGSN to BR 1 in Subnet 1 (with IP address 180.20.55.6).
The IPIP tunnel running over virtual port 1101 connects the GGSN to BR 2 in Subnet 1 (with IP address 180.20.55.124).
The IPIP tunnel running over virtual port 301 connects the GGSN to BR 1 in Subnet 2 (with IP address 180.20.106.6).
The IPIP tunnel running over virtual port 1301 connects the GGSN to BR 2 in Subnet 2 (with IP address 180.20.106.124).

On each subnet, the GGSN uses the IPIP tunnel with the highest priority value to connect to the CEN. If the IPIP tunnel with the highest priority fails, the GGSN switches to the lower-priority tunnel until the higher-priority link is re-established.

Process:

Configuring GGSN 1 in the HA GGSN Sample Topology

1. Include the following configurations in the basic setup for the GGSN by performing the following actions:



NOTE: These commands are included in the `boot.cfg` file of the GGSN 1.

- a. Define `port_up` and `port_down` macros that enable and disable VRRP control on the GGSN Ethernet ports associated with the specified VRIDs by entering:

```
define port_up=(  
  SETDefault !1 -VRRP CONTrol = Enable 83  
  SETDefault !2 -VRRP CONTrol = Enable 73  
)  
define port_down=(  
  SETDefault !1 -VRRP CONTrol = Disable 83  
  SETDefault !2 -VRRP CONTrol = Disable 73  
)
```



NOTE: The VRIDs you specify must match the VRIDs you assign to the ports when you configure VRRP backup functionality (see [step 1d](#) and [step 1e](#)).

- b. As part of the IP address setup, configure a system IP address (SysIP) by entering:

```
SETDefault -IP SystemIP 10.1.225.242
```

- c. Enable the loopback system IP feature for the NTP, FTP, PING, SNMP, Telnet, TFTP, and Syslog services so that the SysIP is used by sessions initiated by the GGSN by entering:

```
SETDefault -SYS UseSystemIP = (NTP, FTP, PING, SNMP, Telnet, TFTP, SYSLOG)
```

- d. To configure VRRP backup router functionality, associate GGSN Ethernet port 1 with a VRID and an IP address on CEN subnet 1 by entering the following commands:

```
ADD !1 -VRRP BackUp AsOwner 83 10.1.225.20  
SETDefault !1 -VRRP AdvertiseInt = 300 83  
SETDefault !1 -VRRP HoldTime = 900 83  
SETDefault !1 -VRRP PRIOrity = 210 83  
SETDefault !1 -VRRP CONTrol = Enable 83
```

- e. To configure VRRP backup router functionality, associate GGSN Ethernet port 2 with a VRID and an IP address on CEN subnet 2 by entering the following commands:

```
ADD !2 -VRRP BackUp AsOwner 73 10.1.223.136  
SETDefault !2 -VRRP AdvertiseInt = 300 73  
SETDefault !2 -VRRP HoldTime = 900 73  
SETDefault !2 -VRRP PRIOrity = 210 73  
SETDefault !2 -VRRP CONTrol = Enable 73
```

- f. Set up the scheduler to invoke the `port_down` and `port_up` macros you defined in [step 1a](#) when GGSN ports 1 and 2 go down or come up by entering the following commands:

```
SETDefault -SCH EbmeCONTrol = Enable  
ADD !1 -SCH EbmeEvent PortDown 450 do port_down  
ADD !1 -SCH EbmeEvent PortUp 2000 do port_up  
ADD !2 -SCH EbmeEvent PortDown 450 do port_down
```

```
ADD !2 -SCH EbmeEvent PortUp 2000 do port_up
```

- g. Specify the HA peer of the GGSN by entering:

```
ADD -GGSN HAStandBy 10.1.225.241 HelloTime 1 3
```

2. Configure GGSN functionality on the device by performing the following actions:



NOTE: These commands comprise the `xsgn.cfg` file of the GGSN 1.

- a. Specify the maximum number of contexts and the type of service (TOS) values for signal and data packets by entering the following commands:

```
SETDefault -GGSN MaxPDPContexts = 20000
```

```
SETDefault -GGSN SignalTos = 160
```

```
SETDefault -GGSN DataTos = 64
```

- b. Configure the four virtual ports over which the IPIP tunnels connecting the GGSN to the CEN run and name the ports (Port_1A, Port_1B, Port_1C, and Port_1D) by entering the following commands:

```
ADD !v101 -Port VirtualPort IPIP 10.1.223.136 10.1.223.158 VRID 83@1, 73@2
```

```
SETDefault !v101 -Port Name = "Port_1A"
```

```
ADD !v1101 -Port VirtualPort IPIP 10.1.223.136 10.1.223.59 VRID 83@1, 73@2
```

```
SETDefault !v1101 -Port Name = "Port_1B"
```

```
ADD !v301 -Port VirtualPort IPIP 10.1.225.20 10.7.223.158 VRID 83@1, 73@2
```

```
SETDefault !v301 -Port Name = "Port_1C"
```

```
ADD !v1301 -Port VirtualPort IPIP 10.1.225.20 10.7.223.59 VRID 83@1, 73@2
```

```
SETDefault !v1301 -Port Name = "Port_1D"
```



NOTE: Each IPIP tunnel corresponds to a BR in the CEN (two BRs per subnet to support the HA feature).

- c. Assign IP addresses to the virtual ports you created in [step 2b](#) by entering the following commands:

```
SETDefault !v101 -IP NETAddr = 180.20.55.5 255.255.255.252
```

```
SETDefault !v1101 -IP NETAddr = 180.20.55.123 255.255.255.252
```

```
SETDefault !v301 -IP NETAddr = 180.20.106.5 255.255.255.252
```

```
SETDefault !v1301 -IP NETAddr = 180.20.106.123 255.255.255.252
```



NOTE: The virtual port IP addresses must be on the same subnet as the BR to which the IPIP tunnel is connected.

- d. Specify the access point name (APN) by which you want to identify the CEN and bind the IPIP tunnels to the APN, associating each IPIP tunnel with a BR in the CEN (two BRs for each subnet) and assigning it a session priority by entering the following commands:

```
SETDefault !v101 -GGSN AccessPointName = "CountyFire.net" 180.20.55.6 4
```

```
SETDefault !v1101 -GGSN AccessPointName = "CountyFire.net" 180.20.55.124 3
```

```
SETDefault !v301 -GGSN AccessPointName = "CountyFire.net" 180.20.106.6 2
```

```
SETDefault !v1301 -GGSN AccessPointName = "CountyFire.net" 180.20.106.124 1
```

- e. Enable the APN by entering:

```
SETDefault !CountyFire.net -GGSN AccessPointCtl = Enable
```

- f. Configure the APN to assign subscriber IP addresses via a local DHCP server and configure the local DHCP server by entering the following commands:

```
SETDefault !CountyFire.net -GGSN IPAssigner = LocalDhcpServer
ADD !CountyFire.net -GGSN AddressPool 10.0.0.1 - 10.0.254.254
SETDefault !CountyFire.net -GGSN DHCPReXmitTime = 2
SETDefault !CountyFire.net -GGSN DHCPReXmitCnt = 3
SETDefault !CountyFire.net -GGSN DHCPSTwitchCnt = 1
```

- g. Optional: Configure GGSN charging. See [Configuring GGSN Charging on page 659](#).
- h. Optional: Configure overload protection. See the [Configuring EOS Devices for Overload Protection on page 752](#).
- i. Enable GGSN functionality on the device by entering:

```
SETDefault -GGSN CONTro1 = Enable
```

Configuring GGSN 2 in the HA GGSN Sample Topology

- 3. Include the following configurations in the basic setup for the GGSN by performing the following actions:



NOTE: These commands are included in the `boot.cfg` file of the GGSN 2.

- a. Define `port_up` and `port_down` macros that enable and disable VRRP control on the GGSN Ethernet ports associated with the specified VRIDs by entering:

```
define port_up=(
SETDefault !1 -VRRP CONTro1 = Enable 83
SETDefault !2 -VRRP CONTro1 = Enable 73
)
define port_down=(
SETDefault !1 -VRRP CONTro1 = Disable 83
SETDefault !2 -VRRP CONTro1 = Disable 73
)
```



NOTE: The VRIDs you specify must match the VRIDs you assign to the ports when you configure VRRP backup functionality (see [step 3d](#) and [step 3e](#)).

- b. As part of the IP address setup, configure a system IP address (SysIP) by entering:

```
SETDefault -IP SystemIP 10.1.225.241
```
- c. Enable the loopback system IP feature for the NTP, FTP, PING, SNMP, Telnet, TFTP, and Syslog services so that the SysIP is used by sessions initiated by the GGSN by entering:

```
SETDefault -SYS UseSystemIP = (NTP, FTP, PING, SNMP, Telnet, TFTP, SYSLOG)
```
- d. To configure VRRP backup router functionality, associate GGSN Ethernet port 1 with a VRID and an IP address on CEN subnet 1 by entering the following commands:

```
ADD !1 -VRRP BackUp AsOwner 83 10.1.225.20
SETDefault !1 -VRRP AdvertiseInt = 300 83
SETDefault !1 -VRRP HoldTime = 900 83
SETDefault !1 -VRRP PRIOrity = 110 83
SETDefault !1 -VRRP CONTro1 = Enable 83
```
- e. To configure VRRP backup router functionality, associate GGSN Ethernet port 2 with a VRID and an IP address on CEN subnet 2 by entering the following commands:

```
ADD !2 -VRRP BackUp AsOwner 73 10.1.223.136
```



```
SETDefault !2 -VRRP AdvertiseInt = 300 73
```

```
SETDefault !2 -VRRP HoldTime = 900 73
```

```
SETDefault !2 -VRRP PRIOrity = 110 73
```

```
SETDefault !2 -VRRP CONTrol = Enable 73
```

- f. Set up the scheduler to invoke the `port_down` and `port_up` macros you defined in [step 3a](#) when GGSN ports 1 and 2 go down or come up by entering the following commands:

```
SETDefault -SCH EbmeCONTrol = Enable
```

```
ADD !1 -SCH EbmeEvent PortDown 450 do port_down
```

```
ADD !1 -SCH EbmeEvent PortUp 2000 do port_up
```

```
ADD !2 -SCH EbmeEvent PortDown 450 do port_down
```

```
ADD !2 -SCH EbmeEvent PortUp 2000 do port_up
```

- g. Specify the HA peer of the GGSN by entering:

```
ADD -GGSN HAsStandBy 10.1.225.242 HelloTime 1 3
```

4. Configure GGSN functionality on the device by performing the following actions:



NOTE: These commands comprise the `xsgn.cfg` file of the GGSN 2.

- a. Specify the maximum number of contexts and the type of service (TOS) values for signal and data packets by entering the following commands:

```
SETDefault -GGSN MaxPDPContexts = 20000
```

```
SETDefault -GGSN SignalTos = 160
```

```
SETDefault -GGSN DataTos = 64
```

- b. Configure the four virtual ports over which the IPIP tunnels connecting the GGSN to the CEN run and name the ports (Port_1A, Port_1B, Port_1C, and Port_1D) by entering the following commands:

```
ADD !v101 -Port VirtualPort IPIP 10.1.223.136 10.1.223.158 VRID 83@1, 73@2
```

```
SETDefault !v101 -Port Name = "Port_1A"
```

```
ADD !v1101 -Port VirtualPort IPIP 10.1.223.136 10.1.223.59 VRID 83@1, 73@2
```

```
SETDefault !v1101 -Port Name = "Port_1B"
```

```
ADD !v301 -Port VirtualPort IPIP 10.1.225.20 10.7.223.158 VRID 83@1, 73@2
```

```
SETDefault !v301 -Port Name = "Port_1C"
```

```
ADD !v1301 -Port VirtualPort IPIP 10.1.225.20 10.7.223.59 VRID 83@1, 73@2
```

```
SETDefault !v1301 -Port Name = "Port_1D"
```



NOTE: Each IPIP tunnel corresponds to a BR in the CEN (two BRs per subnet to support the HA feature).

- c. Assign IP addresses to the virtual ports you created in [step 4b](#) by entering the following commands:

```
SETDefault !v101 -IP NETAddr = 180.20.55.5 255.255.255.252
```

```
SETDefault !v1101 -IP NETAddr = 180.20.55.123 255.255.255.252
```

```
SETDefault !v301 -IP NETAddr = 180.20.106.5 255.255.255.252
```

```
SETDefault !v1301 -IP NETAddr = 180.20.106.123 255.255.255.252
```



NOTE: The virtual port IP addresses must be on the same subnet as the BR to which the IPIP tunnel is connected.

- d. Specify the access point name (APN) by which you want to identify the CEN and bind the IPIP tunnels to the APN, associating each IPIP tunnel with a BR in the CEN (two BRs for each subnet) and assigning it a session priority by entering the following commands:

```
SETDefault !v101 -GGSN AccessPointName = "CountyFire.net" 180.20.55.6 4
SETDefault !v1101 -GGSN AccessPointName = "CountyFire.net" 180.20.55.124 3
SETDefault !v301 -GGSN AccessPointName = "CountyFire.net" 180.20.106.6 2
SETDefault !v1301 -GGSN AccessPointName = "CountyFire.net" 180.20.106.124 1
```

- e. Enable the APN by entering:

```
SETDefault !CountyFire.net -GGSN AccessPointCtl = Enable
```

- f. Configure the APN to assign subscriber IP addresses via a local DHCP server and configure the local DHCP server by entering the following commands:

```
SETDefault !CountyFire.net -GGSN IPAssigner = LocalDhcpServer
ADD !CountyFire.net -GGSN AddressPool 10.0.0.1 - 10.0.254.254
SETDefault !CountyFire.net -GGSN DHCPReXmitTime = 2
SETDefault !CountyFire.net -GGSN DHCPReXmitCnt = 3
SETDefault !CountyFire.net -GGSN DHCPSTwitchCnt = 1
```

- g. Optional: Configure GGSN charging. See [Configuring GGSN Charging on page 659](#).
h. Optional: Configure overload protection. See the [Configuring EOS Devices for Overload Protection on page 752](#).
i. Enable GGSN functionality on the device by entering:

```
SETDefault -GGSN CONTROL = Enable
```

35.14

Configuring Conventional Data Redundancy (HA and DSR) Feature

To configure the conventional data redundancy (HA and DSR) feature you must configure the following devices:

- GGSNs
- Border routers

Process:

1. Configure a GGSN to support conventional data redundancy (HA and DSR). See [Configuring a GGSN to Support Conventional Data Redundancy \(HA and DSR\) on page 622](#).
2. Configure border routers (BRs) to support conventional data redundancy (HA and DSR). See [Configuring Border Routers to Support Conventional Data Redundancy \(HA and DSR\) on page 629](#).

35.14.1

Configuring a GGSN to Support Conventional Data Redundancy (HA and DSR)

You can perform this procedure to configure a GGSN to support a conventional data redundancy (HA and DSR) feature.

You must include the configuration from [step 1](#) in the `boot.cfg` file of the GGSN.

You must include the configurations from [step 2](#) through [step 10](#) in the `xgsn.cfg` file of the GGSN.

Process:

1. Configure the GGSN as an HA peer by performing the following actions:
 - a. Define macros to enable and disable the VRRP control on the GGSN Ethernet ports that connect to the data and zone network management (ZNM) subnets, by entering:

```
define <enable_macro_name>=(
SETDefault !<port> -VRRP CONTrol = Enable <VRID_1>
SETDefault !<port> -VRRP CONTrol = Enable <VRID_2>
)
define <disable_macro_name>=(
SETDefault !<port> -VRRP CONTrol = Disable <VRID_1>
SETDefault !<port> -VRRP CONTrol = Disable <VRID_2>
)
```

where:

<enable_macro_name> is the name that you assign to the macro which enables the VRRP control (for example `vrrpEnable`)
 <disable_macro_name> is the name that you assign to the macro which disables the VRRP control (for example `vrrpDisable`)
 <VRID_1> is the VRID associated with the data subnet
 <VRID_2> is the VRID associated with the ZNM subnet



NOTE: The VRIDs you specify must match the VRIDs you assign to the ports when you configure the backup VRRP functionality (see [Configuring the GGSN as a Backup VRRP Router on page 628](#)).

- b. Define `lanUp` and `lanDown` macros that execute the macros you defined in [step 1a](#), to enable and disable VRRP control by entering:

```
define lanUp=
$global++
if $global==2
do <enable_macro_name>
end
)
define lanDown=
$global++
if $global==1
do <disable_macro_name>
end
)
```

- c. Assign IP addresses on the data and ZNM subnets to the GGSN Ethernet ports that connect to these subnets by entering:

```
SETDefault !<port> -IP NETaddr = <IP_address> <mask>
```

where:

<port> is the Ethernet port number,
<IP_address> is the IP address, on the appropriate subnet,
<mask> is the network mask.

- d. As part of the IP address setup, configure a system IP address (SysIP) on the GGSN by entering:

```
SETDefault -IP SystemIP = <IP_address>
```

where *<IP_address>* is the SysIP.

- e. Enable the loopback system IP feature for the NTP, FTP, PING, SNMP, Telnet, TFTP, and Syslog services so that the SysIP is used by sessions initiated by the GGSN by entering the following command on both HA GGSN peers:

```
SETDefault -SYS UseSystemIP = (NTP, FTP, PING, SNMP, Telnet, TFTP, SYSLOG)
```

- f. Configure the GGSN as a backup VRRP router on the data subnet and on the ZNM subnet. See [Configuring the GGSN as a Backup VRRP Router on page 628](#).

- g. For each GGSN Ethernet port on which VRRP is configured, set up the scheduler for event-based macro execution by entering the following commands:

```
SETDefault -SCH EbmeCONTROL = Enable
```

```
ADD !<port> -SCH EbmeEvent PortDown <DebounceTimer> do lanDown
```

```
ADD !<port> -SCH EbmeEvent PortUp <DebounceTimer> do lanUp
```

where:

<port> is the GGSN Ethernet port number,
<DebounceTimer> is the amount of time (50-3600000 milliseconds) to wait to wait after receiving the port status change notification and before executing the macro.

When the port goes down (the *portDown* event), the scheduler executes the *lanDown* macro (which in turn executes the *vrrpDisable* macro).

When the port comes up (the *portUp* event), the scheduler executes the *lanUp* macro (which in turn executes the *vrrpEnable* macro).

- h. Specify the HA peer of the GGSN by entering:

```
ADD -GGSN HANdStandBy <IP_address> [<Peer_SysName>] [HelloTime <seconds> <count>]
```

where *<IP_address>* is the system IP address of the HA GGSN peer.



NOTE: EOS software supports only one context database synchronization peer per GGSN and only IPv4 address format.

<Peer_SysName> is an optional name used to identify the HA GGSN peer in SNMP traps. If you do not specify a *<SysName>*, the IP address of the GGSN context database synchronization peer is used in SNMP traps.

HelloTime <seconds> <count> is the heartbeat timer for context database synchronization. The specified *<seconds>* and *<count>* are the length of time (in seconds) and number of hellos before the GGSN declares the context database synchronization peer down and restarts the synchronization relationship. If you do not specify a *HelloTime*, the default values (3 seconds and 5 hellos) are used, it means that the GGSN times out the HA peer in 15 seconds (3 seconds per hello x 5 hellos).



WARNING: To avoid unexpected synchronization interruption, you must configure the same *HelloTime* values on both HA GGSN peers.

2. Configure basic GGSN parameters by performing the following actions:

- a. Specify the maximum number of PDP contexts that the GGSN can support by entering:

```
SETDefault -GGSN MaximumPDPContexts = <value>
```

where *<value>* is the maximum number of PDP contexts supported on the GGSN (1-65535). The default *<value>* is 20000.

- b. Set the type of service (TOS) value for the signal packets on the Gn interface (the interface with the Packet Data Gateway (PDG) over which the GTP runs) by entering:

```
SETDefault -GGSN SignalTos = <tos>
```

where *<tos>* is the TOS-byte value (0-255). The default *<tos>* is 0.



NOTE: The TOS value is a one-byte value, with the TOS priority value expressed in the first three bits. For a table that maps *SignalTos* values to the corresponding hexadecimal and TOS priority values, see "GGSN Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

- c. Set the TOS value for the data packets on the Gn interface (the interface with the Packet Data Gateway (PDG) over which GTP runs) by entering:

```
SETDefault -GGSN DataTos = <tos>
```

where *<tos>* is the TOS-byte value (0-255). The default *<tos>* is 0.

3. Configure the VRRP-controlled IPIP tunnels connecting the GGSN to the border routers by performing the following actions:

- a. Create and name the virtual ports over which the IPIP tunnels run (two tunnels to each border router by entering the following commands:

```
ADD !v<port> -Port VirtualPort IPIP <local_ipaddr> <remote_ipaddr> VRID  
<vrid@port-no>, <vrid@port-no>
```

```
SETDefault !v<port> -Port Name = "<port_name>"
```

where:

<local_ipaddr> is the service IP shared by the two GGSNs in each zone (one service IP per CEN subnet),

<remote_ipaddr> is the tunnel endpoint address on the border router.

The VRID keyword indicates that the IPIP tunnel is a VRRP-controlled tunnel, and the *<vrid@port-no>*, *<vrid@port-no>* syntax specifies the virtual router identifier (VRID) and the GGSN port number with which the VRID is associated for each of the two subnets.



NOTE: The "service IP" for the data subnet is the foreign agent IP address.

- b. Assign IP addresses to the virtual ports by entering:

```
SETDefault !v<port> -IP NETaddr = <IP_address> <mask>
```

where:

<IP_address> is the virtual port IP address.

<mask> is the subnet mask.



NOTE: The virtual port IP addresses must be on the same subnet as the border router to which the IPIP tunnel is connected.

- c. Specify that fragmented IPIP packets are reassembled at the receiver end of the IPIP tunnel before forwarding to their destinations by entering:

```
SETDefault !<port> -IP FwdCONTRol = ReassembleFragPkts
```

4. For each IPIP tunnel to the border router, configure a static route to the home agent, with path metrics to prioritize the routes by entering:

```
ADD -IP ROUTe <IP_address> <mask> <gateway> <metric> [Override] BFD
```

where

<IP_address> is the home agent IP address,

<mask> is the subnet mask,
<gateway> is IPIP tunnel endpoint address on the border router,
<metric> is the path metric (the lower the metric, the higher the priority for the route)



NOTE: The GGSN uses the IPIP tunnel with the highest priority value to connect to the CEN. If that tunnel fails, the GGSN switches over to the other IPIP tunnel until the higher priority link is re-established. When a tunnel fails, BFD removes the corresponding static route to the home agent and the route with the next lowest metric becomes active.

The **Override** keyword specifies that a learned route can take precedence over the static route (the device can use a gateway other than the one specified to forward a packet to the destination). Do not include this keyword when configuring the highest priority route.

5. Specify the APN name, bind it to each of the IPIP tunnels, and assign each tunnel a session priority by entering the following command for each IPIP tunnel:

```
SETDefault !<port> -GGSN AccessPointName = "<APN>" <IP_address> <Priority>
```

where:

<APN> is the access point name
<IP_address> is the IPIP tunnel endpoint address on the border router
<Priority> is the session priority (the higher the number, the higher the priority).

6. Configure the mobility agent (foreign agent) functionality on the GGSN by performing the following actions:

- a. Enable foreign agent functionality by entering:

```
SETDefault -MOBILEIP CONTROL = ForeignAgent
```

- b. Configure the foreign agent Care Of Address (COA) by entering:

```
ADD -MOBILEIP ForeignAgent <COAddress> [<HomeAgentAddr>]
```

where:

<COAddress> is the foreign agent COA
<HomeAgentAddr> is the home agent address



NOTE: The two GGSNs in a zone share the same COA, and this address must match the COA configured on the border routers.

- c. Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol uses to replicate the foreign agent visitor tables by entering:

```
SETDefault -MOBILEIP MblReplIP = <IP_address>
```

where *<IP_address>* is the IP address to be used for foreign agent visitor table replication.

7. Enable the APN and configure the home agent IP address associated with this APN by entering the following commands:

```
SETDefault !<APN> -GGSN AccessPointCtl = Enable
```

```
SETDefault !<APN> -GGSN HomeAgentAddr = <IP_address>
```

where *<IP_address>* is the home agent IP address.

8. Configure a multi-hop VRRP-controlled BFD session for the link between the foreign agent and the home agent by performing the following actions:

- a. Configure a multi-hop BFD session on the Ethernet port that connects to the data subnet by entering:

```
ADD !<port> -IP BFDSESSion <session_ID> <peer_IP> [<TTL> (2..255)]  
[<KeepAliveTime> (100 - 100000ms) (Default = 300)] [<TimeoutMultiplier> (3 - 10)  
(Default = 3)]
```

where:

<session_ID> is the session number you want to assign to this session,
<peer_IP> is the home agent IP address,
<TTL> is the time-to-live value (the number of hops before the packet is discarded),
<KeepAliveTime> and <TimeoutMultiplier> are used to calculate the amount of time within which the BFD protocol must receive a reply to its ping before changing the BFD session state to down

There is one multi-hop BFD session per CEN. This enables the MOBILEIP service to switch to the backup GGSN if the active GGSN fails.



NOTE: Setting the TTL value to a number greater than 1 specifies that a multi-hop BFD session is established over UDP port 4784. If you set the TTL value to 1, you configure a single-hop BFD session.

- b. Register the MOBILEIP service with the BFD session you created by entering:

```
ADD -MOBILEIP MblBfdEvent <session_ID>
```

where <session_ID> is the session number you assigned when you created the multi-hop BFD session.

This configures the MOBILEIP service to receive BFD up/down notifications for the link between the foreign agent and the home agent.

- c. Configure a VRRP trigger to stop and start the BFD session (this enables the MOBILEIP service to switch to the backup GGSN if the active GGSN fails) by entering:

```
SETDefault !<port> -IP BFDTrigger = VRRPTrigger <vrid>
```

where:

<port> is the Ethernet port that connects to the data subnet,
<vrid> is the VRID associated with the port (defined after VRRP backup functionality on the port is configured. For more information, see [Configuring the GGSN as a Backup VRRP Router on page 628](#)).

When you configure a VRRP trigger (set the BFDTrigger parameter to VRRPTrigger), BFD registers with VRRP to check if the GGSN is the VRRP master on that interface.

- If the GGSN is the VRRP master on that interface, it starts the BFD session and sends/receives BFD packets.
- If the GGSN is the VRRP backup on that interface it does not send or respond to VRRP packets.

9. Configure subscriber home address allocation through the remote DHCP server configured on the border routers by performing the following actions:

- a. Specify the DHCP server to be the IP assigner by entering:

```
SETDefault !<APN> -GGSN IPAssigner = RemoteDhcpServer
```

- b. Specify IP address and address pool from which subscriber IP addresses are allocated on the primary DHCP server by entering:

```
SETDefault !<APN> -GGSN DHCPPrimSvr = <IP_address> <Gateway> <mask>
```

where:

<IP_address> is the home agent IP address
<Gateway> is the Gateway IP address (the gateway IP address to the DHCP server)
<mask> is the subnet mask for DHCP profile 1.



NOTE: The DHCP profiles and address pools are configured on the border router.

- c. Specify the IP address and address pool from which subscriber IP addresses are allocated on the secondary (backup) DHCP server by entering:

```
SETDefault ! <APN> -GGSN DHCPsecSrvr = <IP_address> <Gateway> <mask>
```

where:

<IP_address> is the home agent IP address

<Gateway> is the Gateway IP address (the gateway IP address to the DHCP server)

<mask> is the subnet mask for DHCP profile 2.

- d. Specify the time interval, in seconds between DHCP message retransmissions by entering:

```
SETDefault ! <APN> -GGSN DHCPReXmitTime = <seconds>
```

The default retransmission interval is 2 seconds.

- e. Specify the maximum number of times the GGSN retransmits mobile subscriber IP address requests by entering:

```
SETDefault ! <APN> -GGSN DHCPReXmitCnt = <value>
```

The default retransmission count is 3.

- f. Set the DHCP switchover count to 0 by entering:

```
SETDefault ! <APN> -GGSN DHCPswitchCnt = 0
```

Setting the DHCPswitchCnt parameter to 0 specifies that when a mobile subscriber IP address request fails, the GGSN does **not** switch from the primary DHCP server to the secondary DHCP server or the secondary DHCP server to the primary DHCP server.

- g. Optional: Configure GGSN charging. See [Configuring GGSN Charging on page 659](#).

- h. Optional: Configure overload protection. See the [Configuring EOS Devices for Overload Protection on page 752](#).

10. Enable GGSN functionality by entering:

```
SETDefault -GGSN CONTrol = Enable
```

35.14.1.1

Configuring the GGSN as a Backup VRRP Router

You can perform this procedure to configure the GGSN as a backup VRRP router on the data subnet and on the ZNM subnet.

You must include the following configurations in the `xgsn.cfg` file of the GGSN.

Procedure:

1. Configure backup VRRP router functionality on the Ethernet port that connects to the datasubnet and to the Ethernet port that connects to the ZNM subnet by entering:

```
ADD ! <port> -VRRP BackUp AsOwner <vrid> <IP_address>
```

where:

<port> is the port on which you configure the VRRP,

AsOwner is a keyword indicating that when the port takes over the IP address, it behaves as if it owns the IP address,

<vrid> is the VRID to be associated with the port,

<IP_address> is the "service IP" shared by the two HA GGSN peers in the zone for that subnet. The GGSN that is the master VRRP owns these addresses. If the master VRRP goes down, the backup VRRP takes over these addresses.



NOTE: On the data subnet, the "service IP" is the foreign agent IP address for the zone.

2. Set the advertisement interval for VRRP packets by entering:

```
SETDefault !<port> -VRRP AdvertiseInt = <adv_time> <vrid>
```

where:

<port> is the port on which you configure the VRRP,
<adv_time> is the advertisement interval, in milliseconds (1 - 100,000),
<vrid> is the VRID associated with the port.



NOTE: You must set the advertisement interval to the same value on both HA GGSN peers.

3. Set the VRRP hold time (the time interval before the backup VRRP declares the master VRRP down) by entering:

```
SETDefault !<port> -VRRP HoldTime = <hold_time> <vrid>
```

where:

<port> is the port on which you are configuring VRRP,
<hold_time> is the hold time in milliseconds (300 – 3000000),
<vrid> is the VRID associated with the port.



NOTE: The hold time should be at least three times the advertisement interval.

4. Set the VRRP priority by entering:

```
SETDefault !<port> -VRRP PRIOrity = <priority_number> <vrid>
```

where:

<port> is the port on which you are configuring VRRP,
<priority_number> is the priority value (1-254),
<vrid> is the VRID associated with the port.



NOTE: You must set the VRRP priority to different values on the two ports and the values you set must match the VRRP priorities on the corresponding ports on the HA GGSN peer (the other GGSN in the zone).

35.14.2

Configuring Border Routers to Support Conventional Data Redundancy (HA and DSR)

You can perform this procedure to configure a border router to support the conventional data redundancy (HA and DSR) feature. It is not a comprehensive border router configuration procedure.

You must include the configurations from this procedure in the `boot.cfg` file of the border router.

Procedure:

1. Configure virtual ports to access the CEN VLANs by performing the following actions:
 - a. Create virtual ports to access the CEN applications and home agent VLANs by entering:

```
ADD !v<port> -Port VirtualPort <path> Vlan <vlid>
```

where:

<vlid> is the path number to associated with this virtual port,
<port> is the unique 802.1Q virtual LAN identifier, in decimal (1-4094).



NOTE: Creating a separate home agent VLAN allows you to separate the mobile IP traffic on the border router from the other CEN applications.

- b. Assign IP addresses to the virtual ports by entering:

```
SETDefault !v<port> -IP NETaddr = <IP_address> <mask>
```

where:

<IP_address> is the virtual port IP address,
<mask> is the subnet mask.



NOTE: Configure the IP address of the virtual port that accesses the CEN applications VLAN as the SysIP. Configure the IP address of the virtual port that accesses the home agent VLAN as an address on the same subnet as the home agent IP address.

2. Configure a system IP (SysIP) address by performing the following actions:

- a. Specify a SysIP address on the border router by entering:

```
SETDefault -IP SystemIP = <IP_address>
```

where <IP_address> is the SysIP address.

- b. Enable the loopback system IP feature for the NTP, FTP, PING, SNMP, Telnet, TFTP, and Syslog services so that the SysIP is used by sessions initiated by the border router by entering:

```
SETDefault -SYS UseSystemIP = (NTP, FTP, PING, SNMP, Telnet, TFTP, SYSLOG)
```

3. Configure the virtual ports over which the IPIP tunnels connecting the border router to the GGSN are run by performing the following actions:

- a. Create the virtual ports by entering:

```
ADD !v<port> -Port VirtualPort IPIP <local_ipaddr> <remote_ipaddr>
```

where:

<local_ipaddr> is the IP address of the virtual port on the border router over which the encrypted tunnel to the firewall runs,
<remote_ipaddr> is the firewall IP address for the GGSN. You must configure one tunnel to each GGSN (two GGSNs per zone).



NOTE: The IPIP tunnels connecting the border router to the GGSNs run inside encrypted tunnels between the border router and the firewall and are unencrypted between the firewall and the GGSN. For this reason, you must configure the local tunnel endpoints on the border router to be the IP addresses of the virtual ports on the border router over which the encrypted tunnels run, and you must configure the remote tunnel endpoints to be the firewall IP addresses for the GGSNs.

- b. Enable TOS priority queuing for the virtual ports by entering:

```
SETDefault !v<port> -IPQOS CONTrol = TosQueue
```

- c. Assign IP addresses to the virtual ports by entering:

```
SETDefault !v<port> -IP NETaddr = <IP_address> <mask> Ones <MTU>
```

where:

<IP_address> is the virtual port IP address,
<mask> is the subnet mask. Ones indicates that for packets using this IP broadcast address, the host portion of the IP address contains all 1 bits.
<MTU> is the maximum transmission unit, in bytes.

4. Configure the BFD-controlled static routes to the foreign agents on the GGSNs in both zones (two routes per zone) with path metrics to prioritize the two routes in each zone by entering:

```
ADD -IP ROUTe <IP_address> <mask> <gateway> <metric> [Override] BFD
```

where:

<IP_address> is the foreign agent COA,
<gateway> is the IP address of the virtual port on the GGSN over which the IPIP tunnel connecting the border router to the GGSN runs.

5. Configure BFD-controlled static routes from the gateway IP addresses of the DHCP servers (from which the subscriber home addresses are allocated) to the GGSNs, with path metrics to prioritize the routes by entering:

```
ADD -IP Route <IP_address> <mask> <gateway> <metric> [Override] BFD
```

where:

<IP_address> is the gateway IP address of the DHCP server configured on the border router,

<mask> is the subnet mask,

<gateway> is the IP address of the virtual port on the GGSN over which the IPIP tunnel connecting the border router to the GGSN runs.



NOTE: The gateway IP address to the DHCP server is also referred to as the *giaddr*.

6. Configure NAT redundancy by performing the following actions:



NOTE: NAT redundancy requires that both border routers in the same CEN be configured with the same NAT replication IP address (NatReplIP) and the same NAT address mapping.

- a. Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol uses to replicate the TCP connections on the border router by entering:

```
SETDefault -NAT NatReplIP = <IP_address>
```

where <IP_address> is the VRRP address shared by the two border routers on the CEN LAN interface. The border router that is the master VRRP owns this address. If the master VRRP goes down, the backup VRRP takes over this address.

- b. Optional: Configure the NAT replication port, by entering:

```
ADD !<localport> -NAT NATReplPort <remoteport>
```

where:

<localport> is the NAT port on the local device,

<remoteport> is the port on the NAT port on the replication peer (the other border router).



NOTE: Specifying a NATReplPort is required only if the border router and its replication peer (the other border router in the CEN) use different NAT ports.

7. Configure OSPF service parameters to support inter-system data functionality by performing the following actions:

- a. Configure the border router to always add a Type 7 LSA route, independent of the P-bit value in the Type 7 LSA by entering:

```
SETDefault -OSPF NSSACont = IgnorePbitAbr
```

- b. Specify that the networks on the virtual ports over which the IPIP tunnels connecting the border router to the GGSNs are run are advertised into the OSPF domain as Type 1 external OSPF routes, and that these routes have a corresponding Type 7 LSA only (no Type 5 LSA is generated) by entering:

```
SETDefault !v<port> -OSPF DirectPolicy = (Advertise, Type1)  
DontAdvertiseLsaType5
```

where <port> is the virtual port over which the IPIP tunnel connecting the border router to the GGSN runs.

- c. Add the foreign agent Care of Address (COA) for each zone to the list of statically-configured routes, by specifying that the metric in the routing table is used as the advertised cost of this route, a link state metric (Type1) is used, and the route has a corresponding Type 7 LSA only (no Type 5 LSA is generated) by entering:

```
ADD -OSPF StaticPolicy <IP_address> 0 Type1 DontAdvertiseLsaType5
```

where *<IP_address>* is the foreign agent Care Of Address (COA).

8. Configure mobility agent (home agent) functionality by performing the following actions:

- a. Configure an address list associated with the home agent subnet, by entering:

```
ADD -IP AddressList <name> (1-15 char) <IP_address> [/ <mask> (0-32)]
```

where:

<name> is the address list name,
<IP_address> is home agent subnet address,
<mask> is the prefix length.

- b. Enable home agent functionality on the border router, by entering:

```
SETDefault -MOBILEIP CONTROL = HomeAgent
```

- c. Configure the home agent IP address on the border router and associate it with the address list you configured in [step 7a](#), by entering:

```
SETDefault -MOBILEIP HomeAgent = <IP_address> <addr_list_name>
```

where:

<IP_address> is the home agent IP address,
<addr_list_name> is the address list name.

- d. Specify the foreign agent COAs (one COA per zone, shared by the two GGSNs in the zone) by entering:

```
ADD -MOBILEIP ForeignAgent <COAddress> [<HomeAgentAddr>]
```

where:

<COAddress> is the COA for the foreign agent
<HomeAgentAddr> is the home agent address.



NOTE: The COAs must match the addresses you configure on the GGSNs.

- e. Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol use to replicate the home agent mobility database by entering:

```
SETDefault -MOBILEIP MblReplIP = <IP_address>
```

where *<IP_address>* is the IP address to be used for home agent mobility database replication.

9. Configure the border router as a backup VRRP router on the CEN applications VLAN and on the home agent VLAN. See [Configure the Border Router as a Backup VRRP Router on page 635](#).

10. Configure multihop BFD sessions for the links between the home agent and the foreign agent in each zone by performing the following actions:

- a. Create two multihop BFD sessions (one session per zone) on the virtual port that accesses the home agent VLAN by entering:

```
ADD !<port> BFDSESSion <session_ID> <peer_IP> [<TTL> (2..255)] [<KeepAliveTime>  
(100 - 100000ms) (Default = 300)] [<TimeOutMultiplier> (3 - 10) (Default = 3)]
```

where:

<session_ID> is the session number you want to assign to this session,
<peer_IP> is the foreign agent COA,
<TTL> is the time-to-live value (the number of hops before the packet is discarded,

<KeepAliveTime> and <TimeoutMultiplier> are values used to calculate the amount of time within which the BFD protocol must receive a reply to its ping before changing the BFD session state to down.



NOTE: Setting the <TTL> value to a number greater than 1 specifies that a multi-hop BFD session is established over UDP port 4784. If you set the TTL value to 1, you configure a single-hop BFD session.

- b. Register the MOBILEIP service with the BFD sessions you created by entering:

```
ADD -MOBILEIP MblBfdEvent <session_ID>
```

where <session_ID> is the session number you assigned when you created the multi-hop BFD session.

This configures the MOBILEIP service to receive BFD up/down notifications for the link between the home agent and the foreign agent.

- c. Configure a VRRP trigger to stop and start the BFD session (enabling the MOBILEIP service to switch to the backup border router if the active border router fails) by entering:

```
SETDefault !v<port> -IP BFDTrigger = VRRPTrigger <vrid>
```

where:

<port> is the virtual port that connects to the home agent VLAN

<vrid> is the VRID associated with the port (defined when you configured VRRP backup functionality on the port. See [Configure the Border Router as a Backup VRRP Router on page 635](#).



NOTE: When you configure a VRRP trigger (set the BFDTrigger parameter to VRRPTrigger), BFD registers with VRRP to check if the border router is the VRRP master on that interface.

If the border router is the VRRP master on that interface, it starts the BFD session and sends/receives BFD packets.

If the border router is the VRRP backup on that interface it does not send or respond to VRRP packets.

11. Configure the border router to execute event-based macros to enable and disable VRRP control on the virtual ports associated with the CEN applications and home agent VLANs by performing the following actions:

- a. Define the lanUp and lanDown macros that are executed when a PortUp or PortDown event occurs by entering:

```
define lanUp=
```

```
SETDefault !v<port> -VRRP CONTROL = Enable <VRID_1>
```

```
SETDefault !v<port> -VRRP CONTROL = Enable <VRID_2>
```

```
)
```

```
define lanDown=(
```

```
SETDefault !v<port> -VRRP CONTROL = Disable <VRID_1>
```

```
SETDefault !v<port> -VRRP CONTROL = Disable <VRID_2>
```

```
)
```

where:

<VRID_1> is the VRID associated with CEN applications VLAN,

<VRID_2> is the VRID associated with the home agent VLAN.

You configure these VRIDs when you configure VRRP control on the VLANs. For more information, see [Configure the Border Router as a Backup VRRP Router on page 635](#).

- b. Enable event-based macro execution and specify that event-based macro execution is enabled and the *PortDown* actions performed when the primary connections fail to establish within 5 minutes after the device boots and that the macro name and a complete/incomplete message is reported in the system log buffer by entering:

```
SETDefault -SCH EbmeCONTRol = (Enabled, RunOnBootFail, Log)
```

- c. Set up the scheduler for event-based macro execution by entering the following commands:

```
SETDefault -SCH EbmeCONTRol = (Enable, RunOnBootFail, Log)
```

```
ADD !<port> -SCH EbmeEvent PortUp <DebounceTimer> do port_up
```

```
ADD !<port> -SCH EbmeEvent PortDown <DebounceTimer> do port_down
```

where:

<port> is the port corresponding to the path you associated with the virtual ports that access the CEN applications and home agent VLANs,

<DebounceTimer> is the amount of time (50-3600000 milliseconds) to wait after receiving the port status change notification and before executing the macro.

When the Ethernet port comes up (the *portUp* event), the scheduler executes the *lanUp* macro.

When the Ethernet port goes down (the *portDown* event), the scheduler executes the *lanDown* macro.

12. Configure DHCP server functionality by performing the following actions:

- a. On the virtual port that accesses the home agent VLAN and on the virtual ports over which the IPIP tunnels connecting the border router to the GGSN, enable DHCP; configure DHCP to use the address pool defined with the *AddressPool* parameter, specify that the DHCP server does not issue an ICMP ECHO request to make sure the address chosen from the pool is not duplicated on the network before allocating the address to a client, and disable subnet mask checking by entering:

```
SETDefault !v<port> -DHCP CONTRol = (Enable, AddressPool, NoIcmpCheck, NoMaskCheck)
```

- b. Configure the address pool from which subscriber home addresses are allocated on the primary DHCP server by entering:

```
ADD -DHCP AddressPool <IPaddr1> - <IPaddr2> !P1 <mask>
```

where:

<IPaddr1> - <IPaddr2> is the address range from which subscriber home addresses is allocated

<mask> is the subnet mask



NOTE: !P1 corresponds to the !<profileid> syntax. EOS creates DHCP profile P1 by default.

- c. Configure the gateway IP addresses (*giaddr* addresses) to the primary DHCP server subnet as reserved addresses which cannot be used for allocation by entering:

```
ADD -DHCP ReserveAddress <IP_address>
```

where <IP_address> is the *giaddr* you want to exclude.

- d. Create a second DHCP profile by entering:

```
ADD -DHCP PProfile <profileid>
```

where <profileid> is the profile ID you want to assign to the second DHCP profile.

- e. Configure the address pool from which subscriber home addresses are allocated on the secondary DHCP server by entering:

```
ADD -DHCP AddressPool <IPaddr1> - <IPaddr2> !P1 <mask>
```

where:

<IPaddr1> - <IPaddr2> is the address range from which subscriber home addresses are allocated,
<mask> is the subnet mask.

- f. Configure the gateway IP addresses (*giaddr* addresses) to the secondary DHCP server subnet as reserved addresses which cannot be used for allocation by entering:

```
ADD -DHCP ReserveAddress <IP_address>
```

where <IP_address> is the *giaddr* you want to exclude.

- g. Configure the IP address of the virtual port that connects to the home agent VLAN as a reserved address which cannot be used for allocation by entering:

```
ADD -DHCP ReserveAddress <IP_address>
```

where <IP_address> is the virtual port IP address you want to exclude.

- h. Configure the reserved IP address which cannot be used for allocation by entering:

```
ADD -DHCP ReserveAddress <IP_address>
```

where <IP_address> is the reserved IP address.

- i. Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol use to replicate DHCP allocations by entering:

```
SETDefault -DHCP DhcpReplIP = <IP_address>
```

where <IP_address> is the IP address to be used for DHCP allocation replication.

- j. Register the DHCP service on the home agent to receive BFD up/down notifications for the links between the home agent and the foreign agent in each zone by entering:

```
ADD -DHCP DhcpBfdEvent <session_ID>
```

where <session_ID> is the session number you assigned when you configured the BFD sessions for the links. For more information, see [Configuring a GGSN to Support Conventional Data Redundancy \(HA and DSR\) on page 622](#).

35.14.2.1

Configure the Border Router as a Backup VRRP Router

You can use this procedure to configure the border router as a backup VRRP router on the CEN applications VLAN and on the home agent VLAN.

Procedure:

1. Configure backup VRRP router functionality on the virtual port that access the CEN application and home agent VLANs by entering:

```
ADD !v<port> -VRRP BackUp AsOwner <vrid> <IP_address>
```

where:

<port> is the virtual port on which you configure VRRP,

AsOwner is a keyword indicating that when the port takes over the IP address, it behaves as if it owns the IP address,

<vrid> is the VRID to be associated with the port,

<IP_address> is the broadcast address associated with the SysIP (for the CEN applications VLAN) or home agent IP address (for the home agent VLAN) shared by the two border routers for that VLAN. The border router that is the master VRRP owns these addresses. If the master VRRP goes down, the backup VRRP takes over these addresses.

2. Set the advertisement interval for VRRP packets by entering:


```
SETDefault !<port> -VRRP AdvertiseInt = <adv_time> <vrid>
```

where:

<port> is the virtual port on which you configure VRRP,
<adv_time> is the advertisement interval, in milliseconds (1 – 100,000),
<vrid> is the VRID associated with the port.



NOTE: You must set the advertisement interval to the same value on both border routers.

3. Set the VRRP hold time (the time interval before the backup VRRP declares the master VRRP down) by entering:

```
SETDefault !<port> -VRRP HoldTime = <hold_time><vrid>
```

where:

<port> is the virtual port on which you configure VRRP,
<hold_time> is the hold time, in milliseconds (300 - 3000000),
<vrid> is the VRID associated with the port.



NOTE: The hold time must be at least three times the advertisement interval.

4. Set the VRRP priority by entering:

```
SETDefault !<port> -VRRP PRIOrity = <priority_number> <vrid>
```

where:

<port> is the virtual port on which you configure VRRP,
<hold_time> is the priority value (1-254),
<vrid> is the VRID associated with the port.



NOTE: You must set the VRRP priority to the same value on the two virtual ports on each border router; however the VRRP priority values on the border router connected to one zone must be different than those on the border router connected to the other zone.

35.15

Conventional Data Redundancy (HA and DSR) Configuration Example

The following figures illustrate a sample conventional data redundancy (HA and DSR) topology with two zones (Zone 1 and Zone 2) and two CENs (CEN 1 and CEN 2) and describes some of the key elements that must be configured in order to create this topology.

To configure the GGSNs in the sample topology, see [Configuring the GGSNs in the Sample Conventional Data Redundancy \(HA and DSR\) Topology on page 643](#).

To configure the border routers in the sample topology, see [Configuring Border Routers in the Sample Conventional Data Redundancy \(HA and DSR\) Topology on page 650](#).

The figure illustrates basic components of the sample conventional data redundancy (HA and DSR) topology, with backup VRRP functionality configured on the GGSNs in the data and ZNM subnets in both zones. The topology features:

Four GGSNs:


- GGSN 1 and GGSN 2 in Zone 1
- GGSN 3 and GGSN 4 in Zone 2

Four border routers (BRs)

- BR 1, in CEN 1 and physically attached to Zone 1

- BR 2, in CEN 1 and physically attached to Zone 2
- BR 3, in CEN 2 and physically attached to Zone 1
- BR 4, in CEN 2 and physically attached to Zone 2

Figure 179: Sample Conventional Data Redundancy (HA and DSR) Topology with Backup VRRP Functionality Configured on the GGSNs in the Data and ZNM Subnets in Both Zones



BR 1 is physically connected to Zone 1. BR 2 is physically connected to Zone 2.

On GGSN 1:
Ethernet port 1 (IP address 10.1.215.63), is configured for backup VRRP functionality on the data subnet with VRID 63.
Ethernet port 2 (IP address 10.1.213.140), is configured for backup VRRP functionality on the ZNM subnet with VRID 53.
GGSN 2, with SysIP 10.1.215.249, is configured as the HA GGSN peer.


On GGSN 2:
Ethernet port 1 (IP address 10.1.215.64) is configured for backup VRRP functionality on the data subnet with VRID 63.
Ethernet port 2 (IP address 10.1.213.141) is configured for backup VRRP functionality on the ZNM subnet with VRID 53.
GGSN 1, with SysIP 10.1.215.250, is configured as the HA GGSN peer.

BR 3 is physically connected to Zone 1. BR 4 is physically connected to Zone 2.

On GGSN 3:
Ethernet port 1 (IP address 10.2.215.63), is configured for backup VRRP functionality on the data subnet with VRID 63.
Ethernet port 2 (IP address 10.2.213.140), is configured for backup VRRP functionality on the ZNM subnet with VRID 53.
GGSN 4, with SysIP 10.2.215.249, is configured as the HA GGSN peer.

On GGSN 4:
Ethernet port 1 (IP address 10.2.215.64) is configured for backup VRRP functionality on the data subnet with VRID 63.
Ethernet port 2 (IP address 10.2.213.141) is configured for backup VRRP functionality on the ZNM subnet with VRID 53.
GGSN 3, with SysIP 10.2.215.250, is configured as the HA GGSN peer.

Figure 180: Sample Conventional Data Redundancy (HA and DSR) Topology: IPIP Tunnels from GGSNs to BRs on Data Subnet




On the Zone 1 GGSNs (GGSN 1 and GGSN 2):

A VRRP-controlled IPIP tunnel running over virtual port 101 (with IP address 180.35.75.11) connects GGSN 1 to BR 1 (virtual port 11 with IP address 180.35.75.12) on the data subnet in CEN 1; if GGSN 1 fails, the same tunnel moves to GGSN 2
A VRRP-controlled IPIP tunnel running over virtual port 1101 (with IP address 180.35.75.143) connects GGSN 1 to BR 2 (virtual port 12 with IP address 180.35.75.144) on the data subnet in CEN 1; if GGSN 1 fails, the same tunnel moves to GGSN 2
A VRRP-controlled IPIP tunnel running over virtual port 102 (with IP address 180.35.75.14) connects GGSN 1 to BR 3 (virtual port 11 with IP address 180.35.75.15) on the data subnet in CEN 2; if GGSN 1 fails, the same tunnel moves to GGSN 2.
A VRRP-controlled IPIP tunnel running over virtual port 1102 (with IP address 180.35.75.147) connects GGSN 1 to BR 4 (virtual port 12 with IP address 180.35.75.148) on the data subnet in CEN 2; if GGSN 1 fails, the same tunnel moves to GGSN 2.

On the Zone 2 GGSNs (GGSN 3 and GGSN 4):

A VRRP-controlled IPIP tunnel running over virtual port 101 (with IP address 180.35.41.11) connects GGSN 3 to BR 1 (virtual port 21 with IP address 180.35.41.12) on the data subnet in CEN 1; if GGSN 3 fails, the same tunnel moves to GGSN 4.
A VRRP-controlled IPIP tunnel running over virtual port 1101 (with IP address 180.35.41.143) connects GGSN 3 to BR 2 (virtual port 22 with IP address 180.35.41.144) on the data subnet in CEN 1; if GGSN 3 fails, the same tunnel moves to GGSN 4.
A VRRP-controlled IPIP tunnel running over virtual port 102 (with IP address 180.35.41.14) connects GGSN 3 to BR 3 (virtual port 21 with IP address 180.35.41.15) on the data subnet in CEN 2; if GGSN 3 fails, the same tunnel moves to GGSN 4.
A VRRP-controlled IPIP tunnel running over virtual port 1102 (with IP address 180.35.41.147) connects GGSN 3 to BR 4 (virtual port 22 with IP address 180.35.41.148) on the data subnet in CEN 2; if GGSN 3 fails, the same tunnel moves to GGSN 4.

Figure 181: Sample Conventional Data Redundancy (HA and DSR) Topology: IPIP Tunnels from GGSNs to BRs on ZNM Subnet



On the Zone 1 GGSNs (GGSN 1 and GGSN 2):

A VRRP-controlled IPIP tunnel running over virtual port 301 (with IP address 180.35.112.11) connects GGSN 1 to BR 1 (virtual port 13 with IP address 180.35.112.12) on the ZNM subnet in CEN 1; if GGSN 1 fails, the same tunnel moves to GGSN 2.

A VRRP-controlled IPIP tunnel running over virtual port 1301 (with IP address 180.35.112.143) connects GGSN 1 to BR 2 (virtual port 14 with IP address 180.35.112.144) on the ZNM subnet in CEN 1; if GGSN 1 fails, the same tunnel moves to GGSN 2.

A VRRP-controlled IPIP tunnel running over virtual port 302 (with IP address 180.35.112.14) connects GGSN 1 to BR 3 (virtual port 13 with IP address 180.35.112.15) on the ZNM subnet in CEN 2; if GGSN 1 fails, the same tunnel moves to GGSN 2.

A VRRP-controlled IPIP tunnel running over virtual port 1302 (with IP address 180.35.112.147) connects GGSN 1 to BR 4 (virtual port 14 with IP address 180.35.112.148) on the ZNM subnet in CEN 2; if GGSN 1 fails, the same tunnel moves to GGSN 2.

On the Zone 2 GGSNs (GGSN 3 and GGSN 4):

A VRRP-controlled IPIP tunnel running over virtual port 301 (with IP address 180.35.106.11) connects GGSN 3 to BR 1 (virtual port 23 with IP address 180.35.106.12) on the ZNM subnet in CEN 1; if GGSN 3 fails, the same tunnel moves to GGSN 4.

A VRRP-controlled IPIP tunnel running over virtual port 1301 (with IP address 180.35.106.143) connects GGSN 3 to BR 2 (virtual port 24 with IP address 180.35.106.144) on the ZNM subnet in CEN 1; if GGSN 3 fails, the same tunnel moves to GGSN 4.

A VRRP-controlled IPIP tunnel running over virtual port 302 (with IP address 180.35.106.14) connects GGSN 3 to BR 3 (virtual port 23 with IP address 180.35.106.15) on the ZNM subnet in CEN 2; if GGSN 3 fails, the same tunnel moves to GGSN 4.

A VRRP-controlled IPIP tunnel running over virtual port 1302 (with IP address 180.35.106.147) connects GGSN 3 to BR 4 (virtual port 24 with IP address 180.35.106.148) on the ZNM subnet in CEN 2; if GGSN 3 fails, the same tunnel moves to GGSN 4.

Figure 182: Sample Conventional Data Redundancy (HA and DSR) Topology: Mobility Agents






Figure 183: Sample Conventional Data Redundancy (HA and DSR) Topology: Roamed Subscriber



When the subscriber roams to another zone, it maintains its IP address on the 21.21.x.x or 21.22.x.x subnet and its association with FA 1 (with COA 10.1.214.35). Traffic to and from the subscriber continues to be tunneled through HA 1, (with address 21.1.2.254 and associated with APN "CEN1APN1").

Figure 184: Sample Conventional Data Redundancy (HA and DSR) Topology: DSR Failure Scenario



35.15.1

Configuring the GGSNs in the Sample Conventional Data Redundancy (HA and DSR) Topology

Procedure:

1. Configure the GGSN as an HA peer by performing the following actions:
 - a. On all GGSNs, define `vrpEnable` and `vrpDisable` macros to enable and disable VRRP control on the GGSN Ethernet ports that connect to the data and zone network management (ZNM) subnets by entering:

```
define vrrpEnable=(
SETDefault !1 -VRRP CONTROL = Enable 63
SETDefault !2 -VRRP CONTROL = Enable 53
)
```

```

define vrrpDisable=(
  SETDefault !1 -VRRP CONTrol = Disable 63
  SETDefault !2 -VRRP CONTrol = Disable 53
)
)

```

- b.** On all GGSNs, define `lanUp` and `lanDown` macros that execute the macros you defined in [step 1a](#) (to enable and disable VRRP control) by entering:

```

define lanUp=
$global++
if $global==2
do vrrpEnable
end
)
define lanDown=
$global++
if $global==1
do vrrpDisable
end
)

```

- c.** To assign IP addresses on the data and ZNM subnets to the GGSN Ethernet ports that connect to these subnets, on the GGSN 1 enter:

```

SETDefault !1 -IP NETaddr = 10.1.215.63 255.255.255.0
SETDefault !2 -IP NETaddr = 10.1.213.140 255.255.255.0

```

- d.** To assign IP addresses on the data and ZNM subnets to the GGSN Ethernet ports that connect to these subnets, on the GGSN 2 enter:

```

SETDefault !1 -IP NETaddr = 10.1.215.64 255.255.255.0
SETDefault !2 -IP NETaddr = 10.1.213.141 255.255.255.0

```

- e.** To assign IP addresses on the data and ZNM subnets to the GGSN Ethernet ports that connect to these subnets, on the GGSN 3 enter:

```

SETDefault !1 -IP NETaddr = 10.2.215.63 255.255.255.0
SETDefault !2 -IP NETaddr = 10.2.213.140 255.255.255.0

```

- f.** To assign IP addresses on the data and ZNM subnets to the GGSN Ethernet ports that connect to these subnets, on the GGSN 4 enter:

```

SETDefault !1 -IP NETaddr = 10.2.215.64 255.255.255.0
SETDefault !2 -IP NETaddr = 10.2.213.141 255.255.255.0

```

- g.** As part of the IP address setup, to configure a system IP address (SysIP) on the GGSN, on the GGSN 1 enter:

```

SETDefault -IP SystemIP = 10.1.215.250

```

- h.** As part of the IP address setup, to configure a system IP address (SysIP) on the GGSN, on the GGSN 2 enter:

```

SETDefault -IP SystemIP = 10.1.215.249

```


- i. As part of the IP address setup, to configure a system IP address (SysIP) on the GGSN, on the GGSN 3 enter:

```
SETDefault -IP SystemIP = 10.2.215.250
```

- j. As part of the IP address setup, to configure a system IP address (SysIP) on the GGSN, on the GGSN 4 enter:

```
SETDefault -IP SystemIP = 10.2.215.249
```

- k. Enable the loopback system IP feature for the NTP, FTP, PING, SNMP, Telnet, TFTP, and Syslog services so that the SysIP is used by sessions initiated by the GGSN.

On all GGSNs, enter:

```
SETDefault -SYS UseSystemIP = (NTP, FTP, PING, SNMP, Telnet, TFTP, SYSLOG)
```

- l. To configure the GGSN as a backup VRRP router on the data subnet and on the ZNM subnet, configure backup VRRP router functionality on the Ethernet port that connects to the data subnet and to the Ethernet port that connects to the ZNM subnet by entering the following commands on GGSN 1 and GGSN 2:

```
ADD !1 -VRRP BackUp AsOwner 63 10.1.215.35
```

```
ADD !2 -VRRP BackUp AsOwner 53 10.1.213.151
```

- m. To configure the GGSN as a backup VRRP router on the data subnet and on the ZNM subnet, configure backup VRRP router functionality on the Ethernet port that connects to the data subnet and to the Ethernet port that connects to the ZNM subnet by entering the following commands on GGSN 3 and GGSN 4:

```
ADD !1 -VRRP BackUp AsOwner 63 10.2.215.35
```

```
ADD !2 -VRRP BackUp AsOwner 53 10.2.213.151
```

- n. To configure the GGSN as a backup VRRP router on the data subnet and on the ZNM subnet, set the advertisement interval for VRRP packets by entering the following commands on all GGSNs:

```
SETDefault !1 -VRRP AdvertiseInt = 300 63
```

```
SETDefault !2 -VRRP AdvertiseInt = 300 53
```

- o. To configure the GGSN as a backup VRRP router on the data subnet and on the ZNM subnet, set the VRRP hold time (the time interval before the backup VRRP declares the master VRRP down) by entering the following commands on all GGSNs:

```
SETDefault !1 -VRRP HoldTime = 900 63
```

```
SETDefault !2 -VRRP HoldTime = 900 53
```

- p. To configure the GGSN as a backup VRRP router on the data subnet and on the ZNM subnet, set the VRRP priority by entering the following commands on all GGSNs:

```
SETDefault !1 -VRRP PRIOrity = 210 63
```

```
SETDefault !2 -VRRP PRIOrity = 210 53
```

- q. For each GGSN Ethernet port on which VRRP is configured, set up the scheduler for event-based macro execution by entering the following commands on all GGSNs:

```
SETDefault -SCH EbmeCONTrol = Enable
```

```
ADD !1 -SCH EbmeEvent PortDown 50 do lanDown
```

```
ADD !1 -SCH EbmeEvent PortUp 5000 do lanUp
```

```
ADD !2 -SCH EbmeEvent PortDown 50 do lanDown
```

```
ADD !3 -SCH EbmeEvent PortUp 5000 do lanUp
```

- r. To specify the HA peer of the GGSN, on GGSN 1 enter:

```
ADD -GGSN HANstandBy 10.1.215.249 z1ggsn2 HelloTime 1 3
```

- s. To specify the HA peer of the GGSN, on GGSN 2 enter:**

```
ADD -GGSN HANstandBy 10.1.215.250 z1ggsn1 HelloTime 1 3
```

- t. To specify the HA peer of the GGSN, on GGSN 3 enter:**

```
ADD -GGSN HANstandBy 10.2.215.249 z2ggsn2 HelloTime 1 3
```

- u. To specify the HA peer of the GGSN, on GGSN 4 enter:**

```
ADD -GGSN HANstandBy 10.2.215.250 z2ggsn1 HelloTime 1 3
```

- 2. Configure basic GGSN parameters by performing the following actions:**

- a. Specify the maximum number of PDP contexts that the GGSN can support by entering on all GGSNs:**

```
SETDefault -GGSN MaximumPDPContexts = 20000
```

- b. Set the TOS (type of service) value for the signal packets on the Gn interface (the interface with the PDG (Packet Data Gateway) over which GTP runs) by entering on all GGSNs:**

```
SETDefault -GGSN SignalTos = 160
```

- c. Set the TOS (type of service) value for the data packets on the Gn interface (the interface with the PDG (Packet Data Gateway) over which GTP runs) by entering on all GGSNs:**

```
SETDefault -GGSN DataTos = 64
```

- 3. Configure the VRRP-controlled IPIP tunnels connecting the GGSN to the border routers by performing the following actions:**

- a. Create and name the eight virtual ports over which the IPIP tunnels run (two tunnels to each border router) by entering the following commands on GGSN 1 and GGSN 2:**

```
ADD !v101 -Port VirtualPort IPIP 10.1.213.151 10.1.213.173 VRID 63@1, 53@2
```

```
SETDefault !v101 -Port Name = "Port_CEN1A"
```

```
ADD !v1101 -Port VirtualPort IPIP 10.1.213.151 10.1.213.74 VRID 63@1, 53@2
```

```
SETDefault !v1101 -Port Name = "Port_CEN1B"
```

```
ADD !v301 -Port VirtualPort IPIP 10.1.215.35 10.2.213.173 VRID 63@1, 53@2
```

```
SETDefault !v101 -Port Name = "Port_CEN1C"
```

```
ADD !v1301 -Port VirtualPort IPIP 10.1.215.35 10.2.213.74 VRID 63@1, 53@2
```

```
SETDefault !v101 -Port Name = "Port_CEN1D"
```

```
ADD !v102 -Port VirtualPort IPIP 10.1.213.151 10.1.213.174 VRID 63@1, 53@2
```

```
SETDefault !v101 -Port Name = "Port_CEN2A"
```

```
ADD !v1102 -Port VirtualPort IPIP 10.1.213.151 10.1.213.75 VRID 63@1, 53@2
```

```
SETDefault !v1101 -Port Name = "Port_CEN2B"
```

```
ADD !v302 -Port VirtualPort IPIP 10.1.215.35 10.2.213.174 VRID 63@1, 53@2
```

```
SETDefault !v101 -Port Name = "Port_CEN2C"
```

```
ADD !v1302 -Port VirtualPort IPIP 10.1.215.35 10.2.213.75 VRID 63@1, 53@2
```

```
SETDefault !v101 -Port Name = "Port_CEN2D"
```

- b. Create and name the eight virtual ports over which the IPIP tunnels run (two tunnels to each border router) by entering the following commands on GGSN 3 and GGSN 4:**

```
ADD !v101 -Port VirtualPort IPIP 10.2.213.151 10.2.213.173 VRID 63@1, 53@2
```

```
SETDefault !v101 -Port Name = "Port_CEN1A"
```

```
ADD !v1101 -Port VirtualPort IPIP 10.2.213.151 10.2.213.74 VRID 63@1, 53@2
SETDefault !v1101 -Port Name = "Port_CEN1B"
ADD !v301 -Port VirtualPort IPIP 10.2.215.35 10.1.213.173 VRID 63@1, 53@2
SETDefault !v101 -Port Name = "Port_CEN1C"
ADD !v1301 -Port VirtualPort IPIP 10.2.215.35 10.1.213.74 VRID 63@1, 53@2
SETDefault !v101 -Port Name = "Port_CEN1D"
ADD !v102 -Port VirtualPort IPIP 10.2.213.151 10.2.213.174 VRID 63@1, 53@2
SETDefault !v101 -Port Name = "Port_CEN2A"
ADD !v1102 -Port VirtualPort IPIP 10.2.213.151 10.2.213.75 VRID 63@1, 53@2
SETDefault !v1101 -Port Name = "Port_CEN2B"
ADD !v302 -Port VirtualPort IPIP 10.2.215.35 10.1.213.174 VRID 63@1, 53@2
SETDefault !v101 -Port Name = "Port_CEN2C"
ADD !v1302 -Port VirtualPort IPIP 10.2.215.35 10.1.213.75 VRID 63@1, 53@2
SETDefault !v101 -Port Name = "Port_CEN2D"
```

- c. Assign IP addresses to the virtual ports by entering the following commands on GGSN 1 and GGSN 2:**

```
SETDefault !v101 -IP NETAddr = 180.35.75.11 255.255.255.252
SETDefault !v1101 -IP NETAddr = 180.35.75.143 255.255.255.252
SETDefault !v301 -IP NETAddr = 180.35.112.11 255.255.255.252
SETDefault !v1301 -IP NETAddr = 180.35.112.143 255.255.255.252
SETDefault !v102 -IP NETAddr = 180.35.75.14 255.255.255.252
SETDefault !v1102 -IP NETAddr = 180.35.75.147 255.255.255.252
SETDefault !v302 -IP NETAddr = 180.35.112.14 255.255.255.252
SETDefault !v1302 -IP NETAddr = 180.35.112.147 255.255.255.252
```

- d. Assign IP addresses to the virtual ports by entering the following commands on GGSN 3 and GGSN 4:**

```
SETDefault !v101 -IP NETAddr = 180.35.41.11 255.255.255.252
SETDefault !v1101 -IP NETAddr = 180.35.41.143 255.255.255.252
SETDefault !v301 -IP NETAddr = 180.35.106.11 255.255.255.252
SETDefault !v1301 -IP NETAddr = 180.35.106.143 255.255.255.252
SETDefault !v102 -IP NETAddr = 180.35.41.14 255.255.255.252
SETDefault !v1102 -IP NETAddr = 180.35.41.147 255.255.255.252
SETDefault !v302 -IP NETAddr = 180.35.106.14 255.255.255.252
SETDefault !v1302 -IP NETAddr = 180.35.106.147 255.255.255.252
```

- e. Specify that fragmented IPIP packets are reassembled at the receiver end of the IPIP tunnel before forwarding to their destinations by entering the following commands on all GGSNs:**

```
SETDefault !101 -IP FwdCONTRol = ReassembleFragPkts
SETDefault !1101 -IP FwdCONTRol = ReassembleFragPkts
SETDefault !301 -IP FwdCONTRol = ReassembleFragPkts
```

```

SETDefault !1301 -IP FwdCONTRol = ReassembleFragPkts
SETDefault !102 -IP FwdCONTRol = ReassembleFragPkts
SETDefault !1102 -IP FwdCONTRol = ReassembleFragPkts
SETDefault !302 -IP FwdCONTRol = ReassembleFragPkts
SETDefault !1302 -IP FwdCONTRol = ReassembleFragPkts

```

4. For each IPIP tunnel to the border router, configure a static route to the home agent by performing the following actions:

a. On GGSN 1 and GGSN 2, enter the following commands:

```

ADD -IP ROUTe 21.1.2.254 255.255.255.255 180.35.75.12 10 BFD
ADD -IP ROUTe 21.1.2.254 255.255.255.255 180.35.75.144 20 Override BFD
ADD -IP ROUTe 21.1.2.254 255.255.255.255 180.35.112.12 30 Override BFD
ADD -IP ROUTe 21.1.2.254 255.255.255.255 180.35.112.144 40 Override BFD
ADD -IP ROUTe 22.1.2.254 255.255.255.255 180.35.75.15 10 BFD
ADD -IP ROUTe 22.1.2.254 255.255.255.255 180.35.75.148 20 Override BFD
ADD -IP ROUTe 22.1.2.254 255.255.255.255 180.35.112.15 30 Override BFD
ADD -IP ROUTe 22.1.2.254 255.255.255.255 180.35.112.148 40 Override BFD

```

b. On GGSN 3 and GGSN 4, enter the following commands:

```

ADD -IP ROUTe 21.1.2.254 255.255.255.255 180.35.41.12 10 BFD
ADD -IP ROUTe 21.1.2.254 255.255.255.255 180.35.41.144 20 Override BFD
ADD -IP ROUTe 21.1.2.254 255.255.255.255 180.35.106.12 30 Override BFD
ADD -IP ROUTe 21.1.2.254 255.255.255.255 180.35.106.144 40 Override BFD
ADD -IP ROUTe 22.1.2.254 255.255.255.255 180.35.41.15 10 BFD
ADD -IP ROUTe 22.1.2.254 255.255.255.255 180.35.41.148 20 Override BFD
ADD -IP ROUTe 22.1.2.254 255.255.255.255 180.35.106.15 30 Override BFD
ADD -IP ROUTe 22.1.2.254 255.255.255.255 180.35.106.148 40 Override BFD

```

5. Specify the APN name, bind it to each of the IPIP tunnels, and assign each tunnel a session priority by performing the following actions:

a. On GGSN 1 and GGSN 2, enter the following commands:

```

SETDefault !v101 -GGSN AccessPointName = "CEN1APN1" 180.35.75.12 4
SETDefault !v1101 -GGSN AccessPointName = "CEN1APN1" 180.35.75.144 3
SETDefault !v301 -GGSN AccessPointName = "CEN1APN1" 180.35.112.12 2
SETDefault !v1301 -GGSN AccessPointName = "CEN1APN1" 180.35.112.144 1
SETDefault !v102 -GGSN AccessPointName = "CEN2APN1" 180.35.75.15 4
SETDefault !v1102 -GGSN AccessPointName = "CEN2APN1" 180.35.75.148 3
SETDefault !v302 -GGSN AccessPointName = "CEN2APN1" 180.35.112.15 2
SETDefault !v1302 -GGSN AccessPointName = "CEN2APN1" 180.35.112.148 1

```

b. On GGSN 3 and GGSN 4, enter the following commands:

```

SETDefault !v101 -GGSN AccessPointName = "CEN1APN1" 180.35.41.12 4
SETDefault !v1101 -GGSN AccessPointName = "CEN1APN1" 180.35.41.144 3
SETDefault !v301 -GGSN AccessPointName = "CEN1APN1" 180.35.106.12 2
SETDefault !v1301 -GGSN AccessPointName = "CEN1APN1" 180.35.106.144 1

```

```
SETDefault !v102 -GGSN AccessPointName = "CEN2APN1" 180.35.41.15 4
SETDefault !v1102 -GGSN AccessPointName = "CEN2APN1" 180.35.41.148 3
SETDefault !v302 -GGSN AccessPointName = "CEN2APN1" 180.35.106.15 2
SETDefault !v1302 -GGSN AccessPointName = "CEN2APN1" 180.35.106.148 1
```

6. Configure mobility agent (foreign agent) functionality on the GGSN by performing the following actions:

a. Enable foreign agent functionality by entering on all GGSNs:

```
SETDefault -MOBILEIP CONTROL = ForeignAgent
```

b. Configure the foreign agent Care Of Address (COA) by entering on GGSN 1 and GGSN 2:

```
ADD -MOBILEIP ForeignAgent 10.1.215.35
```

c. Configure the foreign agent Care Of Address (COA) by entering on GGSN 3 and GGSN 4:

```
ADD -MOBILEIP ForeignAgent 10.2.215.35
```

d. Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol uses to replicate foreign agent visitor tables by entering on GGSN 1 and GGSN 2:

```
SETDefault -MOBILEIP MblReplIP = 10.1.215.35
```

e. Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol uses to replicate foreign agent visitor tables by entering on GGSN 3 and GGSN 4:

```
SETDefault -MOBILEIP MblReplIP = 10.2.215.35
```

7. For each CEN, enable the APN and configure the home agent IP address associated with this APN by entering the following commands on all GGSNs:

```
SETDefault !CEN1APN1 -GGSN AccessPointCtl = Enable
SETDefault !CEN1APN1 -GGSN HomeAgentAddr = 21.1.2.254
SETDefault !CEN2APN1 -GGSN AccessPointCtl = Enable
SETDefault !CEN2APN1 -GGSN HomeAgentAddr = 22.1.2.254
```

8. Configure a multihop VRRP-controlled BFD session for the link between the foreign agent and the home agent by performing the following actions:

a. Configure a multi-hop BFD session on the Ethernet port that connects to the data subnet in each zone by entering on all GGSNs:

```
ADD !1 -IP BFDSESSion 1 21.1.2.254 32 3000 4
ADD !1 -IP BFDSESSion 2 22.1.2.254 32 3000 4
```

b. Register the MOBILEIP service with the BFD sessions you created by entering on all GGSNs:

```
ADD -MOBILEIP MblBfdEvent 1
ADD -MOBILEIP MblBfdEvent 2
```

c. Configure a VRRP trigger to stop and start the BFD session (enabling the MOBILEIP service to switch to the backup GGSN if the active GGSN fails) by entering on all GGSNs:

```
SETDefault !1 -IP BFDTRigger = VRRPTRigger 63
```

There is one multihop BFD session per CEN.

9. Configure subscriber IP address (home address) allocation via the remote DHCP server configured on the border routers by performing the following actions:

a. Specify the DHCP server to be the IP assigner by entering the following commands on all GGSNs:

```
SETDefault !CEN1APN1 -GGSN IPAssigner = RemoteDhcpServer
```

```
SETDefault !CEN2APN1 -GGSN IPAssigner = RemoteDhcpServer
```

- b.** Specify the IP address and address pool from which subscriber home addresses are allocated on the primary DHCP server by entering the following commands on all GGSNs:

```
SETDefault !CEN1APN1 -GGSN DHCPPrimSrvr = 21.1.2.254 21.21.2.1 255.255.0.0
```

```
SETDefault !CEN2APN1 -GGSN DHCPPrimSrvr = 22.1.2.254 22.21.2.1 255.255.0.0
```

- c.** Specify the IP address and address pool from which subscriber home addresses is allocated on the secondary (backup) DHCP server by entering the following commands on all GGSNs:

```
SETDefault !CEN1APN1 -GGSN DHCPSecSrvr = 21.1.2.254 21.22.2.1 255.255.0.0
```

```
SETDefault !CEN2APN1 -GGSN DHCPSecSrvr = 22.1.2.254 22.22.2.1 255.255.0.0
```

- d.** Specify the time interval, in seconds, between DHCP message retransmissions by entering the following commands on all GGSNs:

```
SETDefault !CEN1APN1 -GGSN DHCPReXmitTime = 2
```

```
SETDefault !CEN2APN1 -GGSN DHCPReXmitTime = 2
```

- e.** Specify the maximum number of times the GGSN retransmits mobile subscriber IP address requests by entering the following commands on all GGSNs:

```
SETDefault !CEN1APN1 -GGSN DHCPReXmitCnt = 3
```

```
SETDefault !CEN2APN1 -GGSN DHCPReXmitCnt = 3
```

- f.** Set the DHCP switchover count to 0 by entering the following commands on all GGSNs:

```
SETDefault !CEN1APN1 -GGSN DHCPSwitchCnt = 0
```

```
SETDefault !CEN2APN1 -GGSN DHCPSwitchCnt = 0
```

10. Optional: Configure GGSN charging. See [Configuring GGSN Charging on page 659](#).

11. Optional: Configure overload protection. See the [Configuring EOS Devices for Overload Protection on page 752](#).

12. Enable GGSN functionality by entering on all GGSNs:

```
SETDefault -GGSN CONTro1 = Enable
```

35.15.2

Configuring Border Routers in the Sample Conventional Data Redundancy (HA and DSR) Topology

Procedure:

- 1.** Configure virtual ports to access the CEN VLANs by performing the following actions:
- a.** Create virtual ports to access the CEN applications and home agent VLANs by entering the following commands on both border routers (BRs):

```
ADD !v1 -Port VirtualPort 1 Vlan 134
```

```
ADD !v15 -Port VirtualPort 1 Vlan 135
```

- b.** Assign IP addresses to the virtual ports by entering the following commands on BR 1:

```
SETDefault !v1 -IP NETAddr = 21.1.1.1 255.255.255.0
```

```
SETDefault !v15 -IP NETAddr = 21.1.2.1 255.255.255.0
```

- c.** Assign IP addresses to the virtual ports by entering the following commands on BR 2:

```
SETDefault !v1 -IP NETAddr = 21.1.1.2 255.255.255.0
```

```
SETDefault !v15 -IP NETAddr = 21.1.2.2 255.255.255.0
```

- d. Assign IP addresses to the virtual ports by entering the following commands on BR 3:**

```
SETDefault !v1 -IP NETAddr = 22.1.1.1 255.255.255.0
```

```
SETDefault !v15 -IP NETAddr = 22.1.2.1 255.255.255.0
```

- e. Assign IP addresses to the virtual ports by entering the following commands on BR 4:**

```
SETDefault !v1 -IP NETAddr = 22.1.1.2 255.255.255.0
```

```
SETDefault !v15 -IP NETAddr = 22.1.2.2 255.255.255.0
```

- 2. Configure a system IP (SysIP) address by performing the following actions:**

- a. Specify a SysIP address on the border router by entering on BR 1:**

```
SETDefault -IP SystemIP = 21.1.1.1
```

- b. Specify a SysIP address on the border router by entering on BR 2:**

```
SETDefault -IP SystemIP = 21.1.1.2
```

- c. Specify a SysIP address on the border router by entering on BR 3:**

```
SETDefault -IP SystemIP = 22.1.1.1
```

- d. Specify a SysIP address on the border router by entering on BR 4:**

```
SETDefault -IP SystemIP = 22.1.1.2
```

- e. Enable the loopback system IP feature for the NTP, FTP, PING, SNMP, Telnet, TFTP, and Syslog services so that the SysIP is used by sessions initiated by the border router by entering on all border routers:**

```
SETDefault -SYS UseSystemIP = (NTP, FTP, PING, SNMP, Telnet, TFTP, SYSLOG)
```

- 3. Configure the virtual ports over which the IPIP tunnels connecting the border router to the GGSNs run by performing the following actions:**

- a. Create the virtual ports by entering the following commands on BR 1:**

```
ADD !v11 -Port VirtualPort IPIP 180.35.1.12 192.158.21.176
```

```
ADD !v13 -Port VirtualPort IPIP 180.35.2.12 192.158.22.60
```

```
ADD !v21 -Port VirtualPort IPIP 180.35.2.12 192.158.22.176
```

```
ADD !v23 -Port VirtualPort IPIP 180.35.1.12 192.168.21.30
```

- b. Create the virtual ports by entering the following commands on BR 2:**

```
ADD !v12 -Port VirtualPort IPIP 180.35.1.144 192.158.221.176
```

```
ADD !v14 -Port VirtualPort IPIP 180.35.2.144 192.158.222.60
```

```
ADD !v22 -Port VirtualPort IPIP 180.35.2.144 192.158.222.176
```

```
ADD !v24 -Port VirtualPort IPIP 180.35.1.144 192.158.221.60
```

- c. Create the virtual ports by entering the following commands on BR 3:**

```
ADD !v11 -Port VirtualPort IPIP 180.35.1.15 192.158.41.176
```

```
ADD !v13 -Port VirtualPort IPIP 180.35.2.15 192.158.42.60
```

```
ADD !v21 -Port VirtualPort IPIP 180.35.2.15 192.158.42.176
```

```
ADD !v23 -Port VirtualPort IPIP 180.35.1.15 192.158.41.60
```

- d. Create the virtual ports by entering the following commands on BR 4:**

```
ADD !v12 -Port VirtualPort IPIP 180.35.1.148 192.158.241.176
```

```
ADD !v14 -Port VirtualPort IPIP 180.35.2.148 192.158.242.60
ADD !v22 -Port VirtualPort IPIP 180.35.2.148 192.158.242.176
ADD !v24 -Port VirtualPort IPIP 180.35.1.148 192.158.241.60
```

- e. Enable TOS priority queuing for the virtual ports by entering the following commands on BR 1 and BR 3:**

```
SETDefault !v11 -IPQOS CONTrol = TosQueue
SETDefault !v13 -IPQOS CONTrol = TosQueue
SETDefault !v21 -IPQOS CONTrol = TosQueue
SETDefault !v23 -IPQOS CONTrol = TosQueue
```

- f. Enable TOS priority queuing for the virtual ports by entering the following commands on BR 2 and BR 4:**

```
SETDefault !v12 -IPQOS CONTrol = TosQueue
SETDefault !v14 -IPQOS CONTrol = TosQueue
SETDefault !v22 -IPQOS CONTrol = TosQueue
SETDefault !v24 -IPQOS CONTrol = TosQueue
```

- g. Assign IP addresses to the virtual ports by entering the following commands on BR 1:**

```
SETDefault !v11 -IP NETaddr = 180.35.75.12 255.255.255.252 Ones 1368
SETDefault !v13 -IP NETaddr = 180.35.112.12 255.255.255.252 Ones 1368
SETDefault !v21 -IP NETaddr = 180.35.41.12 255.255.255.252 Ones 1368
SETDefault !v23 -IP NETaddr = 180.35.106.12 255.255.255.252 Ones 1368
```

- h. Assign IP addresses to the virtual ports by entering the following commands on BR 2:**

```
SETDefault !v12 -IP NETaddr = 180.35.75.144 255.255.255.252 Ones 1368
SETDefault !v14 -IP NETaddr = 180.35.112.144 255.255.255.252 Ones 1368
SETDefault !v22 -IP NETaddr = 180.35.41.144 255.255.255.252 Ones 1368
SETDefault !v24 -IP NETaddr = 180.35.106.144 255.255.255.252 Ones 1368
```

- i. Assign IP addresses to the virtual ports by entering the following commands on BR 3:**

```
SETDefault !v11 -IP NETaddr = 180.35.75.15 255.255.255.252 Ones 1368
SETDefault !v13 -IP NETaddr = 180.35.112.15 255.255.255.252 Ones 1368
SETDefault !v21 -IP NETaddr = 180.35.41.15 255.255.255.252 Ones 1368
SETDefault !v23 -IP NETaddr = 180.35.106.15 255.255.255.252 Ones 1368
```

- j. Assign IP addresses to the virtual ports by entering the following commands on BR 4:**

```
SETDefault !v12 -IP NETaddr = 180.35.75.148 255.255.255.252 Ones 1368
SETDefault !v14 -IP NETaddr = 180.35.112.148 255.255.255.252 Ones 1368
SETDefault !v22 -IP NETaddr = 180.35.41.148 255.255.255.252 Ones 1368
SETDefault !v24 -IP NETaddr = 180.35.106.148 255.255.255.252 Ones 1368
```

- 4. Configure BFD-controlled static routes to the foreign agents on the GGSNs in both zones (two routes per zone) with path metrics to prioritize the two routes in each zone by performing the following actions:**

- a. On BR 1, enter the following commands:**

```
ADD -IP ROUTe 10.1.215.35 255.255.255.255 180.35.75.11 10 BFD
ADD -IP ROUTe 10.1.215.35 255.255.255.255 180.35.112.11 30 Override BFD
```



```
ADD -IP ROute 10.2.215.35 255.255.255.255 180.35.41.11 10 BFD
ADD -IP ROute 10.2.215.35 255.255.255.255 180.35.106.11 30 Override BFD
```

b. On BR 2, enter the following commands:

```
ADD -IP ROute 10.1.215.35 255.255.255.255 180.35.75.143 20 Override BFD
ADD -IP ROute 10.1.215.35 255.255.255.255 180.35.112.143 40 Override BFD
ADD -IP ROute 10.2.215.35 255.255.255.255 180.35.41.143 20 Override BFD
ADD -IP ROute 10.2.215.35 255.255.255.255 180.35.106.143 40 Override BFD
```

c. On BR 3, enter the following commands:

```
ADD -IP ROute 10.1.215.35 255.255.255.255 180.35.75.14 10 BFD
ADD -IP ROute 10.1.215.35 255.255.255.255 180.35.112.14 30 Override BFD
ADD -IP ROute 10.2.215.35 255.255.255.255 180.35.41.14 10 BFD
ADD -IP ROute 10.2.215.35 255.255.255.255 180.35.106.14 30 Override BFD
```

d. On BR 4, enter the following commands:

```
ADD -IP ROute 10.1.215.35 255.255.255.255 180.35.75.147 20 Override BFD
ADD -IP ROute 10.1.215.35 255.255.255.255 180.35.112.147 40 Override BFD
ADD -IP ROute 10.2.215.35 255.255.255.255 180.35.41.147 20 Override BFD
ADD -IP ROute 10.2.215.35 255.255.255.255 180.35.106.147 40 Override BFD
```

5. Configure BFD-controlled static routes from the gateway IP addresses of the DHCP servers (from which the IP subscriber home addresses are allocated) to the GGSNs, with path metrics to prioritize the routes by performing the following actions:

a. On BR 1, enter the following commands:

```
ADD -IP ROute 21.21.1.1 255.255.255.255 180.35.75.11 10 BFD
ADD -IP ROute 21.21.1.1 255.255.255.255 180.35.112.11 30 Override BFD
ADD -IP ROute 21.21.1.2 255.255.255.255 180.35.41.11 10 BFD
ADD -IP ROute 21.21.1.2 255.255.255.255 180.35.106.11 30 Override BFD
ADD -IP ROute 21.22.1.1 255.255.255.255 180.35.75.11 10 BFD
ADD -IP ROute 21.22.1.1 255.255.255.255 180.35.112.11 30 Override BFD
ADD -IP ROute 21.22.1.2 255.255.255.255 180.35.41.11 10 BFD
ADD -IP ROute 21.22.1.2 255.255.255.255 180.35.106.11 30 Override BFD
```

b. On BR 2, enter the following commands:

```
ADD -IP ROute 21.21.1.1 255.255.255.255 180.35.75.143 20 Override BFD
ADD -IP ROute 21.21.1.1 255.255.255.255 180.35.112.143 40 Override BFD
ADD -IP ROute 21.21.1.2 255.255.255.255 180.35.41.143 20 Override BFD
ADD -IP ROute 21.21.1.2 255.255.255.255 180.35.106.143 40 Override BFD
ADD -IP ROute 21.22.1.1 255.255.255.255 180.35.75.143 20 Override BFD
ADD -IP ROute 21.22.1.1 255.255.255.255 180.35.112.143 40 Override BFD
ADD -IP ROute 21.22.1.2 255.255.255.255 180.35.41.143 20 Override BFD
ADD -IP ROute 21.22.1.2 255.255.255.255 180.35.106.143 40 Override BFD
```

c. On BR 3, enter the following commands:

```
ADD -IP ROute 22.21.1.1 255.255.255.255 180.35.75.14 10 BFD
ADD -IP ROute 22.21.1.1 255.255.255.255 180.35.112.14 30 Override BFD
ADD -IP ROute 22.21.1.2 255.255.255.255 180.35.41.14 10 BFD
```

```
ADD -IP ROUTe 22.21.1.2 255.255.255.255 180.35.106.14 30 Override BFD
ADD -IP ROUTe 22.22.1.1 255.255.255.255 180.35.75.14 10 BFD
ADD -IP ROUTe 22.22.1.1 255.255.255.255 180.35.112.14 30 Override BFD
ADD -IP ROUTe 22.22.1.2 255.255.255.255 180.35.41.14 10 BFD
ADD -IP ROUTe 22.22.1.2 255.255.255.255 180.35.106.14 30 Override BFD
```

d. On BR 4, enter the following commands:

```
ADD -IP ROUTe 22.21.1.1 255.255.255.255 180.35.75.147 20 Override BFD
ADD -IP ROUTe 22.21.1.1 255.255.255.255 180.35.112.147 40 Override BFD
ADD -IP ROUTe 22.21.1.2 255.255.255.255 180.35.41.147 20 Override BFD
ADD -IP ROUTe 22.21.1.2 255.255.255.255 180.35.106.147 40 Override BFD
ADD -IP ROUTe 22.22.1.1 255.255.255.255 180.35.75.147 20 Override BFD
ADD -IP ROUTe 22.22.1.1 255.255.255.255 180.35.112.147 40 Override BFD
ADD -IP ROUTe 22.22.1.2 255.255.255.255 180.35.41.147 20 Override BFD
ADD -IP ROUTe 22.22.1.2 255.255.255.255 180.35.106.147 40 Override BFD
```

6. Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol uses to replicate the TCP connections on the border router by performing the following actions:

a. On BR 1 and BR 2, enter:

```
SETDefault -NAT NatReplIP = 21.1.1.254
```

b. On BR 3 and BR 4, enter:

```
SETDefault -NAT NatReplIP = 22.1.1.254
```

7. Configure OSPF service parameters to support inter-system data functionality by performing the following actions:

a. Configure the border router to always add a Type 7 LSA route, independent of the P-bit value in the Type 7 LSA by entering on all BRs:

```
SETDefault -OSPF NSSACont = IgnorePbitAbr
```

b. Specify that the networks on the virtual ports over which the IPIP tunnels connecting the border router to the GGSNs run are advertised into the OSPF domain as Type 1 external OSPF routes, and that these routes have a corresponding Type 7 LSA only (no Type 5 LSA is generated) by entering the following commands on BR 1 and BR 3:

```
ADD !v11 -OSPF DirectPolicy = (Advertise, Type1) DontAdvertiseLsaType5
ADD !v13 -OSPF DirectPolicy = (Advertise, Type1) DontAdvertiseLsaType5
ADD !v21 -OSPF DirectPolicy = (Advertise, Type1) DontAdvertiseLsaType5
ADD !v23 -OSPF DirectPolicy = (Advertise, Type1) DontAdvertiseLsaType5
```

c. Specify that the networks on the virtual ports over which the IPIP tunnels connecting the border router to the GGSNs run are advertised into the OSPF domain as Type 1 external OSPF routes, and that these routes have a corresponding Type 7 LSA only (no Type 5 LSA is generated) by entering the following commands on BR 2 and BR 4:

```
ADD !v12 -OSPF DirectPolicy = (Advertise, Type1) DontAdvertiseLsaType5
ADD !v14 -OSPF DirectPolicy = (Advertise, Type1) DontAdvertiseLsaType5
ADD !v22 -OSPF DirectPolicy = (Advertise, Type1) DontAdvertiseLsaType5
ADD !v24 -OSPF DirectPolicy = (Advertise, Type1) DontAdvertiseLsaType5
```

d. Add the foreign agent Care Of Address (COA) for each zone to the list of statically-configured routes, specifying that the metric in the routing table is used as the advertised cost of this route,

a link state metric (Type1) is used, and the route has a corresponding Type 7 LSA only (no Type 5 LSA is generated) by entering the following commands on all BRs:

```
ADD -OSPF StaticPolicy 10.1.215.35 0 Type1 DontAdvertiseLsaType5
```

```
ADD -OSPF StaticPolicy 10.2.215.35 0 Type1 DontAdvertiseLsaType5
```

8. Configure mobility agent (home agent) functionality by performing the following actions:

- a.** Configure an address list associated with the home agent subnet by entering on BR 1 and BR 2:

```
ADD -IP AddressList CEN1APN1 21.1.2.0/24
```

- b.** Configure an address list associated with the home agent subnet by entering on BR 3 and BR 4:

```
ADD -IP AddressList CEN2APN1 22.1.2.0/24
```

- c.** Enable home agent functionality on the border router by entering on all BRs:

```
SETDefault -MOBILEIP CONTROL = HomeAgent
```

- d.** Configure the home agent IP address on the border router and associate it with the address list you configured in [step 8a](#) by entering on BR 1 and BR 2:

```
SETDefault -MOBILEIP HomeAgent = 21.1.2.254 CEN1APN1
```

- e.** Configure the home agent IP address on the border router and associate it with the address list you configured in [step 8a](#) by entering on BR 3 and BR 4:

```
SETDefault -MOBILEIP HomeAgent = 22.1.2.254 CEN2APN1
```

- f.** Specify the foreign agent COAs (one COA per zone, shared by the two GGSNs in the zone) by entering the following commands on all BRs:

```
ADD -MOBILEIP ForeignAgent 10.1.215.35
```

```
ADD -MOBILEIP ForeignAgent 10.2.215.35
```

- g.** Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol uses to replicate the home agent mobility database by entering on BR 1 and BR 2:

```
SETDefault -MOBILEIP MblReplIP = 21.1.2.254
```

- h.** Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol uses to replicate the home agent mobility database by entering on BR 3 and BR 4:

```
SETDefault -MOBILEIP MblReplIP = 22.1.2.254
```

9. To configure VRRP control on the CEN applications and home agent VLANs, configure the border router as a backup VRRP router on the CEN applications VLAN and on the home agent VLAN by performing the following actions:

- a.** Configure backup VRRP router functionality on the virtual port that access the CEN application and home agent VLANs by entering the following commands on BR 1 and BR 2:

```
ADD !v1 -VRRP BackUp AsOwner 1 21.1.1.254
```

```
ADD !v15 -VRRP BackUp AsOwner 15 21.1.2.254
```

- b.** Configure backup VRRP router functionality on the virtual port that access the CEN application and home agent VLANs by entering the following commands on BR 3 and BR 4:

```
ADD !v1 -VRRP BackUp AsOwner 1 22.1.1.254
```

```
ADD !v15 -VRRP BackUp AsOwner 15 22.1.2.254
```

- c.** Set the advertisement interval for VRRP packets by entering the following commands on all border routers:

```
SETDefault !v1 -VRRP AdvertiseInt = 300 1
```

```
SETDefault !v15 -VRRP AdvertiseInt = 300 15
```

- d. Set the VRRP hold time (the time interval before the backup VRRP declares the master VRRP down) by entering the following commands on all border routers:

```
SETDefault !v1 -VRRP HoldTime = 900 1
SETDefault !v15 -VRRP HoldTime = 900 15
```

- e. Set the VRRP priority by entering the following commands on all border routers:

```
SETDefault !v1 -VRRP PRIOrity = 254 1
SETDefault !v15 -VRRP PRIOrity = 254 15
```

10. Configure multihop BFD sessions for the links between the home agent and the foreign agent in each zone by performing the following actions:

- a. Create two multihop BFD sessions (one session per zone) on the virtual port that accesses the home agent VLAN by entering the following commands on all border routers:

```
ADD !v15 -IP BFDSEssion 1 10.1.215.35 32 3000 4
ADD !v15 -IP BFDSEssion 2 10.2.215.37 32 3000 4
```

- b. Register the MOBILEIP service with the BFD sessions you created by entering on all border routers:

```
ADD -MOBILEIP MblBfdEvent 1
ADD -MOBILEIP MblBfdEvent 2
```

- c. Configure a VRRP trigger to stop and start the BFD session by entering on all border routers:

```
SETDefault !v15 -IP BFDTRigger = VRRPTRigger 15
```

11. Configure the border router to execute event-based macros to enable and disable VRRP control on the virtual ports associated with the CEN applications and home agent VLANs by performing the following actions:

- a. Define the lanUp and lanDown macros that are executed when a PortUp or PortDown event occurs by entering on all border routers:

```
define lanUp=
SETDefault !v1 -VRRP CONTrol = Enable 1
SETDefault !v15 -VRRP CONTrol = Enable 15
)
)
define lanDown=(
SETDefault !v1 -VRRP CONTrol = Disable 1
SETDefault !v15 -VRRP CONTrol = Disable 15
)
)
```

- b. Enable event-based macro execution and specify that event-based macro execution are enabled and the PortDown actions performed when the primary connections fail to establish within 5 minutes after the device boots and that the macro name and a complete/incomplete message are reported in the system log buffer by entering on all border routers:

```
SETDefault -SCH EbmeCONTrol = (Enabled, RunOnBootFail, Log)
```

- c. Set up the scheduler for event-based macro execution by entering the following commands on all border routers:

```
ADD !1 -SCH EbmeEvent PortUp 500 do port_up
ADD !1 -SCH EbmeEvent PortDown 500 do port_down
```

12. Configure DHCP server functionality by performing the following actions:

- a. On the virtual port that accesses the home agent VLAN and on the virtual ports over which the IPIP tunnels connecting the border router to the GGSN run: enable DHCP; configure DHCP to use the address pool defined with the `AddressPool` parameter; specify that the DHCP server does not issue an ICMP ECHO request to make sure the address chosen from the pool is not duplicated on the network before allocating the address to a client; and disable subnet mask checking by entering the following commands on BR 1 and BR 3:

```
SETDefault !v15 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
SETDefault !v11 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
SETDefault !v13 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
SETDefault !v21 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
SETDefault !v23 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
```

- b. On the virtual port that accesses the home agent VLAN and on the virtual ports over which the IPIP tunnels connecting the border router to the GGSN run: enable DHCP; configure DHCP to use the address pool defined with the `AddressPool` parameter; specify that the DHCP server does not issue an ICMP ECHO request to make sure the address chosen from the pool is not duplicated on the network before allocating the address to a client; and disable subnet mask checking by entering the following commands on BR 2 and BR 4:

```
SETDefault !v15 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
SETDefault !v12 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
SETDefault !v14 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
SETDefault !v22 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
SETDefault !v24 -DHCP CONTROL = (Enable,AddressPool,NoIcmpCheck,NoMaskCheck)
```

- c. Configure the address pool from which subscriber home addresses are allocated on the primary DHCP server by entering on BR 1 and BR 2:

```
ADD -DHCP AddressPool 21.21.0.0 - 21.21.255.255 !P1 255.255.0.0
```

- d. Configure the address pool from which subscriber home addresses are allocated on the primary DHCP server by entering on BR 3 and BR 4:

```
ADD -DHCP AddressPool 22.21.0.0 - 22.21.255.255 !P1 255.255.0.0
```

- e. Configure the gateway IP addresses (giaddr addresses) to the primary DHCP server subnet as reserved addresses which cannot be used for allocation by entering on BR 1 and BR 2:

```
ADD -DHCP ReserveAddress 21.21.1.1
ADD -DHCP ReserveAddress 21.21.1.2
```

- f. Configure the gateway IP addresses (giaddr addresses) to the primary DHCP server subnet as reserved addresses which cannot be used for allocation by entering on BR 3 and BR 4:

```
ADD -DHCP ReserveAddress 22.21.1.1
ADD -DHCP ReserveAddress 22.21.1.2
```

- g. Create a second DHCP profile by entering on all border routers:

```
ADD -DHCP PRofile P2
```

- h. Configure the address pool from which subscriber home addresses are allocated on the secondary DHCP server by entering on BR 1 and BR 2:

```
ADD -DHCP AddressPool 21.22.0.0 - 1.22.255.255 !P2 255.255.0.0
```

- i. Configure the address pool from which subscriber home addresses are allocated on the secondary DHCP server by entering on BR 3 and BR 4:

```
ADD -DHCP AddressPool 22.22.0.0 - 2.22.255.255 !P2 255.255.0.0
```

- j. Configure the gateway IP addresses (giaddr addresses) to the secondary DHCP server subnet as reserved addresses which cannot be used for allocation by entering on BR 1 and BR 2:

```
ADD -DHCP ReserveAddress 21.22.1.1
```

```
ADD -DHCP ReserveAddress 21.22.1.2
```

- k. Configure the gateway IP addresses (giaddr addresses) to the secondary DHCP server subnet as reserved addresses which cannot be used for allocation by entering on BR 3 and BR 4:

```
ADD -DHCP ReserveAddress 22.22.1.1
```

```
ADD -DHCP ReserveAddress 22.22.1.2
```

- l. Configure the IP address of the virtual port that connects to the home agent VLAN as a reserved addresses which cannot be used for allocation by entering on BR 1:

```
ADD -DHCP ReserveAddress 21.1.2.1
```

- m. Configure the IP address of the virtual port that connects to the home agent VLAN as a reserved addresses which cannot be used for allocation by entering on BR 2:

```
ADD -DHCP ReserveAddress 21.1.2.2
```

- n. Configure the IP address of the virtual port that connects to the home agent VLAN as a reserved addresses which cannot be used for allocation by entering on BR 3:

```
ADD -DHCP ReserveAddress 22.1.2.1
```

- o. Configure the IP address of the virtual port that connects to the home agent VLAN as a reserved addresses which cannot be used for allocation by entering on BR 4:

```
ADD -DHCP ReserveAddress 22.1.2.2
```

- p. Configure the reserved IP address which cannot be used for allocation by entering on BR 1 and BR 2:

```
ADD -DHCP ReserveAddress 21.1.2.254
```

- q. Configure the reserved IP address which cannot be used for allocation by entering on BR 3 and BR 4:

```
ADD -DHCP ReserveAddress 22.1.2.254
```

- r. Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol uses to replicate DHCP allocations by entering on BR 1 and BR 2:

```
SETDefault -DHCP DhcpReplIP = 21.1.2.254
```

- s. Specify the VRRP-controlled IP address shared by the replication peer devices (floating VRRP address) that the replication protocol uses to replicate DHCP allocations by entering on BR 3 and BR 4:

```
SETDefault -DHCP DhcpReplIP = 22.1.2.254
```

- t. Register the DHCP service on the home agent to receive BFD up/down notifications for the links between the home agent and the foreign agent in each zone by entering on all border routers:

```
ADD -DHCP DhcpBfdEvent 1
```

```
ADD -DHCP DhcpBfdEvent 2
```

35.16

Configuring GGSN Charging

The GGSN charging feature allows the GGSN to send billing information about GTP usage to a charging gateway. The GGSN supports one primary and one secondary charging gateway per system. Follow this procedure to configure a GGSN for GGSN charging.

You should add the commands in this procedure to the GGSN's `XGSN.cfg` file so that if the GGSN reboots, its GGSN configuration is maintained.

Procedure:

1. Enable charging functionality (the GTP' protocol) on the GGSN by entering:

```
SETDefault -GGSN CHGCONTRol = Enable
```

You can also enable or disable the GTP' protocol for a particular APN by using the `CHGAPControl` parameter.

2. Identify the primary charging gateway by entering:

```
SETDefault -GGSN CHGPrimGateway = <IP_address>
```

where `<IP_address>` is the IP address of the primary charging gateway.

3. Optional: Identify the secondary charging gateway by entering:

```
SETDefault -GGSN CHGSecGateway = <IP_address>
```

where `<IP_address>` is the IP address of the secondary charging gateway.



NOTE: If you do not specify a secondary charging gateway, the GGSN assumes that there is only a single charging gateway in the system.

4. Optional: Specify the charging source IP address by entering:

```
SETDefault -GGSN CHGSourceAddr = <IP_address>
```

where `<IP_address>` is the source IP address used by the GGSN for communication with the charging gateway. This is the source IP address that is used for packet data billing purposes.



NOTE: If you do not specify a charging source IP address, the GGSN uses the system IP address configured with the IP service `SystemIP` parameter by default.

5. Optional: Specify a name (up to 20 characters) for the node that generates the G-CDR by entering:

```
SETDefault -GGSN CHGNodeId = <Node_Id>
```

where `<Node_Id>` is the name string of the node.

6. Optional: Enable generation of a local record sequence number in each G-CDR by entering:

```
SETDefault -GGSN CHGLocalRecSeq = Enable
```

The local record sequence number is unique to this GGSN and can be used to identify missing records in systems with multiple GGSNs.

7. Configure the maximum number of bytes reported in each G-CDR (1-2000000000) by entering:

```
SETDefault -GGSN CHGRecMaxSize = <bytes>
```

The recommended value for the `CHGRecMaxSize` parameter is the default value (1000000 bytes).

8. Configure the time interval (1-300 seconds) between echo request messages sent from the GGSN to the charging gateway by entering:

```
SETDefault -GGSN CHGEchoTime = <seconds>
```

The recommended value for the `CHGEchoTime` parameter is 10 seconds.

9. Configure the maximum number of times the GGSN retransmits charging echo request messages to the primary charging gateway without receiving a response before switching over to the secondary charging gateway by entering:

```
SETDefault -GGSN CHGEchoRetryCnt = <value>
```

The recommended value for the CHGEchoRetryCnt parameter is 5 seconds.

10. Configure the time interval, in seconds, between charging echo message retransmissions by entering:

```
SETDefault -GGSN CHGEchoRetryTm = <seconds>
```

The recommended value for the CHGEchoRetryTm parameter is 2 seconds.

11. Specify how many seconds (1-600) the GGSN waits before each retransmission of DRT request messages by entering:

```
SETDefault -GGSN CHGReXmitTime = <seconds>
```

The recommended value for the CHGReXmitTime parameter is 10 seconds.

12. Specify the number of times (1-100) the GGSN retransmits DRT request messages by entering:

```
SETDefault -GGSN CHGReXmitCnt = <value>
```

The recommended value for the CHGReXmitCnt parameter is 3 retries.

13. Configure the time interval (5-86400 seconds) between the DRT request messages sent for a certain PDP context by entering:

```
SETDefault -GGSN CHGXferInterval = <seconds>
```

The recommended value for the CHGXferInterval parameter is the default value (100 seconds).

14. Configure the size of the buffer (1-10000 messages) that holds unacknowledged DRT request messages sent to the charging gateway by entering:

```
SETDefault -GGSN CHGXferBuffer = <value>
```

The recommended value for the CHGXferBuffer parameter is 10000 messages.

15. Configure the maximum number of G-CDRs (1-255) that can be put into a single DRT request message by entering:

```
SETDefault -GGSN CHGRecMaxCnt = <value>
```

The recommended value for the CHGRecMaxCnt parameter is the default value (255).

16. Configure the times of day when charging tariffs change by entering:

```
ADD -GGSN CHGTaRiff = <hh:mm[:ss]>
```

You can add up to 12 tariff changes.



NOTE: You cannot add multiple tariff changes in a single command. You must add each tariff change with a separate ADD -GGSN CHGTaRiff command.

35.17

GGSN Configuration for Dynamic DNS

To configure a GGSN for dynamic DNS (DDNS) functionality, you must specify the IP address of the DNS server when you configure each APN on the GGSN. Optionally, you can also specify the domain name to be used for dynamic DNS updates.

The DDNS configuration is included in the `xgsn.cfg` file of the device.

By default, the DDNS domain name is the configured `AccessPointName` for the APN. If you want to change the domain name from its default setting, you must use the GGSN service `DDName` parameter to change it. For details, see the “GGSN Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

35.18

Configuring GGSN Hardware Redundancy for Dimetra Systems

You can use this procedure only for Dimetra system releases that support GGSN hardware redundancy.

Procedure:

1. From the EOS CLI, prepare the GGSNs for a redundant hardware configuration by performing the following actions:
 - a. If you do not use a network management application to configure the IP-IP tunnels, cConfigure the virtual port over which the IP-IP VPN tunnel connecting the CEN and the GGSN run as a synchronized VRRP-controlled IP-IP tunnel.

For details, see [GGSNs Preparation for Hardware Redundancy Configuration Through the EOS CLI on page 661](#).

- b. Configure both GGSNs as VRRP backup routers.



NOTE: The VRRP configuration is included in the `boot.cfg` files of GGSNs.

For details, see [GGSNs Preparation for Hardware Redundancy Configuration Through the EOS CLI on page 661](#).

2. From the network management application UI, configure the GGSN hardware redundancy feature.
For details, see your network management application's documentation.



NOTE: In order for the GGSN hardware redundancy feature to operate, you must manually configure the redundant GGSN pair from your network management application.

35.18.1

GGSNs Preparation for Hardware Redundancy Configuration Through the EOS CLI

The GGSN hardware redundancy feature provides a hot standby for the active GGSN.

To configure an EOS device for GGSN hardware redundancy, you must first do the following from the EOS CLI:

- Configure the IP-IP tunnel connecting the GGSN and the CEN as a VRRP-controlled tunnel.
- Configure the GGSNs as VRRP backup routers.

Configuring a VRRP-Controlled IP-IP Tunnel

To configure the virtual port over which the IP-IP VPN tunnel connecting the CEN and the GGSN runs as a synchronized VRRP-controlled IP-IP tunnel, you must use the following syntax:

```
ADD !v<port> -Port VirtualPort IPIP <src-ip> <dest-ip> [TTL <tll>] [VRID <vrid@port-no>, <vrid@port-no>]
```

where:

`v<port>` is the virtual port number over which the IP-IP tunnel is created.

`<src-ip>` is the source address used to transmit to the CEN. The `<src-ip>` can be the subnet address of an interface on the GGSN or an unused address on that interface.

`<dest-ip>` is the destination address of the tunnel peer. The `<dest-ip>` must be a valid address and should be unique. The IP-IP VPN tunnel peer should have the same configuration, with `<src-ip>` and `<dest-ip>` swapped.

TTL *<ttl>* specifies the number of hops the IP packets travel before dropping. The optional TTL argument is part of the ADD -Port VirtualPort command syntax but is not used in redundant GGSN configuration. VRID specifies the VRRP virtual router identifier (VRID) for a port on the GGSN. *<port-no>* specifies the port on which the VRID is configured.

Configuring GGSNs as Backup VRRP Routers

To configure the GGSNs that makes up the redundant GGSN pair as VRRP backup routers, you must use the following syntax:

```
ADD !<portlist> -VRRP BackUp [AsOwner] <vrid> <ip_address> [, <ip_address>] [<SysName>]
```

where:

<AsOwner> is a Motorola Solutions extension to VRRP that specifies that when this device takes over the IP address(es) it behaves as though it owns the IP address(es): it responds to PINGs, allow TCP connections, and so on.

<vrid> is the number you assigned to the GGSN port using the VRid parameter. The value is between 1 and 255. For details about the VRid parameter, see “VRRP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

<ip_address> is the IP address configured on the owner router. If the owner router fails, the backup router takes over the IP address. Each backup router must have a real IP address already configured in the IP Service. The backup router’s IP address must have the same subnet as the owner router.

<SysName> is the SysName of the VRRP master router. Specifying the master router SysName allows a network management application to discover one peer of the VRRP session.

Viewing GGSN Hardware Redundancy Status

With the EOS CLI, you can view the status of the VRRP-controlled IP-IP tunnel and associated virtual port, as well as the status of the redundant GGSN.

To view the status of the VRRP-controlled IP-IP tunnel, you must issue the SHow -Port VirtualPort command:

- To display the status of the IP-IP tunnel, you can use:
SHow !v<port> -Port VirtualPort
- To display the status of the IP-IP tunnel and the VRIDs, you can use:
SHow !v<port> -Port VirtualPort LOnG
- To display the status of the virtual port, you can use:
SHow !v<port> -Port VirtualPort conf
- To display the status of the virtual port and the VRIDs, you can use:
SHow !v<port> -Port VirtualPort conf LOnG

To display the hardware redundancy status of the GGSN, issue the SHow -GGSN RedundantStatus command. This command returns one of the following values:

Singular – No synchronized VRRP sessions exist on the device. All APNs are in non-redundant configuration.

Active – The device is the master for the synchronized VRRP sessions, and all sessions are UP on this device.

Standby – The device is the backup for the synchronized VRRP sessions, and all sessions are DOWN on this device.

Unconfigured – No GTP APNs are configured.

Misconfigured – Some synchronized VRRP sessions are UP and some are DOWN. This is an error condition that is not likely to occur in a deployed system.

Chapter 36

Network Management Configuration

This chapter describes the management activities in which EOS devices participate. Most management activities require configuration because they are disabled by default. The system manages networks in the following ways:

- Using file service
- Building network maps (netmaps)
- Sending AuditLog messages to a Network Management Station



NOTE: This chapter does not describe the protocols involved in these management activities. For information on the protocols, see the appropriate RFCs. For information on the parameters referenced in this chapter, see the *Enterprise OS Software Version 18.2 Reference Guide*. The network management information in this chapter applies to the system regardless of its functionality unless otherwise specified.

36.1

Simple Network Management Protocol

Simple Network Management Protocol (SNMP) allows you to modify and display some of the EOS device's system parameters from a host; you do not need to attach a terminal to the system console port to change its configuration. The system implementation of SNMPv1/v2c follows the specifications in RFCs 1155, 1157, and 1213. The system parameters and the Motorola Solutions-extended parameters that are described in several RFCs including RFC 1213 can both be accessed from the host.

On an EOS device, SNMPv1/v2c support is enabled by default. If you want to disable SNMP support, you must use:

```
SETDefault -SNMP CONTrol = NoManage
```

36.1.1

Configuring the SNMP Service for SNMPv1/v2c Management

Follow this procedure to configure the SNMP Service for SNMPv1/v2c management. For information about SNMPv3 support, see [SNMPv3 Support on page 668](#).



NOTE: By default, the community name ANYCOM exists with read access to management information base (MIB) variables and allows unrestricted access to the EOS device. To ensure that access is available only to the proper system administrator, Motorola Solutions recommends that you delete the ANYCOM community name, and add the appropriate community string and the manager's IP address.

Procedure:

1. Delete the default community string "ANYCOM".

Step example:

```
DELeTe -SNMP COMMunity "ANYCOM"
```

2. Configure at least one new community string with read/write access.

Step example:

```
ADD -SNMP COMMunity "private" Triv RW AL1
```

3. Add other community strings with read-only access as required.

You can have up to 10 managers for each community. Including ANYCOM, you can have up to 6 communities.



NOTE: When you enter ANYCOM to the list of community names, you must enter it in all uppercase letters.

Step example:

```
ADD -SNMP COMMunity "public" Triv RO AL1
```

4. Configure at least one SNMP manager to the read/write community string.

Step example:

```
ADD -SNMP MANager "private" <IP_addr>
```

5. If SNMP is disabled on your system and you want to enable it, enter:

```
SETDefault -SNMP CONTrol = Manage
```

SNMP is enabled by default.

36.1.2

Request Validation

The following options are available with a request for validation:

- For security purposes, the SNMP agent on the system validates SNMP requests before responding. This prevents unauthorized users from viewing or changing the device configuration.
- You can specify that only the hosts with known community names can send requests. All the community names known to the system are specified by the `COMMunity` parameter. A request cannot be authenticated if its community name is not included in the `COMMunity` parameter. To allow requests from any community, add ANYCOM to the list of community names.
- The information in the ANYCOM entry then processes requests with unmatched community names. When adding a community to the list, you can also specify the level of access to the MIB, read or read/write, and the type of trap sent to managers associated with the community name.
- In addition to specifying community names, you can create a list of managers for each community name. If there is no manager list associated with a community name, the system responds to any request with that community name; otherwise, before an incoming request is processed, it must have a matching IP address for the community name specified by the `MANager` parameter.

36.1.3

Traps and Trap Profiles

Traps are sent by the SNMPv1/v2c agent to alert the network management station of unusual events. For a complete list of the traps supported on EOS devices, see [Traps Supported on EOS Devices on page 871](#).

You can configure different trap profiles in the device to receive different types of traps. To do this, you use the `-SNMP TrapProfile` parameter to create trap profiles that enable specific types of traps. After you create the trap profiles, you can map them to different network management stations using the `-SNMP MANager` parameter. For details about the `MANager`, and `TrapProfile` parameters, see "SNMP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: Creating trap profiles and adding SNMPv1/v2c managers with the existing SNMP service `TrapProfile` and `MANager` parameters applies to SNMPv1/v2c only. Settings that you make for these parameters have no effect on SNMPv3 functionality.

Trap managers are defined by both an IP address and (optionally) a UDP port number, and EOS can send to multiple ports per IP address. As a result, you can place multiple trap managers on a single server box. This provides flexibility if you run multi-vendor network management systems on a single server box.

For example, assume that you have two network management stations in your network:

- NMS1 (with IP address 10.79.135.20)
- NMS2 (with IP address 10.79.136.20)

If you want NMS1 to receive path and port traps, and you want NMS2 to receive PIM (multicast) traps, for the necessary configuration, see [Configuring Trap Profiles – Multiple Network Management Stations on page 667](#).

SNMPv3 and SNMPv1 Trap Generation

In SNMPv3, traps are received through the managers configured with the `ADD -SNMP V3MANager` command. You are not required to delete these managers before changing to SNMPv1/v2c operation. However, when you switch from SNMPv3 to SNMPv1/v2c operation, you must add v1 managers with the `ADD -SNMP MANager` command in order to receive SNMPv1 traps.



NOTE: EOS does not support SNMPv2c traps.

To configure trap profiles (multiple network management stations), see [Configuring Trap Profiles – Multiple Network Management Stations on page 667](#).

To configure trap profiles (single management station), see [Configuring Trap Profiles – Single Management Station on page 667](#).

Trap Throttling

EOS supports numerous traps. To avoid a situation in which these traps overwhelm the network, you may want to throttle certain traps. You can throttle a trap by using the SNMP service `TrapThrottle` parameter to set the throttle time – the minimum number of seconds in which the next trap of the specified type is sent:

```
SETDefault -SNMP TrapThrottle = <OID /MnemonicName> <ThrottleTime>
```

As indicated in the command syntax above, you can identify the trap by its mnemonic name or its OID. For example, to set the throttle time to 60 seconds for the trap with OID 1.3.6.1.4.1.43.2.48.1.1.0.2 and mnemonic name `mnrGtpCreateFail`, you must use:

```
SETDefault -SNMP TrapThrottle = 1.3.6.1.4.1.43.2.48.1.1.0.2 60
```

or

```
SETDefault -SNMP TrapThrottle = mnrGtpCreateFail 60
```

The `TrapThrottle` parameter defines a window during which only one trap with a specific trap type for a given index is sent. By increasing the throttle time, you decrease the number of traps sent on the network. If you do not set a throttle time, a registered default is used. Traps with `ThrottleTime` set to 0 are disabled for throttling. For further details about configuring trap throttling, see the description of the `TrapThrottle` parameter in “SNMP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: Not all EOS traps support throttling by using the `TrapThrottle` parameter. To display a list of the traps that do support throttling via the `TrapThrottle` parameter, you must use:

```
SHow -SNMP TrapThrottle All
```

Trap Delay Queueing and Event-Based Trap Suppression

EOS implements a trap delay queueing mechanism to reduce bursts and duplicates of certain traps, and configure the traps to be dequeued and processed at a regular interval. In addition, the affected traps are suppressed during cooperative WAN routing (CWR) switchover and device bootup, minimizing the performance issues and network congestion that may be caused by the large number of traps triggered

by these events. The trap delay queueing and event-based trap suppression feature allows you to identify more important traps more easily in the Unified Event Manager (UEM).

The trap delay queueing and event-based trap suppression feature applies to both SNMPv1 and SNMPv3 traps and affects the following traps:

- linkDown
- linkUp
- aMotorolaCaPortDown
- aMotorolaCaPortUp
- aMotorolaCaPathDown
- aMotorolaCaPathUp
- mnrdS1E1LineChangeTrap
- audlSyslogMessage
- audlCfgChange
- audlFileAccess
- audlLoginFail



NOTE: All other traps supported by EOS devices and not listed on the previous page are processed as they were prior to the implementation of delay queueing and event-based trap suppression feature. These traps are not processed via the delay queue or suppressed during CWR switchover and device bootup.

Instead of sending immediately to the SNMP managers, the traps listed on the previous page are put in a delay queue for the number of milliseconds configured by the new SNMP service `TrapQueueDelay` parameter. The default `TrapQueueDelay` is 1500 milliseconds (1.5 seconds). At the end of the `TrapQueueDelay`, the traps are dequeued and processed in the order in which they entered the queue at the interval configured with the new SNMP `TrapInterval` parameter (1-10 milliseconds). The default `TrapInterval` is 5 milliseconds.

In addition, EOS suppresses the affected traps during the periods defined by the start and end of the following events:

CWR switchover event

This event starts when the initial `mnrdCWRStatusChangeTrap` is sent and ends approximately 3 seconds later, after the device updated the status for all ports and paths based on the switchover.

Device bootup event

This event starts when hardware initialization occurs and before the `boot.cfg` file is processed. It ends after the `boot.cfg` processing finished, plus an additional delay of approximately 3 minutes to allow the device to bring up all paths, ports, and other services. Console messages similar to the following indicate the start and end of the device bootup event:

System Event - Bootup Start

System Event - Bootup End

Once the CWR switchover or device bootup event ends, the suppressed traps are processed via the delay queue.

Disabling The Trap Delay Queueing And Event-based Trap Suppression Feature

To disable the trap delay queueing and event-based trap suppression feature, set the `TrapQueueDelay` parameter to 0. When `TrapQueueDelay` is set to 0, affected traps are handled as follows:

- Any existing traps in the queue are dequeued and processed first. If a CWR switchover or device bootup event is in progress, the traps are suppressed until the end of the event.

- Any new traps are not placed in the queue or suppressed – even if a CWR switchover or device bootup event is in progress.

36.1.3.1

Configuring Trap Profiles – Multiple Network Management Stations

In this example, since no UDP port number is specified with the `MANager` parameter, the default (162) is used.

Procedure:

1. Add an SNMPv1/2c community name ("private") with all traps enabled by entering:
`ADD -SNMP COMMunity "private" all`
2. Create a trap profile (with profile name "profile1") for path and port traps by entering:
`ADD -SNMP TrapProfile "profile1" PATH PORT`
3. Create a trap profile (with profile name "profile2") for PIM traps by entering:
`ADD -SNMP TrapProfile "profile2" MULTICAST`
4. Add NMS1 to the SNMPv1/2c manager list for community name "private" and map it to the trap profile name "profile1" (so that it receives path and port traps) by entering:
`ADD -SNMP MANager "private" 10.79.135.20 "profile1"`
5. Add NMS2 to the SNMPv1/2c manager list for community name "private" and map it to the trap profile name "profile2" (so that it receives PIM (multicast) traps) by entering:
`ADD -SNMP MANager "private" 10.79.136.20 "profile2"`

36.1.3.2

Configuring Trap Profiles – Single Management Station

Assume that you have a single network management station in your network – NMS1 (with IP address 10.79.135.20) – and you want NMS1 to receive path and port traps on UDP port 50000 and PIM (multicast) traps on UDP port 60000. You can use this procedure to perform necessary configuration.

Procedure:

1. Add an SNMPv1/2c community name ("private") with all traps enabled by entering:
`ADD -SNMP COMMunity "private" all`
2. Create a trap profile (with profile name "profile1") for path and port traps by entering:
`ADD -SNMP TrapProfile "profile1" PATH PORT`
3. Create a trap profile (with profile name "profile2") for PIM traps by entering:
`ADD -SNMP TrapProfile "profile2" MULTICAST`
4. Add NMS1 to the SNMPv1/2c manager list for community name "private" and map it to the trap profile name "profile1" so that it receives path and port traps, by entering:
`ADD -SNMP MANager "private" 10.79.135.20 "profile1" 50000`
5. Add NMS2 to the SNMPv1/2c manager list for community name "private" and map it to the trap profile name "profile2" so that it receives PIM (multicast) traps, by entering:
`ADD -SNMP MANager "private" 10.79.135.20 "profile2" 60000`

36.1.4

Set Operation

You can use the SNMPv1 SetRequest protocol data unit (PDU) to set the objects on an EOS device. For one PDU, you may specify multiple objects, a single object, or a table entry field. If the request does not follow these restrictions, the SNMP agent responds with a general error PDU, and the set operation is disallowed.



NOTE: EOS does not support the SNMPv2 SET operation.

36.1.5

SNMPv2 GETBULK Support

EOS supports SNMPv2 GETBULK operations, as well as SNMPv2 GET and GETNEXT operations. SNMPv2 GETBULK support minimizes interoperability issues for EOS devices managed by InfoVista.



NOTE: EOS does not support SNMPv2 SET operations or notifications (traps).

36.1.6

Reports of Unconfigured Ports

The following changes in how SNMP reports unconfigured (channelized) ports were introduced beginning with EOS version 14.0:

- When you issue an SNMP GETNEXT/BULK command, EOS does not return information from unconfigured (channelized) ports.
- When you issue an SNMP GET command for a port for which no Port Owner was assigned, EOS returns the following values:

`ifType = other(1)`

`ifDescr = ???-<port_number>.<channel_number>/1`

For example:

`???-4F.1/1`

- When you issue an SNMP GET command for an unconfigured CWR port, EOS returns the following value:
`ifType = 0`

36.2

SNMPv3 Support

EOS supports SNMPv3, which introduces security to the SNMP model and supports two security-related capabilities:

User security model (USM)

Provides authentication and privacy (encryption) functions and operates at the message level.

View-based access control model (VACM)

Determines whether a given principal is allowed to access a particular MIB object to perform specific functions and operates at the protocol data unit (PDU) level. For details, see [SNMPv3 Procedure Reference on page 684](#).

In addition, SNMPv3 includes the concept of a **principal**, the entity on whose behalf services are provided or processing takes place. Each SNMP entity includes a single SNMP engine. An SNMP engine implements functions for sending and receiving messages, authenticating and encrypting/decrypting messages, and

controlling access to managed objects. For details about specifying the SNMP engine when managing SNMPv3 USM users, see [SNMPv3 Procedure Reference on page 684](#).

36.2.1

SNMPv3 Configuration on EOS Devices

EOS devices are configured for SNMPv3 management through dynamic registration from the configured network manager.



NOTE: If you configure devices through a `boot.cfg` configuration file, ensure that the following SNMPv1/v2c commands are present in the `boot.cfg` file:

```
ADD -SNMP COMMunity "<RO_community_name>" ro
```

```
ADD -SNMP COMMunity "<RW_community_name>" rw
```

where `<RO_community_name>` and `<RW_community_name>` are the read-only and read-write community strings configured for your system.

36.2.2

SNMPv3 Data Reset

To reset SNMPv3 data, you must issue the `ResetV3` command:

The `ResetV3` command resets the following data:

- SNMPv3 information saved in memory
- SNMPv3 information written to the persistence file
- Information retained by the SNMP Common Agent



IMPORTANT: When a `ResetV3` command is executed, SNMPv3 data is synchronized with the Unified Network Configurator (UNC). You must clear the USM cache after issuing the `ResetV3` command. For details, see “Accessing and Executing Existing Saved Commands” in the *Unified Network Configurator User Guide*.



CAUTION: The `ResetV3` command was developed primarily for use in certain controlled situations and is not intended to be performed as a part of normal system operation. Incorrect use of the `ResetV3` command can potentially result in SNMPv3 protocol authentication failure, which causes the device to become unmanageable from the Unified Network Configurator (UNC). If this occurs, you must delete the device from the UNC and rediscover it.

For details about the `ResetV3` command, see “Commands” in the *Enterprise OS Software Version 18.2 Reference Guide*.

36.2.3

Opening the SNMP V3 User Manager Menu

The starting point for SNMPv3 user management is **SNMP V3 User Manager Menu**, which you access from the `UserManage` command menu.



NOTE: SNMP control must be set to V3 (`SETDefault -SNMP CONTrol = V3`) before you can access **SNMP V3 User Manager Menu**.

Procedure:

1. Issue the `UserManage` command to open the following menu:

```
1 - User Manager Menu
2 - SNMP V3 User Manager Menu
Select (1-2) ... <CR> to Exit ====>
```

2. Select **SNMP V3 User Manager Menu** by entering: 2

The **USM User:** prompt appears.

3. At the **USM User:** prompt, enter: **MotoAdmin**



IMPORTANT: You must log on as the **MotoAdmin** user with the current **MotoAdmin** passphrases. If you attempt to log in as a **non-MotoAdmin** user or enter the incorrect passphrases, the system returns you to the root menu with the error message about invalid SNMPv3 credentials.

The **Authentication PassPhrase:** prompt appears.

4. At the **Authentication PassPhrase:** prompt, enter the authentication passphrase for the **MotoAdmin** user.



NOTE: For security purposes, the system does not display or echo back characters when passphrases are entered.

The **Encryption PassPhrase:** prompt appears.

5. At the **Encryption PassPhrase:** prompt, enter the encryption passphrase for the **MotoAdmin** user.

SNMP V3 User Manager Menu displays:

```
1 - Create USM User
2 - Delete USM User
3 - Change Non-Inform User Authentication PassPhrase
4 - Change Non-Inform User Encryption PassPhrase
5 - Change Inform User Credentials
6 - List All Accounts
Select (1-6) ... <CR> to Exit =====>
```



NOTE: When you log on as the **MotoAdmin** user, the EOS software generates an **aud1CfGChg** trap similar to the following:

"EOS USM user MotoAdmin with engineID 800000a1010a01fd36 Authentication passphrase changed"

This trap is generated because the EOS software sends an **snmpChangeAuthPassphrase** command to the Common Agent as part of the passphrase validation process. The trap is expected and does not require any action on your part.

36.2.4

Predefined SNMPv3 USM Users Management on EOS Devices

EOS supports two predefined USM users:

MotoAdmin

Can issue any command from the SNMP V3 User Manager menu. The **MotoAdmin** user is associated with the **admin_grp** VACM group.

MotoMaster

Can only change its own authentication and encryption passphrases. The **MotoMaster** user is associated with one of the following VACM groups: **MotoMaster_grp_authpriv**, **MotoMaster_grp_authnopriv**, or **MotoMaster_grp_noauthnopriv**. By default, the **MotoMaster** user is associated with the **MotoMaster_grp_noauthnopriv** VACM group and is not secured via either authentication or encryption.

For details about the predefined VACM groups, see [SNMPv3 Procedure Reference on page 684](#).

In addition, EOS supports inform users, which are used solely for the inform traffic supported by the NMA reliable communication feature. Each agent always uses the inform user account (**MotoInformA/B**) to send SNMPv3 inform events to the manager. An inform user is created automatically in the USM table in clear

mode when the Common Agent receives a multiple manager request (MMR) for reliable communication from a permanent manager. The number of permanent managers supported on an EOS device depends on whether or not the device is configured as a conventional channel gateway (CCGW).

- If the EOS device is not configured as a CCGW, it supports up to eight permanent managers (and correspondingly eight inform users).
- If the EOS device is configured as a CCGW, it supports up to four permanent managers (and correspondingly four inform users).

The **Inform** users are identified by their name (**MotoInformA/B**) and the engine ID of the manager that sent the MMR. For further details about the **Inform** users, see [Inform Users on page 675](#).



IMPORTANT: If you do not know the FQDN or Engine ID value that is currently assigned to the USM user you are managing, you must enter 6 from **SNMP V3 User Manager Menu** to select **List All Accounts** and record the value currently assigned to the user. You need this value to complete the management procedures. For details about FQDN and Engine ID values, see [SNMPv3 Procedure Reference on page 684](#).

If you change the system IP address on an EOS device, the SNMPv3 credentials must be re-established with that device. Therefore if SNMPv3 users were configured on the device prior to the system IP address change, you must issue the `ResetV3` command to reset the SNMPv3 data, then reconfigure the SNMPv3 users with the appropriate privilege levels.

36.2.4.1

Changing the MotoAdmin User Password



NOTE: The security level for the **MotoAdmin** user cannot be changed, and the **MotoAdmin** user cannot be deleted.

Procedure:

1. Open **SNMP V3 User Manager Menu**.
2. Change the authentication passphrase for the **MotoAdmin** user by performing the following actions:
 - a. From **SNMP V3 User Manager Menu**, select **Change Non-Inform User Authentication Passphrase** by entering: 3
The **USM User:** prompt appears:
 - b. At the **USM User:** prompt, enter: `MotoAdmin`.
The **Engine ID:** prompt appears.
 - c. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine assigned to the **MotoAdmin** user.
The default engine ID is the local device. If this is the engine ID assigned to the **MotoAdmin** user, you can press ENTER at the **Engine ID:** prompt.
The **Current Authentication PassPhrase:** prompt appears.
 - d. At the **Current Authentication PassPhrase:** prompt, enter the current authentication passphrase for the **MotoAdmin** user.
The **Current Encryption PassPhrase:** prompt appears.
 - e. At the **Current Encryption PassPhrase:** prompt, enter the current encryption passphrase for the **MotoAdmin** user.
The **Enter New Authentication PassPhrase:** prompt appears.

- f. At the **Enter New Authentication PassPhrase:** prompt, enter the new authentication passphrase for the **MotoAdmin** user.

The **Retype New Authentication PassPhrase:** prompt appears.

- g. At the **Retype New Authentication PassPhrase:** prompt, re-enter the new authentication passphrase for verification.

The **MotoAdmin** authentication passphrase is changed, and you are returned to **SNMP V3 User Manager Menu**.

3. Change the encryption passphrase for the **MotoAdmin** user by performing the following actions:

- a. From **SNMP V3 User Manager Menu**, select **Change Non-Inform User Encryption Passphrase** by entering: 4

The **USM User:** prompt appears.

- b. At the **USM User:** prompt, enter: **MotoAdmin**.

The **Engine ID:** prompt appears.

- c. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine assigned to the **MotoAdmin** user.

The default engine ID is the local device. If this is the engine ID assigned to the **MotoAdmin** user, you can press ENTER at the **Engine ID:** prompt.

The **Current Authentication PassPhrase:** prompt appears.

- d. At the **Current Authentication PassPhrase:** prompt, enter the authentication passphrase that you assigned to the **MotoAdmin** user in [step 2e](#).

The **Current Encryption PassPhrase:** prompt appears.

- e. At the **Current Encryption PassPhrase:** prompt, enter the current encryption passphrase for the **MotoAdmin** user.

The **New Encryption PassPhrase:** prompt appears.

- f. At the **New Encryption PassPhrase:** prompt, enter the new encryption passphrase for the **MotoAdmin** user.

The **Retype New Encryption PassPhrase:** prompt appears.

- g. At the **Retype New Encryption PassPhrase:** prompt, re-enter the new encryption passphrase for verification.

The **MotoAdmin** encryption passphrase is changed, and you are returned to **SNMP V3 User Manager Menu**.

36.2.4.2

Changing the Privilege Level for the MotoMaster User

By default, the **MotoMaster** user is associated with the **MotoMaster_grp_noauthnopriv** VACM group and is not secured via either authentication or encryption (NoAuthNoPriv privileges).

Procedure:

1. Open **SNMP V3 User Manager Menu**.
2. Delete the **MotoMaster** user by performing the following actions:

- a. From the **SNMP V3 User Manager** menu, select **Delete USM User** by entering: 2
The **USM User:** prompt appears.
- b. At the **USM User:** prompt, enter: **MotoMaster**
The **Engine ID:** prompt appears.
- c. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine assigned to the **MotoMaster** user.
The default engine ID is the local device. If this is the engine ID assigned to the **MotoAdmin** user, you can press ENTER at the **Engine ID:** prompt.

The **MotoMaster** user is deleted, and you are returned to **SNMP V3 User Manager Menu**.

3. Recreate the **MotoMaster** user with AuthPriv or AuthNoPriv privileges by performing the following actions:

- a. From **SNMP V3 User Manager Menu**, select **Create USM User** by entering: 1
The **USM User:** prompt appears.
- b. At the **USM User:** prompt, enter: **MotoMaster**
The **Engine ID:** prompt appears.
- c. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine you want to assign to the **MotoMaster** user.
The default engine ID is the local device. To use the default value, press ENTER at the **Engine ID:** prompt.

The following menu appears:

```
Security Level
1 - AuthPriv
2 - AuthNoPriv
3 - NoAuthNoPriv
```

- d. Perform one of the following actions.
 - To assign AuthPriv privileges to the **MotoMaster** user, enter: 1
 - To assign AuthNoPriv privileges to the **MotoMaster** user, enter: 2

The **MotoMaster** user is recreated, and you are returned to **SNMP V3 User Manager Menu**.

4. Change the **MotoMaster** user authentication passphrase by performing the following actions:

- a. From **SNMP V3 User Manager Menu**, select **Change Non-Inform User Authentication Passphrase** by entering: 3
The **USM User:** prompt appears.
- b. At the **USM User:** prompt, enter: **MotoMaster**
The **Engine ID:** prompt appears.
- c. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine assigned to the **MotoMaster** user.

The default engine ID is the local device. If this is the engine ID assigned to the **MotoMaster** user, you can press ENTER at the **Engine ID:** prompt.

The **Current Authentication PassPhrase:** prompt appears.

- d. At the **Current Authentication PassPhrase:** prompt, enter the current authentication passphrase for the **MotoAdmin** user.
- e. Perform one of the following actions:
 - If the current security level for the **MotoMaster** user is AuthPriv, at the **Current Encryption PassPhrase:** prompt that appears, enter the current encryption passphrase for the **MotoAdmin** user.
 - If the current security level for the **MotoMaster** user is AuthNoPriv, you are not prompted for the current encryption passphrase.

The **Enter New Authentication PassPhrase:** prompt appears.

- f. At the **Enter New Authentication PassPhrase:** prompt, enter the new authentication passphrase for the **MotoMaster** user.

The **Retype New Authentication PassPhrase:** prompt appears.

- g. Re-enter the new authentication passphrase for verification.

The **MotoMaster** user authentication passphrase is changed, and you are returned to **SNMP V3 User Manager Menu**.

5. Change the **MotoMaster** user encryption passphrase by performing the following actions:

- a. From **SNMP V3 User Manager Menu**, select **Change Non-Inform User Encryption Passphrase** by entering: 4

The **USM User:** prompt appears.

- b. Enter: `MotoMaster`

The **Engine ID:** prompt appears.

- c. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine assigned to the **MotoMaster** user.

The default engine ID is the local device. If the engine ID is assigned to the **MotoMaster** user, you can press ENTER at the **Engine ID:** prompt.

The **Current Authentication PassPhrase:** prompt appears.

- d. At the **Current Authentication PassPhrase:** prompt, enter the current authentication passphrase for the **MotoMaster** user that you assigned in [step 4f](#).

The **Current Encryption PassPhrase:** prompt appears.

- e. At the **Current Encryption PassPhrase:** prompt, enter the encryption passphrase for the **MotoAdmin** user.

The **New Encryption PassPhrase:** prompt appears.

- f. At the **New Encryption PassPhrase:** prompt, enter the new encryption passphrase for the **MotoMaster** user.

The **Retype New Encryption PassPhrase:** prompt appears

- g. Re-enter the new encryption passphrase for verification.

The **MotoMaster** user encryption passphrase is changed, and you are returned to **SNMP V3 User Manager Menu**.

36.2.4.3

Inform Users

EOS devices support an inform user account (**MotoInformA/B**) that sends trap (INFORM) requests to registered managers in support of reliable communications. The number of permanent managers supported on an EOS device depends on whether or not the device is configured as a conventional channel gateway (CCGW).

- If the EOS device is not configured as a CCGW, it supports up to eight permanent managers (and correspondingly eight inform users).
- If the EOS device is configured as a CCGW, it supports up to four permanent managers (and correspondingly four inform users).

The Inform users are identified by their name (**MotoInformA/B**) and the engine ID of the manager that sent the MMR.

The **MotoInformA** user is initially created in one of two ways:

- Manually, via the SNMP V3 User Manager Menu. For details about how to manually create the **MotoInformA** user with NoAuthNoPriv privileges (clear mode), see [Creating MotoInformA \(NoAuthNoPriv\) Users Manually on page 676](#).
- Automatically, when there is no **MotoInform** user on the device, and the device needs to send out a trap (due to an existing or manual registration). In this case, the **MotoInformA** account is created with NoAuthNoPriv privileges (clear mode). The **MotoInformA** account is also automatically created with NoAuthNoPriv privileges if the **MotoInform** user is deleted and a trap needs to be sent after the deletion.

Each **MotoInformA** account has a corresponding **MotoInformB** account – associated with the same engine ID. Only one of these accounts (either **MotoInformA** or **MotoInformB**) is active on the device at a time.

- The active inform user is referred to as the current/used inform user.
- The non-active inform user is referred to as the target/unused inform user.

You cannot change the credentials (security level and passphrases) for the current/used inform user; instead, you must activate the target/unused user and assign the desired privileges and passphrases to that user. You do this using option **5 (Change Inform User Credentials)** from the **SNMP V3 User Manager** menu. When you completed this procedure successfully, the EOS software activates the target/unused user, which then becomes the current/used user, and deletes the current/used user, which then becomes the target/unused user.

For example, if you need to assign AuthPriv or AuthNoPriv privilege levels to a **MotoInformA** account that was automatically created by an MMR for reliable communication, then you must use option **5 (Change Inform User Credentials)** from the **SNMP V3 User Manager** menu to create and activate the corresponding **MotoInformB** user and assign it the appropriate security level and passphrase or passphrases. Once you do this, the initial **MotoInformA** account is deleted, and the device uses the **MotoInformB** account.



NOTE: The **MotoInformA** or **MotoInformB** account always belongs to the **notify_grp** VACM group.

36.2.4.4

Creating MotoInformA (NoAuthNoPriv) Users Manually

You can use this procedure to create a **MotoInformA** user if no inform user was previously created on the device. You can use option **6 (List All Accounts)** from **SNMP V3 User Manager Menu** to check for the existence of an inform user on the device.



IMPORTANT: If you maintain the **MotoInform** user for the Unified Event Manager (UEM), enter the “Service Name” from the UEM interface when prompted for the FQDN in this procedure. The “Service Name” is case-sensitive, so enter it carefully.

Procedure:

1. Open **SNMP V3 User Manager Menu**.
2. From **SNMP V3 User Manager Menu**, select **Create USM User** by entering: 1
The **USM User Name:** prompt appears.
3. At the **USM User Name:** prompt, enter: **MotoInformA**
The following menu appears:

```
1 - FQDN Format
2 - Engine ID Format
```
4. To indicate that you specify the **MotoInformA** user SNMP engine ID in FQDN format, enter: 1
The **FQDN:** prompt appears.
5. At the **FQDN:** prompt, enter the FQDN of the SNMP engine that you want to assign to the **MotoInformA** user.
The FQDN (fully qualified domain name) is the complete domain name of the authorized manager involved in the SNMPv3 message exchange. The authorized manager is the manager that registers with the SNMPv3 engine on the device. The FQDN consists of two parts: the host name and the domain name.
For example, to identify host **z001uem01** within domain **zone1** as the authorized manager involved in the SNMPv3 message exchange, you must enter:
z001uem01.zone1
You must enter a FQDN in response to the **FQDN:** prompt. If you simply press ENTER, the system returns the invalid credentials error message.
The following menu appears:

```
Security Level
1 - AuthPriv
2 - AuthNoPriv
3 - NoAuthNoPriv
```
6. To assign **NoAuthNoPriv** privileges to the **MotoInformA** user, enter: 3
You are returned to **SNMP V3 User Manager Menu**.

36.2.4.5

Changing Credentials for Inform Users

You can use this procedure to change the security level and associated passphrase(s) for an inform user.

Procedure:

1. Open **SNMP V3 User Manager Menu**.
2. To determine the name (**MotoInformA** or **MotoInformB**) of the inform user for which you want to change the credentials, from **SNMP V3 User Manager Menu**, select **List All Accounts** by entering: 6

A display similar to the following appears:

```
USM Name: MotoAdmin
Security Level: AuthPriv
EngineID: 800000A10100000000
Status: Active

USM Name: MotoMaster
Security Level: NoAuthNoPriv
EngineID: 800000A10100000000
Status: Active

USM Name: MotoInformA
Security Level: NoAuthNoPriv
EngineID: 800000A1047A30303175656D30312E7A6F6E6531
FQDN: z001uem01.zone1
Status: Active

USM Name: MotoInformA
Security Level: NoAuthNoPriv
EngineID: 800000A1047A30303175656D30322E7A6F6E6531
FQDN: z001uem02.zone1
Status: Active
```



NOTE: If the display lists two accounts for the same FQDN (one with Status **Active** and one with Status **Not in Service**), you must delete the **Not in Service** account before proceeding. To delete this account, from the **SNMP V3 User Manager** menu, select **Delete USM User** by entering: 2

3. From **SNMP V3 User Manager Menu**, select **Change Inform User Credentials** by entering: 5
The **USM User to Activate:** prompt appears.
4. At the **USM User to Activate:** prompt, enter the name of the target/unused inform user corresponding to the current/used inform user for which you want to change the credentials. (**MotoInformA** or **MotoInformB**).

This name is not listed in the **List All Accounts** display returned in [step 2](#).

For example, if you want to change the credentials for the **MotoInformA** user with FQDN **z001uem01.zone1**, you must enter:

MotoInformB



NOTE: You cannot change the credentials of a current/used inform user. If you enter the name of the current/used inform user, the invalid SNMPv3 credentials error message is returned.

The following menu appears:

```
1 - FQDN Format
2 - Engine ID Format
```

5. To indicate that you specify the inform user's SNMP engine ID in FQDN format, enter: 1
The **FQDN:** prompt appears.

6. At the **FQDN:** prompt, enter the FQDN of the inform user you want to activate.

The FQDN is the same for the current/used **MotoInformA/B** user and its corresponding target/unused **MotoInformA/B** user.

For example, you can enter: `z001uem01.zone1`



NOTE: You must enter a FQDN in response to the **FQDN:** prompt. If you press **Enter**, the system returns the following error message:

`Invalid credential`

The following menu appears:

```
Security Level
1 - AuthPriv
2 - AuthNoPriv
3 - NoAuthNoPriv
```

7. Enter the number corresponding to the security level that you want to assign to the inform user.
8. Perform one of the following actions:
 - If you specified security level 1 (AuthPriv) or 2 (AuthNoPriv), the **Current Authentication PassPhrase:** prompt appears.
 - If you specified security level 3 (NoAuthNoPriv), you are not prompted for passphrases and must skip [step 9](#).
9. At the **Current Authentication PassPhrase:** prompt, enter the current authentication passphrase for the **MotoAdmin** user.



NOTE: When the target/unused inform user is activated, it is cloned from the **MotoAdmin** user; therefore, the current authentication passphrase is that of the **MotoAdmin** user, and not the current/used inform user.

The **Current Encryption PassPhrase:** prompt appears.

10. At the **Current Encryption PassPhrase:** prompt, enter the current encryption passphrase for the **MotoAdmin** user.



NOTE: When the target/unused inform user is activated, it is cloned from the **MotoAdmin** user; therefore, the current encryption passphrase is that of the **MotoAdmin** user, and not the current/used inform user.

The **New Authentication PassPhrase:** prompt appears:

11. At the **New Authentication PassPhrase:** prompt, enter the new authentication passphrase for the inform user.

The **Retype New Authentication PassPhrase:** prompt appears.

12. Re-enter the new authentication passphrase for verification.

If your entry does not match the authentication passphrase you entered in [step 11](#), EOS returns the `Invalid SNMPv3 credentials`

error message. In addition, EOS creates a `Not in Service` inform user and does not activate the target/unused inform user. For this reason, if you see the `Invalid SNMPv3 credentials` error message, you must delete the `Not in Service` user (from **SNMP V3 User Manager Menu**, you must select **Delete USM User** by entering: 2), and return to [step 1](#) to change the inform user credentials. This also applies when you are entering the encryption passphrase for an inform user with security level 1 (AuthPriv).

13. Perform one of the following actions:

- If you assigned the inform user security level 1 (AuthPriv), at the **New Encryption PassPhrase:** prompt, enter the new encryption passphrase for the inform user and retype the new encryption passphrase at the prompt for verification.
- If you assigned the inform user security level 2 (AuthNoPriv), you are not prompted for a new encryption passphrase.

The current/used inform user is deleted (and becomes the target/unused inform user); the target/unused inform user becomes the current/used inform user; and you are returned to the SNMP V3 User Manager menu.

14. To verify that, for the inform user for which you changed the credentials, the target/unused inform user is now the current/used inform user, from **SNMP V3 User Manager Menu** menu, select **List All Accounts** by entering: 6

A display similar to the following appears:

```
USM Name: MotoAdmin
Security Level: AuthPriv
EngineID: 800000A10100000000
Status: Active

USM Name: MotoMaster
Security Level: NoAuthNoPriv
EngineID: 800000A10100000000
Status: Active

USM Name: MotoInformB
Security Level: NoAuthNoPriv
EngineID: 800000A1047A30303175656D30312E7A6F6E6531
FQDN: z001uem01.zone1
Status: Active

USM Name: MotoInformA
Security Level: NoAuthNoPriv
EngineID: 800000A1047A30303175656D30322E7A6F6E6531
FQDN: z001uem02.zone1
Status: Active
```

Note that the USM name for the inform user corresponding to FQDN z001uem01.zone1 is now **MotoInformB**.

15. If you assigned the inform user AuthPriv or AuthNoPriv privileges, verify that authorized and encrypted (AuthPriv) or authorized (AuthNoPriv) informs can be received and read by the network manager.

The newly-activated **MotoInformA/B** account is now used for encrypted and authenticated (AuthPriv) or authenticated (AuthNoPriv) traps.

36.2.5

General SNMPv3 USM Users Management on EOS Devices

When you add a new user to the USM database, it is cloned from the **MotoAdmin** user and assigned to a VACM group based on the security level that you specify when you add the user (AuthPriv, AuthNoPriv, or NoAuthNoPriv).

The new user also inherits the **MotoAdmin** user passphrases:

- If you assign the new user AuthPriv security level, the user inherits the **MotoAdmin** user's current authentication and encryption passphrases. You must change these passphrases after the user is created.
- If you assign the new user AuthNoPriv security level, the user inherits the **MotoAdmin** user's current authentication passphrase. You must change the authentication passphrase after the user is created.

- If you assign the new user NoAuthNoPriv security level, it does not have passphrases.



IMPORTANT: If you do not know the Engine ID value that is currently assigned to the USM user you are managing, from **SNMP V3 User Manager Menu** you must select **List All Accounts** (by entering: 6) and record the Engine ID value currently assigned to the user. You need this value to complete the management procedures. For details about FQDN and Engine ID values, see [SNMPv3 Procedure Reference on page 684](#).

36.2.5.1

Adding SNMPv3 USM Users

Procedure:

1. Open **SNMP V3 User Manager Menu**.
2. Select **Create USM User** by entering: 1
The **USM User:** prompt appears.
3. At the **USM User:** prompt, enter a user name of up to 31 characters.
The **Engine ID:** prompt appears.
4. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine you want to assign to the user.
The default engine ID is the local device. To assign this engine ID to the user, at the **Engine ID:** prompt, you must press ENTER.

The following menu appears:

```
Security Level
1 - AuthPriv
2 - AuthNoPriv
3 - NoAuthNoPriv
```

5. Enter the number corresponding to the security levels you want to assign to the user.
 - 1 - AuthPriv - The user is secured via authentication and privacy (encryption).
 - 2 - AuthNoPriv - The user is secured via authentication but not via encryption.
 - 3 - NoAuthNoPriv - The user is not secured via either authentication or encryption (clear SNMPv3).

The user is created, and you are returned to **SNMP V3 User Manager Menu**.

6. **Users with security level 1 (AuthPriv) or 2 (AuthNoPriv):** Change the user authentication passphrase. See [Changing Authentication Passphrases for SNMPv3 USM Users on page 681](#).
7. **Users with security level 1 (AuthPriv):** Change the user encryption passphrase. See [Changing Encryption Passphrases for SNMPv3 USM Users on page 682](#).

36.2.5.2

Deleting SNMPv3 USM Users

Procedure:

1. Open **SNMP V3 User Manager Menu**.
2. Select **Delete USM User** by entering: 2
The **USM User:** prompt appears.

3. At the **USM User:** prompt, enter the user name of the USM user you want to delete.
The **Engine ID:** prompt appears.
4. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine assigned to the user.
The default engine ID is the local device. If this is the engine ID assigned to the user, at the **Engine ID:** prompt, you can press **Enter**.
The user is deleted, and you are returned to **SNMP V3 User Manager Menu**.

36.2.5.3

Changing Authentication Passphrases for SNMPv3 USM Users

This procedure applies only to users assigned security level 1 (AuthPriv) or security level 2 (AuthNoPriv). Users assigned security level 3 (NoAuthNoPriv) are not associated with authentication or encryption passphrases.

This procedure applies only to non-Inform USM users. To change the authentication passphrase for the **MotoInformA/B** user, from the **SNMP V3 User Manager** menu you must use select **Change Inform User Credentials** by entering: 5 For details, see [Changing Credentials for Inform Users on page 677](#).

Procedure:

1. Open **SNMP V3 User Manager Menu**.
2. Select **Change Non-Inform User Authentication Passphrase** by entering: 3
The **USM User:** prompt appears.
3. At the **USM User:** prompt, enter the user name of the USM user for which you want to change the authentication passphrase.
If you enter the user name of a USM user with security level 3 (NoAuthNoPriv), the `Invalid SNMPv3 credentials` error message is returned.
The **Engine ID:** prompt appears.
4. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine assigned to the user.
The default engine ID is the local device. If this is the engine ID assigned to the user, you can press **ENTER** at the **Engine ID:** prompt.
The **Current Authentication PassPhrase:** prompt appears.
5. At the **Current Authentication PassPhrase:** prompt, enter the current authentication passphrase for the specified user.
If you just created the user and did not change the authentication passphrase, the current authentication passphrase is the authentication passphrase for the **MotoAdmin** user.
6. Perform one of the following actions:
 - If the user has security level 1 (AuthPriv), at the **Current Encryption PassPhrase:** prompt appears, enter the current encryption passphrase for the specified user.
If you just created the user and have not changed the encryption passphrase, the current encryption passphrase is the authentication passphrase for the **MotoAdmin** user.
 - If the user has security level 2 (AuthNoPriv), you are not prompted for the current encryption passphrase.The **New Authentication PassPhrase:** prompt appears.

7. At the **New Authentication PassPhrase:** prompt, enter the new authentication passphrase for the specified user.

The **Retype New Authentication PassPhrase:** prompt appears.

8. At the **Retype New Authentication PassPhrase:** prompt, re-enter the new authentication passphrase for verification.

The user authentication passphrase is changed, and you are returned to **SNMP V3 User Manager Menu**.

36.2.5.4

Changing Encryption Passphrases for SNMPv3 USM Users

This procedure applies only to users assigned security level 1 (AuthPriv). Users assigned security level 2 (AuthNoPriv) or 3 (NoAuthNoPriv) are not associated with encryption passphrases.

This procedure applies only to non-Inform USM users. To change the authentication passphrase for the **MotoInformA/B** user, from **SNMP V3 User Manager Menu**, you must use select **Change Inform User Credentials** by entering: 5 For details, see [Changing Credentials for Inform Users on page 677](#).

Procedure:

1. Open **SNMP V3 User Manager Menu**.
2. Select **Change Non-Inform User Encryption Passphrase** by entering: 4
The **USM User:** prompt appears.
3. At the **USM User:** prompt, enter the user name of the USM user for which you want to change the encryption passphrase.
If you enter the user name of a USM user with security level 2 (AuthNoPriv) or 3 (NoAuthNoPriv), the **Invalid SNMPv3 credentials** error message is returned.
The **Engine ID:** prompt appears.
4. At the **Engine ID:** prompt, enter the engine ID number of the SNMP engine assigned to the user.
The default engine ID is the local device. If this is the engine ID assigned to the user, at the **Engine ID:** prompt, you can press ENTER .
The **Current Authentication PassPhrase:** prompt appears.
5. At the **Current Authentication PassPhrase:** prompt, enter the current authentication passphrase for the specified user.
If you just created the user and did not change the authentication passphrase, the current authentication passphrase is the authentication passphrase for the **MotoAdmin** user.
The **Current Encryption PassPhrase:** prompt appears.
6. At the **Current Encryption PassPhrase:** prompt, enter the current encryption passphrase for the specified user.
If you just created the user and did not change the encryption passphrase, the current encryption passphrase is the passphrase for the **MotoAdmin** user.
The **New Encryption PassPhrase:** prompt appears.
7. At the **New Encryption PassPhrase:** prompt, enter the new encryption passphrase for the specified user.
The **Retype New Encryption PassPhrase:** prompt appears.

8. At the **Retype New Encryption PassPhrase:** prompt, re-enter the new encryption passphrase for verification.

The user encryption passphrase is changed, and you are returned to **SNMP V3 User Manager Menu**.

36.2.5.5

Viewing SNMPv3 USM Users

You can use this procedure to view a list of the currently-configured SNMPv3 USM users.

Procedure:

1. Open **SNMP V3 User Manager Menu**.
2. Enter: 6

A display similar to the following appears:

```
USM Name: MotoAdmin
Security Level: AuthPriv
EngineID: 800000A1010A01FD33
Status: Active

USM Name: MotoMaster
Security Level: NoAuthNoPriv
EngineID: 800000A1010A01FD33
Status: Active

USM Name: MotoInformA
Security Level: NoAuthNoPriv
EngineID: 800000A1047A30303175656D30312E7A6F6E6531
FQDN: z001uem01.zone1
Status: Active

USM Name: MotoInformA
Security Level: NoAuthNoPriv
EngineID: 800000A1047A30303175656D30322E7A6F6E6531
FQDN: z001uem02.zone1
Status: Active
```

This display provides the following information about each user in the USM database:

USM Name

USM user name.

Security Level

Security privilege level for the user (AuthPriv, AuthNoPriv, or NoAuthNoPriv).

EngineID

SNMP engine ID assigned to the user. If the engine ID was assigned as a FQDN, the FQDN is translated into an engine ID value.

FQDN

The fully-qualified domain name of the SNMP engine assigned to the user. Displayed for inform users only.

Status

Status of the USM user:

- 1:Active – The user is created and ready to be used.
- 2:Not in service – Operations on the user were suspended.
- 3:Not Ready – The user is created but not ready to be used.
- 4:Create and go – The user is created and is active as soon as it is created.
- 5:Create and wait – The user is created but is not ready to be used.

6:Destroy – The user is or should be deleted.

36.2.6

SNMPv3 Procedure Reference

SNMPv3 Passphrase Restrictions

The SNMPv3 user authentication or encryption passphrase must be 8-64 characters long and may contain uppercase and lowercase alphabetic characters (A-Z) and (a-z); numeric characters (0-9); and any of the allowed special characters (! " % & " () * + , - . / : ; < = > ?).

Predefined VACM Groups

Table 61: Predefined VACM Groups Supported on EOS Devices

Group Name	Description
MotoMaster_grp_authpriv	For MotoMaster users only; can change its own authentication and encryption passphrases and is allowed access to the entire SNMP-USER-BASED-SM-MIB tree.
MotoMaster_grp_authnopriv	For MotoMaster users only; can change its own authentication and passphrase but not its own encryption passphrase and encryption passphrases and is allowed access to the entire SNMP-USER-BASED-SM-MIB tree.
MotoMaster_grp_noauthnopriv	For MotoMaster users only; cannot change its authentication or encryption passphrases and is allowed access to the entire SNMP-USER-BASED-SM-MIB tree.
notify_grp	For MotoInform user only; cannot change its authentication or encryption passphrases and is allowed only to receive notifications. Users in the notify_grp group have no read/write access to the SNMP-USER-BASED-SM-MIB tree.
admin_grp	For the MotoAdmin user only; can change the authentication and encryption passphrases for any USM user and is allowed access to all MIB trees.

FQDN/Engine ID Value

When an EOS device is first booted, its system IP address is automatically assigned as the FQDN/Engine ID value for the predefined USM users (**MotoAdmin**, **MotoMaster**, and **MotoInform**). You can change this value if you delete and recreate the predefined user.

In some network positions (customer enterprise network (CEN) border routers, for example), the system IP address may be a duplicate address in the radio network interface (RNI) network. To avoid a situation in which nonunique engine IDs would be generated from duplicate system IP addresses, you can use the SNMP service `EngineIDIP` parameter to specify a unique IP address from which to generate the engine ID for a

device. For details, see the description of the `EngineIDIP` parameter in “SNMP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



IMPORTANT: Specifying a unique IP address using the `EngineIDIP` parameter is the only way to generate a unique engine ID. Unless you specify an IP address with the `EngineIDIP` parameter, EOS uses the system IP address to generate engine IDs, and duplicate system IP addresses result in duplicate engine IDs.

EOS generates the engine ID when SNMPv3 management is enabled and does not support dynamic changes of the engine ID. Therefore, the `SETDefault -SNMP EngineIDIP` command must be included before the `SETDefault -SNMP CONTROL = V3` command in the device’s configuration file.

For non-predefined USM users, you assign the FQDN/Engine ID value when you add the user.

If you do not know the FQDN/Engine ID value assigned to a USM user, and you need to enter this value in order to modify the user, you can select option

6 (List All Accounts) from the SNMP V3 User Manager menu to view the currently assigned values for all configured USM users.

The FQDN/Engine ID value identifies the authoritative SNMP engine — the SNMP engine ID of the authorized manager involved in the SNMPv3 message exchange. The authorized manager is the manager that registers with the SNMPv3 engine on the EOS device.

- The authoritative engine is the destination for a message that expects a response; for example a Get, GetNext, GetBulk, Set or Inform. In this case, the remote manager (for example the UEM) performs the registration. You must specify an remote engine ID when you create a USM user. You can enter the remote engine ID as either an FQDN or an engine ID (32-byte hexadecimal digit). If you enter the value as a FQDN, EOS converts the FQDN to an engine ID, which is decoded to a prefix plus an FQDN . For example:

```
800000A1047A30303175656D30312E7A6F6E6531
```

is decoded to:

```
800000A104(prefix) + z001uem01.zone1
```

- The authoritative engine is the source for a message that does not expect a response; for example, a trap, response, or report. In this case, EOS does the registration. The engine ID is local, and you are not required to specify it when you create a USM user. You can press ENTER in response to either the **FQDN:** or the **Engine ID:** prompt to direct the device to use the local engine ID. EOS decodes the local engine ID to the prefix plus the device’s system IP address. For example:

```
800000A (prefix) + 10.1.253.51
```

36.3

Remote Network Monitoring Alarms

To assist network managers in identifying abnormal activity that could adversely affect network performance, the system contains a Remote Monitoring (RMON) agent based on RFC 1271. The Motorola Solutions implementation supports two RMON object groups:

Alarm group

You can set up alarms to monitor the MIB objects of interest through the Alarm group; this group includes an Alarm Table, which you must configure before using alarms. The Motorola Solutions implementation allows you to configure five alarms in the Alarm Table.

Event group

You can set up events to either record the monitoring information or notify the network management station. The Event group includes an Event Table and a Log Table. You must configure the Event Table before using it. The Log Table needs no prior configuration; it is a read-only data table for the network management station. The Motorola Solutions implementation allows you to configure five events in the Event Table.

Because the system contains no user interface to access the RMON agent, you must access it through an SNMP network management station by using SNMP SET/GET/GET-NEXT requests, or by using an SNMP application that can generate SNMP requests. In order to monitor MIB objects of interest according to a set or preset alarms and events, you need to configure the RMON alarms and events with the desired control parameters. The control parameters to be configured include alarm threshold values, which are used to determine if an event needs to be generated.

An event is generated if the newly read MIB value crossed the threshold. The event can take any of the following actions:

- The system sends an SNMP trap to the network management station.
The management station is notified immediately. The management station determines how to react to the SNMP trap.
- The system logs the event into a Log Table in the agent system.
The management station can retrieve the information stored in the Log Table for further analysis. For example, the information collected can be used in selecting proper threshold values.
- The system sends an SNMP trap and logs into the Log Table.

36.4

Network Maps

A network map (netmap) contains the Ethernet and Internet addresses of each EOS device on the attached network and the software version on the device. The EOS device can participate in building the netmap by broadcasting its addresses to the network at regular intervals (defined by `NetMapTime` in the SYS Service).

By default, the value of the `NetMapTime` parameter is set to 0, which means that the system does not broadcast its addresses to the network. To configure the system to broadcast its addresses to the network a specific number of seconds, you must use:

```
SETDefault -SYS NetMapTime = <number>(0 to 120 seconds)
```

The network map can be used as a network management tool, because you can see at a glance which Motorola Solutions devices are on the attached network. To display the network map, you can use:

```
SHOW -SYS NetMAP
```

For the possible variations for this command and what these variations display, see “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. For example, to show all the devices that support the TCP/IP Protocol on the network attached to port 1 of the system, you can use:

```
SHOW !1 -SYS NetMAP tcp
```

A display similar to the following appears:

NETWORK 1 MAP			
0-%08000200F84A	129.213.16.206	1-%08000201136A	129.213.16.208
2-%08000200D26C	129.213.17.152	3-%080002013E86	129.213.18.47
4-%080002013E98	129.213.19.67	5-%080002013E9C	129.213.18.52
6-%0800020128B8	129.213.17.147	7-%08000201ECD3	129.213.16.215

If you want to display the software version running on each Motorola Solutions device, you can use the `Long` value in the command. For example, you can use:

```
SHOW !1 -SYS NetMAP Long
```

Devices that support TCP on the network attached to port 1 are displayed. A display similar to the following appears:

NETWORK 1 MAP		
0-%08000200F84A	129.213.16.206	SW/1-PCS 10000
1-%08000201136A	129.213.16.208	SW/1-PCS 10000
2-%08000200D26C	129.213.17.152	SW/1-TCP 3.0
3-%080002013E86	129.213.18.47	SW/NB-IB-2.0
4-%080002013E98	129.213.19.67	SW/NB-BR-3.0

5-%080002013E9C	129.213.18.52	SW/NB-IB-2.0
6-%0800020128B8	129.213.17.147	SW/200-TCP 3.0
7-%08000201ECD3	129.213.16.215	SW/1-PCS 1001h

36.5

Statistics Retrieved Through SNMP

You can retrieve statistics remotely through SNMP. You can also use SNMP to reset the statistics to zero.

36.5.1

SNMPv2 GETBULK Support

EOS now supports SNMPv2 GETBULK operations, as well as SNMPv2 GET and GETNEXT operations. SNMPv2 GETBULK support minimizes interoperability issues for EOS devices managed by InfoVista.



NOTE: EOS does not support SNMPv2 SET operations or notifications (traps).

36.5.2

Performance Statistics

EOS supports retrieval of T1/E1 physical layer and ifXTable performance statistics (64-bit counters) via SNMP.

Handling of 64-Bit Counter Overflows

EOS currently supports 64-bit counters for the following statistics:

- POrt service statistics (viewed with the `SHoW STATistics -POrt` command)
- PAtH service statistics (viewed with the `SHoW STATistics -PAtH` command)
- DS1/E1 physical layer performance statistics (retrieved via SNMP)
- ifXTable statistics (retrieved via SNMP)

If a 64-bit statistic value overflows, it is set to the maximum unsigned 64-bit integer, ULLONG_MAX:

18,446,744,073,709,551,615

Statistics displays precede this maximum 64-bit value with an asterisk (*) to indicate overflow and append one of the following letters to indicate the units:

- K (Kilo = 10^{**3})
- M (Mega = 10^{**6})
- G (Giga = 10^{**9})
- T (Tera = 10^{**12})
- P (Peta = 10^{**15})
- E (Exa = 10^{**18})

How many digits of the ULLONG_MAX integer are displayed and which unit letter is used depends on the number of columns (and the resulting column width) in the statistics display.

For example, if the display has a single column, ULLONG_MAX is shown as:

*18446744T

where T is an overflow indicator meaning Tera (10^{**12}).

If the display has multiple columns, ULLONG_MAX could also be displayed as:

*18446P

or

*18E

where P and E are overflow characters meaning Peta (10^{15}) and Exa (10^{18}) respectively.



NOTE: When a 64-bit statistic value reaches the maximum value, it has to be flushed for further incrementation by issuing the `FLush -SYS STATistics` command.

T1/E1 Performance Statistics

EOS supports retrieval of the following T1/E1 physical layer performance statistics (64-bit counters) via SNMP:

SEFS

The count of severely errored framing seconds (SEFS). An SEFS is a second with one or more out-of-frame (OOF) events.

ES

The count of errored seconds (ES). An ES is a second with one or more out-of-frame (OOF) events, path code violations (BPV or CRC), controlled slips, or AIS defect events.

SES

The count of severely errored seconds (SES).

- For ESF frame mode, an SES is a second with 320 or more path code violations (BPV or CRC) or OOF events, or a detected AIS event.
- For E1-CRC frame mode, an SES is a second with 832 or more path code violations (BPV or CRC) or OOF events, or a detected AIS event.
- For E1-noCRC frame mode, an SES is a second with 2048 or more line code violations (LCVs).
- For D4 (SF) signals, an SES is a second with OOF or framing error events.

UAS

The count of unavailable seconds (UAS); the total number of seconds that the interface is in an unavailable signal state. An unavailable signal state occurs at the onset of 10 consecutive SESs. The state is cleared at the onset of 10 seconds with no SESs.

BES

The count of bursty errored seconds (BES). This statistic is applicable to ESF frame mode only. A BES is a second with fewer than 320 and more than 1 path coding violation error events, no severely errored frame defects, and no detected incoming AIS defects.

CSS

The count of controlled slip seconds (CSS). A CSS is a second containing one or more controlled slips.

LES

The count of line-errored seconds (LES). An LES is a second containing one or more line errors. Line errors include BPV, excessive zeros (EXZ), and loss of signal (LOS) defects.

TEFS

The count of total error-free seconds (EFS). This is the total number of seconds for which the port was not in error state since it became operational. This counter is not contiguous; it is a cumulative count of all the seconds in which the port's state was error-free.

EFS

The count of error free seconds (EFS). This is the number of seconds for which the port was not in error state since the last errored second occurred; it is a contiguous counter.

FramingErrors

The count of framing errors. This statistic is reported directly from the 8370 chipset in the performance monitoring registers (registers 0x50, 0x51).

BPVErrors

The count of BPV errors. This statistic is reported directly from the 8370 chipset in the performance monitoring registers (registers 0x054, 0x055).

CRCErrors

The count of cyclic redundancy check (CRC) errors. This statistic is reported directly from the 8370 chipset in the performance monitoring registers (registers 0x052, 0x053).

CSErrors

The count of controlled slip (CS) errors. This statistic is reported directly from the 8370 chipset in the performance monitoring registers (register 0x06).

EXZErrors

The count of excessive zeros (EXZ) errors. This statistic is reported directly from the 8370 chipset in the performance monitoring registers (registers 0x54, 0x55). The registers 0x54 and 0x55 show the combined EXZ and BPV errors, provided the EXZ_LCV bit is set in the register 0x45.

TDM

The count of total degraded minutes (TDM) errors. A degraded minute is one in which the estimated error rate exceeded 1E-6 but not 1E-3. This counter is not contiguous; it is a cumulative count of all the minutes in which TDM errors occurred.

DM

The count of degraded minutes (DM) errors. A degraded minute is one in which the estimated error rate exceeded 1E-6 but not 1E-3. This is the total number of minutes that this condition occurred since the last time the port was not in this condition; it is a contiguous counter.

LOSCnt

The total number of times that the line lost signal. The 8370 chipset alarm register (0x004) provides this statistic, which is read by the T1/E1 driver in a polling routine called every 250 milliseconds.

RunningTime

The number of seconds since the path was activated. This is the length of time that the port was set to an operational state.

For further information about these counters, consult the T1/E1 MIB definition file (`t1e1.mib`).

ifXTable Performance Statistics

EOS supports retrieval of the following ifXTable performance statistics (64-bit counters) through SNMP:



NOTE: When a 64-bit statistic value reaches the maximum value, for further incrementation you must issue the `FLush -SYS STATistics` command. For details about the handling of 64-bit counter overflow, see the description of the `FLush` command in "Commands" in the *Enterprise OS Software Version 18.2 Reference Guide*.

ifHCInOctets

The total number of octets received on the interface, including framing characters. This object is a 64-bit version of `ifInOctets`.

ifHCInUcastPkts

The number of packets delivered by this sublayer to a higher sublayer, which were not addressed to a multicast or broadcast address at this sublayer. This object is a 64-bit version of `ifInUcastPkts`.

ifHCInMulticastPkts

The number of packets delivered by this sublayer to a higher sublayer, which were addressed to a multicast address at this sublayer. For a MAC layer protocol, this includes both group and functional addresses. This object is a 64-bit version of `ifInMulticastPkts`.

ifHCInBroadcastPkts

The number of packets delivered by this sublayer to a higher sublayer, which were addressed to a broadcast address at this sublayer. This object is a 64-bit version of `ifInBroadcastPkts`.

ifHCOutOctets

The total number of octets transmitted out of the interface, including framing characters. This object is a 64-bit version of `ifOutOctets`.

ifHCOutUcastPkts

The number of packets that higher-level protocols requested be transmitted and which were not addressed to a multicast or broadcast address at this sublayer, including those that were discarded or not sent. This object is a 64-bit version of `ifOutUcastPkts`.

ifHCOutMulticastPkts

The number of packets that higher-level protocols requested be transmitted and which were addressed to a multicast address at this sublayer, including those that were discarded or not sent. For a MAC layer protocol, this includes both group and functional addresses. This object is a 64-bit version of `ifOutMulticastPkts`.

ifHCOutBroadcastPkts

The number of packets that higher-level protocols requested be transmitted and which were addressed to a broadcast address at this sublayer, including those that were discarded or not sent. This object is a 64-bit version of `ifOutBroadcastPkts`.

For further information about these counters, consult the IF-MIB as defined by RFC2863.

36.5.3

Priority Queuing Statistics

You can retrieve the following priority queuing statistics remotely via SNMP:

Total In Packets

The number of packets input to the priority queue

Fragment Sequence

The accumulated count of the number of fragments generated when large packets are fragmented.

Xmtu

The number of packets sent to the priority queue that exceeded the maximum transfer unit (MTU).

PRQXmitByte

The total number of bytes transmitted from this priority queue.

Total Packets Dropped

The number of packets dropped from this priority queue because the queue was full.

HiQueueDepth

The maximum attained queue depth, in packets, for this priority queue.

QueuedPacketCount

The current queue depth, in packets, for this priority queue.

These statistics are kept for every priority queue. They count the number of packets and the possible fragmentation of these packets, allowing you to monitor the performance of both the TOS-based interleaving and the FRF.12 fragmentation capabilities.

36.5.4

PIM Statistics

You can retrieve the following PIM statistics remotely via SNMP:

MocaPimTotalMulticastRouteGroup

The total number of groups in the PIM multicast table.

MocaPimRPFChanges

The total number of route changes in the unicast route table reported to PIM.

For further information about these PIM statistics, consult the PIM MIB definition file. The PIM MIB is implemented under PIMMIB.

36.5.5

Multicast IP (MIP) Statistics

You can retrieve the following multicast IP (MIP) statistics remotely via SNMP:

MocaMipTotRcvdPkts

The total number of multicast data packets received.

MocaMipTotFwdPkts

The total number of multicast data packets forwarded.

MocaMipTotDiscardPkts

The total number of multicast data packets discarded.

MocaMipForwarding

The total number of connections between two ports that forward multicast IP packets in the EOS device.

MocaMipCacheEntries

The total number of multicast flows active in the EOS device. This is a global MIP statistic.

aMotorolaCaMipTotActiveSessions

The number of sources and receivers which are involved in active sessions.

aMotorolaCaMipTotMaxActiveSessions

The maximum value reached by the aMotorolaCaMipTotActiveSessions counter.

For further information about these MIP statistics, consult the MIP MIB definition file. The MIP MIB is implemented under MIPMIB.

36.5.6

BGP Statistics

You can retrieve the following BGP statistics remotely via SNMP:

MocaBgpRouteChanges

The total number of changes in the BGP routing table.

For further information about these BGP statistics, consult the BGP MIB definition file. The BGP MIB is implemented under BGPMIB.

36.6

Configuration Changes Logged with the AuditLog Service

The AuditLog Service sends event log messages to syslog servers to provide a history of configuration changes and other events useful in monitoring EOS devices. Up to five syslog servers can receive messages from the AuditLog Service.



NOTE: To use the AuditLog Service, you must have IP and User Datagram Protocol (UDP), and a configuration of the UNIX Syslog daemon.

The log messages provided by the AuditLog Service offer information concerning:

- All use of identification and authentication mechanism
- The termination of an interactive session
- Failed login or set privilege attempts

- All configuration changes
- All cryptographic keys operations
- Password changes
- Starting and stopping services
- Invalid SNMP community strings
- SNMP configuration changes
- File operations
- System and dial history messages
- Reboot information
- Unsuccessful x509 certificate validation
- All attempts of firmware update (success and failures)
- Changes to time
- Initiation/termination/failure to establish IPsec and SSH session
- Successful SSH rekey

AuditLog messages are sent to one or more syslog servers, and logged locally in the local audit log buffer.

The following are sample AuditLog display messages:

```
<189> Seq:49 Thu Apr 14 02:42:05 2016 Sev:5 Res:Success From:E0S/SYS/E0S Msg:System  
Initialized and Running
```

```
<188> Seq:67 Thu Apr 14 02:40:55 2016 Sev:4 Res:Success From:root/SYS/CONSOLE Msg:UI cmd  
"rb" executed
```

```
<189> Seq:53 Thu Apr 14 02:42:59 2016 Sev:5 Res:Success From:E0S/AuditLog/E0S Msg:System  
Booted
```

The message format in the local log is as follows:

```
<priority> Seq:SeqNumber Date Time Sev:Severity Res:<Success/Failed> From:<Entity/  
Username>/Service/Source Msg:Text
```

The message form in the syslog server is as follows:

```
<priority> Date Time From:<Hostname/MAC Address>/  
<Entity/Username>)/Service/Source Seq:SeqNumber Sev:Severity Res:<Success/Failed>  
Msg:Text
```

The message includes the following fields:

Table 62: AuditLog Display Message Fields

Value	Description
<priority>	The priority of the message. The priority is a combination of the severity level and the LocalFacility value.
SeqNumber	A number from 0 to 255. This is the unique identifier for the Syslog event.
<Hostname/MAC Address>	The resolved host name or MAC address. When displaying the log files on the syslog server, this field is prepended to each log entry.
Severity	The severity level in numeric form. Severity levels are as follows: 0 = Emergency 1 = Alert

Value	Description
	2 = Critical 3 = Error 4 = Warning 5 = Notice 6 = Info 7 = Debug
<Success/Failed>	The outcome of the audit log event.
<Entity/Username>	The entity or username that initiated the log message. Username specifies the user who initiated the command. The entity EOS specifies the EOS system initiated the log message.
Service	Service of the EOS that initiated the log message.
Source	Source of the log message. Possible sources include: CONSOLE – The console port. EOS – The Enterprise OS system. Telnet xxx.xxx.xxx.xxx – The Telnet session IP address. LoadConfig – The UI LoadConfig command. This source is visible only on locally logged messages. xxx.xxx.xxx.xxx – The IP address of the SNMP management station that initiated an SNMP SET request. This source is visible only on locally logged messages.
Text	A description of the event.

For more information on syslog messages, see [Syslog Messages on page 997](#).

For more information on FRF.17, IPSec, and IKE syslog messages, see [FRF.17, IPSec, and IKE Error Logging on page 516](#).

36.7

Auditlog Files Viewing

To display the contents of the local audit buffer, you must use:

```
SHoW -AuditLog LocalLog
```

To display the contents of a certain log file, you must use:

```
SHoW -AuditLog LocalLog <logfile_name>
```

where <logfile_name> is the name of the log file.

For example, to view the contents of log file Agn, you must use:

```
SHoW -AuditLog LocalLog AUD006.LOG
```

A new log file is created for each reboot, with the AUDxxx.LOG file name, where xxx ranges from 001 to 999. The file name number is incremented from the number of the most recent existing log file at the time that the new file is created. For example, the file created for the first reboot (that is the second bootup after initial powerup boot) when no other log files exist is AUD001.LOG and the file created for the second reboot is AUD002.LOG. A log file list is displayed with the most recent entry at the top of the display, and the oldest entry at the bottom. The contents and syntax of each entry is the same as for the syslog message so that you can compare the two easily.

To display the log file created from the most recent boot, if this file is available, you must use:

```
SHoW -AuditLog LocalLog LastFile
```

A named log file is a text file containing a list of null-terminated ASCII strings with a maximum length of 160 characters per entry. For EOS devices, the maximum size of a log file can be 10240 bytes, or 64 entries (with a maximum of 160 characters each). When a device is rebooted, the in-memory image of the display list is dumped into a named log file in an AUDLOG subdirectory, in the folder for boot configuration files storage. The storage location can be either drive A or drive B depending on the location of the configuration files.

To display the current contents of the display buffer in memory, you must use:

```
SHoW -AuditLog LocalLog Memory
```

36.8

Configuring Remote Logging of Syslog Messages

The AuditLog Service uses the Syslog logging service provided by the syslogd daemon available on most UNIX systems. To use AuditLog, the network must support the IP/ UDP protocol between the EOS device and the network management station. Because delivery of UDP messages is not guaranteed, some Syslog messages may be lost due to network conditions between the EOS device and the network management station.

Before using the AuditLog Service, you must configure the network management station to receive log messages from the EOS device.



NOTE: Various Syslog managers may be configured differently. You must consult your network management station manual for the configuration parameters and settings specific to your server.

Procedure:

1. Configure the network station for AuditLog by performing the following actions:

- a. Log on to your network management station as **root**.
- b. Create an empty log file to receive the Syslog messages.
This must be the same file you specified in `syslog.conf`.

Step example:

```
cat /dev/null > /var/log/auditlog
```

- c. Restart the Syslog daemon by entering:

```
kill -1 'cat /etc/syslog.pid'
```

2. Configure logging to remote syslog servers by performing the following actions:

- a. Configure the first (initial) syslog server by entering:

```
SETDefault -AuditLog LogServerAddr = <ip_address>
```

The initial syslog server uses the local facility level and default behavior that are specified by the `LocalFacility` and `DefAction` parameters. You can specify these parameters by using:

```
SETDefault -AuditLog LocalFacility = <0-7>
```

```
SETDefault -AuditLog DefAction = Include | Exclude
```

The `LocalFacility` parameter specifies the Syslog facility (LOCAL0 (0) through LOCAL7 (7)) to which log messages are directed. This parameter allows you to group all the messages from the EOS devices into one file on the syslog server. The default value is 7.

The `DefAction` parameter configures the default action for logging messages to the Syslog.

`Include` indicates that messages are logged to the Syslog. `Exclude` indicates that messages are not logged to the Syslog.

- b. Specify if the AuditLog service logs events by entering:

```
SETDefault -AuditLog CONTrol = [COnfig | NoCOnfig], [MESSAGES | NoMessages],  
[SEcurity | NoSEcurity]
```

Event categories:

- COnfig – Commands and parameters that successfully changed configurations.
- MESSAGES – System service messages and dial history messages (under the PORT Service).
- SEcurity – Messages concerning failed logon attempts, failed set privilege, or invalid SNMP community strings.



NOTE: In order for a log to be generated, you must use the CONTrol parameter to enable at least one of the three log message categories (set CONTrol to COnfig, MESSAGES, or SEcurity).

Each log message category is assigned a preset syslog priority level as follows:

- COnfig – LogWarning
- MESSAGES – LogNotice
- SEcurity – LogAlert

c. Optional: Add additional syslog servers by entering:

```
ADD -AuditLog LogServerAddr <IP_address> [<local_facility> (0-7)]  
[<default_action> Exclude | Include]
```

You can add up to four additional syslog servers.

You can configure the <local_facility> and <default_action> values to be different than those you configured for the initial syslog server.

d. To delete a logserver, enter:

```
DELeTe -AuditLog LogServerAddr <IP_address> | All
```

e. To display configured logservers, enter:

```
SHow -AuditLog LogServerAddr
```

A display similar to the following appears:

Log Server	Facility	DefAction
-----	-----	-----
100.100.1.1	7	Include (initial)
100.100.1.3	3	Exclude
100.100.1.4	2	Include

36.9

AuditLog Filters Management

You can configure the AuditLog Service to filter based on syslog server address, service, severity, facility, and/or message identifier.

When a filter is added, it takes effect immediately. AuditLog filter rules are executed in the order of their <filterid>. The filter with the lowest <filterid> is applied first, then the filter with the next lowest <filterid>, and so on.

For example, if filter 5 is an include filter, and filter 6 is an exclude filter, the message is sent.

Defining Filters by Using the ADD LogFILTER Command

To define a filter, you must use:

```
ADD -AuditLog LogFilter <filterid> Include | Exclude [Logserver=<log_server_addr_list>]
[Service=<service_list>] [Facility=<facility_list>] [SEverity=<severity_list>]
[Message=<message_id_list>]
```

For detailed information on rule syntax and corresponding values, see “AuditLog Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Displaying Filters

To display all the filters that are currently defined, you must use:

```
SHOW -AuditLog LogFilter
```

Deleting Filters

You must delete filters individually from the system.

To delete a filter, you must use:

```
DELEte -AuditLog LogFILter <filterid>
```

The filter is deleted immediately.

To delete all filters, you must use:

```
DELEte -AuditLog LogFILter ALL
```

AuditLog Filter Examples

The following five examples of AuditLog filtering are provided:

- Logging all events of severity level 5 and above to the default server.
- Suppressing all ISDN up/down events when severity level 3 and above are sent to the default server.
- Sending all VPN-related Syslog messages to server 100.100.1.2 and all other Syslog messages to all servers.
- Suppressing sending all Syslog messages resulting from UI commands to the default server.
- Sending events with severity level 4 to server 100.100.1.1 and events with severity level 5 to server 100.100.1.2

Logging All Events of Severity Level 5 (Notification) and Above to the Default Server

In this example, all events of severity level 5 (Notification) and above are sent to the default syslog server. Events with severity levels lower than 4 are not sent to any syslog servers.

To configure this example, you must use the following commands:

```
SETDefault -AuditLog CONTrol = (COnfig,Messages,SEcurity)
SETDefault -AuditLog LogServerAddr = 100.100.1.1
SETDefault -AuditLog DefaultAction = Exclude
ADD -AuditLog LogFilter 1 Include Logserver=100.100.1.1 Severity=0-5
```

Suppressing All ISDN UP/DOWN Events When Severity Level 3 (Error) and Above Are Sent to the Default Server

In this example, ISDN UP/DOWN events are suppressed. Other events with severity level 3 (Error) and above are sent to the default server. Events with severity levels less than 3 are not sent.

To configure this example, you must use the following commands:

```
SETDefault -AuditLog CONTrol = (COnfig,Messages,SEcurity)
SETDefault -AuditLog LogServerAddr = 100.100.1.1
ADD -AuditLog LogFilter 1 Include Logserver 100.100.1.1 Severity 0-3
SETDefault -AuditLog LogFilter 2 Exclude Logserver 100.100.1.1 Message 1309,1310
```

Sending All VPN-Related Syslog Messages to Server 100.100.1.2 and All Other Syslog Messages to All Servers

In this example, VPN-related events are sent to one syslog server, while all other events are sent to a different syslog server. VPN-related events are those events sent by the RAS services.

The filter causes messages from RAS to be sent to 100.100.1.2 because of the default action, but blocked from being sent to server 100.100.1.1 because of the action specified by the filter `Exclude`. All other messages are handled by the default action, which is to send messages to both syslog servers.

To configure this example, you must use the following commands:

```
SETDefault -AuditLog CONTrol = (COnfig,Messages,SEcurity)
SETDefault -AuditLog DefAction = Include
SETDefault -AuditLog LogServerAddr = 100.100.1.1
ADD -AuditLog LogServerAddr = 100.100.1.2
SETDefault -AuditLog LogFILter 1 Exclude Logserver 100.100.1.1 Service RAS
```

Suppress Sending All Syslog Messages Resulting from UI Commands to the Default Server

In this example, events generated by entering UI commands are not sent to the default syslog server. All other types of events are sent. Events generated by entering UI commands have an identifier of 9999.

To configure this example, you must use the following commands:

```
SETDefault -AuditLog CONTrol = (COnfig,Messages,SEcurity)
SETDefault -AuditLog LogServerAddr = 100.100.1.1
ADD -AuditLog LogFILter 1 Exclude Message=9999
```

Send Events with Severity Level 4 to Server 100.100.1.1 and Events with Severity Level 5 to Server 100.100.1.2

In this example, events with severity level 4 are sent to one syslog server and events with severity level 5 are sent to a different syslog server. All other types of events are not sent. There is no default log server.

To configure this example, you must use the following commands:

```
SETDefault -AuditLog CONTrol = (COnfig,Messages,SEcurity)
SETDefault -AuditLog DefaultAction = Exclude
ADD -AuditLog LogServerAddr = 100.100.1.1
ADD -AuditLog LogServerAddr = 100.100.1.2
ADD -AuditLog LogFILter 1 Include Logserver=100.100.1.1 SEVerity=4
ADD -AuditLog LogFILter 2 Include Logserver=100.100.1.2 SEVerity=5
```

36.10

Syslog Message Throttling

EOS supports a syslog message throttling feature that allows you configure the maximum number of syslog messages that the EOS device is allowed to transmit (`SysLogBurstThr`) within a specified throttle time period (`SysLogThrTime`).

Before transmitting a syslog message with the throttling feature enabled, the EOS software checks if the configured throttle time elapsed since the last message was logged for transmission.

- If so, the software logs the message for transmission.
In addition, if more than one throttle time period passed since the last message was logged for transmission, the software accumulates “credits” to send extra messages. These credits are incremented at a rate of one credit for each throttle time period that elapsed, up to the configured maximum number of syslog messages per throttle time period.

- If not, the software checks to see if it accumulated any credits to send extra syslog messages.
 - If so, the software logs the message for transmission and decreases the accumulated credit count by 1.
 - If not, the software discards the message.



NOTE: If bringing up the device in compliance with Common Criteria, do not use Syslog Messages Throttling. Confirm that this function is disabled by the `SHow -AuditLog SysLogThrottle` command.

36.11

SNMP Event Notification Traps

The AuditLog Service provides notification for tracking and management of EOS device configuration changes. Notification is based on SNMP enterprise-specific trap messages. You can generate traps independently, even if the audit log capability was disabled.

The following types of events can trigger a change or notification trap:

- Configuration Change Trap (Numeric Value 101)
This trap tracks changes originating from the execution of the `SET`, `SETDefault`, `ADD`, and `DELeTe` commands.
- User Authentication Trap (Numeric Value 102)
This trap captures failed logins.
- File Operation Command (Numeric Value 103)
This trap captures file operations from `COPy`, `RemoveFile`, `RemoveDir`, `ReName`, `MakeDir`, `GET`, and `PUT` commands.

Audit log notification requires no specific configuration or parameters, but it is associated with the configuration of the SNMP Service trap generation. To configure SNMP for notification traps, see [Configuring SNMP for Notification Traps on page 698](#)

36.11.1

Configuring SNMP for Notification Traps

Procedure:

1. Enable SNMP trap generation by entering:

```
SETDefault -SNMP CONTRrol = Trap
```

2. Add an SNMP community for your network management station by entering:

```
ADD -SNMP COMMunity <"name"> TRiv RW ALl
```

The value `ALl` specifies that all SNMP traps are enabled. The SNMP trap types are general, authentication, and enterprise-specific. The change traps are enterprise-specific. You can also specify an SNMP trap profile, and use the `AUDitLog` keyword. For more information, see "SNMP Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

3. Specify a destination for traps by entering:

```
ADD -SNMP MANager <community> <IP_Address>[ <mask>]
```

36.12

Root and Admin Passwords Changes

EOS supports the following privileged local logon accounts:

root

Associated with the Network Manager privilege level, which enables all software commands. To log on as **root**, you must enter the Network Manager password.

admin

Associated with the user privilege level, which enables only a subset of the software commands. To log on as **admin**, you must enter the user password.

You can set the Network Manager (**root**) and user (**admin**) passwords by using:

- The SYS service NMPassWord and UMPassWord parameters.
- The SysPassWord command menu.
- The PassWord command, when logged on as **root** or **admin**.



NOTE: To change the privilege level without logging off, you must use:

```
SET PRiVilege = User | NetMgr.
```

If you are logged on at the User privilege level and you issue the SET PRiVilege = NetMgr command to change the privilege level to Network Manager, you are prompted to enter the Network Manager password. You must enter the Network Manager password in order to raise your privilege level.

The default password for both the Network Manager (**root**) and the user (**admin**) privilege is a null string, which you generate by pressing the RETURN key. To set the passwords the first time, you must execute the SETDefault command with an empty string for the old password by using:

```
SETDefault -SYS NMPassWord = " " "<new_password>" "<new_password>"
```

```
SETDefault -SYS UMPassWord = " " "<new_password>" "<new_password>"
```

If you want to a password, you must log on as **root** and define the Network Manager password before you define the user password.

To change both the Network Manager and user passwords, you must use:

```
SysPassWord
```

You must follow the displayed menu to set the Network Manager password and then the user password.



CAUTION: You should set both the Network Manager and the user password. If you do not set the Network Manager password, any unauthorized user can log on with Network Manager privilege. If you do not set the user password, any unauthorized user can log on with user privilege.



IMPORTANT: For security purposes, all default passwords have to be changed prior to operational use. Those include both **root** and **admin** user passwords.

36.12.1

Password Requirements

EOS observes the following password requirements for local logon accounts:

- Passwords are case-sensitive and must be at least 7 characters long. You can configure the minimum required password length by using:
 - The SYS service PwMinLength parameter (if the AC service PWComplexity parameter is set to Standard (the default)).
 - The AC service PWMInchar parameter (if the AC service PWComplexity parameter is set to Enhanced).

- The maximum password length varies, depending on whether or not the account is configured for RADIUS authentication.
 - If the account is configured for RADIUS authentication (the AC service `RESolutionOrder` parameter is set to `Radius` or `Radius Local`), the maximum password length is 31 characters.
 - If the account is configured for local authentication (the AC service `RESolutionOrder` parameter is set to `Local` or `None`), the maximum password length is 128 characters.
- Valid characters are limited to ASCII codes 32 through 126 (alphanumeric characters plus the following special characters):
`! % & ( ) * + : ; < + > ?`
- The string of eleven asterisks (`*****`) is not allowed as a Network Manager (**root**) or User (**admin**) password. This string is reserved for use as a non-operational value when passwords are captured by using the ASCII capture feature.



NOTE: If the Network Manager (**root**) or User (**admin**) password is set to NULL, it is allowed to remain at that value until the password is changed. Once the Network Manager (**root**) or User (**admin**) password is changed from NULL, however, it must comply with the minimum password length and password complexity requirements and cannot be set back to NULL except from the boot monitor or by issuing the `ZEROize` command. For details about password complexity requirements, see [Password Complexity on page 700](#).

36.12.2

Password Complexity

The extent to which you can configure password complexity depends on the setting of the AC service `PWComplexity` parameter:



NOTE: You must be logged on as **root** in order to set password complexity parameters.

When `PWComplexity` is set to `Standard` (default), the following parameters allow you to configure password complexity:

- The SYS service `PwMinLength` parameter specifies the minimum password length (7-12 characters). If a user attempts to set a password that is shorter than the configured minimum length, the system returns an error message.
The `PwMinLength` parameter functions in the same way as the AC service `PWMinchar` parameter, which takes effect when the `PWComplexity` parameter is set to `Enhanced`. The `PwMinLength` parameter allows you to specify a password length of 7-12 characters, while the AC service `PWMinchar` parameter allows you to specify a password length of 7-128 characters.
- The AC service `EXPIrationTimer` parameter sets a password expiration timer.
- The AC service `LockOutTimer` parameter specifies the length of time that the user is locked out after entering the incorrect password the number of times specified by the `TotalPwTrial` parameter.
- The AC service `TotalPwTrial` parameter specifies the number of unsuccessful logon attempts the user is allowed before locking out from the device command line interface for the length of time specified by the AC service `LockOutTimer` parameter.

For example, when `PWComplexity` is set to `Standard`, you must use the following commands to specify that local user logon account passwords must be at least 10 characters long and expire every 30 days, and that users are locked out for 60 minutes after 3 failed logon attempts:

```
SETDefault -SYS PwMinLength = 10
```

```
SETDefault -AC EXPIrationTimer = 30
```

```
SETDefault -AC LockOutTimer = 60
```



```
SETDefault -AC TotalPwTrial = 3
```

When **PWComplexity** is set to **Enhanced**, the following parameters allow you to configure password complexity:

- The AC service **PWMinchar** parameter specifies the minimum password length (7-128 characters). If a user attempts to set a password that is shorter than the configured minimum length, the system returns an error message.
The **PWMinchar** parameter functions in the same way as the SYS service **PwMinLength** parameter, which takes effect when the **PWComplexity** parameter is set to **Standard**. The **PWMinchar** parameter allows you to specify a password length of 7-128 characters, while the SYS service **PwMinLength** parameter allows you to specify a password length of 7-128 characters.
- The AC service **EXPIRATIONTimer** parameter sets a password expiration timer.
- The AC service **TotalPwTrial** parameter specifies the number of unsuccessful logon attempts the user is allowed before locking out from the device command line interface for the length of time specified by the AC service **LockOutTimer** parameter.
- The AC service **LockOutTimer** parameter specifies the length of time that the user is locked out after entering the incorrect password the number of times specified by the **TotalPwTrial** parameter.
- The AC service **PWNumber** parameter specifies the minimum number of digits required in the password.
- The AC service **PWLowercase** parameter specifies the minimum number of lowercase letters required in the password.
- The AC service **PWUppercase** parameter specifies the minimum number of uppercase letters required in the password.
- The AC service **PWSpecial** parameter specifies the minimum number of special characters required in the password.

For example, when **PWComplexity** is set to **Enhanced**, you must use the following commands to specify that local user logon account passwords must be at least 30 characters long, must include at least one uppercase and one lowercase letter, one number, and one special character, expire every 30 days, and that users are locked out for 60 minutes after 3 failed logon attempts:

```
SETDefault -AC PWComplexity = Enhanced
SETDefault -SYS PwMinLength = 30
SETDefault -AC PWUppercase = 1
SETDefault -AC PWLowercase = 1
SETDefault -AC PWNumber = 1
SETDefault -AC PWSpecial = 1
SETDefault -AC EXPIRATIONTimer = 30
SETDefault -AC LockOutTimer = 60
SETDefault -AC TotalPwTrial = 3
```

36.12.3

Password Clearing

You cannot clear the passwords associated with the **root** and **admin** accounts from EOS. These passwords can only be cleared from the boot monitor.

36.13

Executive Logon Timeout

You can log on to the EOS UI in the following ways:

- Console – EOS UI access through a console connection (the console port of the device is connected directly to a PC or workstation).
- Telnet – EOS UI access through a Telnet session.
- SSH – EOS UI access through an SSH session.

The executive logon timeout feature allows you to configure timeout mechanisms (by which idle users are logged out and have to log on again to regain access) for both access methods.

You configure the timeout values by using the SYS service `NetAccessTimer` parameter:

```
SET -SYS NetAccessTimer = {ConsoleTimeout <value> TelnetTimeout <value> SshTimeout <value>} | All <value>
```

```
SETDefault -SYS NetAccessTimer = {ConsoleTimeout <value> TelnetTimeout <value> SshTimeout <value>} | All <value>
```

The `SET` command allows you to set a runtime value for `NetAccessTimer`, thereby allowing you to change the `ConsoleTimeout`, `TelnetTimeout`, or `SshTimeout` value while you are connected. The value configured with the `SET` command remains in effect until the current session is terminated (through logging out, timeout, or device reboot). If you use the `SET` command to set `NetAccessTimer`, your setting is not reflected in a `SHoWDefault` command; it is only reflected in a `SHoW` command.

The `SETDefault` command allows you to set a default value for `NetAccessTimer`. The value configured with the `SETDefault` command takes effect in subsequent sessions and remains in effect until a `SET` or `SETDefault` command (issued either manually or within a `boot.cfg` file) changes it.



NOTE: To change `NetAccessTimer` for both the current and subsequent sessions, you must configure the desired `NetAccessTimer` value with both the `SET` command (to change the current session) and the `SETDefault` command (to change subsequent sessions).

The timers are refreshed as follows:

- The `ConsoleTimeout` timer is refreshed only if a user logged on through a console session sends a command by pressing `ENTER` or if a carriage return is sent to the EOS UI.
- The `TelnetTimeout` timer is refreshed any time a user logged on through a Telnet session presses any key.
- The `SshTimeout` timer is refreshed any time a user logged on through an SSH session presses any key.



NOTE: `SshTimeout` and `TelnetTimeout` are synonyms affecting both SSH and Telnet (incoming) sessions. If you modify one of these values, the other is automatically modified to the same value.

If the configured timeout period expires with no commands issued, keys pressed, or HTTP requests sent, the idle user is automatically logged out and must log on again to access the EOS UI.

If you set `NetAccessTimer` to 0, EOS UI sessions accessed with the specified method never time out. This setting is useful if you want to prevent Telnet-based session timeouts from prematurely ending long-running macro command sequences. Since commands issued from within a macro do not refresh the `TelnetTimeout` timer, a Telnet connection may time out during execution of a long-running macro, even though commands are issued from within the macro. To prevent this, you must set `NetAccessTimer` to never time out Telnet sessions, as follows:

```
SETDefault -SYS NetAccessTimer = tt 0
```



NOTE: You must be logged on at the `NetMgr` privilege level to set any of the `NetAccessTimer` timeout values to 0.

For details about the `NetAccessTimer` parameter, see “SYS Service Commands” in the *Enterprise OS Software Version 18.2 Reference Guide*.

36.14

System Remote Access

You can access your system remotely to perform network management operations. These operations can be completed through a remote station, such as a Sun workstation, which does not need to be physically connected to the console port of your system. Remote access is accomplished by using the proprietary `REMOte` command from another device or by using the Telnet Protocol from a Telnet client.

Depending on your reasons for accessing the system remotely, you may want to use `TELnet` instead of `REMOte`. These commands differ in the following ways:

- The `REMOte` command provides access to a subset of the EOS device commands, is UDP-based, and can truncate long displays.
- The `TELnet` command provides access to all EOS device commands and is TCP-based.

After the connection was established through remote mode or the Telnet Protocol, you can change your privilege level to Network Manager, provide the correct password, and perform network management operations or configuration procedures.

You can also prevent unauthorized users from making remote connections to your system by configuring the `-SYS NetAccess` parameter, and you can restrict remote access to specific users by configuring the `-SYS RemoteManager` parameter.

The `REMOte` command allows you to execute commands on your EOS device from a remote terminal. After you enter the `REMOte` command followed by an IP address or name, you are in remote mode. Remote mode is indicated on your remote terminal by the appearance of the remote prompt (`Remote:`). In remote mode, all commands entered affect your device until you exit remote mode. To exit remote mode, you must press the **Break** key and enter `TELnet` to return to Telnet mode.


 **NOTE:** The `REMOte` command is not subject to a configured password. To restrict access to the system, you must disable `REMOte` and use the `TELnet` command instead.

Some EOS device commands cannot be used in remote mode. For a list of these commands, see “Commands” in the *Enterprise OS Software Version 18.2 Reference Guide*.

36.14.1

Displaying the Software Version with the REMote Command

Figure 185: Accessing the System Remotely



The Telnet Protocol also can be used on a remote terminal to access an EOS device. In this situation, your system functions as the Telnet server (the destination), and the remote terminal functions as the Telnet client (the initiator).

To access a device called "gateway," on the remote terminal (RBCS or Sun workstation), you must use: `TELnet gateway`. The user level command prompt of the device named "gateway" appears on the remote terminal. After you change the privilege level to Network Manager and enter the password, you can perform network management procedures. After your management activities are complete, you must use the `LISten` command to disconnect the session and place the port in listen mode.

Procedure:

1. On the remote RBCS terminal, enter:

```
REMote 129.3.4.2
```

2. At the remote prompt, enter:

```
SHow -SYS VERSion
```

The version information of the software running on the system is displayed.

3. To return to the command prompt of your remote terminal, press the **BREAK** key.

You can also enter the `REMote` command, followed by the address of the EOS device, and then followed by a command to be executed.

For example, the following command displays the IP Routing Table of the system that has the address of 129.3.2.4:

```
REMote 129.3.2.4 SHow -IP AllRoutes
```

36.14.2

Preventing Remote Access

By default, you can access your system remotely by using the SSH protocol. You can regulate access from remote devices by using the `SETDefault -SYS NetAccess` command.



CAUTION: The software allows the EOS device to be disabled without giving any warning messages. After assigning `NoRemote` or `NoTelnet`, you can no longer access the system parameters to perform software configuration. You must boot the system from a software image that contains an enabled `NetAccess` parameter before you can regain access.

Procedure:

- To disable access with the REMote command, enter:
`SETDefault -SYS NetAccess = NoRemote`
- To disable access with the Telnet Protocol, enter:
`SETDefault -SYS NetAccess = NoTelnet`
- To disable access with SSH, enter:
`SETDefault -SYS NetAccess = NoSsh`

36.14.3

Restricting Remote Access

You can use this procedure to specify that only certain devices can access your system.

You can configure a maximum of three RemoteManager addresses. For additional information on these commands, see “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Prerequisites:

Obtain the Network Manager privilege.

Procedure:

1. Remove remote access to your system from all devices by entering:
`DELeTe -SYS RemoteManager *.*.*.*`
2. Specify the IP address of the device that only can access your system by entering:
`ADD -SYS RemoteManager <IP_address>`

Step example:

To allow only the device with the address 129.98.96.99 to access your system, enter:

```
ADD -SYS RemoteManager 129.98.96.99
```

36.14.4

Restricting Telnet Access

You can use this procedure to specify that only certain devices can access your system through the Telnet Protocol.

You can configure a maximum of 64 `TelnetManager` addresses. For more information, see the description of the `TelnetManager` parameter in “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: The `TelnetManager` parameter and the `SSHManager` parameter modify the same list; therefore, modifying one list affects the other.

Prerequisites:

Obtain the Network Manager privilege.

Procedure:

1. Remove Telnet access to your system from all devices by entering:
2. Specify that only the devices with specific IP address(es) can access your system via Telnet by entering:

```
DELeTe -SYS TelnetManager *.*.*.*
```

```
ADD -SYS TelnetManager <IP_address>
```

Step example:

To specify that only the devices with the addresses 129.213.48.18 and 139.87.180.* can access your system through Telnet, enter:

```
ADD -SYS TelnetManager 129.213.48.18
```

```
ADD -SYS TelnetManager 139.87.180.*
```

36.14.5

Restricting SSH Access

You can use this procedure to specify that only certain devices can access your system through SSH.

You can configure a maximum of 64 `SSHManager` addresses. For more information, see the description of the `SSHManager` parameter in “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: The `SSHManager` parameter and the `TelnetManager` parameter modify the same list; therefore, modifying one list affects the other.

Prerequisites:

Obtain the Network Manager privilege.

Procedure:

1. Remove SSH access to your system from all devices by entering:
2. Specify that only the devices with specific IP address(es) can access your system through SSH by entering:

```
DELeTe -SYS SSHManager *.*.*.*
```

```
ADD -SYS SSHManager <IP address>
```

Step example:

To specify that only the devices with the addresses 129.213.48.18 and 139.87.180.* can access your system through SSH, enter:

```
ADD -SYS SSHManager 129.213.48.18  
ADD -SYS SSHManager 139.87.180.*
```

36.15

Resynchronization Feature for Encryption Devices

You can connect KG-81/KG-94 encryption devices to WAN ports (for example, ports 3 and 4, which have an RS-449 connector). The KG-81 and KG-94 encryptors are available for use by U.S. government installations only.

After you establish the connection between the EOS devices and one of these encryption devices, you must enable the resynchronization feature on the EOS device by using:

```
SETDefault !<path> -PATH CONTrol = Crypto
```

For example, to enable the resynchronization feature on path 3, you must use:

```
SETDefault !3 -PATH CONTrol = Crypto
```

When this feature is enabled, if a path goes down, a pulse that attempts to resynchronize the devices is generated on one of the signal lines (DTR) on the RS-449 connector. Each pulse is 10,000 microseconds long and is sent approximately every 10 seconds. Pulses are sent until resynchronization occurs.

You should disable the resynchronization feature only when you disconnect the EOS device from the encryption device. For example, to disable the resynchronization feature on path 3, you must use:

```
SETDefault !3 -PATH CONTrol = NoCrypto
```

To display the current settings of the CONTrol parameter in the PATH Service, you must use:

```
SHow -PATH CONTrol
```



NOTE: This feature is supported only on ports running PPP.

Chapter 37

Remote Polling Protocol Configuration

This chapter describes how to configure the Remote Polling Protocol (REMP). REMP monitors the reachability of target network devices and collects network performance data for a list of configured remote targets by periodically polling the targets for an echo response.

37.1

REMP Configuration

You can use REMP anytime you want to monitor the reachability of any target from an EOS device. When a target cannot be reached, you can be notified by an SNMP trap (if configured) or REMP.

Targets can be added to the remote-polling list in one of two ways: statically or dynamically.

Static Targets

You configure static targets manually by using SNMP or the `LoadConfig` command. Static targets are saved to the remote polling configuration file, and are preserved across reboots. This group of targets is always polled.

In this example, a static target is added with an IP address of 129.213.10.1 and default settings. A static target with IP address of 129.213.20.2 is added with specific settings (ATtempts=8, Mode=RetryOnFailure, Wait=20 seconds).

To set up this example with the command-line interface, for a target with an IP address of 129.213.10.1, you must enter:

```
ADD -REMPolling RemoteTarget 129.213.10.1
ADD -REMPolling RemoteTarget 129.213.20.2 A 8 M rof W 20
```

To set up this example with SNMP for a target with an IP address of 129.213.10.1, you must:

1. Send an SNMP “createRequest” command to create a `remPollEntry` for 129.213.10.1.
2. Send an SNMP “valid” command to set the entry just created for 129.213.10.1 to valid.

To set up this example with SNMP for a target with an IP address of 129.213.20.2, you must:

1. Send an SNMP “createRequest” command to create a `remPollEntry` for 129.213.20.2.
2. Send SNMP set request to set the following objects for 129.213.20.2:
`remPollMode` to `RetryOnFailure`
`remPollAttempts` to 8
`remPollTimeout` to 2000
3. Send an SNMP “valid” command to set the entry just created for 129.213.20.2 to validate.

Dynamic Targets

Dynamic targets are added to the REMP Service automatically. Dynamic targets are not saved to the REMP configuration files, and they are not preserved across reboots.

The following three dynamic target groups are supported:

Remote Authentication Service (RAS)

RAS group targets are obtained when a RAS session is established. Each RAS entry is removed from the polling list when the RAS session is torn down.

Virtual Leased Line (VLL)

VLL group targets are obtained when you configure a VLL port by using:

```
ADD -l2t VLeasedLine <Dest_IP_Address> {Protocol = <PPTP / I2TP>}
```

VLL targets are removed from the polling list when the VLL configuration is deleted.

Tunnel Peer (TP)

TP group targets are obtained when you configure a port by using one of the following commands:

- ADD <!port> -Port DNL "ip address"
- ADD <!port> -Port VirtualPort IPIP remote_ipaddr
- ADD <!port> -Port VirtualPort IPIP p2mp
followed by
ADD -ip addr <ip1> IPIP <ip2>

TP targets are removed from the polling list when the port configuration is deleted.

37.2

Target Group Priority

REMP support a maximum of 100 targets in the polling list. Because of this, a priority is used to determine which target(s) are added to the list when the list is full.

By default, the priority order is as follows:

- Static group
- RAS group
- TP group
- VLL group

For example, if the polling list is at capacity with 55 static targets, 15 RAS targets, and 30 VLL targets, and you want to add an additional static target, one of the VLL targets is deleted from the list to make room for the new static target.

In the case when a lower priority target needs to be deleted to make room for a higher priority target, a target in the lowest priority group that was added first (first-in-first-out) is removed.

You can assign more than one group the same priority. In this case, the FIFO deletion rule is also applied.

To configure the group priority, you must use:

```
SETDefault -REMPolling PRIOrity = <priority_level> (1-4) <Static / RAS / TP / VLL>
```

The REM Polling Service creates a log file which you can configure to be either ASCII or binary format. The AuditLog Service controls the distribution of messages to the Syslog server(s).

Chapter 38

Virtual Router Redundancy Protocol Configuration

This chapter describes how to configure the Virtual Router Redundancy Protocol (VRRP) on EOS devices. VRRP allows a backup router to immediately take the place of a failed master router using the same IP and MAC addresses. A master router is a default router explicitly specified by end hosts.

VRRP is a protocol between routers that allows backup routers to monitor the status of a master router. When the master router fails, the backup router can take over the function of the master router. The new master router keeps the IP and MAC address of the original master, so that hosts that are configured with a single default gateway do not have their network connectivity disrupted if the gateway fails.

With VRRP enabled, the master router sends out regular VRRP packets to indicate that it is alive. If the VRRP packets stop, the backup router adopts the IP and MAC address of the master, in addition to its own IP and MAC addresses. If you have more than one backup router, the router with the highest priority becomes the master router.

Each router running VRRP is either the owner or the backup router for a VRID. There can be only one owner for each VRID. The owner router owns the IP address, configured in the IP service, that is used by the backup router if the owner fails. The owner router is always the master router if it is active, because it has a higher priority than any backup router.

VRRP supports the following media:

- Ethernet
- Fast Ethernet

VRRP can be used with IPv4 only.

MAC Address

The MAC address used by the master router is 00-00-5E-00-01-*<vrid>*. Each VRID has one MAC address associated with it, so a router that has multiple VRIDs has a different MAC address for each VRRP interface.

38.1

Configuring the VRRP

Prerequisites:

Ensure that:

- You are logged on to the system with Network Manager privilege.
- Ports and paths of your EOS device are set up. See [Basic Ports and Paths Configuration on page 50](#).
- IP routing is configured on each EOS device. See [IP Routing Configuration on page 119](#).

Process:

1. Enable the owner router. See [Owner Router Enabling on page 711](#).
2. Enable the backup routers. See [Backup Routers Enabling on page 711](#).
3. Set up priorities for multiple backup routers. See [Priority Settings for Multiple Backup Routers on page 712](#).
4. Enable the VRRP. See [VRRP Enabling on page 714](#).

5. Register the VRRP service. See [VRRP Registration on page 714](#).
6. To disable/delete the VRRP from the owner router. See [VRRP Disabling/Deleting on page 715](#).

38.1.1

Owner Router Enabling

The owner router is always the master router when it is active. You can assign a virtual router identifier (VRID) to the owner router by using:

```
ADD !<port> -VRRP VRid <vrid> <IP_address> [ , <IP_address>]
```

where:

<port> is the port on which you define VRRP. You must specify the port that is connected to the LAN.

<vrid> is any number between 1 and 255 that you want to use to identify the VRRP owner IP addresses. If you have more than one owner router on a LAN, you must use unique VRIDs.

<IP_address> is the IP address or addresses that you assigned to the port in the IP service. You can specify up to five addresses. This IP address is used by the backup router when the owner router goes down.

If you have multiple subnets on the port, you can add more than one VRID if you want a separate backup router for each subnet. You can add multiple VRIDs per router as long as each VRID is unique.

38.1.2

Backup Routers Enabling

You can assign a backup router to the owner router by using:

```
ADD !<portlist> -VRRP BackUp [AsOwner] <vrid> <ip_address> [ , <ip_address>]  
[ <MasterSysName>]
```

where:

<portlist> is the port on which you define VRRP. Specify the port that is connected to the LAN. The VRRP port must be an Ethernet or Fast Ethernet interface.

<vrid> is the VRID defined on the owner router. The *<vrid>* value must match on the owner and all backup routers.

<ip_address> is the owner router IP address or addresses. The address must match the owner router IP address exactly. The backup router's own IP address, configured in the IP service, must have the same subnet as the owner router.

<MasterSysName> is the SysName of the VRRP master router. If specified, this parameter is sent in the SNMP trap when the master router goes down or comes up.

You can use the BackUp parameter on the backup router only. A router can be configured as both an owner router and a backup router, so long as each vrid is configured for only one function.

The BackUp parameter identifies one or more IP addresses, the associated *<vrid>*, and the port where they should be backed up. The owning router uses the VRID parameter, not the BackUp parameter.

[AsOwner] is an optional Motorola Solutions extension to VRRP that specifies that when this router takes over the IP addresses it behaves as though it owns the IP address: it responds to Ping, allow TCP connections, and so on. This behavior is contrary to the RFC, but is useful for some applications.

Often in this type of application, there is no real owner of the IP address – it moves among the backup routers.

38.1.3

Priority Settings for Multiple Backup Routers

When the master router goes down, it stops sending VRRP packets, enabling the backup router to become the master router.

If you have more than one backup router, which router becomes the master is determined by the following settings:


- **Preempt** – This option, set in the `CONTROL` parameter, allows a router that has a higher priority than the current master to assume control as the master router. `NoPreempt` reduces network instability caused by master router changes, so only use `Preempt` if you must ensure the primary backup router.
- **HoldTime and skew time** – The `HOLDTime` parameter sets the amount of time before the backup router declares the master router to be down. The total time consists of the hold time plus the skew time. The skew time depends on the priority of the router, and is either 500 ms or 1 second. The router with the lowest [hold time + skew time] becomes master router first.
- **PRIOrity** – If more than one router has the same `Preempt` setting, the router with the highest priority becomes master. Also, the priority determines the skew time, either 500 ms or 1 second.

The owner router is **always** master when it is active, because it is set for `Preempt` and has the highest priority (255).

If you have more than one backup router, we recommend that you use the `PRIOrity` field to establish an order in which each backup router would become the master should the master router fail.

The following figure illustrates the process by which one router becomes master when there are multiple backups.

Figure 186: VRRP Priorities



*You cannot control which router becomes master if they have the same characteristics. Because of the timer on the NETBuilder bridge/router, the same [hold time + skew time] can vary by as much as 500 ms.

Setting the Hold Time

The highest priority backup router should have the lowest `HOLDtime`. Set the `HOLDtime` for the highest priority router, then set the `HOLDtime` to be one second longer for the next priority router. Add one second for each succeeding router, in order of priority.

You can set the amount of time before the backup router declares the master router to be down by using:

```
SETDefault !<port> -VRRP HOLDtime = <hold_time>(3 to 255 seconds) <vrid>
```

where `<hold_time>` is a number from 3 to 255 in seconds. `<hold_time>` should be at least three times the advertisement interval (see [VRRP Customization on page 715](#)). The default is 3.

If you want a backup router to become the master no matter what the `HOLDtime` is, you can specify `Preempt` in the `CONTROL` parameter and assign it the highest priority using the `PRIORITY` parameter.

Setting the Priority

If you have more than one `Preempt` router, the `PRIORITY` determines which router becomes master. The value of 255 is reserved for the router that owns the IP address associated with the virtual router. You can use the VRRP service `PRIORITY` parameter to configure the backup routers with any value from 1 through 254.

The `PRIOrity` also determines the length of the skew time, which in conjunction with the `HOLdTime` determines when the backup router can declare the master router down. The skew time is defined as the time to skew the master down interval, in seconds, calculated as: $(256 - \text{PRIOrity}) / 256$.

Thus if the `PRIOrity` is between 1 and 128, the skew time is 1 second. If the `PRIOrity` is between 129 and 254, the skew time is 500 ms.

If the `HOLdTime` is the same on two routers, but the priorities fall in different ranges, then the higher priority router becomes master. Because [hold time + skew time] on a router can vary as much as 500 ms, you cannot determine which router becomes master if you have more than one router within a range. Motorola Solutions recommends setting the `HOLdTime` itself at least one second apart on the different routers to guarantee priority.

You can set the priority of a backup router to any number between 1 and 254 by using:

```
SETDefault !<port> -VRRP PRIOrity = <priority_number> <vrid>
```

38.1.4

VRRP Enabling

You can enable VRRP on each router by using:

```
SETDefault !<port> -VRRP CONTrol = (Enable, [PreEmpt | NoPreEmpt]) <vrid>
```

`PreEmpt` allows a router that has a higher priority than the current master router to assume control as the master router. Because the owner router has the maximum priority, it is always the master router when it is active. You cannot set the owner router to `NoPreEmpt`. If you have more than one backup router, and you use `NoPreEmpt` (the default), the router assumes control only when the current master router fails, even if it has a higher priority. If you know that one backup router should always be the primary backup, you should specify `PreEmpt` in addition to giving it a high priority. If you do not have a strong preference, keep `NoPreEmpt` on all backup routers. `NoPreEmpt` reduces network instability caused by master router changes.

If the owner router goes down, and then comes up later, it becomes the master router again without waiting for the backup to fail.

38.1.5

VRRP Registration

Any service can register with the VRRP service. When VRRP discovers that a given router has a failure, it sends messages to the registered clients, which can act accordingly. Clients register for the VRRP service on a given port (allowed only if VRRP is enabled on that port). When the master router fails on that port, VRRP notifies all the registered clients and includes the IP address of the failed router in the notification. Similarly, when the master router comes back up, VRRP notifies all registered clients.

VRRP for IGMP

The VRRP for IGMP feature was designed to address a particular dual site router architecture scenario. It can also be used in configurations with more than one IGMP querier on a LAN so that if one IGMP querier goes down, the other router can take over as IGMP querier using VRRP convergence times.

You can use the MIP service `VRRPNotify` parameter to enable or disable IGMP registration with VRRP on a given port. When MIP receives a notification stating that a VRRP Master router is down, it elects itself querier and immediately begins to query, providing the following conditions are true:

- IGMP and VRRP Notify are enabled on that port (the MIP service `IgmpControl` and `VRRPNotify` parameters are set to `Enable`).
- The IP address of the current IGMP querier is the same as the IP address going down.

VRRP for PIM

The VRRP for PIM feature was designed to detect PIM failovers at faster convergence times. Use the PIM service `VRRPNotify` parameter to enable or disable PIM-SM registration with VRRP.

VRRP for BFD

The VRRP for BFD feature was designed to ensure that BFD sessions configured on a VRRP-controlled address start only when the EOS device is the VRRP master on that interface. Use the IP service `BFDTrigger` parameter to register BFD with VRRP to start and stop the BFD session.

38.1.6

Ping/Telnet Virtual Router IP (VIP)

If the master of the virtual router is not the owner router (the owner router is down), attempts to Ping or Telnet to the Virtual Router IP (VIP) get no response. This is the default behavior as specified in the VRRP specification. By adding the `AsOwner` keyword to the `VRRP BackUp` command, you can direct the backup router to treat the IP address as if it was its owner. For example, the backup router responds to pings, and can be used as the destination of a TCP connection for Telnet.

38.1.7

VRRP Disabling/Deleting

Before disabling/deleting a VRRP from the owner router, you must disable/delete the VRID from the backup routers. Failing to do so may cause the same IP address to be mapped to different MAC addresses on different nodes. This happens because the new master and the owner are both responding to the ARP request for the VIP. You may also see "Local-Col" displayed in response to a `SHOW -IP` address command. The "Local-Col" is cleared two hours after the problem is corrected (after the vrid backup routers are disabled/deleted).



NOTE: The System IP address of a router should not be used as a VRRP Virtual Router IP (VIP). Doing so may cause two routers to respond to the same IP address. When a request (such as SNMP) is sent to the VIP, the response can be returned from either of the routers.

38.2

VRRP Customization

To customize your VRRP configuration, you must set the interval between VRRP packets sent by the master router by using:

```
SETDefault !<port> -VRRP AdvertisementInt = <adv_time> (100 to 100000 milliseconds)
```

where `<adv_time>` is a number from 100 to 100000 milliseconds. The default is 1 second (10000 milliseconds). You must set the advertisement interval to the same value on all routers associated with a VRID.

Associated with the VRRP advertisement interval is the master down interval, which is calculated as follows:

master down interval = 3 * VRRP advertisement interval + skew time

where skew time is the time to skew the master down interval, in seconds, calculated as: $(256 - \text{Priority}) / 256$.

A lower skew time is preferred to minimize the amount of time required to learn that a master router went down.

38.3

VRRP Scenarios

The following VRRP scenarios can be used:


- Gateway to WAN
- Connecting two LANS
- Load sharing with redundancy

38.3.1

Configuring VRRP – Gateway to a WAN

The following figure shows a basic VRRP setup with one gateway router connecting a LAN to the Internet. If port 1 goes down, the backup router becomes master.

Figure 187: Owner Router with One Backup Router



Procedure:

1. On the owner router, perform the following actions:

- a. Create the VRID by entering:

```
ADD !1 -VRRP VRid 1 144.232.29.1
```

- b. Enable VRRP by entering:

```
SETDefault !1 -VRRP CONTrol = Enable 1
```

2. On the backup router, perform the following actions:

- a. Back up the VRID 1 by entering:

```
ADD !3 -VRRP BackUp 1 144.232.29.1
```

- b. Enable VRRP by entering:

```
SETDefault !3 -VRRP CONTrol = Enable 1
```


If the owner router fails, the backup router automatically takes over using the IP address 144.232.29.1 and the MAC address 00-00-5E-00-01-01.

38.3.2

Configuring VRRP – One Master Router Connecting Two LANs

The following figure shows two LANs connected by a master router, with one backup router.

Figure 188: Two VRIDs on One Router Connecting Two LANs



Procedure:

1. On the master router, perform the following actions:

a. Add VRID 1 to port 1 by entering:

```
ADD !1 -VRRP VRid 1 144.232.29.1
```

b. Add VRID 2 to port 2 by entering:

```
ADD !2 -VRRP VRid 2 139.215.65.3
```

c. Enable VRRP on both ports by entering:

```
SETDefault !1 -VRRP CONTrol = Enable 1
```

```
SETDefault !2 -VRRP CONTrol = Enable 2
```

2. On the backup router, perform the following actions:

a. Back up VRID 1 by entering:

```
ADD !3 -VRRP BackUp 1 144.232.29.1
```

b. Back up VRID 2 by entering:

```
ADD !5 -VRRP BackUp 2 139.215.65.3
```

c. Enable VRRP on both ports by entering the following commands:

```
SETDefault !3 -VRRP CONTrol = Enable 1
```

```
SETDefault !5 -VRRP CONTrol = Enable 2
```


If the master router fails, the backup router automatically takes over. Port 3 uses IP address 144.232.29.1 and MAC address 00-00-5E-00-01-01. Port 5 uses IP address 139.215.65.3 and MAC address 00-00-5E-00-01-02.

38.3.3

Configuring VRR – Load Sharing with Redundancy

The following figure shows two routers sharing the network traffic on subnet 144.232.29.x. Half of the hosts have Owner router VRID 1 configured as the default gateway and the other half have Owner router VRID 2 as the default gateway. Each owner router is configured to back up the other if one fails.

Figure 189: Load Sharing with Redundancy



Procedure:

1. On the owner router VRID 1, perform the following actions:
 - a. Add VRID 1 to port 1 by entering:
`ADD !1 -VRRP VRid 1 144.232.29.1`
 - b. Back up VRID 2 by entering:
`ADD !1 -VRRP BackUp 2 144.232.29.23`
 - c. Enable VRRP for each VRID by entering the following commands:
`SETDefault !1 -VRRP CONTrol = Enable 1`
`SETDefault !1 -VRRP CONTrol = Enable 2`
2. On the owner router VRID 2, perform the following actions:
 - a. Add VRID 2 to port 3 by entering:
`ADD !3 -VRRP VRid 2 144.232.29.23`
 - b. Back up VRID 1 by entering:
`ADD !3 -VRRP BackUp 1 144.232.29.1`
 - c. Enable VRRP for each VRID by entering the following commands:
`SETDefault !3 -VRRP CONTrol = Enable 2`
`SETDefault !3 -VRRP CONTrol = Enable 1`

If Owner router VRID 1 fails, Owner router VRID 2 automatically takes over. Port 3 uses IP address 144.232.29.1 and MAC address 00-00-5E-00-01-01, in addition to its own IP and MAC addresses, 144.232.29.23 and 00-00-5E-00-01-02.

If Owner router VRID 2 fails, Owner router VRID 1 automatically takes over. Port 1 uses IP address 144.232.29.23 and MAC address 00-00-5E-00-01-02, in addition to its own IP and MAC addresses, 144.232.29.1 and 00-00-5E-00-01-01.

38.4

VRRP Traps Enabling and Disabling

EOS supports two VRRP-specific traps. These traps are generated:

- When the VRRP master router goes down

- When the VRRP master router comes back up

You can enable and disable the VRRP SNMP traps by using the VRRP `SnmpTrap` parameter. For details, see “VRRP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. For details about the traps, see [Traps Supported on EOS Devices on page 871](#).



NOTE: If you specify the `SysName` of the master router with the `ADD -VRRP BackUp` command, this information is also included in the VRRP SNMP traps. For details, see the description of the `BackUp` parameter in “VRRP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Chapter 39

Scheduling and Event-Based Macro Execution

This chapter describes how to schedule repetitive events to occur on specific days or dates using the SCHEDuling Service and how to set up an automatic back-up for a port using Event-based Macros Execution (EBME).

SCHEDuling Service

The SCHEDuling Service allows you to schedule repetitive events to be executed on a specified calendar date each year or on a specified day of each week. Scheduled events can consist of commands or macros. Used with the dial-up feature, the scheduler allows you to perform many useful tasks, including:

- Matching traffic prioritization with work schedules.
- Updating configuration and test-booting off-hours.
- Allowing for time zone and work habit differences, and giving several remote sites scheduled access to a single WAN port at the central site.
- Creating a more secure Internet with remote sites having scheduled access and only the central site originating calls.
- Synchronizing the mail server and WAN link to control cost of electronic mail distribution.

The scheduler allows you to define up to 12 **daily schedules** and to assign each daily schedule to one or more calendar dates or days of the week. Each daily schedule contains one or more events, each of which consists of a time of day and a command or macro to be executed.

An **active schedule** is the combination of a calendar day or day of the week with a daily schedule. On any given day, only one active schedule can be in use.

If an active schedule exists for today's calendar date (for example, March 31) and an active schedule exists for today's day of the week (for example, Wednesday), the active schedule for the calendar date takes precedence.

EBME

EBME is a platform routing protocol and provides a backup link when the primary link between two sites fails. It also provides port loopback detection; EBME supports a port-level backup to a primary link failure and loopback detection and recovery.

You can use EBME to configure commands or macros for execution when the status of a given port changes from UP to DOWN, from DOWN to UP or to LOOPBACK condition. You can configure commands and macros containing instructions to bring back the primary connection if an UP to DOWN event occurs on the port. You can also configure commands and macros to be executed when a loopback event occurs.

EBME consists of event generators, a port-based event responder, and an event and action database. The EBME Service provides the interface for these components and for the EOS software.

When the EBME Service is enabled to log all events, the sequence of events to execute the action you configured is as follows:

- When the port is configured with some EBME events, its status is monitored by EBME.
- The port status – either UP, DOWN or LOOPBACK – is recognized by the port-based event responder from the event generator. Each port may have no more than three events: UP, DOWN or LOOPBACK.
- The port UP/DOWN event becomes valid when the debounce timer runs out.

- When the event responder determines that the event occurred, it queues the event and selects the appropriate action from the event and action database.
- The action is scheduled to take place.
- When the action completes, a message is written to the log file stating that the action is complete.



NOTE: Only the first 80 characters of the output from the UI command are printed in the log file, and only the macro name and a complete or incomplete message are printed.

39.1

Schedules Creation

Defining a Daily Schedule

You must define at least one daily schedule before you can create an active schedule.

You can create a daily schedule and define its first event in one step by using:

```
ADD -SCH Event <daily_schedule> <hh:mm> <command-string>
```

You can add additional events to the daily schedule using the same syntax, where *<daily_schedule>* is the name of a daily schedule for which at least one step was already defined.

Creating an Active Schedule

After the daily schedule is complete, you must assign it to a calendar date or to a day of the week by using:

```
ADD -SCH ActiveSchedule <mm/dd / SUN / MON / TUE / WED / THU / FRI / SAT>  
<daily_schedule>
```

The scheduler is only active when two conditions are met:

- The `-SCH CONTROL` parameter is set to Enabled.
- The `-SCH CONTROL` parameter is set to RealTimeClock.

The `-SCH CONTROL` parameter also allows you to choose the hardware clock or software clock and select or deselect the logging function. For a complete summary of settings for the `-SCH CONTROL` parameter, see “SCH Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

You can enable the scheduler by using:

```
SETDefault -SCH CONTROL = Enabled
```

Executing Macros with the Scheduler

The scheduler is a batch-oriented utility. Always test macros thoroughly before submitting them to the scheduler. The scheduler submits the macro to the system without evaluating it and does not report success or failure – even if logging is enabled, the scheduler reports only that the macro was submitted.

Ensure that any macros submitted to the scheduler contain:

- No infinite loops.
- No input variables.
- No illegal commands.
- Minimal output message.



CAUTION: The only way to stop a macro that contains an infinite loop is to reboot the system.

39.2

Scheduling WAN Connections

You can create a daily schedule to establish a dial-up connection and hang up the line at a specified times, and assign this daily schedule to one or more calendar dates or days of the week.

Procedure:

1. Create a daily schedule that establishes a dial-up connection by entering:

```
ADD -SCH EVent <daily_schedule> <hh:mm> DIal !<port>  
" <telephone_number>"
```

Step example:

To define the daily schedule, "Daily," that establishes a connection to the telephone number 555-1212 on port 3 at 11 a.m., enter:

```
ADD -SC EVent daily 11:00 DIal !3 "5551212"
```

2. Add to the daily schedule an event that hangs up the connection by entering:

```
ADD -SCH EVent <daily_schedule> <hh:mm> HangUp !<port>
```

Step example:

To add to the daily schedule, "Daily," a hangup event at 3 p.m., enter:

```
ADD -SCH EVent daily 15:00 HangUp !3
```

3. Assign the daily schedule, "Daily," to each weekday by entering:

```
ADD -SCH ActiveSCHedule MON DAILY
```

4. Repeat [step 3](#) four times for each of the remaining weekday designators: (TUE, WED, THU, and FRI) in the place of MON.

39.3

Event-Based Commands/Macros

Event-based Command/Macro Execution (EBME) provides automatic back-up if a connection fails between two sites when a primary link goes down. EBME also provides loopback detection and recovery; it is a port-based service.

EBME provides the following features:

- User-defined command or macro configuration.
- A backup action that is executed when the status of a port changes. EBME can back up any kind of port including static, virtual, parent, and dial ports.
- An action that is executed when a port is in a loopback condition. EBME can disable the port and then re-enable the port after a defined delay.
- A user-configurable debounce timer that provides a delay before the command or macro is executed. This timer is provided to prevent the software from reacting to transient changes in the port status.
- A log to track system command or macro execution.
- User control to enable or disable the service.

EBME is also a batch-oriented utility. Always test macros thoroughly before submitting them to EBME. EBME submits the macro to the system without evaluating it and does not report success or failure – even if logging is enabled, the scheduler reports only that the macro was submitted.

You must ensure that any macros submitted to the EBME contain:

- No infinite loops.

- No input variables.
- No illegal commands.
- Minimal output message.



CAUTION: The only way to stop a macro that contains an infinite loop is to reboot the system.

39.3.1

Setting Up a Backup Port

You can use this procedure to configure EBME, to bring up a port when a primary link fails.

Procedure:

1. Configure the command or macro for the backed up port by entering:

```
ADD !2 -SCH EbmeEvent PortDown 30 D0 port2down_macro
```

When port 2 fails, this command causes the EBME Service to execute the `port2down_macro` file and sets the debounce timer to 30 seconds. The `port2down_macro` brings up the backup port you specify in the macro.

2. Enable the EBME Service by entering:

```
SETDefault -SCH EbmeCONTRol = Enable
```

3. Set the `Log` option to record the commands or macros that are executed in the system log buffer by entering:

```
SETDefault -SCH EbmeCONTRol = Log
```

39.3.2

Hanging Up a Port

Procedure:

To configure EBME to hang up a backup port when the primary link becomes active again, enter:

```
ADD !2 -SCH EbmeEvent PortUp HangUp !3
```

When port 2 comes up, this command causes the EBME Service to execute the `HangUp` command for port 3.

39.3.3

Recovering from Port Loopback

The Spanning Tree Protocol (STP) is designed to prevent bridges from forming a loop in active paths. However, STP does not prevent data from being looped back into the port that emitted it. This situation may cause an incorrect station hop and/or broadcast storm problems.

You can use EBME to detect a port loopback condition by configuring EBME to execute a user-defined command or macro when a loopback condition occurs.

Procedure:

1. Configure EBME to recover from a loopback condition on port 2 by defining the port recovery macro, `port2LBmacro`, by entering:

```
(  
define port2LBmacro =
```

```

setd !2 -port CONT=Disable
pause 300
setd !2 -port CONT=Enable
)

```

This macro checks every five minutes (300 seconds) to see if the loopback condition is still happening. When the loopback is gone, the port is returned to the forwarding state.

2. Enable the Spanning Tree Protocol by entering:

```
SETDefault -STP CONTrol = Enable
```

3. Configure the macro for the port by entering:

```
ADD !2 -SCH EbmeEvent LoopBack DO port2LBmacro
```

4. Enable the EBME Service by entering:

```
SETDefault -SCH EbmeCONTrol = Enable
```



CAUTION: The event command string is executed sequentially. A `Pause` command delays another event from execution even if the other event occurs concurrently. Motorola Solutions recommends that you use the `Pause` command only in loopback events.

39.3.4

VRRP Event-Triggered Macros Configuration

One way to achieve redundancy and fast fault recovery is to set the timers to very low values for unicast protocols. However it is not always possible to have hosts run the same protocols with the same speed. In such cases, VRRP is used to provide redundancy.

VRRP is a protocol used between devices that share a common IP address on a broadcast LAN. One device is the master (or owner) router and own this address. The master router responds to ARPs for this IP address. The other devices are backup routers. All these devices use VRRP keepalives to signal that they are still active on that interface. If any backup router notices that the master router stopped sending keepalives, the backup router internally generates a `Master Down` event, and the backup router with the highest priority becomes the master. When the master router recovers, the backup router sees the keepalives from the master, internally generates a `Master Up` event, and relinquishes control over the shared IP address.

The VRRP event-triggered macros feature allows you to specify an EOS command or macro to be executed when a VRRP `Master Down` or `Master Up` event occurs. You add, delete, and display VRRP event-triggered macros using the SCH service `VrrpEvent` parameter.

For example, to specify that a macro named `master_down` be run with a debounce timer of 50 milliseconds when the VRRP master router on port 1 at IP address 10.69.121.30 goes down, you must use:

```
ADD !1 -SCH VrrpEvent 10.69.121.30 50 MasterDown do master_down
```

For details about the `VrrpEvent` parameter, see “SCH Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

39.3.5

Configuring Bidirectional Forwarding Detection (BFD) with Event-Based Macros

You can use single-hop or multi-hop bidirectional forwarding detection (BFD) to trigger an event-based macro. Follow this procedure to configure EBME to execute a user-defined command or macro when the BFD session state changes from up to down or from down to up.

For further details about BFD, see [Bidirectional Forwarding Detection \(BFD\) Configuration on page 138](#).

Procedure:

1. Configure the command or macro for the BFD event by entering:

```
ADD !<portlist> -SCH BfdEvent <IP_Address> [TTL <t1>] <PeerUp / PeerDown>  
[ <milliseconds>(50-3600000)] <Any_UI_command/Macro_(1-80_chars)>
```

Step example:

Single-hop BFD:

```
ADD !2 -SCH BfdEvent 20.20.20.202.43 PeerDown 50 D0 BFDdown_macro
```

When a failure occurs on the single-hop static route to IP address 20.20.20.202, triggering a BFD Peer Down event, this command causes the EBME Service to execute the BFDdown_macro file and sets the debounce timer to 50 milliseconds.

Multi-hop BFD:

```
ADD !2 -SCH BfdEvent 20.20.20.202.43 TTL 9 PeerDown 50 D0 BFDdown_macro
```

When a failure occurs on the multi-hop static route (with TTL value 9) to IP address 20.20.20.202, triggering a BFD Peer Down event, this command causes the EBME Service to execute the BFDdown_macro file and sets the debounce timer to 50 milliseconds.



NOTE: If the BFD session starts from a VRRP-controlled address (VRRP is enabled on the interface), you must also set the IP service BFDTRigger parameter to configure BFD to register with VRRP on that interface to start and stop the BFD session. For details, see the description of the BFDTRigger parameter in “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

2. Enable the EBME Service by entering:

```
SETDefault -SCH EbmeCONTrol = Enable
```

3. Set the Log option to record the commands or macros that are executed in the system log buffer by entering:

```
SETDefault -SCH EbmeCONTrol = Log
```

Chapter 40

CCGW Functionality Configuration

This chapter describes how to configure the MNR S25000, MCG 8000 and GGM 8000 as a conventional channel gateway (CCGW).

The term *CCGW* is used generically in this chapter to refer to the MNR S2500 or the GGM 8000 or the MCG 8000 gateway configured as conventional channel gateway.

The CCGW SNMP Agent performs fault management for the CCGW. For details about supported CCGW MIBs and traps, see [MNR CCGW Private MIBs on page 867](#) and [CCGW Fault Traps on page 880](#).

40.1

Supported CCGW Features

EOS software supports the following CCGW features:

Conventional integrated voice and data (IV&D)

Supports voice and data on the same digital channel. For more information, see [Conventional IV&D on page 739](#).

IP conventional

Supports IP connectivity between the CCGW and conventional IP base stations. For more information, see [IP Conventional on page 740](#).

Mixed mode

Supports inbound and outbound mixed mode calls, when audio is processed over the analog interface, and signaling related to the call is processed over the digital (V.24 or IP) interface. For more information, see [Mixed Mode on page 740](#).

MDC 1200 signaling

Supports additional MDC 1200 message processing over analog channels. For more information, see [MDC 1200 Signaling on page 742](#).

ASTRO console interface module (ACIM) link support

Enables connection of the Motorola Solutions consolette to the CCGW through the ACIM link protocol interface. For more information, see [ACIM Link Channel Support on page 742](#).

Advanced SECURENET

Supports the key selection and key selection messaging between the MCC 7500 and the CCGW. For more information, see [Advanced SECURENET Support on page 742](#).

LDAP server backwards compatibility

Supports pre-ASTRO® 25 7.9 LDAP server releases on the CCGW LDAP client. You can specify which LDAP server release to use from the EOS UI. For more information, see “CCGW Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Conventional data DSR support

Provides a complete DSR solution for conventional systems: conventional data DSR as well as conventional voice DSR. For details, see [CCGW Support for Conventional Data DSR on page 744](#).

Conventional talkgroups

Provides separation of digital voice communications on conventional channels. Subscribers and console operators who use the same talkgroup, can communicate with each other, and users of other talkgroups cannot hear them. Conventional talkgroups also provide the separation of emergency alarms. For more information, see [Conventional Talkgroups on page 744](#).

Console repeat

Enables dispatch console users to force a set of scanners to detect the carrier on a particular channel and remain on that channel as long as the carrier is present. For more information, see [Console Repeat on page 746](#).

Awareness of Main/Alternate channel pair status

Prevents the CCGW from oversending call requests for digital channels and undersending call requests for analog channels. For more information, see [Main/Alternate Channel Status Awareness on page 746](#).

Duplex patch with two-party conference

Provides the ASTRO® 25 system support for the duplex patch feature already supported on existing CENTRACOM Gold series systems. When two analog conventional or MDC 1200 channels are in a duplex patch, the non-transmitting users on both channels hear the sum of the audio from the transmitting user on their channel and the transmitting user on the other channel. For more information, see [Duplex Patch with Two-Party Conference on page 747](#).

Ethernet/4-Wire Hybrid Interface support

Provides Analog and MDC1200 channels in configuration with G-Series Analog Comparator an option to use a channel link with additional connectivity failure detection facility. For more information, see [CCGW Functionality on page 736](#).

IP Analog

Supported only by the MCG 8000. IP Analog channel enables analog conventional operation over an IP interface between the CCGW and the base station. Physical analog interfaces are not used for this channel type.

IP MDC 1200

Supported only by the MCG 8000. IP MDC1200, like the IP Analog, uses an IP interface between the CCGW and the base station. Apart from the analog conventional audio, it features the MDC1200 signaling.

40.2

CCGW Hardware Requirements

CCGW functionality requires either an MNR S2500 or a GGM 8000 gateway.

The following table describes the maximum number of channels supported per CCGW, supported channel types, and allowable channel IDs for each supported hardware platform/Site Type combination.

Table 63: Allowable Channel IDs per CCGW Hardware Platform, Site Type, and Channel Type

CCGW Hardware Platform	Maximum # of Channels per CCGW	Site Type	Channel Type	Allowable Channel IDs
High Density Enhanced Conventional Gateway (GGM 8000 with High Density Enhanced Conventional Gateway module)	32 ¹	Combination-HD	Any non-data capable channel type (analog conventional, MDC 1200)	1-32
			Any data-capable channel type (digital conventional, IP conventional, mixed mode ² , or ACIM)	1-30
Low Density Enhanced Conventional Gateway (GGM 8000 with Low	24 ³	Combination-HD	Any non-data capable channel type (analog conventional, MDC 1200)	1-32

CCGW Hardware Platform	Maximum # of Channels per CCGW	Site Type	Channel Type	Allowable Channel IDs
Density Enhanced Conventional Gateway module)			Any data-capable channel type (digital conventional, IP conventional, mixed mode ² , or ACIM)	1-30
GGM 8000 with expansion module and analog/V.24 interface kit	14	Combination ⁴	Analog conventional	1-4
			MDC 1200	1-4
			Digital conventional	1-4
			IP conventional	21-30
			Mixed mode ²	1-4
			ACIM	1-4
GGM 8000 Base Unit	16	Combination-HD	IP conventional	1-30
	10	Combination	IP conventional	21-30
MNR S2500 with conventional-to-IP interface kit ⁵	4	Analog	Analog conventional	1-4
MNR S2500 with V.24 module	2	Digital	Digital conventional	1-4
MCG 8000 with High Density Enhanced Conventional Gateway module	32 ⁶	Combination-MCG	Any non-data capable channel type (analog conventional, MDC 1200)	1-32
			Any data-capable channel type (digital conventional, IP conventional, mixed mode ⁷ , or ACIM)	1-30
MCG 8000 Base Unit	16	Combination-MCG	IP conventional	1-30

¹ 8 channels using an analog interface, plus 8 channels using a V.24 interface, plus 16 IP conventional channels.

² A mixed mode channel can be either analog + V.24 (V.24 mixed mode) or analog + IP (IP mixed mode).

³ 4 channels using an analog interface, plus 4 channels using a V.24 interface, plus 16 IP conventional channels.

⁴ The GGM 8000 with analog/V.24 interface kit does not support the Combination-HD site type.

⁵ The S2500 conventional-to-IP interface kit consists of an E&M module and DSP SIMM.

⁶ 8 channels using an analog interface, plus 8 channels using a V.24 interface, plus 16 IP conventional channels.

⁷ A mixed mode channel can be either analog + V.24 (V.24 mixed mode) or analog + IP (IP mixed mode).

If the CCGW receives a configuration from the LDAP server, and the required hardware does not exist on the box, the CCGW either accepts the configuration or sends a state change trap (ccgwChannelStateChangeTrap) with channel state DOWN, and fault cause 6 (User Configured Channel).

40.3

CCGW Port-to-Channel Mapping

IP conventional channels are not tied to physical ports, so port-to-channel mapping is not applicable for these channels.

Allowable analog or MDC 1200 channel IDs are 1-32, but only one channel can be configured at a time on each analog interface.

Allowable digital channel IDs are 1-30, but only one channel can be configured at a time on each digital interface.

Allowable mixed mode or ACIM channel IDs are 1-30, but only one channel can be configured at a time on each analog/digital interface pair.




NOTE: Non-data capable channels (analog, MDC 1200) can have channel IDs in the range 1-32 and data capable channels (digital, IP conventional, mixed mode, or ACIM) can have IDs in the range 1-30.

Table 64: MCG 8000 Port-To-Channel Mapping

Port/Path Number	Analog or MDC 1200 Channel Number	Digital Channel Number	Mixed Mode or ACIM Channel Number
Low Density Enhanced Conventional Gateway			
9D + 8D	1-32	N/A	N/A
9C + 8C	1-32	N/A	N/A
9B + 8B	1-32	N/A	N/A
9A + 8A	1-32	N/A	N/A
7B	N/A	1-30	N/A
7A	N/A	1-30	N/A
6B	N/A	1-30	N/A
6A	N/A	1-30	N/A
9D + 8D + 7B	N/A	N/A	1-30
9C + 8C + 7A	N/A	N/A	1-30
9B + 8B + 6B	N/A	N/A	1-30
9A + 8A + 6A	N/A	N/A	1-30
High Density Enhanced Conventional Gateway			
9D + 8D	1-32	N/A	N/A
9C + 8C	1-32	N/A	N/A
9B + 8B	1-32	N/A	N/A
9A + 8A	1-32	N/A	N/A
7B	N/A	1-30	N/A
7A	N/A	1-30	N/A
6B	N/A	1-30	N/A
6A	N/A	1-30	N/A
9D + 8D + 7B	N/A	N/A	1-30
9C + 8C + 7A	N/A	N/A	1-30

Port/Path Number	Analog or MDC 1200 Channel Number	Digital Channel Number	Mixed Mode or ACIM Channel Number
9B + 8B + 6B	N/A	N/A	1-30
9A + 8A + 6A	N/A	N/A	1-30
13D + 12D	1-32	N/A	N/A
13C + 12C	1-32	N/A	N/A
13B + 12B	1-32	N/A	N/A
13A + 12A	1-32	N/A	N/A
11B	N/A	1-30	N/A
11A	N/A	1-30	N/A
10B	N/A	1-30	N/A
10A	N/A	1-30	N/A
13D + 12D + 11B	N/A	N/A	1-30
13C + 12C + 11A	N/A	N/A	1-30
13B + 12B + 10B	N/A	N/A	1-30
13A + 12A + 10A	N/A	N/A	1-30

Figure 190: MCG 8000 CCGW Port Numbering and Port-to-Channel Mapping



HD_ECCGW_port_to_channel_mapping_B

Figure 191: GGM 8000 Low Density CCGW Port Numbering and Port-to-Channel Mapping






Figure 192: GGM 8000 High Density CCGW Port Numbering and Port-to-Channel Mapping




HD_ECCGW_port_to_channel_mapping_B

Figure 193: MNR S2500 CCGW Port Numbering and Port-to-Channel Mapping (Analog)




Annotation	Description
1	E&M port 4 maps to analog channel 1
2	E&M port 5 maps to analog channel 2
3	E&M port 6 maps to analog channel 3
4	E&M port 7 maps to analog channel 4

Figure 194: MNR S2500 CCGW Port Numbering and Port-to-Channel Mapping (Digital with V.24 in Slot A)



Annotation	Description
1	V.24 port 2A maps to digital channel 1
2	V.24 port 2B maps to digital channel 2

Figure 195: MNR S2500 CCGW Port Numbering and Port-to-Channel Mapping (Digital with V.24 in Slot B)



Annotaiton	Description
1	V.24 port 3A maps to digital channel 1
2	V.24 port 3A maps to digital channel 2

40.4 CCGW LEDs

The following table describes how the S2500, the GGM 8000 and the MCG 8000 E&M (analog) port and V.24 port LEDs appear if all supported channels are in-service, depending on the channel type.

Table 65: CCGW LED Status Indicators (All Supported Channels In Service)

Channel Type	E&M (analog) LEDs	V.24 LEDs
Analog conventional	Green	Off
Digital conventional	Off	Green
IP conventional ¹²	Off	Off
V.24 Mixed mode ²	Green ³	Green
IP mixed mode ²	Green	Off
MDC 1200 ²	Green	Off
ACIM ²	Green	Green

¹ IP conventional channels are not associated with a physical E&M or V.24 port; therefore the E&M and V.24 LED status does not indicate the status of these channels.

² IP conventional, V.24 mixed mode, IP mixed mode, MDC 1200, and ACIM channels are supported only on the GGM 8000 and the MCG 8000.

³ If an E&M LED is off, but the corresponding V.24 LED is on, the V.24 mixed mode channel is partially in-service: in analog mode the channel is incapable of transmitting audio, but in digital mode the channel can function properly.

40.5 Configuring EOS Devices as a CCGW

You can perform this procedure to configure the CCGW functionality on the MNR S2500, MCG 8000 or the GGM 8000.

For more information about the EOS CCGW service parameters, see “CCGW Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

For more information about the CCGW data flow and the interaction between the EOS software and the LDAP server, see [CCGW Functionality on page 736](#).

Procedure:

1. Specify the IP address and TCP port of the LDAP server and specify the CCGW as a service for which this server IP address must be used when sending out a request by entering:

```
ADD -LDAP ServerAddress <IP_address> [TCP Port] "ccgw"
```

where:

<IP_address> is the IP address of the LDAP server,

TCP Port is the TCP port at the LDAP server where the connection must be made.

2. Specify the DN value for the LDAP bind request and specify CCGW as the service configuration that should be downloaded from the LDAP server by entering:

```
ADD -LDAP BindDN "<DN_string>" "<password>" "ccgw" [Encrypted]
```

where:

<DN_string> is the DN value, in ASCII text format,

<password > is a string of 0-16 characters.

The DN value is used to authenticate an area in the tree where the ADD/DELEte/Search operations can be performed. If you want to enter the password in an encrypted form, you must use the optional Encrypted keyword.

3. Specify the UDP port to be used by the Notifier to send notifications to the CCGW by entering:

```
SETDefault -CCGW NMUDPPort = <UDP_Port_#>
```

where <UDP_Port_#> is the UDP port number.

4. Specify the zone controller's Netcom1 and Netcom 2 IP addresses by entering the following commands:

```
SETDefault -CCGW ZCLink1Address = <IP_address_1>
```

```
SETDefault -CCGW ZCLink2Address = <IP_address_2>
```

where:

<IP_address_1> is the zone controller's Netcom1 IP address,

<IP_address_2> is the zone controller's Netcom2 IP address.

5. If your system implements the DSR voice or conventional data redundancy (HA and DSR) feature, specify the backup zone controller's Netcom1 and Netcom2 IP addresses by entering the following commands:

```
SETDefault -CCGW BkupZCLink1Addr = <IP_address_1>
```

```
SETDefault -CCGW BkupZCLink2Addr = <IP_address_2>
```

where:

<IP_address_1> is the backup zone controller's Netcom1 IP address,

<IP_address_2> is the backup zone controller's Netcom2 IP address.

6. Enable CCGW functionality on the CCGW by entering:

```
SETDefault -CCGW CONTrol = Enable
```

40.6

CCGW Functionality

The CCGW supports the following channel types:

Analog conventional

Interfaces 4-wire analog audio connections (across WAN or LAN links) with certain types of Motorola Solutions public safety network equipment that support digital audio connections.

Digital conventional (digital over V.24 interface)

Interfaces V.24 ports carrying digital audio (across WAN or LAN links) with certain types of Motorola Solutions public safety network equipment that support digital audio connections.

IP conventional (digital over IP interface)

Supports the IP connectivity between the CCGW and conventional IP base stations and comparators (in IP simulcast systems). An IP conventional channel functions as an IP interface between the base station and the console, creating an IP link between the CCGW and the base station using Motorola Solutions proprietary link management protocol. Typically, IP conventional runs over the Ethernet interface.

Mixed mode

Processes audio over the analog interface and the signaling related to the call over the digital interface. The CCGW supports mixed mode channels over V.24 and IP digital interfaces.

Ethernet/4-Wire Hybrid Interface Support

The Ethernet/4-Wire Hybrid Interface feature supports a method to detect the failure of the wireline between the CCGW and the G-Series Analog Comparator for Analog and MDC1200 channels. The feature requires a GGM 8000 gateway equipped with either a Low Density Enhanced Conventional Gateway module, or a High Density Enhanced Conventional Gateway module. The interface binds an analog interface and an IP digital interface in a way similar to IP Mixed Mode channels, with the following differences:

- Analog Link Monitor Tone – Ethernet/4-Wire Hybrid Interface maintains a signal of 220 Hz over analog interface in both directions to monitor connection status. It is generated continuously, also during a call, mixed with the call audio, but it's filtered at the receiving endpoint so that the tone is not heard by the user.
- IP Interface – IP interface has the following functions:
 - Conveys link or RF site equipment failure conditions detected at the comparator side. An example can be a TX pair failure on the analog interface while RX pair is ok. This type of failure cannot be detected by the CCGW, because the monitor tone would be still received over RX pair.
 - Detects the inbound call start or stop from the station
 - Is used for keying the station for outbound calls if the channel is configured for digital keying. Another option is keying with tone sequence sent out the analog interface, IP interface is not used for outbound call maintenance in that case.
 - Carries main/alt status of the channel to G-Series Analog Comparator (applicable to main/alt configurations only).

MDC 1200

Integrates MDC 1200 and analog conventional signaling and makes digital calls with the zone controller. An MDC channel is a channel where audio and call control are processed over the analog interface. However, the call control messages between the CCGW and the zone controller are in digital conventional format.

ACIM

Enables connection of the Motorola Solutions console to the CCGW via the ACIM link serial control protocol. This topology allows a console site to have access to talkgroups when its link to the zone core failed. It also enables wireless connection to systems utilizing a variety of over-the-air protocols.

IP Analog

Supported only on the MCG 8000. IP Analog channel enables analog conventional operation over IP interface between the CCGW and the base station. Physical analog interfaces are not used for this channel type.

IP MDC 1200

Supported only on the MCG 8000. IP MDC1200 like IP Analog uses IP interface between the CCGW and the base station. Apart from analog conventional audio it features MDC1200 signaling.



NOTE:

Digital conventional (digital over V.24) and IP conventional (digital over IP) channels can also be configured as conventional talkgroup channels.

IP conventional, mixed mode, MDC 1200, and ACIM channels are supported only on the GGM 8000 and MCG 8000.

Analog and MDC1200 channels in configuration with G-Series Analog Comparator may be set up for Ethernet/4-Wire Hybrid Interface, which adds IP connection to existing 4-wire analog audio connection.

IP conventional, mixed mode, MDC 1200, and ACIM channels are supported only on the GGM 8000 and MCG 8000.

The channel type is configured in the LDAP database, and the `ChType` parameter allows you to view the channel type configured on the LDAP server. In addition, the `GroupID` parameter allows you to view the group to which a channel is assigned.

The CCGW supports the following site types:

Analog

Specifies an MNR S2500 CCGW configured as an analog CCGW, which allows channels to be configured only as analog conventional channels.

Digital

Specifies an MNR S2500 CCGW configured as a digital CCGW, which allows channels to be configured only as V.24 digital channels.

Combination

Specifies a GGM 8000 CCGW, which allows channels to be configured as analog, V.24 digital, IP conventional, mixed mode, MDC 1200, or ACIM channels.

Combination-HD

Specifies a GGM 8000 CCGW configured with the Enhanced Conventional Gateway option, which allows channels to be configured as analog, V.24 digital, IP conventional, mixed mode, MDC 1200, or ACIM channels.

Combination-MCG

Specifies a MCG 8000 which allows channels to be configured as analog, V.24 digital, IP conventional, IP analog, IP MDC 1200, mixed mode, MDC 1200, or ACIM channels.

The site type is configured in the LDAP database, and the `EOS SiteType` parameter allows you to view the site type configured on the LDAP server. The `SiteType` parameter corresponds to the **Conventional Gateway Type** field in the Provisioning Manager.

CCGW commands are executed in the device `boot.cfg` file upon bootup. After the `boot.cfg` file was executed, EOS initializes CCGW, which sets up communication with the network management system's LDAP server. EOS then listens to the UDP notification port and configures the CCGW based on the network management system's LDAP configuration.

A CCGW configuration file (CCGWDB) is created on the EOS flash. This file contains the latest CCGW configuration from the network management LDAP server. It is updated when changes are received from the

network management LDAP server after the CCGW was rebooted and is used for bootup persistence. If the CCGW cannot communicate with the LDAP server, it reads the configuration from the EOS flash.



IMPORTANT:

If a change is made to the `CCGWDBVersion`, `DBSchemaVer`, `DSR`, `SiteID`, `SiteType`, or `ZoneID` parameter from the LDAP network management server (to correct a misconfiguration, for example), you must perform one of the following actions after the change is made:

- Remove the CCGWDB file from flash, then reboot the CCGW.
- Resynchronize the CCGWDB file with the LDAP server. To do this, you must issue the following commands to disable then re-enable the CCGW service:

```
SETDefault -CCGW CONTROL = Disable
```

```
SETDefault -CCGW CONTROL = Enable
```

The CCGWDB file is refreshed from the LDAP server. When a CCGW is replaced, the LDAP server is the source for refreshing the CCGWDB file.



CAUTION: When you upgrade or migrate a CCGW, you must remove the CCGWDB file from flash before you upgrade the software on the CCGW.

40.6.1

CCGW Identification

The `mnrSysServiceID` MIB variable identifies whether or not a device is a CCGW.

- If the `SETDefault -CCGW CONTROL = Enabled` command is read from the `boot.cfg` file when the device boots, the device is identified as a CCGW, and it remains identified as such until the next reboot. The device remains identified as a CCGW even if the `SETDefault -CCGW CONTROL = Disabled` command is executed.
- In order to not be identified as a CCGW, the device must be rebooted with a `boot.cfg` file that includes the `SETDefault -CCGW CONTROL = Disabled` command or no `SETDefault -CCGW CONTROL` command (because the `CCGW CONTROL` parameter is set to `Disabled` by default).

You can also use the absence or presence of the `mnrWanSlotObj` variable to determine if the device is a CCGW only or both a CCGW and a site router/gateway.

- If the `mnrWanSlotObj` variable does not exist, and the `mnrSysServiceID` variable returns a value of 2 (CCGW), the device is a CCGW only.
- If the `mnrWanSlotObj` variable exists, and the `mnrSysServiceID` variable returns a value of 2 (CCGW), the device is both a CCGW and a site router/ gateway.

40.6.2

Conventional IV&D

Conventional IV&D adds data service to the existing voice-only analog conventional and digital conventional architecture. The CCGW provisions and integrates voice calls and data transport simultaneously.

The CCGW call manager interacts with the CCGW data manager to mitigate voice and data contention. To ensure that voice calls are not compromised while data transmission is in progress, a single outbound call is buffered up to 1200 ms for a channel. Upon finishing the transmission of an outbound packet PDU, the buffered audio is played out.

40.6.3

IP Analog

IP Analog channel enables analog conventional operation over IP interface between the CCGW and the base station. Physical analog interfaces are not used for this channel type.



NOTE: IP analog is supported only on the MCG 8000.

40.6.4

IP MDC1200

IP MDC1200 like IP Analog uses IP interface between the CCGW and the base station. Apart from analog conventional audio it features MDC1200 signaling.



NOTE: IP MDC1200 is supported only on the MCG 8000.

40.6.5

IP Conventional

The IP conventional feature supports IP connectivity between the CCGW and conventional IP base stations and comparators (in IP simulcast systems).

Before the introduction of the IP conventional feature, CCGW functionality was available only over E&M (analog) ports or V.24 ports. An IP conventional channel functions as an IP interface between the base station and the console, creating an IP link between the CCGW and the base station using Motorola Solutions proprietary link management protocol. Typically, IP conventional runs over the Ethernet interface.



NOTE: IP conventional is supported only on the GGM 8000 and the MCG 8000.

With IP conventional, the conventional channels are not bound to any of the physical ports/paths on the CCGW. Instead, a transparent port (created for internal mapping purposes only and not bound to any physical path) is created for each IP conventional channel, and voice/data/control packets are transported over a physical port determined by IP routing. IP conventional supports up to 10 channels, as opposed to the 4 channels supported for analog conventional or digital conventional operation. The 10 channels correlate 10 IP addresses to 10 base stations.

40.6.6

Mixed Mode

The CCGW supports mixed mode channels over V.24 and IP digital interfaces.

V.24 mixed mode

Binds an analog interface and a V.24 digital interface as a single channel. The CCGW associates the signaling information processed over the digital V.24 interface with the voice processed over the corresponding analog interface.

IP mixed mode

Binds an analog interface and an IP digital interface as a single channel. The IP interface is used for call signaling and control and digital audio for digital calls, while the analog interface is used for call signaling and control and analog audio for analog calls. Each direction is independent of the other.

V.24 mixed mode and IP mixed mode channels use the same channel type (mixed mode); the difference is that for V.24 mixed mode channels the digital interface type (DCCHtype) is V.24, while for IP mixed mode channels the digital interface type (DCCHtype) is Ethernet (IP).

V.24 Mixed Mode Support

A conventional mixed mode channel can run either in pure digital mode (audio and signaling is processed over the V.24 digital interface) or in mixed mode. In mixed mode, audio is processed over the analog E&M interface and the signaling related to this call is processed over the digital V.24 interface. The CCGW associates the signaling information processed over the digital V.24 interface with the voice processed over the corresponding analog E&M interface.

For mixed mode operation, a V.24 port is associated with an E&M port via a static mapping, done by the Call Manager during initialization. For more information on how the CCGW V.24 ports are associated with E&M ports, see [CCGW Port-to-Channel Mapping on page 729](#).

Analog Link Monitor Tone Support

To maintain the knowledge of the operability of the analog portion of the V.24 links when they are used to support signaling for digital conventional mixed mode, the CCGW maintains a signal of 1004 ± 2 Hz on its outbound analog lines while the link is idle (no call activity is occurring on the analog side of a functional established link). This signal is called the Analog Link Monitor Tone (ALMT). It is generated to indicate to a connected device that the link is up.

When the CCGW is supporting mixed mode operation, the analog 4-wire interface is idle whenever a pure digital (as opposed to a mixed mode) call is supported. In these cases, the CCGW provides the ALMT to the connected device (base station or comparator).



NOTE: The ALMT is supported only on V.24 mixed mode channels. It is not supported on IP mixed mode channels that use the IP (Ethernet) interface or on analog conventional channels.

IP Mixed Mode Support

The IP mixed mode feature (digital and analog on the same channel) supports mixed mode comparators, which allow you to consolidate digital channels with existing analog channels without purchasing additional channels.

The GGM 8000 CCGW and the MCG 8000 CCGW support mixed mode conventional channels, which bind an analog interface and an IP interface as a single channel. The IP interface is used for call signaling and control and digital audio for digital calls, while the analog interface is used for call signaling and control and analog audio for analog calls. Each direction is independent of the other. IP mixed mode functionality allows the CCGW to multiplex a variety of analog/digital air interface devices through a combined channel and resolve channel contention in both data and voice traffic. Unlike a V.24 mixed mode channel, an IP mixed mode channel does not have to buffer audio packets while a data session is ongoing.

For inbound calls on an IP mixed mode channel, the CCGW:

- Detects the inbound call from both the IP interface and the E&M interface.
- Initiates a radio call request with the zone controller.
- Detects the end of call from both the IP interface and the E&M interface.
- Sends a radio dekey ICP to the zone controller.
- Converts the digital audio to XIS IMBE packets and the analog audio to G.728 packets for transmission.

For outgoing calls on an IP mixed mode channel, the CCGW:

- Receives the outbound digital or analog call from the console.
- Keys the channel over the IP interface for digital calls and over both the IP interface and the E&M interface for analog calls.
- Receives XIS IMBE audio packets for digital calls and G.728 audio packets for analog calls from the multicast address, then sends the digital audio over the IP interface and sends the audio packets to the DSP vocoder for conversion to analog audio for outbound transmission over the 4-wire interface.
- Receives the outbound console end of transmission and dekeys the channel over the IP interface for digital calls or over the E&M interface for analog calls.

40.6.7

MDC 1200 Signaling

MDC signaling integrates MDC 1200 and analog conventional signaling and makes digital calls with the zone controller. An MDC channel is a channel where audio and call control are processed over the analog E&M interface. However, the call control messages between the CCGW and the zone controller are in digital conventional format.

40.6.8

ACIM Link Channel Support

The ACIM link feature enables connection of the Motorola Solutions consolette to the CCGW through the ACIM link serial control protocol. This topology allows a console site to have access to talkgroups when its link to the zone core fails. It also enables wireless connection to systems that utilize a variety of over-the-air protocols, such as:

- Analog
- ASTRO conventional
- 9600 trunking
- 3600 trunking digital
- 3600 trunking analog

ACIM link channels:

- Vocode all audio with G.728, even when the over-the-air mode is digital or encrypted.
- Do not perform encryption or decryption for encrypted audio. The consolette performs these functions.



NOTE: ACIM link channels are supported only on the GGM 8000 and the MCG 8000. The MNR S2500 does not support the ACIM link feature.

40.6.9

Advanced SECURENET Support

The Advanced SECURENET (ASN) feature supports key selection and key selection messaging between the MCC 7500 and the CCGW. The console operator selects the key and frequency on which to talk. Bundled with each transmission, the console sends the proper key and frequency selection to the CCGW. The CCGW uses the frequency and key selection information to look up the proper tone sequence (which may be one or two tones). The tone sequence is then sent out the analog 4W port before the analog audio sourced by the console is sent.

When an inbound call starts, ends, or is taken over by another call, the CIU coded/clear state changes. The CCGW detects when the inbound audio is encrypted (coded) and sends the coded indication message to all subscribed consoles. The consoles then use the coded indication message to display to the user that the channel is currently in coded mode. When the audio is not encrypted (clear), the CIU does not provide an indication to the CCGW E-lead.



NOTE: The Advanced SECURENET (ASN) feature is supported only on the GGM 8000 and the MCG 8000. The MNR S2500 does not support the ASN feature.

As part of the secure communication mode feature, the CCGW supports:

- The following configuration parameters from the LDAP server:
 - ChFeedback
 - SecureMode

For details about these parameters, see “CCGW Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.



NOTE: The CCGW does not support the console parameter that enables/disables the ASN feature. If ASN is On, the console selects which key is to be used by the CIU. If ASN is Off, the CIU uses the default key. In either case, the CCGW simply accepts the outbound instruction from the console, accesses the tone table, and provides that appropriate station control command in terms of the outbound tone sequence to the CIU.

- A direct connection to the CIU. For details, see “Connecting the GGM 8000 or the MCG 8000 E&M Interface to a CIU for Advanced SECURENET Support” in the *GGM 8000 Hardware User Guide*.
- Positive mode control (PMC) tone tables and dual tone generation
The CCGW accesses the dual tone table provided by the LDAP server and passes the dual tone sequence to the DSP. The DSP then generates the dual tone sequence using the analog interface, and the CCGW generates dual tone commands for keying the channel or executing station control commands over the analog interface.
- Coded/clear input from the CIU/DIU
The CCGW detects the secure status (clear/coded) from the CIU via the E-lead state change and keeps this status in its channel profile. Whenever a digital conventional radio request ICP is sent to the zone controller, the CCGW retrieves the clear/coded state from its channel profile and includes it in the ICP message to the zone controller. When the channel is not configured for secure operation (the SecureMode parameter is set to Clear), the CCGW does not interpret the E-lead transition.



NOTE: How the GGM 8000 or the MCG 8000 CCGW detects the secure status (coded/clear) from the CIU depends on whether the CCGW is equipped with an analog/V.24 interface kit or the Enhanced Conventional Gateway option.

- GGM 8000s equipped with the analog/V.24 interface kit detect the secure status from the CIU via the E-lead state change.
- GGM 8000s equipped with the analog/V.24 interface kit detect the secure status from the CIU via the E-lead state change.
- GGM 8000s equipped with the Enhanced Conventional Gateway option detect the secure status from the CIU via pins 7 and 8 on the Audio, E&M connector.
- MCG 8000 detect the secure status from the CIU via pins 7 and 8 on the Audio, E&M connector.

40.6.10

Relay Keying

The relay keying feature allows you to use the M-lead on the CCGW E&M card as a diagnostic tool to monitor the base station or the console “keyed” state. You can do this by configuring the CCGW to set the M-lead (close the relay) on the E&M card to indicate that the base station or console transmitter is keyed. This configuration is performed on the LDAP server by using the “digital and relay” setting for the “method to key the transmitter” parameter. Relay keying is supported on digital conventional, mixed mode, and ACIM links. For these three channel types, the E&M card is used for transmission and reception of audio, but the E-lead and M-lead signals on each port are unused since all control signaling uses the digital command link, and keying is performed through digital messaging across the V.24 interface.

40.6.11

CCGW Support for Conventional Data DSR

The CCGW supports a complete DSR solution for conventional systems: conventional data DSR support and the conventional voice DSR. This solution provides links to two radio network gateways (RNGs), with different multicast addresses for broadcast on each RNG.

The LDAP database schema distinguishes LDAP server queries from previous releases. The CCGW no longer queries for radio network gateway (RNG) and multicast IP addresses in the Zones table. Instead, the CCGW locally calculates all IP addresses based on a calculation formula and populates the addresses in the respective tables. If the `DataDSRCapable` parameter is set to `Enable`, multicast, backup multicast, and RNG IP addresses are also calculated (and populated in the respective tables) for the backup zone core in a DSR system.

The CCGW calculates IP addresses which are used to establish LAPD connections to the RNG. If the `DataDSRCapable` parameter is set to `Enable`, the CCGW retries the LAPD connection with the primary and backup zone RNG/multicast IP addresses at 30-second intervals until an LAPD connection is established.

40.6.12

Conventional Talkgroups

A digital conventional (digital over V.24) or IP conventional (digital over IP) channel that is configured for conventional talkgroup operation is identified as a *talkgroup channel*.



NOTE: Conventional talkgroups are supported only on digital conventional and IP conventional channels on the GGM 8000 and the MCG 8000 hardware platform.

On a particular talkgroup channel, consoles affiliate to the talkgroup for which they want to receive audio and associate to the talkgroups for which they want to be notified about a call activity, but do not want to receive the audio. The zone controller is aware of which consoles are affiliated and associated with each conventional talkgroup and which conventional talkgroups are affiliated and associated with each console.

The CCGW is responsible for talkgroup call setup for site and audio routing. The ASTRO system does not know which of the conventional talkgroups assigned to a particular channel a subscriber wants to use until the subscriber transmits. Therefore, when the zone controller link to the CCGW is established, the CCGW affiliates all the conventional talkgroups for that channel to the zone controller, and the zone controller is then able to route audio to the correct channel for that conventional talkgroup.

The zone controller determines that a talkgroup is a conventional talkgroup when the CCGW affiliates the talkgroup. If the CCGW never affiliated a talkgroup, the zone controller assumes it is a trunking talkgroup. If the console affiliates a talkgroup before the CCGW does, the zone controller assumes it is a trunking talkgroup until the CCGW affiliates it, at which point the zone controller cleans up any parameters that are specific to trunking talkgroups and maintains parameters that are shared by trunking and conventional talkgroups. If a conventional talkgroup is deaffiliated by the CCGW, it is changed to a trunking talkgroup.

The current infrastructure supports up to 20 talkgroups (each identified by a talkgroup ID) per talkgroup channel. Once a talkgroup is assigned to a channel, it cannot be assigned to another channel, and it cannot be used for trunking. Among the 20 talkgroups supported per talkgroup channel, the following special talkgroups are also supported:

Everyone talkgroup

P25 standards-based talkgroup (\$FFFF) of which all talkgroups are a member. All subscribers unmute to the Everyone talkgroup. The console does not use the Everyone talkgroup directly; instead the console uses the Everyone talkgroup indirectly via the Channel Wide talkgroup.

Channel Wide talkgroup

Talkgroup ID that is designated to be converted (by the CCGW) to the Everyone talkgroup on the air interface. A Channel Wide talkgroup allows the console operator to talk to all subscribers on a channel regardless of what talkgroups the subscribers are using. Every channel uses a different talkgroup ID for the Channel Wide talkgroup. This enables a console operator to talk to everyone on a channel instead

of everyone on every channel. A console receiving on the Channel Wide talkgroup does not hear all the talkgroups on the channel; instead it hears all subscribers using the Everyone talkgroup as well as other consoles using the Channel Wide talkgroup for that channel.

 **NOTE:** Subscribers are not intended to directly use the talkgroup ID allocated to a Channel Wide talkgroup. If a subscriber does use that talkgroup ID, the talkgroup ID is repeated and a network management event is triggered.

The CCGW:

- Receives the Channel Wide talkgroup ID
- Converts the embedded “\$FFFF” Everyone talkgroup ID to the configured Channel Wide talkgroup ID for that channel and sends it to the controller and console; also converts the outbound Channel Wide talkgroup ID to the Everyone talkgroup
- Performs talkgroup affiliation processing
 - Affiliation processing begins when a talkgroup channel is created with talkgroup IDs, regardless of whether the channel is in service or out of service. Affiliations and deaffiliations are considered channel provisioning information and do not affect call processing.
 - When the Netcom link is established, the CCGW sends a Conventional Talkgroup Channel mobile synchronization key (MSK) Request inbound control packet (ICP) for each talkgroup channel to the zone controller or conventional site controller to check if the MSKs on the CCGW and the zone controller or conventional site controller are synchronized.
 - The zone controller or conventional site controller sends a Conventional Talkgroup Channel MSK Request acknowledgement outbound control packet (OCP) indicating whether or not the MSK is matched. If the MSK is matched, the CCGW does nothing (the talkgroups are already affiliated); if the MSK is not matched, the CCGW starts the affiliation processing for that channel.
 - When a talkgroup ID is deleted from a talkgroup channel, the CCGW performs deaffiliation processing by sending a Conventional Talkgroup Deaffiliation Request ICP for each talkgroup ID to the zone controller or conventional site controller.

EOS CCGW Service Support For Conventional Talkgroups

The following CCGW service parameters support the conventional talkgroups feature:

- The `ActiveChConf` and `CONFIguration` parameter displays include the configured `TalkgroupMode` (Talkgroup or Non Talkgroup).
- The `CMStatusdisplay` indicates whether an IP digital or V.24 digital channel is an IP digital talkgroup channel (TG_IP) or a V.24 digital talkgroup channel (TG_V24).
- The `CCGWLog` parameter identifies talkgroup events in the CCGW event log with the notation TG_IP or TG_V24.

For example:

```
04/11 20:04:45 TG_V24 Channel In Service - Channel 1
04/11 20:04:45 TG_IP Channel In Service - Channel 2
04/11 20:04:45 TG_IP Channel In Service - Channel 3
04/11 20:06:26 TG_V24 Audio Inactive Error - Channel 1
```

- The `AckDelayTimer` parameter displays the configured message Ack delay timer for talkgroup channels.
- The `TalkgroupMode` parameter indicates whether or not the specified channel is configured as a talkgroup channel.
- The `TGInfo` parameter provides a table of configured talkgroup parameters for all talkgroup channels.

For more information about these parameters, see “CCGW Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

40.6.13

Console Repeat

The console repeat feature is designed to enable dispatch console users to force a set of scanners to detect the carrier on a particular channel and remain on that channel as long as the carrier is present.

This allows users who are near the scanner to hear all the activity on that channel without interruption caused by scanning other channels. To accomplish this, console users activate the console repeat feature on the desired analog or MDC 1200 channel from their dispatch position's GUI, causing the transmitter on the channel to key up continuously. When the incident is over, the dispatch console user deactivates the console repeat feature from their dispatch position's GUI, which causes the transmitter on the channel to de-key. When the channel's carrier disappears, the scanners resume their normal scanning. The console repeat feature also includes a deactivation timer that automatically deactivates the feature after it was active for a configurable amount of time.

When a channel is configured for console repeat, the controlling console keys the station whenever there is an inbound subscriber call and causes the subscriber's inbound audio to be routed to the transmitter (repeated) through the CCGW. In addition, an MDC 1200 channel that is configured for console repeat can be configured to repeat or not repeat received MDC messages (the default is to repeat them).

EOS CCGW Service Support for Console Repeat

The following CCGW service parameters display the current console repeat configuration for analog and MDC 1200 channels:

MixerOnOffState

Indicates whether or not the channel is operationally in a console repeat call.

RepeatRcvdMDC

Indicates whether or not an MDC1200 channel configured for console repeat is configured to repeat received MDC messages.

For more information about these parameters, see "CCGW Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

40.6.14

Main/Alternate Channel Status Awareness

To prevent the CCGW from oversending call requests for digital channels and undersending call requests for analog channels:

- The CCGW reads the `Channel Role` parameter (which indicates whether or not the channel is in a Main/Alternate configuration) from the conventional channel LDAP table.
- The zone controller or conventional site controller sends a Main Alt State Command message to the CCGW to indicate the current state of the Main/Alternate channel (active or inactive).



NOTE: The zone controller or conventional site controller does not send the Main Alt State Command message when the channel is in a standalone configuration (the "Channel Role" is "Standalone"). A standalone channel is always active.

EOS CCGW Service Support for Main/Alternate Channel Status

The following CCGW service parameters enable you to view Main/Alternate channel pair information:

- The `ChannelRole` parameter indicates whether the channel is in a standalone or a Main/Alternate configuration.
- The `ChannelStatus` parameter indicates the Main/Alternate channel status (Active or Inactive).

- The `ChnlState` field on the `CMStatus` parameter display and the `MainAlt Status` field on the `DMStatus` parameter display indicate the Main/Alternate status (Active or Inactive) for the channels configured on the CCGW.

For more information about these parameters, see “CCGW Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

40.6.15

Duplex Patch with Two-Party Conference

The duplex patch with two-party conference (duplex patch) feature provides ASTRO® 25 system support for the duplex patch feature already supported on existing CENTRACOM Gold series systems. When two analog conventional or MDC 1200 channels are in a duplex patch, the non-transmitting users on both channels hear the sum of the audio from the transmitting user on their channel as well as the transmitting user on the other channel.

 **NOTE:** The duplex patch feature requires the MCG 8000 platform or the GGM 8000 gateway equipped with either a Low Density Enhanced Conventional Gateway module, a High Density Enhanced Conventional Gateway module, or an analog/V.24 interface kit. The MNR S2500 does not support the duplex patch feature.

The only allowable combinations for a duplex patch are two duplex patch capable analog or MDC 1200 conventional channels or a console telephone resource and a duplex patch capable analog or MDC 1200 capable channel. The CCGW that interfaces to the base station sums the transmit and receive audio for the analog or MDC 1200 conventional channels in the duplex patch.

During an incident that requires duplex patch:

- The dispatch console operator uses the console GUI to put the two duplex patch capable channels into a patch group.
- When audio is received on one of the channels, the console begins a patch transmit on both channels, which causes the transmitters on both channels to begin transmitting.
- The transmit audio on each channel becomes the sum of the received audio on that channel and the other channel in the duplex patch.
- The non-transmitting users on both channels hear the summation of the audio from the transmitting subscribers on both channels.
- Both channels continued to transmit until the duplex drop out delay timer expires or the dispatch console operator takes one or both of the patch group members out of the patch group. Both events cause both channels to dekey.

 **NOTE:** The duplex drop out delay timer starts when inbound activity ends on both channels in the duplex patch.
Only two resources can be in a duplex patch. If a third member is added, the patch changes from duplex to simplex mode.

EOS CCGW Service Support for Duplex Patch

The following CCGW service parameters support the duplex patch feature:

- The `SHOW -CCGW StatisticsCM` display includes the following information for duplex patch capable analog or MDC 1200 channels:
 - `Duplex Mode Start` – the number of times a duplex mode call was started on that channel.
 - `Duplex Mode End` – the number of times a duplex mode call was ended on that channel.
- The `CCGWLog` parameter includes an entry in a format similar to the following each time a duplex patch call starts or ends:

```
<Date> <Time> <Analog_Channel|Mdc_channel> Duplex Mode <Start/End> -Channel <id>
```

where:

- <Date> and <Time> are the date and time at which the event occurred,
- <Analog_Channel|Mdc_channel> specifies whether the call was on an analog or MDC 1200 channel,
- <Start/End> indicates whether the event was a duplex patch call start or end,
- <id> is the channel ID of the channel on which the call occurred.

For example:

```
04/11 20:04:45 Analog Channel Duplex Mode Start -Channel 2  
04/11 20:05:56 Analog Channel Duplex Mode End -Channel 2
```

40.7

CCGW Troubleshooting

This section provides information on how to use the CCGW service `IMBEETerr` and `UnkCallStartCnt` parameters and the inbound call event statistics on the `StatisticsCM` display for troubleshooting and system fine tuning.

For information about hardware troubleshooting, see the *MNR S2500 Hardware User Guide*, the *GGM 8000 Hardware User Guide*, or the *MCG 8000 System Gateway Feature Guide*.

40.7.1

Inbound Emergency Call Error Detection Threshold Configuration

An *emergency alarm* is a non-voice message sent by the subscriber unit to the fixed network equipment. It typically is used to inform dispatch of a life-threatening situation. The emergency alarm message is sent to a group, but it is also acknowledged by the fixed network equipment.

An *emergency call* provides a visual indication of the elevated call priority, displayed on the console and PTT log, and is used as a basis for preemption of low priority audio at the console to accommodate audio for the emergency call. Emergency calls are identified by a state in over-the-air protocol. This emergency state is included repeatedly (every 360 milliseconds) in emergency calls.

If the bit error rate of a call is high enough, in which case the audio is usually undecodable or unintelligible, the call can be misinterpreted as an emergency call. To ensure that the emergency state of a call is reliable and to avoid cases in which a call with a high error rate is misinterpreted as an emergency call, a configurable error detection threshold was implemented on the MNR S2500 router, the GGM 8000 gateway and the MCG 8000 platform.

The `IMBEETerr` parameter configures a total improved multiband excitation error threshold (IMBE ET) error count threshold for each V.24 digital, IP conventional, or mixed mode channel configured on the MNR S2500, GGM 8000 or MCG 8000. This threshold defines the maximum number of IMBE errors that can be present in an inbound PTT in order for an emergency state received in that PTT to be considered reliable and the call to be interpreted as an emergency call. If the error rate indicates that the call control information is unreliable, the emergency indication is not passed to the console, until the call error rate indicates that the call control information is reliable.

The `IMBEETerr` parameter is configured with a default value of 50, which should be appropriate for most systems. In general, an inbound PTT with an IMBE error count greater than 50 contains the audio that is unintelligible or undecodable. If necessary, you can use the CCGW service `IMBEETerr` parameter to configure the IMBE ET error count threshold for a different value. For more information, see "CCGW Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

40.7.2

Inbound Call Statistics to Console Unknowns Comparison

In normal operation, "Unknowns" (the calls with the source ID "Unknown") displayed on the console for a particular channel should be calls that start with an undetectable call control information. This is caused by RF conditions or data integrity issues between the RF equipment and the digital CCGW (MNR S2500) or site gateway (conventional channel interface) (GGM 8000 or MCG 8000 platform). This section provides procedures that you can perform to compare the Unknowns displayed on the MCC 7500 console to MNR S2500, GGM 8000 or the MCG 8000 statistics for digital and MDC 1200 conventional channels. These comparisons can be used to determine if the "Unknown" status is displayed appropriately and under the right conditions.



NOTE: An "Unknown" displayed on the console can also be caused by site link failures not measured by the digital CCGW or site gateway (conventional channel interface) and, at a very low level, by normal network congestion.

40.7.2.1

Comparing the Inbound Call Unknown Start Count to Console Unknowns – Digital Channels



CAUTION: This procedure was developed for system fine-tuning and troubleshooting purposes only and is not intended to be performed as a part of normal system operation.

For calls with very low signal strength or very high interference, call control information may not be recoverable due to excessive errors in the signal. The source IDs for these calls may be displayed on the MCC 7500 channel resource display as "Unknown" or as a random source ID. The type of call (group or selective), destination ID, secure mode, and other control information may be in error for as long as the call experiences a high error rate. On V.24 digital conventional channels, IP conventional channels, and mixed mode channels with digital calls, the digital CCGW or site gateway (conventional channel interface) counts the number of calls per channel that start with undetectable call control information. This count is called the inbound call unknown start count.



NOTE: Inbound analog voice calls received on mixed mode channels are also displayed with source ID Unknown on the MCC 7500 channel resource display and are included in the inbound call unknown start count.

Procedure:

1. Reset the inbound call unknown start count for the desirable channel by entering:

```
FLush !c<channel_ID> -CCGW UnkCallStartCnt
```

where <channel_ID> is the channel identifier (1-4 for V.24 digital or mixed mode channels; 21-30 for the IP conventional channels) of the channel for which you want to view the inbound call unknown start count.



NOTE: The inbound call unknown start count does not wrap and can be high if it was not reset recently.

Mixed mode and IP conventional channels are only supported on the GGM 8000 gateway or the MCG 8000 platform. The MNR S2500 only supports the V.24 digital channels.

If you do not know the channel identifier for a channel, you must issue the `SHow -CCGW ChID` command to display a list of channel identifiers for all conventional channels configured on the device.

Step example:

To reset the inbound call unknown start count for the V.24 digital conventional channel 2, enter:

```
FLush !c2 -CCGW UnkCallStartCnt
```

To reset the inbound call unknown start count for the IP conventional channel 23, enter:

```
FLush !c23 -CCGW UnkCallStartCnt
```

2. After the data collection is complete, to view the inbound call unknown start count for the desirable channel, enter:

```
SHow !c<channel_ID> -CCGW UnkCallStartCnt
```

where <channel_ID> is the channel identifier (1-4 for V.24 digital or mixed mode channels, 21-30 for IP conventional channels) of the channel for which you want to view the inbound call unknown start count.



NOTE: In a properly installed system, statistically accurate data requires days of data collection at average call rates. Heavily used channels require hours to days of data collection. Lightly used channels require weeks of data collection.

Mixed mode and IP conventional channels are only supported on the GGM 8000 gateway or the MCG 8000 platform. The MNR S2500 only supports the V.24 digital channels.

If you do not know the channel identifier for a channel, you must issue the `SHow -CCGW ChID` command to display a list of channel identifiers for all conventional channels configured on the device.

Step example:

To view the inbound call unknown start count for the V.24 digital conventional channel 2, enter:

```
SHow !c2 -CCGW UnkCallStartCnt
```

To view the inbound call unknown start count for the IP conventional channel 23, enter:

```
SHow !c23 -CCGW UnkCallStartCnt
```

3. Compare the inbound call unknown start count to the number of "Unknowns" displayed on the MCC 7500 console since the inbound call unknown start count was reset.
4. If necessary, repeat this procedure for each channel that you want to investigate.

40.7.2.2

Comparing the Inbound Call Start Event Counts to Console Unknowns (MDC 1200 Channels)



CAUTION: Consult a Motorola Solutions field service engineer before performing this procedure. The procedure was developed for system fine-tuning and troubleshooting purposes only and is not intended to be performed as a part of normal system operation.

The CCGW `SHow -CCGW StatisticsCM` display reports the following inbound call start event counts per MDC 1200 channel:

- `LOBL CallStart Cnt` – The number of calls for which a LOBL signal was detected.
- `UnkID CallStart Cnt` – The number of calls for which no source ID was detected.
- `Decode Error Cnt` – The number of calls for which MDC 1200 signaling was detected but the signaling was undecodable plus the number of undecodable supplementary signaling messages.



NOTE: MDC 1200 channels are supported only on the GGM 8000 gateway or the MCG 8000 platform. The MNR S2500 does not support the MDC 1200 channels.

Procedure:

1. Reset the call manager statistics by entering:

```
FLush -CCGW StatisticsCM
```

The inbound call event counts can be high if they were not reset recently. You should flush the counts before you begin the data collection.

2. After the data collection is complete, view the call manager statistics by entering:

SHoW -CCGW StatisticsCM



NOTE: In a properly installed system, statistically accurate data requires days of data collection at average call rates. Heavily used channels require hours to days of data collection. Lightly used channels require weeks of data collection.

The MDC 1200 inbound call event counts are listed at the end of the call manager statistics display, in a format similar to the following:

INBOUND CALL EVENTS				
LOBL CallStart Cnt	0	0	0	0
UnkID CallStart Cnt	0	0	0	0
Decode Error Cnt	0	0	0	0

The four columns provide the inbound call event counts for channel IDs 1, 2, 3, and 4 respectively.

3. Compare the inbound call event counts on the channel you are investigating to the number of "Unknowns" displayed on the MCC 7500 console for that channel since the call manager statistics were reset.

A high LOBL CallStart Cnt may indicate unexpected interference on the channel.

A high UnkID CallStart Cnt or Decode Error Cnt may indicate analog calls without signaling, trailing MDC 1200 calls, RF site problems that are reducing the coverage or signal strength, or unexpected interference on the channel.

The Decode Error Cnt statistic counts both MDC 1200 voice calls for which MDC 1200 signaling was detected but the PTT ID fails to decode and undecodable supplementary signaling messages. MDC 1200 voice calls for which the PTT ID fails to decode are logged as Unknowns on the MCC 7500 console. Undecodable supplementary signaling messages, on the other hand are not logged on the console since they are not passed up to the zone core by the site gateway (conventional channel interface). You are comparing the Decode Error Cnt to the console Unknowns.

The MCC 7500 console logs an Unknown for each voice call initiated by a radio subscriber that does not have the PTT ID enabled. However, these calls are not included in the Decode Error Cnt because no MDC 1200 signaling is transmitted by the radio in these cases.

Chapter 41

Overload Protection Configuration

This chapter describes how to configure overload protection on EOS devices. The EOS Overload Protection Management (OPM) feature monitors device CPU utilization, memory usage, and queue drops, and informs registered applications (OPM clients) when any of these parameters meets or exceeds a configured threshold.

EOS software applications are designed such that, under most circumstances, memory usage is not an issue, and CPU utilization is within acceptable margins based on the required messages-per-hour rate. For example, the GGSN application performs adequately by limiting the number of simultaneous contexts open at any one point in time. However if a packet or data storm occurs, the packet rate may increase above the stated required maximum. In such cases, limited data loss can occur in arbitrary places in the data pipeline of the EOS device.

The purpose of the overload protection management (OPM) feature is to monitor the following parameters and to inform registered applications (OPM clients) when any of the parameters meets or exceeds a configured threshold:

- CPU utilization
- Memory utilization
- Queue drops

Based on this information, the registered applications can limit their activity during overload conditions, thereby preventing potential data loss.



NOTE: Queue drops are accumulated for Ethernet ports only.

OPM currently supports the following two clients: GGSN and SNMP.

You set the OPM thresholds for each of the monitored parameters to “high”, “medium,” or “low” using either a network management application or the EOS CLI. Each threshold level (high, medium or low) is associated with internally-configured high-water and low-water marks. OPM clients are called when either of the following events occurs:

- **Overload** – The parameter meets or exceeds the high-water mark associated with the specified threshold level (high, medium, or low).
- **Normal** – The parameter falls back below the low-water mark associated with the specified threshold level (high, medium, or low).



NOTE: The EOS implementation of overload protection management incorporates a 30-second Holddown timer which prevents the EOS device from entering or leaving Overload within that time period. For example, if a device enters Overload state at 3:30.15, the Holddown timer starts, and the device cannot exit Overload state until the 30-second Holddown period expires — in this example, at 3:30:45.

41.1

Configuring EOS Devices for Overload Protection

You can use this procedure to configure overload protection on EOS devices by using the EOS CLI.

For details about the EOS parameters used to configure overload protection, see “SYStem Service Parameters” and “GGSN Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Procedure:

1. Enable overload protection and set the overload protection threshold level for the parameters for which you enabled overload protection by entering one of all the following commands:

- `ADD -SYS OverLoadProtection CpuUtilization [HIgh | MEdition | LOw]`
- `ADD -SYS OverLoadProtection MemoryUsage [HIgh | MEdition | LOw]`
- `ADD -SYS OverLoadProtection QueueDrops [HIgh | MEdition | LOw]`

Step example:

To enable overload protection and set a low overload protection threshold for CPU utilization and to enable overload protection and set a medium overload protection threshold for memory usage, enter:

```
ADD -SYS OverLoadProtection CpuUtilization LOw
```

```
ADD -SYS OverLoadProtection MemoryUsage MEdition
```

2. Specify the overload protection parameters for which you want the GGSN client to be registered by entering:

```
SETDefault -GGSN OverloadProtection = ([ALL | None], [CpuUtilization |  
NoCpuUtilization] [MemoryUsage | NoMemoryUsage] [QueueDrops | NoQueueDrops])
```

It is not necessary to register the SNMP client for overload protection events. When overload protection is enabled on an EOS device, the device sends a trap whenever it enters an overload condition (as defined by the configured thresholds) or returns to normal (non-overload) state. For details, see [Supported MIBs and Traps on page 845](#).

Step example:

To register GGSN for CPU utilization and memory usage overload protection, enter:

```
SETDefault -GGSN OverloadProtection = (CpuUtilization, MemoryUsage)
```

To register GGSN for all three overload protection parameters, enter:

```
SETDefault -GGSN OverloadProtection = ALL
```

41.2

Overload Protection Statistics

You can retrieve the following overload protection statistics:

- CPU Utilization – The current value for CPU utilization.
- Memory Usage – The current value for memory usage.
- Queue Drops – The current value for queue drops.
- The number of times the device went into Overload state.
- A log of the last 5 times the device went into Overload state, including the following information:
 - Start Time – The time at which the device went into Overload state.
 - End Time – The time at which the device exited from Overload state.
 - Duration – The amount of time the device was in Overload state.
 - Cause In – The reason that the device went into Overload state (CPU utilization, memory usage, or queue drops).
 - Cause Out – The reason that the device exited from Overload state (CPU utilization, memory usage, or queue drops).

To view overload protection statistics, you must use:

SHow -SYS OverLoadProtection Stats

Chapter 42

IP Link Converter (IPLC) Functionality Configuration

The IPLC feature allows users with fielded CENTRACOM Gold Elite and MCC 5500 consoles to replace circuit connectivity with IP connectivity and transition legacy digital/mixed mode RF channels to G series RF channels. Users can also retain existing circuit backhaul networks (T1s/E1s).

The IPLC feature is supported on GGM 8000 gateways running Enterprise OS (EOS) software version 16.6 or higher (which introduces an IPCc service) and equipped with one of the following optional modules:

- High Density Enhanced Conventional Gateway module (provides 8 analog ports and 8 V.24 ports)
- Low Density Enhanced Conventional Gateway module (provides 4 analog ports and 4 V.24 ports)

The IPLC feature uses the IP transport capabilities of the GGM 8000 gateway to position users to migrate to ASTRO® 25 IP-based systems. In addition, the IPLC feature utilizes the following conventional channel gateway (CCGW) functionality supported on the GGM 8000:

- V.24/AIS (ASTRO® 25 infrastructure signaling) protocols plus 4-wire interface to the digital interface unit (DIU)
- AIS/XIS (X-Zone infrastructure signaling) conversion
- Support for IP and IP/XIS transport
- Conversion from 4-wire analog to G.728 packets to transport audio over IP
- Support for IP and analog comparators in a mixed mode configuration
- Integrated voice and data (IV&D) to support AIS/XIS data applications

42.1

IPLC Feature Overview

A GGM 8000 with IPLC functionality connects legacy consoles to new base stations and/or comparators to allow migration to:

- Voted channels.
- IP transport between the DIU and the RF site.

GGM 8000s with IPLC functionality may be deployed as standalone boxes or in pairs.

When deployed as a standalone box, the GGM 8000 with IPLC functionality is located at either the site or the core and singlehandedly connects the DIU at the core to the IP base station. If located at the core, the GGM 8000 with IPLC functionality is connected to the base station over IP backhaul.

In a paired IPLC topology:

- One GGM 8000 with IPLC functionality is co-located with the console at the core. The core GGM 8000 with IPLC functionality is connected to the DIU.
- One GGM 8000 with IPLC functionality is located at the site. The site GGM 8000 with IPLC functionality is connected to the base station or comparator.

- The two GGM 8000s with IPLC functionality (IPLC peers) are connected over IP backhaul.



NOTE: The IPCc service `BAckhaul` parameter is a per-channel parameter that specifies whether or not the channel is connected over backhaul to a peer GGM 8000 with IPLC functionality:

- `BAckhaul = yes` means that this channel is part of an IPLC topology that consists of two GGM 8000s with IPLC functionality, connected through an IP link and operating as a pair to provide the IPLC solution (paired IPLC topology).
- `BAckhaul = no` means that this channel is part of an IPLC topology that consists of a standalone GGM 8000 with IPLC functionality providing the IPLC solution (standalone topology).

A single GGM 8000 can be configured with multiple channels, some configured with `BAckhaul = yes` and some configured with `BAckhaul = no`.

The key components of the IPLC feature in a standalone topology are an *end-to-end channel* that connects the DIU to the base station, and spans a *single GGM 8000 with IPLC functionality* containing *upstream and downstream CPorts*.

The key components of the IPLC feature for a paired topology are an *end-to-end channel* that connects the DIU to the base station and spans *two GGM 8000s with IPLC functionality*, each of which supports an *upstream and a downstream channel* and an *upstream and a downstream CPort*.

End-to-end channel

End-to-end, fixed source/destination traffic stream between a legacy DIU and a base station or a comparator. An end-to-end channel may pass through one or two GGM 8000s with IPLC functionality.

An end-to-end channel is associated with one of four end-to-end channel types:

IP-Digital

Purely digital flow connecting a legacy DIU to an IP base station

V24-Digital

Purely digital flow connecting a legacy DIU to a V.24 base station

IP-Mixed

Mixed mode (digital and analog) flow connecting a legacy DIU to a mixed mode base station

V24-Mixed

Mixed mode (digital and analog) flow connecting a legacy DIU to a mixed mode base station



NOTE:

When used in reference to an end-to-end channel, the term *channel type* represents the type of system-wide, end-to-end traffic flow connecting a legacy console subsystem and a site – not a segment of this flow (comprised of an upstream and downstream CPort) that physically resides in the GGM 8000 with IPLC functionality.

Channel

Segment in an end-to-end channel that resides on a single GGM 8000 with IPLC functionality; it is created by linking two CPorts.

A channel may have one of three roles:

Upstream channel

Exists on the GGM 8000 with IPLC functionality co-located at the core with the console; it is designed to be paired with a peer channel (its downstream counterpart) that exists on another GGM 8000 with IPLC functionality co-located with the base station or comparator at the site.

Downstream channel

Exists on the GGM 8000 with IPLC functionality co-located at the site with the base station or comparator. It is designed to be paired with a peer channel (its upstream counterpart) that exists on another GGM 8000 with IPLC functionality co-located with the console at the core.

Standalone channel

Single-handedly connects the DIU at the core to an IP base station.

CPort

Logical port residing in an GGM 8000 with IPLC functionality.

A CPort may have one of two roles:

Upstream CPort

Always core-facing. It receives upstream (toward core) traffic from the downstream CPort of the same channel (in the same GGM 8000 with IPLC functionality), and egresses it toward the console. The upstream CPort also ingresses downstream (toward site) traffic originated from the console and forwards it to the downstream port of the same channel.

Downstream CPort

Always site-facing. It receives downstream (toward site) traffic from the upstream CPort of the same channel (in the same GGM 8000 with IPLC functionality) and egresses toward the base station. It also ingresses upstream (toward core) traffic originated from the base station and forwards it to the upstream CPort on the same channel.

An upstream CPort linked to a downstream CPort defines a channel.

The upstream CPort is always core-facing, and the downstream CPort is always site-facing, regardless of the GGM 8000 with IPLC functionality or channel in which the CPort resides.

A CPort may be one of four types: V24, IP, V24-Mixed, or IP-Mixed. Between the pair of CPorts that constitutes a channel in the GGM 8000 with IPLC functionality, the format of the traffic may be changed between ASTRO® 25 infrastructure signaling (AIS) and X-Zone infrastructure signaling (XIS), depending on the types of the two ports.

V24

Physical V.24 port

IP

Virtual IP port

V24-Mixed

Union of a physical V.24 port and a physical analog 4-wire port

IP-Mixed

Union of a virtual IP port and a physical analog 4-wire port

Figure 196: Standalone IPLC Topology: IP-Digital (GGM 8000 with IPLC functionality located at core)

Standalone IPLC topology with end-to-end channel type

IP-Digital and the GGM 8000 with IPLC functionality located at the core. In this topology, the end-to-end channel type must be configured as IP-Digital.

For this channel, the GGM 8000 with IPLC functionality operates as a standalone box (IPCC service `BAckhaul` parameter is set to `no`). The GGM 8000 with IPLC functionality is connected to a DIU through a V.24 link. The GGM 8000 with IPLC functionality is connected to an IP digital base station over IP backhaul.

This topology is also referred to as a *standalone IP backhaul topology*.

In a voting scenario, the topology also includes a GCM 8000 comparator connected to the IP base station.




Figure 197: Standalone IPLC Topology: IP-Mixed

Standalone IPLC topology with end-to-end channel type

IP-Mixed. In this topology, the GGM 8000 with IPLC functionality must be located at the site.

For this channel, the GGM 8000 with IPLC functionality operates as a standalone box (IPCc service BACKhaul parameter is set to no).

The GGM 8000 with IPLC functionality is connected to a DIU through a circuit backhaul connection (typically through a channel bank).

The GGM 8000 with IPLC functionality box is connected to an MLC 8000 analog comparator through an analog link and to a GCM 8000 comparator through an IP link. An IP link connects the comparator to an MLC 8000 link converter, which in turn is connected to a mixed mode base station through a V.24 link paired with an analog link.

 **NOTE:** The IPLC topology with end-to-end channel type IP-Mixed is supported only when the GGM 8000 with IPLC functionality is located at the site. If the GGM 8000 with IPLC functionality is located at the core, the end-to-end channel type for a standalone topology must be IP-Digital.





Figure 198: Paired IPLC Topology: IP-Mixed

Paired IPLC topology with end-to-end channel type IP-Mixed.
For this channel, the two GGM 8000s with IPLC functionality are connected through an IP link (IPCc service BACkhaul parameter is set to yes).
The GGM 8000 with IPLC functionality at the core is connected to a DIU through a V.24 link paired with an analog link.

The GGM 8000 with IPLC functionality at the site is connected to an MLC analog comparator through an analog link and to a GCM 8000 comparator through an IP link. An IP link connects the comparator to an MLC 8000 link converter, which in turn is connected to a mixed mode base station through a V.24 link paired with an analog link.



* Configured on GGM 8000 with IPLC functionality located at core
** Configured on GGM 8000 with IPLC functionality located at site

Figure 199: Paired IPLC Topology: V24-Digital

Paired IPLC topology with end-to-end channel type V24-Digital.

For this channel, the two GGM 8000s with IPLC functionality are connected through an IP link (IPCc service BAckhaul parameter is set to yes).

The GGM 8000 with IPLC functionality at the core is connected to a DIU through a V.24 link.

The GGM 8000 with IPLC functionality at the site is connected to a V.24 digital base station through a V.24 link.





Figure 200: Paired IPLC Topology: V24-Mixed

Paired IPLC topology with end-to-end channel type V24-Mixed.
For this channel, the two GGM 8000s with IPLC functionality are connected through an IP link (IPCC service BAckhaul parameter is set to yes).
The GGM 8000 with IPLC functionality at the core is connected to a DIU through a V.24 link paired with an analog link.
The GGM 8000 with IPLC functionality at the site is connected to a mixed mode base station through a V.24 link paired with an analog link.



42.1.1 Channel Type-to-CPort Type Mapping

Standalone channels do not support V24-Digital and V24-Mixed end-to-end channel types.
Standalone channels configured on GGM 8000s located at the core do not support the IP-Mixed end-to-end channel type because, in this case, the downstream CPort on the standalone channel is connected via the IP network to a GTR 8000 base station or GCM 8000 comparator at the site. Since the GGM 8000 is not physically located at the site, you cannot also make a direct connection from the analog port on the downstream CPort to an MLC 8000 analog comparator at the site. When the GGM 8000 is located at the site, however, the downstream CPort connects directly to a GTR 8000 base station or GCM 8000 comparator at the site via the IP (Ethernet) port, and a direct connection from the analog port to the MLC 8000 analog comparator at the site is also possible.

Table 66: Channel Type-to-CPort Type Mapping (Standalone Channel at Core)

Standalone Channel		
End-to-End Channel Type	Upstream CPort Type	Downstream CPort Type
IP-Digital	V24	IP

Table 67: Channel Type-to-CPort Type Mapping (Standalone Channel at Site)

Standalone Channel		
End-to-End Channel Type	Upstream CPort Type	Downstream CPort Type
IP-Digital	V24	IP
IP-Mixed	V24-Mixed	IP-Mixed

Table 68: Channel Type to Upstream and Downstream CPort Type Mapping (Paired IPLC Topologies)

End-to-End Channel Type	Upstream Channel		Downstream Channel	
	Upstream CPort Type	Downstream CPort Type	Upstream CPort Type	Downstream CPort Type
IP-Mixed	V24-Mixed	IP	IP	IP-Mixed
V24-Digital	V24	IP	IP	V24
V24-Mixed	V24-Mixed	IP	IP	V24-Mixed

42.1.2

V.24 and Analog Port Pairing

The following table lists the V.24 and analog port pairs used in the CPort configurations for the following end-to-end channel types:

- IP-Mixed (standalone topology; upstream CPort)
- IP-Mixed (paired topology; upstream CPort on upstream channel)
- V24-Mixed (paired topology; upstream CPort on upstream channel and downstream CPort on downstream channel)

Table 69: V.24 and Analog Port Pairing

V.24 Port	Analog Port
6A	8A
6B	8B
7A	8C
7B	8D
10A	12A
10B	12B
11A	12C
11B	12D

42.1.3 IPLC LEDs

Enhanced Conventional Gateway Module Analog Interface Status LEDs

Each analog interface on the Enhanced Conventional Gateway module features two bi-color status LEDs located in the lower left and lower right corners of the “Audio, E&M” (yellow) RJ45 connectors. Only the right-hand LED is operational. The left-hand LED is always off.

 **NOTE:** The IPLC feature does not use the GGM 8000 Enhanced Conventional Gateway module “Analog I/O” (white auxiliary) ports (9D through 9A and 13D through 13A).

Figure 201: Low Density Enhanced Conventional Gateway Module Analog LEDs





Figure 202: High Density Enhanced Conventional Gateway Module Analog LEDs



Each analog interface status LED indicates the status of one of the four analog interfaces as follows:

Steady Green

The analog interface is connected, the link to the DIU, MLC 8000 analog comparator, or base station is up, and the channel corresponding to this analog interface is enabled and idle.

Blinking Green

The analog interface is connected, the link to the DIU, MLC analog comparator, or base station is up, the channel corresponding to this analog interface is enabled, and audio packets are flowing.

Yellow

IPLC functionality is not enabled; the DSP is out of service; or a fault was detected on the channel corresponding to this analog interface.



NOTE: When the GGM 8000 boots up, the analog interface status LEDs are amber until the first channel using an analog interface is configured. Once a channel using an analog interface is configured, the analog interface LEDs are off and turn amber if an enabled/active channel fails. The GGM 8000 with IPLC functionality does not support the “partially available” status supported by the conventional channel gateway (CCGW). If the channel is configured as mixed mode and either the GGM 8000 analog interface link or the GGM 8000 V.24 interface link is down, then the mixed mode channel is considered down.

Off

IPLC functionality is enabled, but the channel corresponding to this analog interface is disabled or is not configured to use an analog interface (for example it is a V.24 digital channel).

Enhanced Conventional Gateway Module V.24 Interface Status LEDs

Each V.24 interface on the Enhanced Conventional Gateway module features two bi-color (green/amber) status LEDs in the lower left and right corners of the “V.24” (blue) RJ45 connectors. Only the right-hand LED is operational. The left-hand LED is always off.

Figure 203: Low Density Enhanced Conventional Gateway Module V.24 LEDs





Figure 204: High Density Enhanced Conventional Gateway Module V.24 LEDs



Each V.24 interface status LED indicates the status of one of the V.24 interfaces as follows:

Steady Green

The V.24 interface is connected, the link to the DIU, base station, or comparator is up, and the channel corresponding to this V.24 interface is enabled and idle.

Blinking Green

The V.24 interface is connected, the link to the DIU, base station, or comparator, is up, the channel corresponding to this V.24 interface is enabled, and audio packets are flowing.

Yellow

The IPCc service is enabled, and the corresponding channel is configured to use the V.24 interface (it is a V.24 digital or mixed mode channel), but either the V.24 interface is not connected or the V.24 interface is connected and the link is down.

Off

The IPCc service is disabled or no channel was configured to use this V.24 interface.

42.1.4

IPLC Alarms

The GGM 8000 with IPLC functionality provides alarm status to the MOSCAD RTU by using an E-lead relay on one of the analog ports to signal a digital input port on a MOSCAD RTU unit. A relay closure occurs and a LED lights under the following conditions:

- Any ALMT link is out of service
- Any DIU link is out of service
- Any backhaul link is out of service
- Any base station link is out of service



NOTE: Which relay is turned on/off is hard-coded in the order the ports initially come up:

- 8A - ALMT link
- 8B - DIU link
- 8C - backhaul link
- 8D - base station link

To connect the GGM 8000 to a MC-EDGE NFM RTU, you must use the pinouts listed in the following table.

Table 70: GGM 8000-to-MOSCAD NFM XC RTU Pinouts

GGM 8000		MOSCAD NFM XC RTU	
Pin	Signal Name	Pin	Signal Name
3	E-lead	(RJ45 ports 45-48) pin 1	DI 48
6	Signal Ground	(RJ45 ports 45-48) pin 2	Common 45-48

To connect the GGM 8000 to a MOSCAD SDM 3000 RTU, you must use the pinouts listed in the following table.

Table 71: GGM 8000-to-MOSCAD SDM 3000 RTU Pinouts

GGM 8000		MOSCAD NFM XC RTU	
Pin	Signal Name	Pin	Signal Name
3	E-lead	(RJ45 ports 1-4) pin 8	DI 4
6	Signal Ground	(RJ45 ports 1-4) pin 7	Common 1-4



IMPORTANT: You must set the RTU up as a wet input (providing loop current). The wet configuration is critical since the GGM 8000 does not provide loop current on the E-lead relay output.

RemotePeerInfo Display

When you issue the `SHoW -IPCc RemotePeerInfo` command, EOS displays the status of the “echo” exchange and prints the remote channel configuration locally if possible, in a display similar to the following.

```
Remote Peer Info: Packet contents unequal.
Retnote Peer Info: Packet comparison failed.
Remote Peer Info: Non-backhaul link on remote peer channel is down.
```

```

-----Remote Channel Info: -----
Channel-ID Index EndToEndChType Role Enabled State Upstream-Cport Downstream-Cport
-----
1 0 Ip-Digital US Yes Down 2 3
-----Upstream Cport-----
--<CR> to continue, Q to quit--
CportID: 2
CportIdx: 0
CportType: IP-mode
CportRole: Upstream Cport
Parent ChannelID: 1
Digital Path Label: N/A
Digital Path Number: 4094
Local IPAddr: 10.22.3.30
Peer IPAddr: 10.22.3.10
-----Downstream Cport-----
CportID: 3
CportIdx: 1
CportType: IP-mode
CportRole: Downstream Cport
Parent ChannelID: 1
Digital Path Label: N/A
Digital Path Number: 4093

```

```
Local IPAddr:      10.22.4.30
Peer IPAddr:       10.22.4.10
```



NOTE: The RemotePeerInfo command does not apply to standalone IPLC topologies.

You can use the information returned by the `SHoW RemotePeerInfo` command to troubleshoot as follows:

- Check the `SHoW -IPCc RemotePeerInfo` display for data integrity issues; for example:
Packet contents unequal
Packet comparison failed
If the `SHoW -IPCc RemotePeerInfo` display does not indicate data integrity issues, the backhaul-facing CPorts in the channel at both the remote and the local GGM 8000s with IPLC functionality are fully functional.
- Issue the `SHoW -IPCc CHannel` command and determine if the local channel is in the UP state.
- If the local channel is in the UP state, check the `SHoW -IPCc RemotePeerInfo` display to see if the remote channel is in the UP state.
If the local channel is in the UP state, but the remote channel is in the DOWN state, then the non-backhaul-facing CPort in the remote channel is not functioning. Check the connectivity of the site link (on the non-backhaul CPort) and of the base station across the site link.

The `SHoW -IPCc RemotePeerInfo` display includes the following information:

- “echo” exchange data, displayed in the following way:
 - A byte in the local buffer is printed in hex, followed by a byte in the same position from the reply buffer, which is followed by a STATUS byte.
If the two bytes match, the STATUS byte is printed as a space.
If the two bytes do not match, the STATUS byte is printed as an asterisk (*).
 - 17 bytes in both buffers are printed in each line.
- Information about the remote peer, including:
 - Data integrity check – whether or not the contents of the packet returned by the peer were the same as the contents of the packet sent to the peer and whether the packet comparison succeeded or failed
 - Status of the backhaul or non-backhaul link on the remote peer channel
- Configuration information about the channel (and its child CPorts) configured on the remote peer, including:
 - Channel configuration information:
 - Channel ID** – The channel ID
 - Index** – The 0-based index of the channel in the channel table
 - End-to-end channel type** – IP-Digital, IP-Mixed, V24-Digital, or V24-Mixed.
 - Role** – US (upstream) or DS (downstream)
 - Enabled** – Yes or No
 - State** – Up or Down
 - Upstream-Cport** – CPort ID of the upstream CPort
 - Downstream-Cport** – CPort ID of the downstream CPort
 - CPort configuration information (for both the upstream and downstream CPort):
 - CportID** – The CPort ID
 - Index** – The 0-based index of the CPort in the CPort table
 - CportType** – IP, IP-Mixed, V24, or V24-Mixed.
 - CportRole** – Upstream or Downstream
 - Parent ChannelID** – The channel ID of the parent channel

Digital Path Label – The path label associated with the V.24 interface on which the channel is configured. For IP-Digital and IP-Mixed end-to-end channel types, this field lists “N/A”.

Digital Path Number – The internal representation of the digital path label for a V.24 digital port - or the IP digital virtual port number

Local IPAddr – The local IP address configured on the remote peer. This should be the same as the remote peer IP address configured on the local GGM 8000 with IPLC functionality.

Peer IPAddr – The remote peer IP address configured on the remote peer. This should be the same as the local IP address configured on the local GGM 8000 with IPLC functionality.

42.2

Configuring IPLC Functionality

You can use this procedure to manually configure the GGM 8000 to support IPLC functionality by entering CLI commands line by line. You can also create and load an ASCII configuration (`boot.cfg`) file. For details, see [IPLC Configuration Templates on page 790](#).

The IPCc service is implemented with no persistence over reboots. If you configure the GGM 8000 to support IPLC functionality by creating and loading an ASCII configuration (`boot.cfg`) file, you can recover the configuration by reloading the `boot.cfg` file. However by default, any additional configuration changes you make to the GGM 8000 after it is operational are not added to the `boot.cfg` file. This means that when the GGM 8000 is rebooted or power cycled, the additional configurations changes are not effective. You can recover the IPLC configuration after a reboot or power cycle using the ASCII Capture feature, which is enabled by default. This feature captures all the successful configuration changes that you make on the GGM 8000 using the `SETDefault`, `ADD`, and `DELeTe` commands and saves them in an ASCII text file (`capture.cfg`) in the configuration directory.

Procedure:

1. Configure box-level parameters, by performing the following actions:

- a. Configure the device to display a logon banner with required notice and consent information by entering:

```
SETDefault -SYS BannerString = "Use of this device and any related service
is your consent to all associated terms and conditions, including consent to
monitoring and disclosure provisions."
```



NOTE: Do not include any carriage returns when entering the banner string.

- b. Enable IPLC functionality on the GGM 8000 by entering:

```
SETDefault -IPCc CONTrol = Enable
```

- c. If the GGM 8000 with IPLC functionality is located at the core, specify the physical location of the GGM 8000 with IPLC functionality by entering:

```
SETDefault -IPCc Location = Core
```

- d. If the GGM 8000 with IPLC functionality is located at the site, specify the physical location of the GGM 8000 with IPLC functionality by entering:

```
SETDefault -IPCc Location = Site
```



NOTE: You must configure the box-level parameters before proceeding to configure the channels and CPorts.

2. Perform one of the following actions:

- If you configure a paired IPLC topology, repeat [step 1](#) for the second GGM 8000 with IPLC functionality.

- If you do not configure a paired IPLC topology, go to [step 3](#).



NOTE: For paired IPLC topologies, you must configure the box-level parameters on both GGM 8000s with IPLC functionality before proceeding to configure the channels and CPorts on the two boxes.

3. Add and configure channels and CPorts by performing one of the following procedures:
 - [Adding and Configuring Channels for IPLC Standalone Topology – IP-Digital \(GGM 8000 with IPLC Functionality at Core\) on page 770](#)
 - [Adding and Configuring Channels for IPLC Standalone Topology – IP-Mixed on page 772](#)
 - [Adding and Configuring Channels for IPLC Paired Topology – IP-Mixed on page 775](#)
 - [Adding and Configuring Channels for IPLC Paired Topology – V24-Digital on page 779](#)
 - [Adding and Configuring Channels for IPLC Paired Topology – V24-Mixed on page 784](#)
 - To add additional channels and CPorts, repeat one or more procedures.
4. To add additional channels and CPorts, repeat one or more procedures specified in [step 3](#).



NOTE: If you want to make changes to a channel configuration after the channel was added and configured, you must disable the channel (`SETDefault -IPCc CHControl = Disable`), make the desired changes, then re-enable the channel (`SETDefault -IPCc CHControl = Enable`).

42.2.1

Adding and Configuring Channels for IPLC Standalone Topology – IP-Digital (GGM 8000 with IPLC Functionality at Core)

For illustration of an IPLC standalone topology with end-to-end channel type IP-Digital and the GGM 8000 with IPLC functionality located at the core (also referred to as a standalone IP backhaul topology), see [Figure 196: Standalone IPLC Topology: IP-Digital \(GGM 8000 with IPLC functionality located at core\) on page 758](#).

The IPCc service `BAckhaul` parameter does not indicate whether or not the GGM 8000 with IPLC functionality is connected to the base station via IP backhaul; instead it indicates whether or not a channel is connected over backhaul to a peer GGM 8000 with IPLC functionality:

- `BAckhaul = yes` – the channel is part of an IPLC topology that consists of two GGM 8000s with IPLC functionality, connected via an IP link and operating as a pair to provide the IPLC solution (paired IPLC topology).
- `BAckhaul = no` – the channel is part of an IPLC topology that consists of a standalone GGM 8000 with IPLC functionality providing the IPLC solution (standalone topology).

For this standalone IP backhaul topology, the IPCc service `BAckhaul` parameter is set to `no`.

For a standalone topology, if the GGM 8000 with IPLC functionality is located at the core, the end-to-end channel type must be IP-Digital.

In a voting scenario, the topology also includes a GCM 8000 comparator connected to the IP base station.

You can also configure the CPorts first, then add the already-configured CPorts to the channel with the `ADD -IPCc CHannel` command.

Procedure:

1. Add the channel and its upstream and downstream CPorts by entering:

```
ADD -IPCc CHannel <Channel_ID> <Upstream_CPortID> <Downstream_CPortID>
```

where `<Channel_ID>`, `<Upstream_CPortID>`, and `<Downstream_CPortID>` are positive integers between 1 and 32767.

Step example:

To add a channel with channel ID 10, upstream CPort ID 126, and downstream CPort ID 226, enter:

```
ADD -IPCc CHannel 10 126 226
```

**IMPORTANT:**

You must include both the *<Upstream_CPortID>* and the *<Downstream_CPortID>* in the ADD -IPCc CHannel command in that order.

The numbers 1280 and 2500 are reserved and are **not** allowed as the *<Channel_ID>*, *<Upstream_CPortID>*, or *<Downstream_CPortID>*. If you attempt to set the *<Channel_ID>*, *<Upstream_CPortID>*, or *<Downstream_CPortID>* to 1280 or 2500, the ADD -IPCc CHannel command is rejected.

If the CPorts you specify do not already exist, the ADD -IPCc CHannel command creates them. If you specify a channel that already exists, the system returns an error message.

2. Indicate that the end-to-end channel type is IP-Digital by entering:

```
SETDefault !C<Channel_ID> -IPCc EndToEndChType = IP-Digital
```

3. Indicate that the channel is used in a standalone IPLC topology by entering:

```
SETDefault !C<Channel_ID> -IPCc Backhaul = No
```

4. Configure the upstream CPort by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = V24
```

where *<CPort_ID>* is the ID of the upstream CPort that owns the digital path.

- b. Specify the digital path name associated with the V.24 port on the GGM 8000 that is connected to the DIU by entering:

```
SETDefault !C<CPort_ID> -IPCc DigitalPath = <Digital_Path_Label>
```

For path-to-port pairing, see [GGM 8000 Port and Path Numbering on page 65](#).

- c. Specify the digital clock source by entering:

```
SETDefault !C<CPort_ID> -IPCc V24ClockSource = External
```

- d. Configure the digital jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc V24JitterBuffer = <Jitter_Time>
```

Step example:

To configure CPort 126 on V.24 port 6A with a 100-millisecond digital jitter buffer, enter the following commands:

```
SETDefault !c126 -IPCc CPTYPE = V24
```

```
SETDefault !c126 -IPCc DigitalPath 6A
```

```
SETDefault !c126 -IPCc V24ClockSource = External
```

```
SETDefault !c126 -IPCc V24JitterBuffer = 100
```

5. Configure the downstream CPort by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = IP
```

- b. Specify the local IP address to be bound to the CPort by entering:

```
SETDefault !C<CPort_ID> -IPCc IPAddr = <Local_IP_Addr>
```



NOTE: You must configure the downstream CPort to use the System IP address (configured with the IP service SystemIP parameter and routable) as its local IP address.

- c. Specify the IP address of the remote peer (the IP base station or comparator) by entering:

```
SETDefault !C<CPort_ID> -IPCc PeerIPAddr = <Remote_IP_Addr>
```

Step example:

To configure CPort 226 with IP address 10.1.123.10 and remote peer IP address 20.1.125.10, enter the following commands:

```
SETDefault !c226 -IPCc CPTYPE = IP
```

```
SETDefault !c226 -IPCc IPAddr = 10.1.123.10
```

```
SETDefault !c226 -IPCc PeerIPAddr = 20.1.125.10
```

6. Enable the channel by entering:

```
SETDefault !C<Channel_ID> -IPCc CHControl = Enable
```

This command directs the underlying IPLC software to activate the channel for operation.



NOTE: Before activating the channel, the software performs comprehensive sanity checks to make sure that the channel and CPort configuration is consistent with the role and type of the channel and its child CPorts. If inconsistencies are detected, an error message is returned. In this case, you must adjust the configuration and attempt to enable the channel again.

For details about channel to CPort role/type mapping, see [Channel Type-to-CPort Type Mapping on page 762](#).

For details about which IPCc service parameters are supported for each CPort type, see “IPCc Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

To enable channel 10, enter:

```
SETDefault !c10 -IPCc CHControl = Enable
```

42.2.2

Adding and Configuring Channels for IPLC Standalone Topology – IP-Mixed

For the illustration of an IPLC standalone topology with end-to-end channel type IP-Mixed, see [Figure 197: Standalone IPLC Topology: IP-Mixed on page 759](#).



NOTE:

The IPLC standalone topology with end-to-end channel type IP-Mixed is supported only if the GGM 8000 with IPLC functionality is located at the site. If the GGM 8000 with IPLC functionality is located at the core, the end-to-end channel type must be IP-Digital.

You can also configure the CPorts first, then add the already-configured CPorts to the channel with the `ADD -IPCc Channel` command.

Procedure:

1. Add the channel and its upstream and downstream CPorts by entering:

```
ADD -IPCc Channel <Channel_ID> <Upstream_CPortID> <Downstream_CPortID>
```


where *<Channel_ID>*, *<Upstream_CPortID>*, and *<Downstream_CPortID>* are positive integers between 1 and 32767.



IMPORTANT: You must include both the *<Upstream_CPortID>* and the *<Downstream_CPortID>* in the `ADD -IPCc CHannel` command in that order.

The numbers 1280 and 2500 are reserved and are not allowed as the *<Channel_ID>*, *<Upstream_CPortID>*, or *<Downstream_CPortID>*. If you attempt to set the *<Channel_ID>*, *<Upstream_CPortID>*, or *<Downstream_CPortID>* to 1280 or 2500, the `ADD -IPCc CHannel` command is rejected.

If the CPorts you specify do not already exist, the `ADD -IPCc CHannel` command creates them. If you specify a channel that already exists, the system returns an error message.

Step example:

To add a channel with channel ID 10, upstream CPort ID 126, and downstream CPort ID 226, enter:

```
ADD -IPCc CHannel 10 126 226
```

2. Indicate that the end-to-end channel type is IP-Mixed by entering:

```
SETDefault !C<Channel_ID> -IPCc EndToEndChType = IP-Mixed
```

3. Indicate that the channel is used in a standalone IPLC topology by entering:

```
SETDefault !C<Channel_ID> -IPCc Backhaul = No
```

4. Configure the upstream CPort by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = V24-Mixed
```

- b. Specify the digital path name associated with the V.24 port on the GGM 8000 that is connected to the DIU by entering:

```
SETDefault !C<CPort_ID> -IPCc DigitalPath = <Digital_Path_Label>
```



NOTE: For path-to-port pairing, see [GGM 8000 Port and Path Numbering on page 65](#).

- c. Specify the analog path name associated with the analog port on the GGM 8000 that is connected to the DIU by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogPath = <Analog_Path_Label>
```



NOTE: For details about how the V.24 and analog ports are paired for mixed mode operation, see [V.24 and Analog Port Pairing on page 763](#).

- d. Specify the digital clock source by entering:

```
SETDefault !C<CPort_ID> -IPCc V24ClockSource = External
```

- e. Specify the analog impedance by entering:

```
SETDefault !C<CPort_ID> -IPCc IMPedance [1 | 2]
```

To specify 600Ω impedance, you must enter: 1. To specify 10KΩ impedance, you must enter 2.

- f. Configure the digital jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc V24JitterBuffer = <Jitter_Time>
```

- g. Configure the analog jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogJitterBuffer = <Jitter_Time>
```

Step example:

To configure CPort 126 on V.24 port 6A and analog port 8A, with 600Ω analog impedance, a 100 millisecond digital jitter buffer, and a 30-millisecond analog jitter buffer, enter the following commands:

```

SETDefault !c126 -IPCc CPTYPE = V24-Mixed
SETDefault !c126 -IPCc DigitalPath = 6A
SETDefault !c126 -IPCc AnalogPath = 8A
SETDefault !c126 -IPCc V24ClockSource = External
SETDefault !c126 -IPCc IMPedance = 1
SETDefault !c126 -IPCc V24JitterBuffer = 100
SETDefault !c126 -IPCc AnalogJitterBuffer = 30

```

5. Configure the downstream CPort by performing the following actions:

a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = IP-Mixed
```

b. Specify the analog path name associated with the analog port on the GGM 8000 that is connected to the IP base station or comparator by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogPath = <Analog_Path_Label>
```



NOTE: For path-to-port pairing, see [GGM 8000 Port and Path Numbering on page 65](#).

c. Specify the local IP address to be bound to the CPort by entering:

```
SETDefault !C<CPort_ID> -IPCc IPAddr = <Local_IP_Addr>
```



NOTE: You must configure the downstream CPort to use the System IP address (configured with the IP service SystemIP parameter and routable) as its local IP address.

d. Specify the IP address of the remote peer (the IP base station or comparator) by entering:

```
SETDefault !C<CPort_ID> -IPCc PeerIPAddr = <Remote_IP_Addr>
```

e. Specify the analog impedance by using:

```
SETDefault !C<CPort_ID> -IPCc IMPedance = [1 | 2]
```

To specify 600Ω impedance, you must enter: 1. To specify 10KΩ impedance, you must enter 2.

f. Configure the analog jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogJitterBuffer = <Jitter_Time>
```

Step example:

To configure CPort 226 on analog port 8B with IP address 10.1.123.10, remote peer IP address 20.1.125.10, 600Ω analog impedance, and a 30-millisecond analog jitter buffer, enter the following commands:

```

SETDefault !c226 -IPCc CPTYPE = IP-Mixed
SETDefault !c226 -IPCc AnalogPath = 8B
SETDefault !c226 -IPCc IPAddr = 10.1.123.10
SETDefault !c226 -IPCc PeerIPAddr = 20.1.125.10
SETDefault !c226 -IPCc IMPedance = 1
SETDefault !c226 -IPCc AnalogJitterBuffer = 30

```

6. Enable the channel by entering:

```
SETDefault !C<Channel_ID> -IPCc CHControl = Enable
```

This command directs the underlying IPLC software to activate the channel for operation.



NOTE:

Before activating the channel, the software performs comprehensive sanity checks to make sure that the channel and CPort configuration is consistent with the role and type of the channel and its child CPorts. If inconsistencies are detected, an error message is returned. In this case, you must adjust the configuration and attempt to enable the channel again.

For details about channel to CPort role/type mapping, see [Channel Type-to-CPort Type Mapping on page 762](#).

For details about which IPCc service parameters are supported for each CPort type, see “IPCc Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

To enable channel 10, enter:

```
SETDefault !c10 -IPCc CHControl = Enable
```

42.2.3

Adding and Configuring Channels for IPLC Paired Topology – IP-Mixed

For the illustration of an IPLC paired topology with end-to-end channel type IP-Mixed, see [Figure 198: Paired IPLC Topology: IP-Mixed on page 760](#).



NOTE: You can also configure the CPorts first, then add the already-configured CPorts to the channel with the `ADD -IPCc CHannel` command.

Procedure:

1. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), add the upstream channel and its upstream and downstream CPorts by entering:

```
ADD -IPCc CHannel <Channel_ID> <Upstream_CPortID> <Downstream_CPortID>
```

where `<Channel_ID>`, `<Upstream_CPortID>`, and `<Downstream_CPortID>` are positive integers between 1 and 32767.

For an IPLC paired topology, with an end-to-end channel spanning two GGM 8000s with IPLC functionality, we recommend that you set the `<Channel_ID>` to the same value on both boxes to help identify the channel across the two boxes.



IMPORTANT:

You must include both the `<Upstream_CPortID>` and the `<Downstream_CPortID>` in the `ADD -IPCc CHannel` command in that order.

The numbers 1280 and 2500 are reserved and are not allowed as the `<Channel_ID>`, `<Upstream_CPortID>`, or `<Downstream_CPortID>`. If you attempt to set the `<Channel_ID>`, `<Upstream_CPortID>`, or `<Downstream_CPortID>` to 1280 or 2500, the

`ADD -IPCc CHannel` command is rejected.

If the CPorts you specify do not already exist, the `ADD -IPCc CHannel` command creates them. If you specify a channel that already exists, the system returns an error message.

Step example:

To add a channel with channel ID 10, upstream CPort ID 126, and downstream CPort ID 226, enter:

```
ADD -IPCc CHannel 10 126 226
```

2. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), indicate that the end-to-end channel type is IP-Mixed by entering:

```
SETDefault !C<Channel_ID> -IPCc EndToEndChType = IP-Mixed
```

3. On the GGM 8000 located at the core (on which you set the Location to Core in [Configuring IPLC Functionality on page 769](#)), indicate that the channel is used in a paired IPLC topology by entering:

```
SETDefault !C<Channel_ID> -IPCc Backhaul = Yes
```

4. On the GGM 8000 located at the core (on which you set Location to Core in [Configuring IPLC Functionality on page 769](#)), configure the upstream CPort on the upstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPortType = V24-Mixed
```

- b. Specify the digital path name associated GGM 8000 V.24 port that is connected to the DIU by entering:

```
SETDefault !C<CPort_ID> -IPCc DigitalPath = <Digital_Path_Label>
```



NOTE: For path-to-port pairing, see [GGM 8000 Port and Path Numbering on page 65](#).

- c. Specify the analog path name associated with the GGM 8000 analog port that is connected to the DIU by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogPath = <Analog_Path_Label>
```

- d. Specify the digital clock source entering:

```
SETDefault !C<CPort_ID> -IPCc V24ClockSource = External
```

- e. Specify the analog impedance by entering:

```
SETDefault !C<CPort_ID> -IPCc IMPedance = [1 | 2]
```

To specify 600Ω impedance, you must enter: 1. To specify 10KΩ impedance, you must enter 2.

- f. Configure the digital jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc V24JitterBuffer = <Jitter_Time>
```

- g. Configure the analog jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogJitterBuffer = <Jitter_Time>
```

Step example:

To configure CPort 126 on V.24 port 6A and analog port 8A, with 600Ω analog impedance, a 100-millisecond digital jitter buffer, and a 30-millisecond analog jitter buffer, enter the following commands:

```
SETDefault !c126 -IPCc CPortType = V24-Mixed
```

```
SETDefault !c126 -IPCc DigitalPath = 6A
```

```
SETDefault !c126 -IPCc AnalogPath = 8A
```

```
SETDefault !c126 -IPCc V24ClockSource = External
```

```
SETDefault !c126 -IPCc IMPedance = 1
```

```
SETDefault !c126 -IPCc V24JitterBuffer = 100
```

```
SETDefault !c126 -IPCc AnalogJitterBuffer = 30
```

5. On the GGM 8000 located at the core (on which you set Location to Core in [Configuring IPLC Functionality on page 769](#)), configure the downstream CPort on the upstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPortType = IP
```

- b. Specify the local IP address to be bound to the CPort by entering:

```
SETDefault !C<CPort_ID> -IPCc IPAddr = <Local_IP_Addr>
```



NOTE: You must configure the downstream CPort to use the System IP address (configured with the IP service SystemIP parameter and routable) as its local IP address.

- c. Specify the IP address of the remote peer (the GGM 8000 with IPLC functionality located at the site) by entering:

```
SETDefault !C<CPort_ID> -IPCc PeerIPAddr <Remote_IP_Addr>
```

Step example:

To configure IP CPort 226 with IP address 10.1.123.10 and remote peer IP address 10.25.84.20, enter the following commands:

```
SETDefault !c226 -IPCc CPTYPE = IP
```

```
SETDefault !c226 -IPCc IPAddr = 10.1.123.10
```

```
SETDefault !c226 -IPCc PeerIPAddr = 10.25.84.20
```

6. On the GGM 8000 located at the core (on which you set Location to Core in [Configuring IPLC Functionality on page 769](#)), enable the upstream channel by entering:

```
SETDefault !C<Channel_ID> -IPCc CHControl = Enable
```

This command directs the underlying IPLC software to activate the channel for operation.



NOTE: Before activating the channel, the software performs comprehensive sanity checks to make sure that the channel and CPort configuration is consistent with the role and type of the channel and its child CPorts. If inconsistencies are detected, an error message is returned. In this case, you must adjust the configuration and attempt to enable the channel again.

For details about channel to CPort role/type mapping, see [Channel Type-to-CPort Type Mapping on page 762](#).

For details about which IPCc service parameters are supported for each CPort type, see “IPCc Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

To enable channel 10, enter:

```
SETDefault !c10 -IPCc CHControl = Enable
```

7. On the GGM 8000 located at the site (on which you set Location to Site in [Configuring IPLC Functionality on page 769](#)), add the downstream channel and its upstream and downstream CPorts by entering:

```
ADD -IPCc CHannel <Channel_ID> <Upstream_CPortID> <Downstream_CPortID>
```

where <Channel_ID>, <Upstream_CPortID>, and <Downstream_CPortID> are positive integers between 1 and 32767.

For an IPLC paired topology, with an end-to-end channel spanning two GGM 8000s with IPLC functionality, we recommend that you set the <Channel_ID> to the same value on both boxes to help identify the channel across the two boxes.

If the CPorts you specify do not already exist, the ADD -IPCc CHannel command creates them. If you specify a channel that already exists, the system returns an error message.



IMPORTANT:

You must include both the <Upstream_CPortID> and the <Downstream_CPortID> in the ADD -IPCc CHannel command in that order.

The numbers 1280 and 2500 are reserved and are not allowed as the <Channel_ID>, <Upstream_CPortID>, or <Downstream_CPortID>. If you attempt to set the <Channel_ID>, <Upstream_CPortID>, or <Downstream_CPortID> to 1280 or 2500, the ADD -IPCc CHannel command is rejected.

Step example:

To add a channel with channel ID 10, upstream CPort ID 326, and downstream CPort ID 426, enter:

```
ADD -IPCc CHannel 10 326 426
```

8. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), indicate that the end-to-end channel type is IP-Mixed by entering:

```
SETDefault !C<Channel_ID> -IPCc EndToEndChType = IP-Mixed
```

9. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), indicate that the channel is used in a paired IPLC topology by entering:

```
SETDefault !!C<Channel_ID> -IPCc Backhaul = Yes
```

10. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), configure the upstream CPort on the downstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = IP
```

- b. Specify the local IP address to be bound to the CPort by entering:

```
SETDefault !C<CPort_ID> -IPCc IPAddr = <Local_IP_Addr>
```



NOTE: This address should match the address you configured as the remote IP address when you configured the downstream CPort on the upstream channel on the GGM 8000 with IPLC functionality at the core.

- c. Specify the IP address of the remote peer (the GGM 8000 with IPLC functionality located at the core) by entering:

```
SETDefault !C<CPort_ID> -IPCc PeerIPAddr = <Remote_IP_Addr>
```



NOTE: This address should match the address you configured as the local IP address when you configured the downstream CPort on the upstream channel on the GGM 8000 with IPLC functionality at the core.

Step example:

To configure IP CPort 326 with IP address 10.25.84.20 and remote peer IP address 10.1.123.10, enter the following commands:

```
SETDefault !c326 -IPCc CPTYPE = IP
```

```
SETDefault !c326 -IPCc IPAddr = 10.25.84.20
```

```
SETDefault !c326 -IPCc PeerIPAddr = 10.1.123.10
```

11. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), configure the downstream CPort on the downstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = IP-Mixed
```

- b. Specify the analog path name associated with the GGM 8000 analog port that is connected to the base station or comparator by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogPath = <Analog_Path_Label>
```



NOTE: For path-to-port pairing, see [GGM 8000 Port and Path Numbering on page 65](#).

- c. Specify the local IP address to be bound to the CPort by entering:

```
SETDefault !C<CPort_ID> -IPCc IPAddr = <Local_IP_Addr>
```



NOTE: You must configure the downstream CPort to use the System IP address (configured with the IP service `SystemIP` parameter and routable) as its local IP address.

- d. Specify the IP address of the remote peer (the IP base station or comparator) by entering:

```
SETDefault !C<CPort_ID> -IPCc PeerIPAddr = <Remote_IP_Addr>
```

- e. Specify the analog impedance by entering:

```
SETDefault !C<CPort_ID> -IPCc IMPedance = [ 1 | 2 ]
```

To specify 600Ω impedance, you must enter: 1. To specify 10KΩ impedance, you must enter 2.

Step example:

To configure CPort 426 on analog port 8D with IP address 10.25.84.20 and remote peer IP address 10.1.123.10, and 600Ω analog impedance, enter the following commands:

```
SETDefault !c426 -IPCc CPTYPE = IP-Mixed
SETDefault !c426 -IPCc AnalogPAth = 8D
SETDefault !c426 -IPCc IPAddr = 10.25.84.20
SETDefault !c426 -IPCc PeerIPAddr = 10.1.123.10
SETDefault !c426 -IPCc IMPedance = 1
```

12. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), enable the downstream channel by entering:

```
SETDefault !C<Channel_ID> -IPCc CHControl = Enable
```

This command directs the underlying IPLC software to activate the channel for operation.



NOTE:

Before activating the channel, the software performs comprehensive sanity checks to make sure that the channel and CPort configuration is consistent with the role and type of the channel and its child CPorts. If inconsistencies are detected, an error message is returned. In this case, you must adjust the configuration and attempt to enable the channel again.

For details about channel to CPort role/type mapping, see [Channel Type-to-CPort Type Mapping on page 762](#).

For details about which IPCc service parameters are supported for each CPort type, see “IPCc Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

To enable channel 10, enter:

```
SETDefault !c10 -IPCc CHControl = Enable
```

42.2.4

Adding and Configuring Channels for IPLC Paired Topology – V24-Digital

For the illustration of an IPLC paired topology with end-to-end channel type V24-Digital, see [Figure 199: Paired IPLC Topology: V24-Digital on page 761](#).



NOTE: You can also configure the CPorts first, then add the already-configured CPorts to the channel with the `ADD -IPCc CHannel` command.

Procedure:

1. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), add the upstream channel and its upstream and downstream CPorts by entering:

```
ADD -IPCc CHannel <Channel_ID> <Upstream_CPortID> <Downstream_CPortID>
```


where *<Channel_ID>*, *<Upstream_CPortID>*, and *<Downstream_CPortID>* are positive integers between 1 and 32767.

For an IPLC paired topology, with an end-to-end channel spanning two GGM 8000s with IPLC functionality, we recommend that you set the *<Channel_ID>* to the same value on both boxes to help identify the channel across the two boxes.

If the CPorts you specify do not already exist, the `ADD -IPCc CHannel` command creates them. If you specify a channel that already exists, the system returns an error message.



IMPORTANT:

You must include both the *<Upstream_CPortID>* and the *<Downstream_CPortID>* in the `ADD -IPCc CHannel` command in that order.

The numbers 1280 and 2500 are reserved and are not allowed as the *<Channel_ID>*, *<Upstream_CPortID>*, or *<Downstream_CPortID>*. If you attempt to set the *<Channel_ID>*, *<Upstream_CPortID>*, or *<Downstream_CPortID>* to 1280 or 2500, the `ADD -IPCc CHannel` command is rejected.

Step example:

To add a channel with channel ID 10, upstream CPort ID 126, and downstream CPort ID 226, enter:

```
ADD -IPCc CHannel 10 126 226
```

2. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), indicate that the end-to-end channel type is V24-Digital by entering:

```
SETDefault !C<Channel_ID> -IPCc EndToEndChType = V24-Digital
```

3. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), indicate that the channel is used in a paired IPLC topology by entering:

```
SETDefault !C<Channel_ID> -IPCc Backhaul = Yes
```

4. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), configure the upstream CPort on the upstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = V24
```

- b. Specify the digital path name associated with the V.24 port on the GGM 8000 that is connected to the DIU by entering:

```
SETDefault !C<CPort_ID> -IPCc DigitalPath = <Digital_Path_Label>
```



NOTE: For path-to-port pairing, see [GGM 8000 Port and Path Numbering on page 65](#).

- c. Specify the digital clock source by entering:

```
SETDefault !C<CPort_ID> -IPCc V24ClockSource = External
```

- d. Configure the digital jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc V24JitterBuffer = <Jitter_Time>
```

Step example:

To configure CPort 126 on V.24 port 6A with a 100-millisecond digital jitter buffer, enter the following commands:

```
SETDefault !c126 -IPCc CPTYPE = V24
```

```
SETDefault !c126 -IPCc DigitalPath 6A
```

```
SETDefault !c126 -IPCc V24ClockSource = External
```

```
SETDefault !c126 -IPCc V24JitterBuffer = 100
```


5. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), configure the downstream CPort on the upstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = IP
```

- b. Specify the local IP address to be bound to the CPort by entering:

```
SETDefault !C<CPort_ID> -IPCc IPAddr = <Local_IP_Addr>
```



NOTE: You must configure the downstream CPort to use the System IP address (configured with the IP service `SystemIP` parameter and routable) as its local IP address.

- c. Specify the IP address of the remote peer (the GGM 8000 with IPLC functionality located at the site) by entering:

```
SETDefault !C<CPort_ID> -IPCc PeerIPAddr <Remote_IP_Addr>
```

Step example:

To configure IP CPort 126 with IP address 10.1.123.10 and remote peer IP address 10.25.84.20, enter the following commands:

```
SETDefault !c126 -IPCc CPTYPE = IP
```

```
SETDefault !c126 -IPCc IPAddr = 10.1.123.10
```

```
SETDefault !c126 -IPCc PeerIPAddr = 10.25.84.20
```

6. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), enable the upstream channel by entering:

```
SETDefault !C<Channel_ID> -IPCc CHControl = Enable
```

This command directs the underlying IPLC software to activate the channel for operation.



NOTE: Before activating the channel, the software performs comprehensive sanity checks to make sure that the channel and CPort configuration is consistent with the role and type of the channel and its child CPorts. If inconsistencies are detected, an error message is returned. In this case, you must adjust the configuration and attempt to enable the channel again.

For details about channel to CPort role/type mapping, see [Channel Type-to-CPort Type Mapping on page 762](#).

For details about which IPCc service parameters are supported for each CPort type, see "IPCc Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

To enable channel 10, enter:

```
SETDefault !c10 -IPCc CHControl = Enable
```

7. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), add the downstream channel and its upstream and downstream CPorts by entering:

```
ADD -IPCc CHannel <Channel_ID> <Upstream_CPortID> <Downstream_CPortID>
```

where `<Channel_ID>`, `<Upstream_CPortID>`, and `<Downstream_CPortID>` are positive integers between 1 and 32767.

For an IPLC paired topology, with an end-to-end channel spanning two GGM 8000s with IPLC functionality, we recommend that you set the `<Channel_ID>` to the same value on both boxes to help identify the channel across the two boxes.

If the CPorts you specify do not already exist, the `ADD -IPCc CHannel` command creates them. If you specify a channel that already exists, the system returns an error message.



IMPORTANT:

You must include both the `<Upstream_CPortID>` and the `<Downstream_CPortID>` in the `ADD -IPCc CHannel` command in that order.

The numbers 1280 and 2500 are reserved and are not allowed as the `<Channel_ID>`, `<Upstream_CPortID>`, or `<Downstream_CPortID>`. If you attempt to set the `<Channel_ID>`, `<Upstream_CPortID>`, or `<Downstream_CPortID>` to 1280 or 2500, the `ADD -IPCc CHannel` command is rejected.

Step example:

To add a channel with channel ID 10, upstream CPort ID 326, and downstream CPort ID 426, enter:

```
ADD -IPCc CHannel 10 326 426
```

8. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), indicate that the end-to-end channel type is V24-Digital by entering:

```
SETDefault !C<Channel_ID> -IPCc EndToEndChType = V24-Digital
```

9. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), indicate that the channel is used in a paired IPLC topology entering:

```
SETDefault !!C<Channel_ID> -IPCc Backhaul = Yes
```

10. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), configure the upstream CPort on the downstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = IP
```

- b. Specify the local IP address to be bound to the CPort by entering:

```
SETDefault !C<CPort_ID> -IPCc IPAddr = <Local_IP_Addr>
```



NOTE: This address should match the address you configured as the remote IP address when you configured the downstream CPort on the upstream channel on the GGM 8000 with IPLC functionality at the core.

- c. Specify the IP address of the remote peer (the GGM 8000 with IPLC functionality located at the core) by entering:

```
SETDefault !C<CPort_ID> -IPCc PeerIPAddr = <Remote_IP_Addr>
```



NOTE: This address should match the address you configured as the local IP address when you configured the downstream CPort on the upstream channel on the GGM 8000 with IPLC functionality at the core.

Step example:

To configure IP CPort 326 with IP address 10.25.84.20 and remote peer IP address 10.1.123.10, enter the following commands:

```
SETDefault !c326 -IPCc CPTYPE = IP
```

```
SETDefault !c326 -IPCc IPAddr = 10.25.84.20
```

```
SETDefault !c326 -IPCc PeerIPAddr = 10.1.123.10
```

11. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), configure the downstream CPort on the downstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = V24
```

- b. Specify the digital path name associated with the V.24 port on the GGM 8000 that is connected to the base station or comparator by entering:

```
SETDefault !C<CPort_ID> -IPCc DigitalPath = <Digital_Path_Label>
```



NOTE: For path-to-port pairing, see [GGM 8000 Port and Path Numbering on page 65](#).

- c. Specify the digital clock source by entering:

```
SETDefault !C<CPort_ID> -IPCc V24ClockSource = External
```

- d. Configure the digital jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc V24JitterBuffer = <Jitter_Time>
```

Step example:

To configure CPort 426 on V.24 port 6A with a 100-millisecond digital jitter buffer, enter the following commands:

```
SETDefault !c426 -IPCc CPTYPE = V24
```

```
SETDefault !c426 -IPCc DigitalPath 6A
```

```
SETDefault !c426 -IPCc V24ClockSource = External
```

```
SETDefault !c426 -IPCc V24JitterBuffer = 100
```

12. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), enable the downstream channel by entering:

```
SETDefault !C<Channel_ID> -IPCc CHControl = Enable
```

This command directs the underlying IPLC software to activate the channel for operation.



NOTE:

Before activating the channel, the software performs comprehensive sanity checks to make sure that the channel and CPort configuration is consistent with the role and type of the channel and its child CPorts. If inconsistencies are detected, an error message is returned. In this case, you must adjust the configuration and attempt to enable the channel again.

For details about channel to CPort role/type mapping, see [Channel Type-to-CPort Type Mapping on page 762](#).

For details about which IPCc service parameters are supported for each CPort type, see "IPCc Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

To enable channel 10, enter:

```
SETDefault !c10 -IPCc CHControl = Enable
```

42.2.5

Adding and Configuring Channels for IPLC Paired Topology – V24-Mixed

For the illustration of an IPLC paired topology with end-to-end channel type V24-Mixed, see [Figure 200: Paired IPLC Topology: V24-Mixed on page 762](#).



NOTE: You can also configure the CPorts first, then add the already-configured CPorts to the channel with the `ADD -IPCc CHannel` command.

Procedure:

1. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), add the upstream channel and its upstream and downstream CPorts by entering:

```
ADD -IPCc CHannel <Channel_ID> <Upstream_CPortID> <Downstream_CPortID>
```

where `<Channel_ID>`, `<Upstream_CPortID>`, and `<Downstream_CPortID>` are positive integers between 1 and 32767.

For an IPLC paired topology, with an end-to-end channel spanning two GGM 8000s with IPLC functionality, we recommend that you set the `<Channel_ID>` to the same value on both boxes to help identify the channel across the two boxes.

If the CPorts you specify do not already exist, the `ADD -IPCc CHannel` command creates them. If you specify a channel that already exists, the system returns an error message.



IMPORTANT:

You must include both the `<Upstream_CPortID>` and the `<Downstream_CPortID>` in the `ADD -IPCc CHannel` command in that order.

The numbers 1280 and 2500 are reserved and are not allowed as the `<Channel_ID>`, `<Upstream_CPortID>`, or `<Downstream_CPortID>`. If you attempt to set the `<Channel_ID>`, `<Upstream_CPortID>`, or `<Downstream_CPortID>` to 1280 or 2500, the `ADD -IPCc CHannel` command is rejected.

Step example:

To add a channel with channel ID 10, upstream CPort ID 126, and downstream CPort ID 226, enter:

```
ADD -IPCc CHannel 10 126 226
```

2. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), indicate that the end-to-end channel type is V24-Mixed by entering:

```
SETDefault !C<Channel_ID> -IPCc EndToEndChType = V24-Mixed
```
3. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), indicate that the channel is used in a paired IPLC topology by entering:

```
SETDefault !C<Channel_ID> -IPCc Backhaul = Yes
```
4. On the GGM 8000 located at the core (on which you set `Location` to `Core` in [Configuring IPLC Functionality on page 769](#)), configure the upstream CPort on the upstream channel by performing the following actions:
 - a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = V24-Mixed
```
 - b. Specify the digital path name associated with the V.24 port on the GGM 8000 that is connected to the DIU by entering:

```
SETDefault !C<CPort_ID> -IPCc DigitalPath = <Digital_Path_Label>
```



NOTE: For path-to-port pairing, see [GGM 8000 Port and Path Numbering on page 65](#).

- c. Specify the analog path name associated with the GGM 8000 analog port that is connected to the DIU by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogPath = <Analog_Path_Label>
```

- d. Specify the digital clock source by entering:

```
SETDefault !C<CPort_ID> -IPCc V24ClockSource = External
```

- e. Specify the analog impedance by entering:

```
SETDefault !C<CPort_ID> -IPCc IMPedance = [1 | 2]
```

To specify 600Ω impedance, you must enter: 1. To specify 10KΩ impedance, you must enter 2.

- f. Configure the digital jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc V24JitterBuffer = <Jitter_Time>
```

- g. Configure the analog jitter buffer by using:

```
SETDefault !C<CPort_ID> -IPCc AnalogJitterBuffer = <Jitter_Time>
```

Step example:

To configure CPort 126 on V.24 port 6A with a 100-millisecond digital jitter buffer, enter the following commands:

```
SETDefault !c126 -IPCc CPTYPE = V24
```

```
SETDefault !c126 -IPCc DigitalPath 6A
```

```
SETDefault !c126 -IPCc V24ClockSource = External
```

```
SETDefault !c126 -IPCc V24JitterBuffer = 100
```

5. On the GGM 8000 located at the core (on which you set Location to Core in [Configuring IPLC Functionality on page 769](#)), configure the downstream CPort on the upstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = IP
```

- b. Specify the local IP address to be bound to the CPort by entering:

```
SETDefault !C<CPort_ID> -IPCc IPAddr = <Local_IP_Addr>
```



NOTE: You must configure the downstream CPort to use the System IP address (configured with the IP service SystemIP parameter and routable) as its local IP address.

- c. Specify the IP address of the remote peer (the GGM 8000 with IPLC functionality located at the site) by entering:

```
SETDefault !C<CPort_ID> -IPCc PeerIPAddr <Remote_IP_Addr>
```

Step example:

To configure IP CPort 126 with IP address 10.1.123.10 and remote peer IP address 10.25.84.20, enter the following commands:

```
SETDefault !c126 -IPCc CPTYPE = IP
```

```
SETDefault !c126 -IPCc IPAddr = 10.1.123.10
```

```
SETDefault !c126 -IPCc PeerIPAddr = 10.25.84.20
```

6. On the GGM 8000 located at the core (on which you set Location to Core in [Configuring IPLC Functionality on page 769](#)), enable the upstream channel by entering:

```
SETDefault !C<Channel_ID> -IPCc CHControl = Enable
```

This command directs the underlying IPLC software to activate the channel for operation.



NOTE: Before activating the channel, the software performs comprehensive sanity checks to make sure that the channel and CPort configuration is consistent with the role and type of the channel and its child CPorts. If inconsistencies are detected, an error message is returned. In this case, you must adjust the configuration and attempt to enable the channel again.

For details about channel to CPort role/type mapping, see [Channel Type-to-CPort Type Mapping on page 762](#).

For details about which IPCc service parameters are supported for each CPort type, see “IPCc Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

To enable channel 10, enter:

```
SETDefault !c10 -IPCc CHControl = Enable
```

7. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), add the downstream channel and its upstream and downstream CPorts by entering:

```
ADD -IPCc CHannel <Channel_ID> <Upstream_CPortID> <Downstream_CPortID>
```

where `<Channel_ID>`, `<Upstream_CPortID>`, and `<Downstream_CPortID>` are positive integers between 1 and 32767.

For an IPLC paired topology, with an end-to-end channel spanning two GGM 8000s with IPLC functionality, we recommend that you set the `<Channel_ID>` to the same value on both boxes to help identify the channel across the two boxes.

If the CPorts you specify do not already exist, the `ADD -IPCc CHannel` command creates them. If you specify a channel that already exists, the system returns an error message.



IMPORTANT:

You must include both the `<Upstream_CPortID>` and the `<Downstream_CPortID>` in the `ADD -IPCc CHannel` command in that order.

The numbers 1280 and 2500 are reserved and are not allowed as the `<Channel_ID>`, `<Upstream_CPortID>`, or `<Downstream_CPortID>`. If you attempt to set the `<Channel_ID>`, `<Upstream_CPortID>`, or `<Downstream_CPortID>` to 1280 or 2500, the `ADD -IPCc CHannel` command is rejected.

Step example:

To add a channel with channel ID 10, upstream CPort ID 326, and downstream CPort ID 426, enter:

```
ADD -IPCc CHannel 10 326 426
```

8. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), indicate that the end-to-end channel type is V24-Digital by entering:

```
SETDefault !C<Channel_ID> -IPCc EndToEndChType = V24-Digital
```

9. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), indicate that the channel is used in a paired IPLC topology entering:

```
SETDefault !!C<Channel_ID> -IPCc Backhaul = Yes
```

10. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), configure the upstream CPort on the downstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = IP
```

- b. Specify the local IP address to be bound to the CPort by entering:

```
SETDefault !C<CPort_ID> -IPCc IPAddr = <Local_IP_Addr>
```



NOTE: This address should match the address you configured as the remote IP address when you configured the downstream CPort on the upstream channel on the GGM 8000 with IPLC functionality at the core.

- c. Specify the IP address of the remote peer (the GGM 8000 with IPLC functionality located at the core) by entering:

```
SETDefault !C<CPort_ID> -IPCc PeerIPAddr = <Remote_IP_Addr>
```



NOTE: This address should match the address you configured as the local IP address when you configured the downstream CPort on the upstream channel on the GGM 8000 with IPLC functionality at the core.

Step example:

To configure IP CPort 326 with IP address 10.25.84.20 and remote peer IP address 10.1.123.10, enter the following commands:

```
SETDefault !c326 -IPCc CPTYPE = IP
```

```
SETDefault !c326 -IPCc IPAddr = 10.25.84.20
```

```
SETDefault !c326 -IPCc PeerIPAddr = 10.1.123.10
```

11. On the GGM 8000 located at the site (on which you set Location to Site in [Configuring IPLC Functionality on page 769](#)), configure the downstream CPort on the downstream channel by performing the following actions:

- a. Specify the CPort type by entering:

```
SETDefault !C<CPort_ID> -IPCc CPTYPE = V24-Mixed
```

- b. Specify the digital path name associated with the V.24 port on the GGM 8000 that is connected to the base station or comparator by entering:

```
SETDefault !C<CPort_ID> -IPCc DigitalPAth = <Digital_Path_Label>
```



NOTE: For path-to-port pairing, see [GGM 8000 Port and Path Numbering on page 65](#).

- c. Specify the analog path name associated with the GGM 8000 analog port that is connected to the base station or comparator by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogPAth = <Analog_Path_Label>
```

- d. Specify the digital clock source by entering:

```
SETDefault !C<CPort_ID> -IPCc V24ClockSource = Internal
```

- e. Specify the analog impedance by entering:

```
SETDefault !C<CPort_ID> -IPCc IMPedance = [1 | 2]
```

To specify 600Ω impedance, you must enter: 1. To specify 10KΩ impedance, you must enter 2.

- f. Configure the digital jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc V24JitterBuffer = <Jitter_Time>
```

- g. Configure the analog jitter buffer by entering:

```
SETDefault !C<CPort_ID> -IPCc AnalogJitterBuffer = <Jitter_Time>
```

Step example:

To configure CPort 426 on V.24 port 7A and analog port 8A, with 600Ω analog impedance, a 100-millisecond digital jitter buffer, and a 30-millisecond analog jitter buffer, enter the following commands:

```
SETDefault !c426 -IPCc CPTYPE = V24-Mixed
```

```
SETDefault !c426 -IPCc DigitalPAth 6A
```



```
SETDefault !c426 -IPCc AnalogPath 8A
SETDefault !c426 -IPCc V24ClockSource = Internal
SETDefault !c426 -IPCc IMPedance 1
SETDefault !c426 -IPCc V24JitterBuffer = 100
SETDefault !c426 -IPCc AnalogJitterBuffer = 30
```

12. On the GGM 8000 located at the site (on which you set `Location` to `Site` in [Configuring IPLC Functionality on page 769](#)), enable the downstream channel by entering:

```
SETDefault !C<Channel1_ID> -IPCc CHControl = Enable
```

This command directs the underlying IPLC software to activate the channel for operation.



NOTE:

Before activating the channel, the software performs comprehensive sanity checks to make sure that the channel and CPort configuration is consistent with the role and type of the channel and its child CPorts. If inconsistencies are detected, an error message is returned. In this case, you must adjust the configuration and attempt to enable the channel again.

For details about channel to CPort role/type mapping, see [Channel Type-to-CPort Type Mapping on page 762](#).

For details about which IPCc service parameters are supported for each CPort type, see “IPCc Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Step example:

To enable channel 10, enter:

```
SETDefault !c10 -IPCc CHControl = Enable
```

42.2.6

Creating a Configuration File with an IPLC Configuration Template

You can use this procedure to create an ASCII configuration file based on an IPLC configuration template.

Procedure:

1. Copy the appropriate template into any simple text editor on a PC or UNIX workstation.
Do not copy the template into a word processing application that embeds text and format processing command codes in the file.
2. Edit the variables in the template (IP addresses, channel IDs, CPort IDs, and so on) to match the specifics of your system.
3. Save the file into the appropriate directory on the PC from which you load the file to the GGM 8000, assigning it a name that is appropriate for your system plan.

For example, you may give the name `coreIPLC01.cfg` to the configuration file for the GGM 8000 with IPLC functionality at the core in a paired topology.

42.2.7

Downloading a Stored Configuration File to the GGM 8000

Procedure:

1. Assign the following IP address and subnet mask to the LAN card on the PC that you use to perform the configuration:

IP Address: 20.0.0.1

Subnet Mask: 255.255.255.0

2. Connect an Ethernet crossover cable between the LAN card on the PC and Ethernet port 1 on the GGM 8000.

The crossover cable crosses over pins 1 and 2 to pins 3 and 6.

3. Connect a null modem cable between the serial port on the PC and the console port on the GGM 8000.
4. Power up the GGM 8000 and connect to it by using a terminal emulation program.
5. Enter the following console port settings in response to the prompts:

Baud rate: 9600

Data bits: 8

Parity: No

Stops bits: 1

6. At the **NetLogin:** prompt, enter: root
7. At the **Password:** prompt, press ENTER.
8. Verify that the GGM 8000 is unconfigured (no IP addresses are assigned to any of the ports) by performing the following actions:

- a. Enter:

```
SHow -IP NETAddr
```

- b. If any IP addresses are listed, delete them by entering:

```
DELeTe !<port> -IP NETAddr <IP_address>
```

where:

<port> is the port to which the IP address is assigned

<IP_address> is the assigned IP address

9. Configure the IP address for Ethernet port 1 on the GGM 8000 by entering:

```
SEtDefault !1 -IP NETAddr = 20.0.0.2 255.255.255.0
```

The IP address must be on the same subnet as the IP address you assigned to the LAN card on the PC in [step 1](#).

10. Select and run a TFTP application.
11. Verify that the IPLC configuration file is present on the TFTP server and that you know its location.
12. Browse to the TFTP directory in which the IPLC configuration file is located.
13. Transfer the configuration file to the GGM 8000 by performing the following actions:

- a. Return to the terminal emulation program and locate the **EnterpriseOS#** prompt.
- b. Copy the stored configuration file to the GGM 8000's primary directory by entering:

```
copy 20.0.0.1: <cfg_filename> a:/primary/boot.cfg
```

where <cfg_filename> is the configuration file name specific to the GGM 8000 with IPLC functionality that you are configuring.

Step example:

To copy the stored configuration file coreipcc01.cfg to the primary directory, enter:

```
copy 20.0.0.1: coreipcc01.cfg a:/primary/boot.cfg
```

The configuration file is transferred to the GGM 8000 and renamed as boot.cfg.

14. Reboot the GGM 8000 by performing the following actions:

- a. Return to the terminal emulation program.
- b. At the **EnterpriseOS#** prompt, enter:

```
ReBoot
```

The GGM 8000 reboots and processes the configuration file. The IP address you assigned to the GGM 8000 is replaced with the IP address configured in the configuration file (and specific to that GGM 8000 in the system). Once the processing is complete, the message appears informing that the system is initialized and running.

15. Verify that the GGM 8000 is running the new configuration by performing the following actions:

- a. Log on to the GGM 8000.

The GGM 8000 prompt now displays the system name of the GGM 8000 (as configured in the `boot.cfg` file) rather than **EnterpriseOS#**.

- b. Enter:

```
cd
```

- c. Enter:

```
cat boot.cfg
```

- d. Compare the `Timestamp` and `Config Summary` sections to the original configuration file stored on the PC and perform one of the following actions:
 - If they match, go to [step 16](#).
 - If they do not match, ensure that the correct configuration file is loaded on the PC and that you select the correct file to load, and repeat [step 1](#) to [step 15](#).

16. Power down the GGM 8000, disconnect the TFTP PC, and connect all system communication cables to the GGM 8000.**17. Power up the GGM 8000.**

The GGM 8000 reboots using the newly-loaded configuration file.

42.3

IPLC Configuration Templates

This section provides IPLC configuration templates and describes how to use the templates to create and load an ASCII configuration file onto the GGM 8000.

The following IPLC topology examples are available:

- Standalone - IP-Digital
- Standalone - IP-Mixed
- Paired - IP-Mixed
- Paired - V24-Digital
- Paired - V24-Mixed
- Multiple Channel Types on Core GGM 8000

You can enter the commands provided in the configuration templates manually, or insert them into an ASCII configuration file and place them in the `a:/primary` directory of the flash file system on the GGM 8000. Then,

each time the GGM 8000 is booted with a `boot.cfg` file, these commands are executed, configuring IPLC functionality on the box without further user intervention.

 **NOTE:** The templates provide the commands that should be included in the configuration file for a GGM 8000 with IPLC functionality. Specific settings contained in the templates (IP addresses, port and channel numbers, and so on) are examples. The actual values vary from system to system.

42.3.1

Configuring a Sample IPLC Standalone Topology – IP-Digital

This procedure provides an example of how to configure an IPLC standalone topology, with end-to-end channel type IP-Digital.

 **NOTE:** The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering:
`SETDefault -IPCc CONTrol = Enable`
2. If the GGM 8000 with IPLC functionality is located at the core, enter:
`SETDefault -IPCc Location = Core`
3. If the GGM 8000 with IPLC functionality is located at the site, enter:
`SETDefault -IPCc Location = Site`
4. Configure the system IP address by entering:
`SETDefault -IP SystemIP = 10.1.2.85`
5. Assign the system IP address to the Ethernet/IP port by entering:
`SETDefault !1 -IP NETaddr = 10.1.2.85 255.255.255.0`
6. Add a default route and enable IP routing by entering the following commands:
`ADD -IP ROute 0.0.0.0 0.0.0.0 10.1.2.254 1`
`SETDefault -IP CONTrol = ROute`
7. Enable the loopback system IP feature for PING by entering:
`SETDefault -SYS UseSystemIP = PING`
8. Add and configure the channel and its upstream and downstream CPorts by entering the following commands:
`ADD -IPCc CHannel 1 2 3`
`SETDefault !c1 -IPCc EndToEndChType = IP-Digital`
`SETDefault !c1 -IPCc BACkhaul = no`
`SETDefault !c2 -IPCc CPType = V24`
`SETDefault !c2 -IPCc DigitalPAth = 6A`
`SETDefault !c2 -IPCc V24ClockSource = Internal`
`SETDefault !c2 -IPCc V24JitterBuffer = 100`
`SETDefault !c3 -IPCc CPType = IP`
`SETDefault !c3 -IPCc IPAddr = 10.1.2.85`

```
SETDefault !c3 -IPCc PeerIPAddr = 10.1.2.166  
SETDefault !c1 -IPCc CHControl = Enable
```

42.3.2

Configuring a Sample IPLC Standalone Topology – IP-Mixed

This procedure provides an example of how to configure an IPLC standalone topology, with end-to-end channel type IP-Mixed.



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```
SETDefault -IPCc CONTrol = Enable  
SETDefault -IPCc Location = Site
```

2. Configure the system IP address by entering:

```
SETDefault -IP SystemIP = 10.1.2.85
```

3. Assign the system IP address to the Ethernet/IP port by entering:

```
SETDefault !1 -IP NETaddr = 10.1.2.85 255.255.255.0
```

4. Add a default route and Enable IP routing by entering the following commands:

```
ADD -IP ROUTe 0.0.0.0 0.0.0.0 10.1.2.254 1  
SETDefault -IP CONTrol = ROute
```

5. Enable the loopback system IP feature for PING by entering:

```
SETDefault -SYS UseSystemIP = PING
```

6. Add and configure the channel and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 1 2 3  
SETDefault !c1 -IPCc EndToEndChType = IP-Mixed  
SETDefault !c1 -IPCc BAcKhaul = no  
SETDefault !c2 -IPCc CPType = V24-Mixed  
SETDefault !c2 -IPCc DigitalPAth = 6A  
SETDefault !c2 -IPCc V24ClockSource = Internal  
SETDefault !c2 -IPCc V24JitterBuffer = 100  
SETDefault !c2 -IPCc AnalogPAth = 8A  
SETDefault !c2 -IPCc AnalogJitterBuffer = 30  
SETDefault !c2 -IPCc Impedance = 1  
SETDefault !c3 -IPCc CPType = IP-Mixed  
SETDefault !c3 -IPCc AnalogPAth = 8B  
SETDefault !c3 -IPCc AnalogJitterBuffer = 30  
SETDefault !c3 -IPCc Impedance = 1
```

```
SETDefault !c3 -IPCc IPAddr = 10.1.2.85
SETDefault !c3 -IPCc PeerIPAddr = 10.1.2.166
SETDefault !c1 -IPCc CHControl = Enable
```

42.3.3

Configuring a Sample IPLC Paired Topology – IP-Mixed (Core GGM 8000)

This procedure provides an example of how to configure the GGM 8000 with IPLC functionality located at the core in an IPLC IP-Mixed Paired Topology with 8 Channels (for GGM 8000 at the core).



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```
SETDefault -IPCc CONTROL = Enable
SETDefault -IPCc Location = Core
```

2. Configure the system IP address by entering:

```
SETDefault -IP SystemIP = 10.1.3.2
```

3. Assign the system IP address to the backhaul port by entering:

```
SETDefault !2 -IP NETAddr = 10.1.3.2 255.255.255.0
```

4. Add a default route and enable IP routing by entering the following commands:

```
ADD -IP ROUTe 0.0.0.0 0.0.0.0 10.1.3.254 1
SETDefault -IP CONTROL = ROUTe
```

5. Enable the loopback system IP feature for PING by entering:

```
SETDefault -SYS UseSystemIP = PING
```

6. Add and configure channel 1 and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 1 21 31
SETDefault !c1 -IPCc EndToEndChType = IP-Mixed
SETDefault !c1 -IPCc BACKhaul = yes
SETDefault !c21 -IPCc CPTYPE = V24-Mixed
SETDefault !c21 -IPCc DigitalPAtH = 6A
SETDefault !c21 -IPCc V24ClockSource = Internal
SETDefault !c21 -IPCc V24JitterBuffer = 100
SETDefault !c21 -IPCc AnalogPAtH = 8A
SETDefault !c21 -IPCc AnalogJitterBuffer = 30
SETDefault !c21 -IPCc Impedance = 1
SETDefault !c31 -IPCc CPTYPE = IP
SETDefault !c31 -IPCc IPAddr = 10.1.3.2
```

```
SETDefault !c31 -IPCc PeerIPAddr = 10.1.3.3
```

```
SETDefault !c1 -IPCc CHControl = Enable
```

7. Add and configure channel 2 and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 2 22 32
```

```
SETDefault !c2 -IPCc EndToEndChType = IP-Mixed
```

```
SETDefault !c2 -IPCc BAckhaul = yes
```

```
SETDefault !c22 -IPCc CPTYPE = V24-Mixed
```

```
SETDefault !c22 -IPCc DigitalPAth = 6B
```

```
SETDefault !c22 -IPCc V24ClockSource = Internal
```

```
SETDefault !c22 -IPCc V24JitterBuffer = 100
```

```
SETDefault !c22 -IPCc AnalogPAth = 8B
```

```
SETDefault !c22 -IPCc AnalogJitterBuffer = 30
```

```
SETDefault !c22 -IPCc Impedance = 1
```

```
SETDefault !c32 -IPCc CPTYPE = IP
```

```
SETDefault !c32 -IPCc IPAddr = 10.1.3.2
```

```
SETDefault !c32 -IPCc PeerIPAddr = 10.1.3.3
```

8. Add and configure channel 3 and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 3 23 33
```

```
SETDefault !c3 -IPCc EndToEndChType = IP-Mixed
```

```
SETDefault !c3 -IPCc BAckhaul = yes
```

```
SETDefault !c23 -IPCc CPTYPE = V24-Mixed
```

```
SETDefault !c23 -IPCc DigitalPAth = 7A
```

```
SETDefault !c23 -IPCc V24ClockSource = Internal
```

```
SETDefault !c23 -IPCc V24JitterBuffer = 100
```

```
SETDefault !c23 -IPCc AnalogPAth = 8C
```

```
SETDefault !c23 -IPCc AnalogJitterBuffer = 30
```

```
SETDefault !c23 -IPCc Impedance = 1
```

```
SETDefault !c33 -IPCc CPTYPE = IP
```

```
SETDefault !c33 -IPCc IPAddr = 10.1.3.2
```

```
SETDefault !c33 -IPCc PeerIPAddr = 10.1.3.3
```

9. Add and configure channel 4 and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 4 24 34
```

```
SETDefault !c4 -IPCc EndToEndChType = IP-Mixed
```

```
SETDefault !c4 -IPCc BAckhaul = yes
```

```
SETDefault !c24 -IPCc CPTYPE = V24-Mixed
```

```
SETDefault !c24 -IPCc DigitalPAth = 7B
```

```

SETDefault !c24 -IPCc V24ClockSource = Internal
SETDefault !c24 -IPCc V24JitterBuffer = 100
SETDefault !c24 -IPCc AnalogPAth = 8D
SETDefault !c24 -IPCc AnalogJitterBuffer = 30
SETDefault !c24 -IPCc Impedance = 1
SETDefault !c34 -IPCc CPTYPE = IP
SETDefault !c34 -IPCc IPAddr = 10.1.3.2
SETDefault !c34 -IPCc PeerIPAddr = 10.1.3.3

```

10. Add and configure channel 5 and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 5 25 35
SETDefault !c5 -IPCc EndToEndChType = IP-Mixed
SETDefault !c5 -IPCc BACkhaul = yes
SETDefault !c25 -IPCc CPTYPE = V24-Mixed
SETDefault !c25 -IPCc DigitalPAth = 10A
SETDefault !c25 -IPCc V24ClockSource = Internal
SETDefault !c25 -IPCc V24JitterBuffer = 100
SETDefault !c25 -IPCc AnalogPAth = 12A
SETDefault !c25 -IPCc AnalogJitterBuffer = 30
SETDefault !c25 -IPCc Impedance = 1
SETDefault !c35 -IPCc CPTYPE = IP
SETDefault !c35 -IPCc IPAddr = 10.1.3.2
SETDefault !c35 -IPCc PeerIPAddr = 10.1.3.3

```

11. Add and configure channel 6 and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 6 26 36
SETDefault !c6 -IPCc EndToEndChType = IP-Mixed
SETDefault !c6 -IPCc BACkhaul = yes
SETDefault !c26 -IPCc CPTYPE = V24-Mixed
SETDefault !c26 -IPCc DigitalPAth = 10B
SETDefault !c26 -IPCc V24ClockSource = Internal
SETDefault !c26 -IPCc V24JitterBuffer = 100
SETDefault !c26 -IPCc AnalogPAth = 12B
SETDefault !c26 -IPCc AnalogJitterBuffer = 30
SETDefault !c26 -IPCc Impedance = 1
SETDefault !c36 -IPCc CPTYPE = IP
SETDefault !c36 -IPCc IPAddr = 10.1.3.2
SETDefault !c36 -IPCc PeerIPAddr = 10.1.3.3

```

12. Add and configure channel 7 and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 7 27 37
SETDefault !c7 -IPCc EndToEndChType = IP-Mixed
SETDefault !c7 -IPCc BAckhaul = yes
SETDefault !c27 -IPCc CPTYPE = V24-Mixed
SETDefault !c27 -IPCc DigitalPAth = 11A
SETDefault !c27 -IPCc V24ClockSource = Internal
SETDefault !c27 -IPCc V24JitterBuffer = 100
SETDefault !c27 -IPCc AnalogPAth = 12C
SETDefault !c27 -IPCc AnalogJitterBuffer = 30
SETDefault !c27 -IPCc Impedance = 1
SETDefault !c37 -IPCc CPTYPE = IP
SETDefault !c37 -IPCc IPAddr = 10.1.3.2
SETDefault !c37 -IPCc PeerIPAddr = 10.1.3.3
```

13. Add and configure channel 8 and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 8 28 38
SETDefault !c8 -IPCc EndToEndChType = IP-Mixed
SETDefault !c8 -IPCc BAckhaul = yes
SETDefault !c28 -IPCc CPTYPE = V24-Mixed
SETDefault !c28 -IPCc DigitalPAth = 11B
SETDefault !c28 -IPCc V24ClockSource = Internal
SETDefault !c28 -IPCc V24JitterBuffer = 100
SETDefault !c28 -IPCc AnalogPAth = 12D
SETDefault !c28 -IPCc AnalogJitterBuffer = 30
SETDefault !c28 -IPCc Impedance = 1
SETDefault !c38 -IPCc CPTYPE = IP
SETDefault !c38 -IPCc IPAddr = 10.1.3.2
SETDefault !c38 -IPCc PeerIPAddr = 10.1.3.3
```


42.3.4

Configuring a Sample IPLC Paired Topology – IP-Mixed (Site GGM 8000)

This procedure provides an example of how to configure the GGM 8000 with IPLC functionality located at the site in an IPLC IP-Mixed Paired Topology with 8 Channels (Site GGM 8000).



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```
SETDefault -IPCc CONTrol = Enable
SETDefault -IPCc Location = Site
```

2. Configure the system IP address by entering:

```
SETDefault -IP SystemIP = 10.1.3.3
```

3. Assign IP addresses to the Ethernet/IP ports by entering:

```
SETDefault !1 -IP NETaddr = 10.1.2.85 255.255.255.0
SETDefault !2 -IP NETaddr = 10.1.3.3 255.255.255.0
```

4. Add a default route and enable IP routing by entering the following commands:

```
ADD -IP ROUTe 0.0.0.0 0.0.0.0 10.1.3.254 1
SETDefault -IP CONTrol = ROUTe
```

5. Enable the loopback system IP feature for PING by entering:

```
SETDefault -SYS UseSystemIP = PING
```

6. Add and configure channel 1 and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 1 21 31
SETDefault !c1 -IPCc EndToEndChType = IP-Mixed
SETDefault !c1 -IPCc BAckhaul = yes
SETDefault !c21 -IPCc CPTYPE = IP
SETDefault !c21 -IPCc IPAddr = 10.1.3.3
SETDefault !c21 -IPCc PeerIPAddr = 10.1.3.2
SETDefault !c31 -IPCc CPTYPE = IP-Mixed
SETDefault !c31 -IPCc AnalogPAth = 8A
SETDefault !c31 -IPCc AnalogJitterBuffer = 30
SETDefault !c31 -IPCc Impedance = 1
SETDefault !c31 -IPCc IPAddr = 10.1.3.3
SETDefault !c31 -IPCc PeerIPAddr = 10.1.2.166
SETDefault !c1 -IPCc CHControl = Enable
```

7. Add and configure channel 2 and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 2 22 32
SETDefault !c2 -IPCc EndToEndChType = IP-Mixed
```

```

SETDefault !c2 -IPCc BAckhaul = yes
SETDefault !c22 -IPCc CPTYPE = IP
SETDefault !c22 -IPCc IPAddr = 10.1.3.3
SETDefault !c22 -IPCc PeerIPAddr = 10.1.3.2
SETDefault !c32 -IPCc CPTYPE = IP-Mixed
SETDefault !c32 -IPCc AnalogPAth = 8B
SETDefault !c32 -IPCc AnalogJitterBuffer = 30
SETDefault !c32 -IPCc Impedance = 1
SETDefault !c32 -IPCc IPAddr = 10.1.3.3
SETDefault !c32 -IPCc PeerIPAddr = 10.1.2.166

```

8. Add and configure channel 3 and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 3 23 33
SETDefault !c3 -IPCc EndToEndChType = IP-Mixed
SETDefault !c3 -IPCc BAckhaul = yes
SETDefault !c23 -IPCc CPTYPE = IP
SETDefault !c23 -IPCc IPAddr = 10.1.3.3
SETDefault !c23 -IPCc PeerIPAddr = 10.1.3.2
SETDefault !c33 -IPCc CPTYPE = IP-Mixed
SETDefault !c33 -IPCc AnalogPAth = 8C
SETDefault !c33 -IPCc AnalogJitterBuffer = 30
SETDefault !c33 -IPCc Impedance = 1
SETDefault !c33 -IPCc IPAddr = 10.1.3.3
SETDefault !c33 -IPCc PeerIPAddr = 10.1.2.166

```

9. Add and configure channel 4 and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 4 24 34
SETDefault !c4 -IPCc EndToEndChType = IP-Mixed
SETDefault !c4 -IPCc BAckhaul = yes
SETDefault !c24 -IPCc CPTYPE = IP
SETDefault !c24 -IPCc IPAddr = 10.1.3.3
SETDefault !c24 -IPCc PeerIPAddr = 10.1.3.2
SETDefault !c34 -IPCc CPTYPE = IP-Mixed
SETDefault !c34 -IPCc AnalogPAth = 8D
SETDefault !c34 -IPCc AnalogJitterBuffer = 30
SETDefault !c34 -IPCc Impedance = 1
SETDefault !c34 -IPCc IPAddr = 10.1.3.3
SETDefault !c34 -IPCc PeerIPAddr = 10.1.2.166

```

10. Add and configure channel 5 and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 5 25 35
SETDefault !c5 -IPCc EndToEndChType = IP-Mixed

```

```

SETDefault !c5 -IPCc BACkhaul = yes
SETDefault !c25 -IPCc CPTYPE = IP
SETDefault !c25 -IPCc IPAddr = 10.1.3.3
SETDefault !c25 -IPCc PeerIPAddr = 10.1.3.2
SETDefault !c35 -IPCc CPTYPE = IP-Mixed
SETDefault !c35 -IPCc AnalogPAth = 12A
SETDefault !c35 -IPCc AnalogJitterBuffer = 30
SETDefault !c35 -IPCc Impedance = 1
SETDefault !c35 -IPCc IPAddr = 10.1.3.3
SETDefault !c35 -IPCc PeerIPAddr = 10.1.2.166

```

11. Add and configure channel 6 and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 6 26 36
SETDefault !c6 -IPCc EndToEndChType = IP-Mixed
SETDefault !c6 -IPCc BACkhaul = yes
SETDefault !c26 -IPCc CPTYPE = IP
SETDefault !c26 -IPCc IPAddr = 10.1.3.3
SETDefault !c26 -IPCc PeerIPAddr = 10.1.3.2
SETDefault !c36 -IPCc CPTYPE = IP-Mixed
SETDefault !c36 -IPCc AnalogPAth = 12B
SETDefault !c36 -IPCc AnalogJitterBuffer = 30
SETDefault !c36 -IPCc Impedance = 1
SETDefault !c36 -IPCc IPAddr = 10.1.3.3
SETDefault !c36 -IPCc PeerIPAddr = 10.1.2.166

```

12. Add and configure channel 7 and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 7 27 37
SETDefault !c7 -IPCc EndToEndChType = IP-Mixed
SETDefault !c7 -IPCc BACkhaul = yes
SETDefault !c27 -IPCc CPTYPE = IP
SETDefault !c27 -IPCc IPAddr = 10.1.3.3
SETDefault !c27 -IPCc PeerIPAddr = 10.1.3.2
SETDefault !c37 -IPCc CPTYPE = IP-Mixed
SETDefault !c37 -IPCc AnalogPAth = 12C
SETDefault !c37 -IPCc AnalogJitterBuffer = 30
SETDefault !c37 -IPCc Impedance = 1
SETDefault !c37 -IPCc IPAddr = 10.1.3.3
SETDefault !c37 -IPCc PeerIPAddr = 10.1.2.166

```

13. Add and configure channel 8 and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 8 28 38
SETDefault !c8 -IPCc EndToEndChType = IP-Mixed

```

```

SETDefault !c8 -IPCc EndToEndChType = IP-Mixed
SETDefault !c8 -IPCc BAckhaul = yes
SETDefault !c28 -IPCc CPTYPE = IP
SETDefault !c28 -IPCc IPAddr = 10.1.3.3
SETDefault !c28 -IPCc PeerIPAddr = 10.1.3.2
SETDefault !c38 -IPCc CPTYPE = IP-Mixed
SETDefault !c38 -IPCc AnalogPAth = 12D
SETDefault !c38 -IPCc AnalogJitterBuffer = 30
SETDefault !c38 -IPCc Impedance = 1
SETDefault !c38 -IPCc IPAddr = 10.1.3.3
SETDefault !c38 -IPCc PeerIPAddr = 10.1.2.166

```

42.3.5

Configuring a Sample IPLC Paired Topology – V24-Digital (Core GGM 8000)

This procedure provides an example of how to configure the GGM 8000 with IPLC functionality located at the core in an IPLC V24-Digital Paired Topology (for GGM 8000 at the core).



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```

SETDefault -IPCc CONTrol = Enable
SETDefault -IPCc Location = Core

```

2. Configure the system IP address by entering:

```
SETDefault -IP SystemIP = 10.1.3.2
```

3. Assign the system IP address to the backhaul port by entering:

```
SETDefault !2 -IP NETAddr = 10.1.3.2 255.255.255.0
```

4. Add a default route and enable IP routing by entering the following commands:

```

ADD -IP ROUTe 0.0.0.0 0.0.0.0 10.1.3.2.254 1
SETDefault -IP CONTrol = ROUTe

```

5. Enable the loopback system IP feature for PING by entering:

```
SETDefault -SYS UseSystemIP = PING
```

6. Add and configure the channel and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 1 2 3
SETDefault !c1 -IPCc EndToEndChType = V24-Digital
SETDefault !c1 -IPCc BAckhaul = yes
SETDefault !c2 -IPCc CPTYPE = V24
SETDefault !c2 -IPCc DigitalPAth = 6A

```

```

SETDefault !c2 -IPCc V24ClockSource = Internal
SETDefault !c2 -IPCc V24JitterBuffer = 100
SETDefault !c3 -IPCc CPTYPE = IP
SETDefault !c3 -IPCc IPAddr = 10.1.3.2
SETDefault !c3 -IPCc PeerIPAddr = 10.1.3.3
SETDefault !c1 -IPCc CHControl = Enable

```

42.3.6

Configuring a Sample IPLC Paired Topology – V24-Digital (Site GGM 8000)

This procedure provides an example of how to configure the GGM 8000 with IPLC functionality located at the site in an IPLC V24-Digital Paired Topology (Site GGM 8000).



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```

SETDefault -IPCc CONTROL = Enable
SETDefault -IPCc Location = Site

```

2. Configure the system IP by entering:

```

SETDefault -IP SystemIP = 10.1.3.3

```

3. Assign IP addresses to the Ethernet/IP ports by entering the following commands:

```

SETDefault !1 -IP NETAddr = 10.1.2.85 255.255.255.0
SETDefault !2 -IP NETAddr = 10.1.3.3 255.255.255.0

```

4. Add a default route and enable IP routing by entering the following commands:

```

ADD -IP ROUTe 0.0.0.0 0.0.0.0 10.1.3.254 1
SETDefault -IP CONTROL = ROUTe

```

5. Enable the loopback system IP feature for PING by entering:

```

SETDefault -SYS UseSystemIP = PING

```

6. Add and configure the channel and its upstream and downstream CPorts by entering:

```

ADD -IPCc CHannel 1 2 3
SETDefault !c1 -IPCc EndToEndChType = V24-Digital
SETDefault !c1 -IPCc BACKhaul = yes
SETDefault !c2 -IPCc CPTYPE = IP
SETDefault !c2 -IPCc IPAddr = 10.1.3.3
SETDefault !c2 -IPCc PeerIPAddr = 10.1.3.2
SETDefault !c3 -IPCc CPTYPE = V24
SETDefault !c3 -IPCc DigitalPAth = 6A
SETDefault !c3 -IPCc V24ClockSource = Internal

```

```
SETDefault !c3 -IPCc V24JitterBuffer = 100  
SETDefault !c1 -IPCc CHControl = Enable
```

42.3.7

Configuring a Sample IPLC Paired Topology – V24-Mixed (Core GGM 8000)

This procedure provides an example of how to configure the GGM 8000 with IPLC functionality located at the core in an IPLC V24-Mixed Paired Topology.



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```
SETDefault -IPCc CONTrol = Enable  
SETDefault -IPCc Location = Core
```

2. Configure the system IP address by entering:

```
SETDefault -IP SystemIP = 10.1.3.2
```

3. Assign the system IP address to the backhaul port by entering:

```
SETDefault !2 -IP NETaddr = 10.1.3.2 255.255.255.0
```

4. Add a default route and enable IP routing by entering the following commands:

```
ADD -IP ROUTe 0.0.0.0 0.0.0.0 10.1.3.254 1  
SETDefault -IP CONTrol = ROUTe
```

5. Enable the loopback system IP feature for PING by entering:

```
SETDefault -SYS UseSystemIP = PING
```

6. Add and configure the channel and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 1 2 3  
SETDefault !c1 -IPCc EndToEndChType = V24-Mixed  
SETDefault !c1 -IPCc BAckhaul = yes  
SETDefault !c2 -IPCc CPType = V24-Mixed  
SETDefault !c2 -IPCc DigitalPAtH = 6A  
SETDefault !c2 -IPCc V24ClockSource = Internal  
SETDefault !c2 -IPCc V24JitterBuffer = 100  
SETDefault !c2 -IPCc AnalogPAtH = 8A  
SETDefault !c2 -IPCc AnalogJitterBuffer = 30  
SETDefault !c2 -IPCc Impedance = 1  
SETDefault !c3 -IPCc CPType = IP  
SETDefault !c3 -IPCc IPAddr = 10.1.3.2  
SETDefault !c3 -IPCc PeerIPAddr = 10.1.3.3
```

```
SETDefault !c1 -IPCc CHControl = Enable
```

42.3.8

Configuring a Sample IPLC Paired Topology – V24-Mixed (Site GGM 8000)

This procedure provides an example of how to configure the GGM 8000 with IPLC functionality located at the site in an IPLC V24-Mixed Paired Topology.



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```
SETDefault -IPCc CONTROL = Enable
```

```
SETDefault -IPCc Location = Site
```

2. Configure the system IP address by entering:

```
SETDefault -IP SystemIP = 10.1.3.3
```

3. Assign IP addresses to the Ethernet/IP ports by entering the following commands:

```
SETDefault !1 -IP NETAddr = 10.1.2.85 255.255.255.0
```

```
SETDefault !2 -IP NETAddr = 10.1.3.3 255.255.255.0
```

4. Add a default route and enable IP routing by entering the following commands:

```
ADD -IP ROUTe 0.0.0.0 0.0.0.0 10.1.3.254 1
```

```
SETDefault -IP CONTROL = ROUTe
```

5. Enable the loopback system IP feature for PING by entering:

```
SETDefault -SYS UseSystemIP = PING
```

6. Add and configure the channel and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 1 2 3
```

```
SETDefault !c1 -IPCc EndToEndChType = V24-Mixed
```

```
SETDefault !c1 -IPCc BACKhaul = yes
```

```
SETDefault !c2 -IPCc CPTYPE = IP
```

```
SETDefault !c2 -IPCc IPAddr = 10.1.3.3
```

```
SETDefault !c2 -IPCc PeerIPAddr = 10.1.3.2
```

```
SETDefault !c3 -IPCc CPTYPE = V24-Mixed
```

```
SETDefault !c3 -IPCc DigitalPAtH = 6A
```

```
SETDefault !c3 -IPCc V24ClockSource = Internal
```

```
SETDefault !c3 -IPCc V24JitterBuffer = 100
```

```
SETDefault !c3 -IPCc AnalogPAtH = 8A
```

```
SETDefault !c3 -IPCc AnalogJitterBuffer = 30
```

```
SETDefault !c3 -IPCc Impedance = 1
```

```
SETDefault !c1 -IPCc CHControl = Enable
```

42.3.9

Configuring a Sample IPLC Paired and Standalone Topology – Core GGM 8000 Supporting Multiple Channel Types

This procedure provides an example of how to configure the GGM 8000 with IPLC functionality located at the core which supports V24-Mixed, IP-Digital, and IP-Mixed (MLC 8000) channels.



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```
SETDefault -IPCc CONTro1 = Enable
```

```
SETDefault -IPCc Location = Core
```

2. Configure the system IP address by entering:

```
SETDefault -IP SystemIP = 10.205.1.10
```

3. Assign the system IP address to the backhaul port by entering:

```
SETDefault !2 -IP NETaddr = 10.205.1.10 255.255.255.0
```

4. Add a default route and enable IP routing

```
ADD -IP ROute 0.0.0.0 0.0.0.0 10.205.1.254 1
```

```
SETDefault -IP CONTro1 = R0ute
```

5. Enable the loopback system IP feature for PING by entering:

```
SETDefault -SYS UseSystemIP = PING
```

6. Add and configure a V24 mixed mode channel and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 1 2 3
```

```
SETDefault !c1 -IPCc EndToEndChType = V24-Mixed
```

```
SETDefault !c1 -IPCc BAckhaul = yes
```

```
SETDefault !c2 -IPCc CPType = V24-Mixed
```

```
SETDefault !c2 -IPCc DigitalPAth = 6A
```

```
SETDefault !c2 -IPCc V24ClockSource = Internal
```

```
SETDefault !c2 -IPCc V24JitterBuffer = 100
```

```
SETDefault !c2 -IPCc AnalogPAth = 8A
```

```
SETDefault !c2 -IPCc AnalogJitterBuffer = 30
```

```
SETDefault !c2 -IPCc Impedance = 1
```

```
SETDefault !c3 -IPCc CPType = IP
```

```
SETDefault !c3 -IPCc IPAddr = 10.205.1.10
```

```
SETDefault !c3 -IPCc PeerIPAddr = 10.205.212.10
```

```
SETDefault !c1 -IPCc CHControl = Enable
```


This channel is paired with a V.24 mixed mode channel configured on the IPLC peer GGM 8000, located at the site.

7. Add and configure one IP digital channel and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 11 12 13
SETDefault !c11 -IPCc EndToEndChType = IP-Digital
SETDefault !c11 -IPCc BAckhaul = no
SETDefault !c12 -IPCc CPTYPE = V24
SETDefault !c12 -IPCc DigitalPAtH = 7A
SETDefault !c12 -IPCc V24ClockSource = Internal
SETDefault !c12 -IPCc V24JitterBuffer = 100
SETDefault !c13 -IPCc CPTYPE = IP
SETDefault !c13 -IPCc IPAddr = 10.205.1.10
SETDefault !c23 -IPCc PeerIPAddr = 10.205.3.2
SETDefault !c11 -IPCc CHControl = Enable
```

This channel is NOT paired with a channel on an IPLC peer GGM 8000.

8. Add and configure a second IP digital channel and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 71 72 73
SETDefault !c71 -IPCc EndToEndChType = IP-Digital
SETDefault !c71 -IPCc BAckhaul = no
SETDefault !c72 -IPCc CPTYPE = V24
SETDefault !c72 -IPCc DigitalPAtH = 7B
SETDefault !c72 -IPCc V24ClockSource = Internal
SETDefault !c72 -IPCc V24JitterBuffer = 100
SETDefault !c73 -IPCc CPTYPE = IP
SETDefault !c73 -IPCc IPAddr = 10.205.1.10
SETDefault !c73 -IPCc PeerIPAddr = 10.205.211.14
SETDefault !c17 -IPCc CHControl = Enable
```

This channel is NOT paired with a channel on an IPLC peer GGM 8000.

9. Add and configure an IP mixed mode channel and its upstream and downstream CPorts by entering the following commands:

```
ADD -IPCc CHannel 41 42 43
SETDefault !c41 -IPCc EndToEndChType = IP-Mixed
SETDefault !c41 -IPCc BAckhaul = yes
SETDefault !c42 -IPCc CPTYPE = V24-Mixed
SETDefault !c42 -IPCc DigitalPAtH = 6B
SETDefault !c42 -IPCc V24ClockSource = Internal
SETDefault !c42 -IPCc V24JitterBuffer = 100
```

```

SETDefault !c42 -IPCc AnalogPath = 8B
SETDefault !c42 -IPCc AnalogJitterBuffer = 30
SETDefault !c42 -IPCc Impedance = 1
SETDefault !c43 -IPCc CPTYPE = IP
SETDefault !c43 -IPCc IPAddr = 10.205.1.10
SETDefault !c43 -IPCc PeerIPAddr = 10.205.3.10
SETDefault !c41 -IPCc CHControl = Enable

```

This channel is paired with an IP mixed mode channel configured on the IPLC peer GGM 8000, located at the site.

42.3.10

Configuring a Sample IPLC Paired Topology – V.24-Mixed (Site GGM 8000)

This procedure provides an example of how to configure the GGM 8000 with IPLC functionality that is connected to V.24 mixed mode channel.



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```

SETDefault -IPCc CONTrol = Enable
SETDefault -IPCc Location = Site

```

2. Configure the system IP address by entering:

```
SETDefault -IP SystemIP = 10.205.212.10
```

3. Assign the system IP address to the backhaul port by entering:

```
SETDefault !2 -IP NETAddr = 10.205.212.10 255.255.255.0
```

4. Add a default route and enable IP routing by entering the following commands:

```

ADD -IP ROUTe 0.0.0.0 0.0.0.0 10.205.212.254 1
SETDefault -IP CONTrol = ROute

```

5. Enable the loopback system IP feature for PING by entering:

```
SETDefault -SYS UseSystemIP = PING
```

6. Add and configure a V.24 mixed mode channel and its upstream and downstream CPorts by entering the following commands:

```

ADD -IPCc CHannel 1 2 3
SETDefault !c1 -IPCc EndToEndChType = V24-Mixed
SETDefault !c1 -IPCc BACkhaul = yes
SETDefault !c2 -IPCc CPTYPE = IP
SETDefault !c2 -IPCc IPAddr = 10.205.212.10
SETDefault !c2 -IPCc PeerIPAddr = 10.205.1.10

```

```
SETDefault !c3 -IPCc CPTYPE = V24-Mixed
SETDefault !c3 -IPCc DigitalPAth = 6A
SETDefault !c3 -IPCc V24ClockSource = Internal
SETDefault !c3 -IPCc V24JitterBuffer = 100
SETDefault !c3 -IPCc AnalogPAth = 8A
SETDefault !c3 -IPCc AnalogJitterBuffer = 30
SETDefault !c3 -IPCc Impedance = 1
SETDefault !c1 -IPCc CHControl = Enable
```

This channel is paired with a V.24 mixed mode channel configured on the IPLC peer GGM 8000, located at the core.

42.3.11

Configuring a Sample IPLC Paired Topology – IP-Mixed (Site GGM 8000)

This procedure provides an example of how to configure the GGM 8000 with IPLC functionality that is connected to an MLC 8000 IP mixed mode channel.



NOTE: The IP addresses, channel numbers, CPort, and port numbers used in this template are for example purposes only. Replace them with the IP addresses, channel numbers, and CPort numbers appropriate for your network topology.

Procedure:

1. Configure the box-level parameters by entering the following commands:

```
SETDefault -IPCc CONTROL = Enable
SETDefault -IPCc Location = Site
```

2. Configure the system IP address by entering:

```
SETDefault -IP SystemIP = 10.205.3.10
```

3. Assign the system IP address to the backhaul port by entering:

```
SETDefault !2 -IP NETAddr = 10.205.3.10 255.255.255.0
```

4. Add a default route and enable IP routing by entering the following commands:

```
ADD -IP ROUTe 0.0.0.0 0.0.0.0 10.205.3.254 1
SETDefault -IP CONTROL = ROUTe
```

5. Enable the loopback system IP feature for PING by entering:

```
SETDefault -SYS UseSystemIP = PING
```

6. Add and configure an IP mixed mode channel and its upstream and downstream CPorts:

```
ADD -IPCc CHannel 41 42 31
SETDefault !c41 -IPCc EndToEndChType = IP-Mixed
SETDefault !c41 -IPCc BACKhaul = yes
SETDefault !c42 -IPCc CPTYPE = IP
SETDefault !c42 -IPCc IPAddr = 10.205.3.10
SETDefault !c42 -IPCc PeerIPAddr = 10.205.1.10
```

```
SETDefault !c43 -IPCc CPTYPE = IP-Mixed
SETDefault !c43 -IPCc AnalogPath = 8B
SETDefault !c43 -IPCc AnalogJitterBuffer = 30
SETDefault !c43 -IPCc Impedance = 1
SETDefault !c43 -IPCc PeerIPAddr = 10.205.3.1
SETDefault !c41 -IPCc CHControl = Enable
```

This channel is paired with an IP mixed mode channel configured on the IPLC peer GGM 8000, located at the core.

Appendix A

Dial-Up Progress and Error Messages

This appendix provides dial-up progress and error messages for modems and integrated services digital network (ISDN) terminal adapters (TAs).

A.1

Dial-up Progress and Error Messages for Modems

The EOS device sends the following messages to indicate call progress, failure on lines configured with modems, or the presence of an incoming call:

INCOMING CALL ON PATH *<path>*, PORT *<port>*
 CALL ON PATH *<path>*, PORT *<port>* REJECTED
 INCOMING CALL ON PATH *<path>*, PORT *<port>* CONNECTED
 CALL ON PATH *<path>*, PORT *<port>* CONNECTED
 CALL ON PATH *<path>*, PORT *<port>* REJECTED, NO CARRIER
 DISCONNECT ON PATH *<path>*, PORT *<port>*
 PATH NOT COMING UP, INITIATING HANGUP ON PATH *<path>*
 PRIMARY IS UP, INITIATING HANGUP ON PATH *<path>*
 SECONDARY IS IDLE, INITIATING HANGUP ON PATH *<path>*
 PATH IS IDLE, INITIATING HANGUP ON PATH *<path>*
 AUTODIAL INITIATING CALL ON PATH *<path>*, PORT *<port>*
 USER INITIATING CALL ON PATH *<path>*, PORT *<port>*
 BOD FEATURE INITIATING CALL ON PATH *<path>*, PORT *<port>*
 DR FEATURE INITIATING CALL ON PATH *<path>*, PORT *<port>*
 RETRY INITIATING CALL ON PATH *<path>*, PORT *<port>*
 DIALNO IS REQUIRED FOR V.25BIS CALLS ON PATH *<path>*
 NO CALL ATTEMPTED, NO-ORIGINATE SET ON PATH *<path>*
 CALL ON PATH *<path>*, PORT *<port>* REJECTED, CODE = *<xx>*
 DOD INITIATING CALL ON PATH *<path><path>*, PORT *<port>*
 DOD RETRY INITIATING CALL ON PATH *<path>*, PORT *<port>*

Modem Response Codes

In response to a failed attempt, a V.25bis modem may return one of the error codes listed in the following table. These codes are associated with the following messages:

DIALNO IS REQUIRED FOR V.25BIS CALLS ON PATH *<path>*
 CALL ON PATH *<path>*, PORT *<port>* REJECTED, CODE = *<xx>*

Table 72: Modem Response Codes

Response Code	Meaning
AB	Abort call
CB	Local DCE busy
ET	Engaged Tone
FC	Forbidden call
NS	Number not stored
NT	Answer tone not detected
RT	Ring tone

A.2

Dial-up Progress and Error Messages for EOS Devices

The following table lists error codes and messages that EOS devices may return in response to failed attempts to communicate.

Table 73: ISDN Dial Failure Cause Codes

Cause Number ¹	Cause	Definition ²
1	Unallocated (unassigned) number	Indicates that the destination requested by the calling user cannot be reached because, although the number is in a valid format, it is not currently assigned (allocated).
2	No route to specified transit network	Indicates that the equipment sending this cause received a request to route the call through a particular transit network, which it does not recognize. The equipment sending this cause does not recognize the transit network either because the transit network does not exist or because that particular transit network, while it does exist, does not service the equipment that is sending this cause. This cause is supported on a network-dependent basis.
3	No route to destination	Indicates that the called user cannot be reached because the network through which the call was routed does not serve the destination desired. This cause is supported on a network-dependent basis.
6	Channel unacceptable	Indicates the channel most recently identified is not acceptable to the sending entity for use in this call.
7	Call awarded and delivered in an established channel	Indicates that the user was awarded the incoming call, and that the incoming call is connected to a channel already established to that user for similar calls.
16	Normal call clearing	Indicates that the call is cleared because one of the users involved in the call requested that the call be cleared. Under normal situations, the source of this cause is not the network.

Cause Number ¹	Cause	Definition ²
17	User busy	Used when the called user indicated the inability to accept another call. The user equipment is compatible with the call.
18	No user responding	Used when a user does not respond to a call establishment message with either an alerting or connect indication within the prescribed period of time allocated.
19	No answer from user (user alerted)	Used when a user provided an alerting indication but was not provided a connect indication within a prescribed period of time.
21	Call rejected	Indicates that the equipment sending this cause does not wish to accept this call, although it could have accepted the call because the equipment sending this cause is neither busy nor incompatible.
22	Number changed	Returned to a calling user when the called party number indicated by the calling user is no longer assigned. If a network does not support this capability, cause #1 "unallocated (unassigned) number" is used.
26	Non-selected user clearing	Indicates that the user was not awarded the incoming call.
27	Destination out of order	Indicates that the destination indicated by the user cannot be reached because the interface to the destination is not functioning correctly. The term "not functioning correctly" indicates that a signaling message was unable to be delivered to the remote user; for example, a physical layer or data link layer failure at the remote user, user equipment off-line, and so forth.
28	Invalid number format	Indicates that the called user cannot be reached because the called party number is not in a valid format or is not complete.
29	Facility rejected	Returned when a facility requested by the user cannot be provided by the network.
31	Normal, unspecified	Used to report a normal event only when no other cause in the normal class applies.
34	No circuit/channel available	Indicates that there is no circuit/channel presently available to handle the call.
38	Network out of order	Indicates that the network is not functioning correctly and that the condition is likely to last a relatively long period of time; for example, immediately retrying the call is not likely to be successful.
41	Temporary failure	Indicates that the network is not functioning correctly and that the condition is not likely to last a long period of time; for example, the user can try another call attempt almost immediately.
42	Switching equipment congestion	Indicates that the switching equipment generating this cause is experiencing a period of high traffic.

Cause Number ¹	Cause	Definition ²
43	Access information discarded	Indicates that the network could not deliver access information to the remote user as requested; for example, a user-to-user information, low layer compatibility, high layer compatibility, or subaddress.
44	Requested circuit/channel not available	Returned when the circuit or channel indicated by the requesting entity cannot be provided by the other side of the interface.
47	Resources unavailable, unspecified	Used to report a resource unavailable event only when no other cause in the resource unavailable class applies.
49	Quality of service not available	Used to report that the requested quality of service, as defined in CCITT Recommendation X.213, cannot be provided (for example, throughput or transit delay cannot be supported).
50	Requested facility not subscribed	Indicates that the requested supplementary service could not be provided by the network because the user did not complete the necessary administrative arrangements with its supporting network.
57	Bearer capability not authorized	Indicates that the user requested a bearer capability that is implemented by the equipment, which generated this cause, but the user is not authorized to use.
58	Bearer capability not presently available	Indicates that the user has requested a bearer capability that is implemented by the equipment, which generated this cause, but which is not available at this time.
63	Service or option not available, unspecified	Used to report a service or option not available event only when no other cause in the service or option not available class applies.
65	Bearer capability not implemented	Indicates that the equipment sending this cause does not support the bearer capability requested.
66	Channel type not implemented	Indicates that the equipment sending this cause does not support the channel type requested.
69	Requested facility not implemented	Indicates that the equipment sending this cause does not support the requested supplementary service.
70	Only restricted digital information bearer capability is available	Indicates that an equipment requested an unrestricted bearer service but that the equipment sending this cause only supports the restricted version of the requested bearer capability.
79	Service or option not implemented, unspecified	Used to report a service or option not implemented event only when no other cause in the service or option not implemented class applies.
81	Invalid call reference value	Indicates that the equipment sending this cause received a message with a call reference that is not currently in use on the user-network interface.

Cause Number ¹	Cause	Definition ²
82	Identified channel does not exist	Indicates that the equipment sending this cause received a request to use a channel not activated on the interface for a call. For example, if a user subscribed to those channels on a primary rate interface numbered from 1 to 12 and the user equipment or the network attempts to use channels 13 through 23, this cause is generated.
83	A suspended call exists, but this call identity does not	Indicates that a call resume was attempted with a call identity, which differs from that in use for any presently suspended call(s).
84	Call identity in use	Indicates that the network received a call suspend request. The call suspend request contained a call identity (including the null call identity) that is already in use for a suspended call within the domain of interfaces over which the call might be resumed.
85	No call suspended	Indicates that the network received a call resume request. The call resume request contained a call identity information element that presently does not indicate any suspended call within the domain of interfaces over which calls may be resumed.
86	Call having the requested call identity was cleared	Indicates that the network received a call resume request. The call resume request contained a call identity information element that once indicated a suspended call; however, that suspended call was cleared while suspended (either by network timeout or by the remote user).
88	Incompatible destination	Indicates that the equipment sending this cause received a request to establish a call that has low layer compatibility, high layer compatibility, or other compatibility attributes (for example, data rate), which cannot be accommodated.
91	Invalid transit network selection	Indicates that a transit network identification was received that is of an incorrect format.
95	Invalid message, unspecified	Used to report an invalid message event only when no other cause in the invalid message class applies. Used to report a resource unavailable event only when no other cause in the resource unavailable class applies.
96	Mandatory information element is missing	Indicates that the equipment sending this cause received a message that is missing an information element, which must be present in the message before that message can be processed.
97	Message type non-existent or not implemented	Indicates that the equipment sending this cause received a message with a message type it does not recognize either because this is a message not defined or defined but not implemented by the equipment sending this cause.

Cause Number ¹	Cause	Definition ²
98	Message not compatible with call state or message type non-existent or not implemented	Indicates that the equipment sending this cause received a message such that the procedures do not indicate that this is a permissible message to receive while in the call state, or a STATUS message was received indicating an incompatible call state.
99	Information element non-existent or not implemented	Indicates that the equipment sending this cause received a message, which includes information elements not recognized because the information element identifier is not defined or it is defined but not implemented by the equipment sending the cause. However, the information element is not required to be present in the message in order for the equipment sending the cause to process the message.
100	Invalid information element contents	Indicates that the equipment sending this cause received an information element that it has implemented; however, one or more of the fields in the information element are coded in such a way that was not implemented by the equipment sending this cause. One reason this cause may occur is that the EOS device requested 56K rate adaption in auto-rate mode, but the switch does not support 56K rate adaption.
101	Message not compatible with call state	Indicates that a message was received, which is incompatible with the call state.
102	Recovery on timer expiry	Indicates that a procedure was initiated by the expiry of a timer in association with ETS 300 102-1 error handling procedures.
111	Protocol error, unspecified	Used to report a protocol error event only when no other cause in the protocol error class applies.
127	Interworking, unspecified	Indicates that there was interworking with a network that does not provide causes for the actions it takes; the exact cause for a message that is sent cannot be determined.

¹ These cause numbers and definitions are found in ETS300 102-1, CCITT Q931 specification.

² In response to the error code, the dial-up software may select a new path, may try another phone number, or may keep the same path and number when retrying the call.

Appendix B

Loopback Testing

This appendix describes how to set up loopback tests.




NOTE: For information about setting up loopback tests for a specific EOS hardware platform, see the appropriate hardware user guide.

B.1

Dial-up Loopback Testing Using Modems

Very few modems or terminal adapters support loopbacks initiated by controlling the LL and RL pins on the DTE/DCE interface. You can perform these loopback functions by setting the `DLTest TestMode` to `Loopback` and manually configuring the modems for the desired loopback type.

Figure 205: Typical Dial or Leased Connection



To successfully troubleshoot your network, you must determine the procedure your DCE uses to perform DCE loopbacks. There are several different ways DCEs use loopbacks. Check your DCE vendor manual, and identify the ways you can activate the DCE to perform the following functions:

- **Local loopback** – A loopback is made at the local DCE with the transmit data returned on the receive data to the DTE (EOS device).
- **Call or connection** – A connection is made to a remote DCE.
- **Remote loopback** – A loopback is made by a far-end DCE where the received data is returned on the transmit path, back to the near-end DCE.

You can configure a DCE in one or more of the following ways:

- Front panel switches
- Console control through an asynchronous terminal connected to a separate control port
- Signal line control of V.24 circuits
- AT commands sent over a combination data and control port

Most DCEs do not support all of these modes of operation. To make sure which type of control your DCE offers, see your DCE vendor's documentation.

For the following information, see "Commands" in the *Enterprise OS Software Version 18.2 Reference Guide*.

- `DLTest` command
- High-speed serial (HSS) software settings
- Internal clock setting
- Leased-line setting

For specific settings, see the DCE manufacturer's documentation.

Before conducting a loopback test, you must check the basic network connectivity:

1. Check the connections between the DCE and the telephone network.
2. Verify both DCE configuration settings.

Both DCEs should support common communication standards.

If you still experience problems, troubleshoot the connection to isolate the problem.




NOTE: When performing the following tests, use commands or methods that are specific to each DCE.

B.1.1

Performing a Local Loopback Test

Figure 206: HSS V.35 and RS-232 Local Loopback Flowchart



Procedure:

1. Check basic configurations at the local site.
2. At the local DCE, initiate the loopback mode by performing the following actions:
 - a. Set the HSS board to Leased mode.
 - b. Run the loopback mode on the HSS board using the `DLTest` command.

- c. Perform one of the following actions:
 - If this loopback test passes, perform [Performing a Remote Loopback Test on page 819](#).
 - If this loopback test fails and you use a V.35 interface, go to [step 7](#).
 - If this loopback test fails and you use an RS-232 interface, go to [step 3](#).
3. Disconnect the RS-232 user cable from the HSS board.
4. Connect the RS-232 loopback fixture to the HSS board.

Table 74: RS-232 Loopback Pin Assignments

Name	Pin	Name	Pin
TD	2	RD	3
RTS	4	CTS	5
DSR	6	DCD	8
		DTR	20
RXC	17	TT	24

5. Set the HSS board to internal clock source and leased-line mode.
6. Run the loopback mode on the HSS board using the `DLTest` command.
 - If the test passes, the DCE or the user cables are faulty. Try using a known good DCE or user cable. After you isolate the connection problem, remember to change back the HSS board settings.
 - If the test fails, try using a known good HSS board or contact the Centralized Managed Support Operations (CMSO).
7. Disconnect the V.35 user cable from the HSS board.
8. Connect the V.35 loopback fixture to the HSS board.

Table 75: V.35 Loopback Pin Assignments

Name	Pin	Name	Pin
SCTEA	U	SCRA	V
SCTEB	W	SCRB	X
RTS	C	CTS	D
DTR	H	DSR	E
		DCD	F
SCTA	Y		
SCTB	AA		


9. Set the HSS board to Test Mode and leased-line mode.
10. Run the loopback mode on the HSS board by using the `DLTest` command.
11. Perform one of the following actions:
 - If the test passes, the DCE or the user cables are faulty. Try using a known good DCE or user cable. After you isolate the connection problem, remember to change back the HSS board settings.
 - If the test fails, try using a known good HSS board or contact the Centralized Managed Support Operations (CMSO).

B.1.2

Performing a Remote Loopback Test

 **NOTE:** This test is conducted between modem and modem.

Figure 207: HSS V.35 and RS-232 Remote Loopback Flowchart



For DCE configuration information, see your vendor documentation.

Procedure:

1. At both sites, turn the local loopback feature off.
2. At the local site, make the DCE call the remote DCE.
3. At the local DCE, start a remote loopback test using a self-generated pattern.
4. Perform one of the following actions:
 - If the test passes, the line and remote DCE are functional. Run the local test procedure at the remote site. After you isolate the connection problem, remember to change back the HSS board settings.
 - If the test fails, the line or remote DCE, or phone connections at either the local or remote site are faulty; recheck the phone, cabling, DCE phone line settings, and/or contact your carrier.

B.2

Performing a Loopback Diagnostic Test for Built-In ISDN Ports

You can use this procedure to perform a loopback test using two B channels on one Integrated Services Digital Network (ISDN) line. The following figure illustrates the data flow in this type of a loopback test. Both of these channels occupy the same physical connector and no modem is required in this configuration; however, the unit must be attached to an ISDN line.

Figure 208: ISDN Loopback Testing



Both channels (B1 and B2) occupy the same physical connector.

Procedure:

1. Connect a console running at 9,600 baud to your EOS devices.
2. Set the path line type to Dialup by entering:

```
SETDefault !<path> -PATH LineType = Dialup
```

Step example:

```
SETDefault !2.1 -PATH LineType = Dialup
```

```
SETDefault !2.2 -PATH LineType = Dialup
```

3. Set the rate adaption parameter to automatically detect the speed of the interface using:

```
SETDefault !<path> -PATH RateAdaption = Auto
```

In [step 1](#), the test originates from path 2.2 and targets path 2.1. To specify this for path 2.2, you must enter:

```
SETDefault !2.2 -PATH RateAdaption = Auto
```

4. Set the switch type by entering:

```
SETDefault !<path> -PATH SwitchType = ETSI | NTT | ATT5ESS | NT1 |  
DMS100 | KDD
```

Step example:

To set the switch type to ETSI, you must enter:

```
SETDefault !2 -PATH SwitchType = ETSI
```

5. Establish the local dial numbers for the bearer channels by entering:

```
SETDefault !<port> -PATH LocalDialNo = "<string>"
```

Step example:

To establish the local dial numbers for the two bearer channels, you must enter:

```
SETDefault !2.1 -PATH LocalDialNo = "4962124"
```

```
SETDefault !2.2 -PATH LocalDialNo = "4962125"
```

6. Configure the ports for loopback testing by entering:


```
SETDefault !<port> -PORT OWNEr = Loopback
```

Step example:

To specify Loopback as the owner on both the sending and receiving ports, enter:

```
SETDefault !2 -PORT OWNEr = Loopback
```

```
SETDefault !3 -PORT OWNEr = Loopback
```

7. Establish a connection between the two bearer channels by dialing out on one channel and dialing into the other by entering:

```
DIal !<path> "<string>"
```

Step example:

To dial port 2 from port 3, enter:

```
DIal !2.2 "4962124"
```

Path 2.2 places a call to the specified number, which is the number for path 2.1. It is not important which port originates or answers the call as long as the port does not try to call itself.

8. When the connection is successfully established, select the loopback testing mode by entering:

```
DLTest TestMode Loopback
```

You can specify the number of seconds the test should run. You can enter this value any time before entering the `DLTest START` command. If you do not specify a test duration, an infinite time duration is assumed. To run the test for a specific number of seconds, you must use:

```
DLTest TestDuration <seconds>
```



CAUTION: Use caution when running the test for a specified test duration. The test ends abruptly as soon as the time duration expires, and a discrepancy between the number of packets transmitted and the number received may result.

9. Start the DLTest by entering:

```
DLTest Start <sendingport>, <receivingport>
```

The loopback test is successful when the number of received packets equals or approximately equals the number of transmitted packets. If the test is not successful, verify that your system is cabled correctly and your model is installed correctly. You can check the number of packets transmitted and the number of errors by using `DLTest Stat`

Step example:

To start the DLTest and designate port 2 to send the DLTest data and port 3 to receive and loop back the data, enter:

```
DLtest START 2,3
```

10. Stop the DLTest by entering:

```
DLTest Abort
```

11. Disconnect the call by entering:

```
HangUp !2.2
```

12. Change the port owner from Loopback to the original owner by entering:

```
SETDefault !<port> -PORT OWNEr = PPP
```

Appendix C

Internet Addressing

This appendix provides information about internet addressing, including dotted decimal notation, addressing rules, and subnet addressing and masks.

C.1

Internet Addresses and Classes

Any universal communications system requires a globally accepted method of identifying individual computers; one globally accepted method is to have the network manager assign unique Internet addresses to devices, or *hosts*, on the Internet. These hosts can be personal computers, communications servers, ports on a communications server, internetwork bridges, network control servers, or UNIX hosts. The Internet uses these assigned addresses when sending or receiving packets.



NOTE: You can obtain valid and unique Internet addresses through the InterNIC Registration Services.

The Internet Protocol (IP) uses Internet addresses. Internet addressing uses a 32-bit address field numbered 0 to 31. This address field is composed of two parts:

- One part identifies the network on which the host resides.
- The second part identifies the host itself.

Hosts attached to the same network must share a common prefix designating their network number. Conceptually, each address is a pair (*net#*, *host#*) where *net#* identifies the network, and *host#* identifies a host on that network.

Internet addressing is divided into four classes: A, B, C, and D. Each address class begins with a unique bit pattern that is used by the Internet software residing on network hosts to identify the address class. After the software identified the address class from the leading bits of the Internet address, it can determine which bits are used to represent the network number and which bits are used to identify the host portion of the address. The next four sections describe these four Internet address classes.

Class A Address Format

Class A addresses have a 7-bit network field and a 24-bit local address. The highest-order bit is set to 0. This allows 127 Class A networks to be defined (network number 0 is not allowed).

The format of a Class A address with the bit numbers on the first line:

0	1	7	8	31
0	Network			Local Host Address

Class B Address Format

Class B addresses have a 14-bit network field and a 16-bit local address. The two highest-order bits are set to 1 and 0. This allows 16,383 Class B networks to be defined.

The format of a Class B address with the bit numbers on the first line:

0	1	2	15	16	31
1	0	Network			Local Host Address

Class C Address Format

Class C addresses have a 21-bit network field and an 8-bit local address. The three highest-order bits are set to 1, 1, and 0. This allows 2,097,151 Class C networks to be defined.,

The format of a Class C address, with the bit numbers on the first line:

0	1	2	3	23	24	31
1	1	0	Network			Local Host Address

Class D Address Format

Class D addresses have a 23-bit multicast address field. The four highest-order bits are set to 1, 1, 1, and 0. This allows 8,000,000 multicast addresses to be defined.

The the format of a Class D address, with the bit numbers on the first line:

0	1	2	3	4 ¹	5	6	31
1	1	1	0	0	0	Multicast Address	

¹ Not used



NOTE: No addresses are allowed with the four highest-order bits set to 1-1-1-1. These addresses are reserved.

C.1.1

Dotted Decimal Notation

An Internet address is specified as four decimal numbers, each separated by a dot. This format is called *dotted decimal notation*. The 32-bit Internet address is divided into four 8-bit fields, called octets; the value of each field is specified as a decimal number with the fields separated by periods.

For example, the Internet address of USC-ISIB.ARPA in binary octets and dotted decimal notation is as follows:

Binary octets: 00001010 00000011 00000000 00110100

Dotted decimal notations: 010.003.000.052 or 10.3.0.52

Valid network numbers for each address class are provided. The “nnn” represents the network portion of the address, which is assigned by the InterNIC. The “hhh” represents the host portion of the address, which is assigned by the network manager.

- Class A networks: (nnn.hhh.hhh.hhh):001.hhh.hhh.hhh through 126.hhh.hhh.hhh
- Class B networks: (nnn.nnn.hhh.hhh):128.001.hhh.hhh through 191.254.hhh.hhh
- Class C networks: (nnn.nnn.nnn.hhh):192.000.001.hhh through 223.255.254.hhh
- Class D networks: 224.000.000.000 through 239.255.255.255



NOTE: The bits defining the local address portion of an Internet address cannot be all zero bits or all 1 bits. These are special addresses, as described in [Addressing Rules on page 824](#).

C.1.2

Addressing Rules

You should observe these general guidelines when assigning Internet addresses:


- The bits used to define the host portion of an Internet address should not be all one bits.
According to the standard, any Internet address with the host portion consisting of all ones is interpreted as meaning “all,” as in “all hosts.” For example, the address 128.1.255.255 is interpreted as meaning all hosts on network 128.1 and is reserved for directed broadcast addressing.
- The bits used to define the network portion of an Internet address should not be all zero bits.
According to the standard, any Internet address with the network portion consisting of all zeros is interpreted as meaning “this,” as in “this network.” For example, the address 0.0.0.63 is interpreted as meaning host 63 on this network.
- The class A network number 127 is assigned the “loopback” function.
The loopback function allows a datagram to be sent by a higher-level protocol to network 127 to loopback inside the host. For example, in a Berkeley Software Distribution (BSD) UNIX environment, if program A needs to communicate with program B, which is running on the same machine, they can do this with IP network number 127. However, no datagram should ever appear on any network with a source or destination network address of 127.
- No addresses are allowed with the four highest-order bits set to 1111.
These addresses are reserved for Class E networks.

C.1.3

Sample Network Using the Class B Address Format

The following figure shows a sample network using the Class B address format. The segments connected by internetwork bridges share the same network fields while having different host fields. For example, Router A segments Network A into two LANs, both of which have the network address of 128.001.

Figure 209: Sample Network Using the Class B Address Format



Segments interconnected by routers must have different network fields to be physically separate networks. For example, Router B physically separates Network A from Network B as indicated by the different network addresses (Network A has a network address of 128.001 and Network B has a network address of 128.002).

C.2

Subnet Addresses and Subnet Masks

The original interpretation of Internet addresses (described in the previous sections) was based on a two-level hierarchy. In this model, each host sees its network as a single entity.

A number of organizations have added a third level to the interpretation of Internet addresses. In this structure, a given Internet network is divided into a collection of subnets. The three-level model is useful in networks in moderately large organizations, where it is often necessary to use more than one LAN cable to cover a local area. Each LAN can then be treated as a subnet belonging to a given main Internet network number. These independent networks are connected by routers. However, each organization that wants to connect to the Internet can usually obtain only a single Internet number.


C.2.1

Subnet Addressing

If multiple Transmission Control Protocol/Internet Protocol (TCP/IP) networks are interconnected across routers, you must assign a different network field to each network. (However, if the network is part of the Internet, you cannot use different network fields because the network field must be assigned by the InterNIC.) Subnet addressing allows an organization to use a single Internet network number for multiple physical networks. You can use subnets with any class of Internet addressing except Class D.

In the following figure, a site uses subnet addressing to span two physical networks with a single class B network address. The router accepts all traffic for network 128.005.000.000 and chooses a physical network based on the third octet of the address.

Figure 210: Subnet Addressing



The format of a regular Internet address and an Internet address with subnet mask are as follows:

Regular Internet Address Format

Network Number	Host Number

Subnet Address Format

Network Number	Subnet Number	Host Number

The network field is already defined in the previous section. The width of the subnet field is constant for a given network number.

For example, on a Class B network with a 6-bit-wide subnet field, an Internet address can be broken down as follows:

0	1	2	15	16	21	22	31
1	0	Network			Subnet		Local Host Address



NOTE: The host portion of an Internet address with the preceding definition cannot be defined as all 1 bits, but the subnet portion of an Internet address can be defined as all 1 bits.

In the preceding example, the subnet field can have any value between 0 and 63, and the host field can have any value between 1 and 1022 (all numbers are decimal). A typical class B Internet address that fits the requirements of the preceding example is the Internet address 128.5.61.100 with a subnet mask of 255.255.252.0.

C.2.2

Subnet Masks

The subnet mask allows the host portion of an Internet address to be divided into two parts. One part is used to identify a physical subnet, and the second part is used to identify a host on that subnet.

Bits in the subnet mask are set to 1 if the network treats the corresponding bit in the Internet address as part of the network address. Bits in the subnet mask are set to 0 if it treats the bit as part of the host identifier.

Subnet Address Format

Network Number	Subnet Number	Host Number

Subnet Mask

11111111 11111111	11111111	00000000

The subnet mask is also defined in the dotted decimal notation.

For example, with a Class B address of 128.121.61.100, the subnet mask is as follows:

10000000 01111001	00111101	01100100

The subnet mask for the preceding example would then be 255.255.255.0.

Adhering to the preceding constant width requirement, when using RIP the value of the subnet mask should be the same on all subnets defined for a given network number. Motorola Solutions routers support variable length subnet masks when using OSPF. For more information, see [Variable Length Subnet Masks on page 833](#).

You can assign subnet masks by using the `NETaddr` parameter in the IP Service. For example, to assign the IP address of 128.005.001.001 with a subnet mask of 255.255.255.000 to port 1 of a router, you must use:

```
SETDefault !1 -IP NETaddr = 128.005.001.001 255.255.255.000
```

For more information on the `NETaddr` parameter, see the “IP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

C.2.3

Establishing Subnets: Example 1

The InterNIC assigns you Class B Internet address 128.001.000.000. You need to establish 256 subnets with each subnet capable of supporting 254 hosts. This is the simplest form of subnetting. The first and second octets of the IP address identify the network, the third octet identifies the subnet, and the fourth octet identifies a host on the subnet. You can use this procedure to determine and assign your subnet addresses.

Procedure:

1. Convert the address assigned by the InterNIC to binary format.

Step example:

128.001.000.000 = 10000000 00000001 00000000 00000000

The underlined binary digits represent the network portion of the Internet address assigned by the InterNIC.

2. Determine the number of binary digits you need to represent 256 subnets.

Eight binary digits are required to define 256 subnets ($2^8 = 256$). You can use the binary values of all zeros (decimal value 0) and all ones (decimal value 255) as subnets. The subnets are numbered 0 through 255. The following table lists these subnets and their binary and decimal equivalents.

Table 76: Sample Subnets with Binary and Decimal Equivalents (Example 1)

Subnet #	Binary	Decimal
0	00000000	0
1	00000001	1
2	00000010	2
-	--	-
254	11111110	254
255	11111111	255

3. Select the eight most significant bits of the host portion of the Internet address to define the subnets.

These bits are displayed in bold text:

128.001.000.000 = 10000000.00000001.**00000000**.00000000

4. Define a subnet mask so that all bits of the network and future subnet fields are set to 1, and all bits of the future host field are set to 0.

Network #:

10000000.00000001.00000000.00000000 = 128.001.000.000

Subnet Mask:

11111111.11111111.11111111.00000000 = 255.255.255.000

You must configure this subnet mask (255.255.255.000) on each host and defined it for each router. Use the same subnet mask for devices on the same subnetted subnet that share the same Internet address.

5. Determine the subnet address for each host.

The 256 subnets have the following addresses:

Subnet #0:

10000000.00000001.00000000.00000000 = 128.001.000.000

Subnet #1:

10000000.00000001.00000001.00000000 = 128.001.001.000

Subnet #2:

10000000.00000001.00000010.00000000 = 128.001.002.000

Subnet #254:

10000000.00000001.11111110.00000000 = 128.001.254.000

Subnet #255:

10000000.00000001.11111111.00000000 = 128.001.255.000

The range of addresses that you can assign for subnet #1 are as follows:

Subnet #1:

10000000.00000001.00000001.00000000 = 128.001.001.000

Low Address:

10000000.00000001.00000001.00000001 = 128.001.001.001

High Address:

10000000.00000001.00000001.11111110 = 128.001.001.254

The range of addresses that you can assign for subnet #35 are as follows:

Subnet #35:

10000000.00000001.00000001.00000000 = 128.001.035.000

Low Address:

10000000.00000001.00000001.00000001 = 128.001.035.001

High Address:

10000000.00000001.00000001.11111110 = 128.001.035.254

The range of addresses that you can assign for subnet #129 are as follows:

Subnet #129:

10000000.00000001.10000001.00000000 = 128.001.129.000

Low Address:

10000000.00000001.10000001.00000001 = 128.001.129.001

High Address:

10000000.00000001.10000001.11111110 = 128.001.129.254

The range of addresses that you can assign for subnet #255 are as follows:

Subnet #255:

10000000.00000001.11111111.00000000 = 128.001.255.000

Low Address:

10000000.00000001.11111111.00000001 = 128.001.255.001

High Address:

10000000.00000001.11111111.11111110 = 128.001.255.254

6. Assign the Internet address to the router.

Step example:

If subnet #1 is connected to router port 1, enter the following command to assign the Internet address:

```
SETDefault !1 -IP NETaddr = 128.001.001.001 255.255.255.000
```

C.2.4

Establishing Subnets: Example 2

The InterNIC assigns you a Class B Internet address of 128.001.000.000. You need to establish four subnets with each subnet capable of supporting up to 16,381 hosts. You can use this procedure to determine and assign your subnet addresses.

Procedure:

1. Convert the address assigned by the InterNIC to binary format:

Step example:

128.001.000.000 = 10000000.00000001.00000000.00000000

The underlined binary digits represent the network portion of the Internet address assigned by InterNIC.

2. Determine the number of binary digits you need to represent four subnets.

Two binary digits are required to define 4 subnets ($2^2 = 4$). You can use the binary values of all zeros (decimal value 0) and all ones (decimal value 255) as subnets. For example, the four subnets you select to use can be numbered 0 through 3 and you can select the two most significant bits of the host portion of the Internet address. The following table shows subnet numbers 0 through 3 and their binary and decimal equivalents. The two most significant bits selected are shown in bold.

Table 77: Sample Subnets with Binary and Decimal Equivalents (Example 2)

Subnet #	Binary	Decimal
0	0000000	000
1	01000000	064
2	10000000	128
3	11000000	192

3. Select the two most significant bits of the host portion of the Internet address to define the subnets.

These bits are displayed in bold text:

128.001.000.000 = 10000000.00000001.00000000.00000000

4. Define a subnet mask so that all bits of the network and future subnet fields are set to 1, and all bits of the future host are set to 0.

Network #:

10000000.00000001.00000000.00000000 = 128.001.000.000

Subnet Mask:

11111111.11111111.11000000.00000000 = 255.255.192.000

You must configure this subnet mask (255.255.192.000) on each host and defined it for each router. Use the same subnet mask for devices on the same subnetted subnet that share the same Internet address.

5. Determine the subnet address for each host.

The four subnets have the following addresses:

Subnet #0:

10000000.00000001.00000000.00000000 = 128.001.000.000

Subnet #1:

10000000.00000001.01000000.00000000 = 128.001.064.000

Subnet #2:

10000000.00000001.10000000.00000000 = 128.001.128.000

Subnet #3:

10000000.00000001.11000000.00000000 = 128.001.192.000

The range of addresses that you can assign for subnet #0 are as follows:

Subnet #0:

10000000.00000001.00000000.00000000 = 128.001.000.000

Low Address:

10000000.00000001.00000000.00000001 = 128.001.000.001

High Address:

10000000.00000001.00111111.11111110 = 128.001.063.254

The range of addresses that you can assign for subnet #1 are as follows:

Subnet #1:

10000000.00000001.01000000.00000000 = 128.001.064.000

Low Address:

10000000.00000001.01000000.00000001 = 128.001.064.001

High Address:

10000000.00000001.01111111.11111110 = 128.001.127.254

The range of addresses that you can assign for subnet #2 are as follows:

Subnet #2:

10000000.00000001.10000000.00000000 = 128.001.128.000

Low Address:

10000000.00000001.10000000.00000001 = 128.001.128.001

High Address:

10000000.00000001.10111111.11111110 = 128.001.191.254

The range of addresses that you can assign for subnet #3 are as follows:

Subnet #3:

10000000.00000001.11000000.00000000 = 128.001.192.000

Low Address:

10000000.00000001.11000000.00000001 = 128.001.192.001

High Address:

10000000.00000001.11111111.11111110 = 128.001.255.254

6. Assign the Internet address to the router.

Step example:

If subnet #1 is connected to router port #2, enter the following command to assign the Internet address:

```
SETDefault !2 -IP NETaddr = 128.001.064.001 255.255.192.000
```

C.2.5

Establishing Subnets: Example 3

The InterNIC assigns you a Class B Internet address of 128.001.000.000. You need to establish 8 subnets with each subnet capable of supporting up to 8,190 hosts. You can use this procedure to determine and assign your subnet addresses.

Procedure:

1. Convert the address assigned by the InterNIC to binary format:

Step example:

128.001.000.000 = 10000000.00000001.00000000.00000000

The underlined binary digits represent the network portion of the Internet address assigned by InterNIC.

2. Determine the number of binary digits you need to represent six subnets.

Three binary digits are required to define 8 subnets ($2^3 = 8$). You can use the binary values of all zeros (decimal value 0) and all ones (decimal value 255) as subnets. For example, you can number the 8 subnets 0 through 7, and you can select the three most significant bits of the host portion of the Internet address to define the subnets.

The following table lists the subnets 0 through 7 and their binary and decimal equivalents. The three most significant bits are shown in bold.

Table 78: Sample Subnets with Binary and Decimal Equivalents (Example 3)

Subnet #	Binary	Decimal
0	00000000	000
1	00100000	032
2	01000000	064
3	01100000	096
4	10000000	128
5	10100000	160
6	11000000	192
7	11100000	224

3. Select the three most significant bits of the host portion of the Internet address to define the subnets.

These bits are displayed in bold text:

128.001.000.000 = 10000000.00000001.**000**00000.00000000

4. Define a subnet mask so that all bits of the network and future subnet fields are set to 1, and all bits of the future host are set to 0.

Network #:

10000000.00000001.**000**00000.00000000 = 128.001.000.000

Subnet Mask:

11111111.11111111.11100000.00000000 = 255.255.224.000

You must configure this subnet mask (255.255.224.000) on each host and defined it for each router. You should use the same subnet mask for devices on the same subnetted subnet that share the same Internet address.

5. Determine the subnet address for each host.

The eight subnets have the following addresses:

Subnet #0:

10000000.00000001.00000000.00000000 = 128.001.000.000

Subnet #1:

10000000.00000001.00100001.00000000 = 128.001.032.000

Subnet #2:

10000000.00000001.01000010.00000000 = 128.001.064.000

Subnet #3:

10000000.00000001.01111110.00000000 = 128.001.096.000

Subnet #4:

10000000.00000001.10011111.00000000 = 128.001.128.000

Subnet #5:

10000000.00000001.10111111.00000000 = 128.001.160.000

Subnet #6:

10000000.00000001.11011111.00000000 = 128.001.192.000

Subnet #7:

10000000.00000001.11111111.00000000 = 128.001.224.000

The range of addresses that you can assign for subnet #3 are as follows:

Subnet #3:

10000000.00000001.01100001.00000000 = 128.001.096.000

Low Address:

10000000.00000001.01100001.00000001 = 128.001.096.001

High Address:

10000000.00000001.01100001.11111110 = 128.001.127.254

The range of addresses that you can assign for subnet #5 are as follows:

Subnet #5:

10000000.00000001.10100001.00000000 = 128.001.160.000

Low Address:

10000000.00000001.10100001.00000001 = 128.001.160.001

High Address:

10000000.00000001.10100001.11111110 = 128.001.191.254

6. Assign the Internet address to the router.

Step example:

If subnet #3 is connected to router port 1, you can enter the following command to assign the Internet address:

SETDefault !1 -IP NETaddr = 128.001.096.001 255.255.224.000

C.3

Variable Length Subnet Masks

Motorola Solutions routers support variable length subnet masks; you can configure more than one subnet mask (of different lengths) for a given network address. In this case, the router supports up to five different subnet masks for a given network address.

For example, on a Class B network, you can assign the following four subnet masks for Network 128.1.0.0:

- 255.255.240.0
FF.FF.F0.00 supports 128.1.16.0 – 128.1.224.0
- 255.255.255.0
FF.FF.FF.00 supports 128.1.1.0 – 128.1.15.0
- 255.255.255.252
FF.FF.FF.FC supports 128.1.0.4 – 128.1.0.12
- 255.255.255.240
FF.FF.FF.F0 supports 128.1.0.16 – 128.1.0.240

The subnet masks are stored in the routing tables so that the longest subnet mask takes precedence over the shortest subnet mask.

Appendix D

Macro Features

This appendix provides information about macro conventions and macros with conditional statements.

D.1

Macro Conventions

The EOS software supports the following macro conventions:

- Macro contents must begin with a left parenthesis. If the definition requires more than one line, you must press the [Return] key after the opening left parenthesis. The **macro** prompt then appears as a Locator for you. Nested parentheses in balanced pairs are allowed.

- Comments are allowed in macro definitions. To identify a comment, use a hash mark (#) as the first character in the line.

For example:

```
define bringuphssi=(  
#This macro will bring up the HSSI link  
SETDefault !$1 -PAth CONTrol = Enable  
)
```

- When you end the macro with the closing right parenthesis, the normal server prompt returns.
- A single macro cannot contain more than 1023 characters.
- Macro names must follow the DOS file naming conventions:
 - The macro name cannot contain more than 14 characters.
 - The macro name extension, if any, cannot contain more than three characters.
- If an error is detected in the macro, the macro stops executing, and an error message appears. The error message includes the macro name and a short explanation. After the message appears, you are returned to command mode.
- A macro can embed the **D0** command to call another macro (nesting). Because of the large amount of memory required to keep track of the calling history and variables, the limit for nested macros is 10.



NOTE: If you use command substitution in macros, the user interface may hang. This is because of the echoing of flow-control characters between the quotation marks in the string sent from the EOS device.

D.2

Macros With Conditional Statements

Macros with conditional statements contain variables (such as arguments and return codes) and control structures (such as “if-else-end” and “switch-case-end”). Control structures instruct the macro to test conditions or make comparisons from which execution decisions can be made. Variables contain the values from which these comparisons are made. The execution decisions affect the final macro output.

D.2.1

Macro Variables

Macro variables store values in memory, which can be evaluated by a macro during execution. Values can be either numeric or strings, and are represented by a variable name such as \$1 or \$rc. Examples of variable types are arguments, input or output requests, return codes, or global variables. A variable can be readable, writable, or both readable and writable by the macro.

The following table lists all possible variables within macros. The table breaks down the variables by type (argument, input, output, return code, asynchronous event, global, system/user), value (numeric or string), and actions (readable or writable). Each variable name begins with a dollar sign (\$).

Table 79: Variables within Macros

Type	Name	Value Numeric (N)/String (S)	Readable	Writable
Argument	#	N	Yes	Yes
	1	S	Yes	Yes
	2	S	Yes	Yes
	3	S	Yes	Yes
	4	S	Yes	Yes
	5	S	Yes	Yes
Input	<	S	Yes	No
Output	>	S	No	Yes
Return Code	rc	N	Yes	No
Asynchronous	error	S	Yes	Yes
Event	brk	S	Yes	Yes
Global	global	N	Yes	Yes
	lpw	S	Yes	No
	sess	N	Yes	No
	prompt	S	Yes	No
	portid	N	Yes	No
	priv	N	Yes	No
	user	N	Yes	No
	nm	N	Yes	No
	lnm	N	Yes	No
	gnm	N	Yes	No
	eth_add	S	Yes	No
	lcerrs	N	Yes	No

D.2.1.1

Macro Variable Types

The Enterprise OS software supports the following nine types of variables:

- Argument (local)
- Input
- Output
- Return code
- Asynchronous
- Global
- System/user information
- Numeric and String
- Readable and Writable

Argument (Local)

Argument (local) variables (\$#, \$1, \$2, \$3, \$4, and \$5) provide the option of passing up to five arguments to a macro. Within the macro, these arguments can be referenced by \$1 through \$5. The \$# variable contains the actual number of arguments passed. Local argument variables apply only to a particular macro. For example:

	\$1	\$2	\$3	\$4	\$5 (empty)
do <i><macro_name></i>	call	me	at	370-6610	

The \$1 variable contains `call`, \$2 contains `me`, \$3 contains `at`, \$4 contains `370-6610`, and \$5 is empty. The \$# variable is 4.

Similar to a C language procedure, argument variable values exist only within the macro, and these values disappear when the macro terminates. You can execute the same macro with different argument variables assigned, giving it a completely independent value. For nested macros, the called macro has its own set of argument variables, independent of the calling macro.

Input

Input variables (\$<) cause the macro to stop executing and wait for your input. \$< is then substituted by your input, and the macro continues executing based on your input. Input variables are illegal in macros that are submitted to the Scheduling Service.

Output

Output variables (\$>) cause any string of characters assigned to this variable to be displayed on your terminal. Output variables can generate all 127 characters on the terminal screen, as does the `Echo` command. For example, both of the following lines generate a bell ([Ctrl]+G) to the terminal when they are executed:

```
echo "^G"  
$> = ^G
```

However, \$> does not operate exactly the same as the `Echo` command. The `Echo` command automatically appends a CR-LF after the string is echoed, and the \$> variable does not. Therefore, \$> is more convenient to use for controlling screen layout.

Return Code

Return code variables (\$rc) contain the return status of the last executed user interface command. \$rc is always 0 (no errors) if the last command executed successfully. When a called macro returns to the calling macro, the \$rc variable is not affected by the return operation.

Asynchronous Event

Asynchronous event variables (\$error and \$brk) handle unexpected conditions that cause a macro to abort. A macro aborts under one of two conditions: when an internal error is detected or when you press the BREAK key.

The \$error variable is used to recover from an error. With \$error defined, if an error occurs, the macro stops executing, and a new macro, as specified by the \$error variable, automatically begins executing to clean up or recover from the error. Without \$error defined, an error detected in a macro stops the execution of the macro and returns to command mode.

The \$break variable defines a macro that begins executing when you press the BREAK key. For example, you can define a macro called "recover" that puts the user into Listen mode, then assign the macro "recover" to \$break so that recover is executed when the [Break] key is pressed.

Without the \$break variable defined, you can exit a macro while it is executing by pressing the BREAK key (unless the NoMacroBreak option is set in the InterAction parameter). The BREAK key exits the macro and returns to command mode.

The \$error and \$brk variables do not cancel the effects of errors or breaks. They restart a new macro service in order to handle the error or break signal. If you do not define \$error or \$brk, errors and break signals force the macro to stop executing, and return to command mode.

In addition to a macro name, \$error and \$brk can contain up to five arguments. For example:

```
$error = <macro_name_for_handling_error> arg1 arg2 ...
```

Global

Global variables (\$global) provide another way to pass information between macros when calling a macro. A global variable is a variable that is globally shared among all macros executed from the same user port.

You can use \$global to test the return status of a macro.

System/User Information

System/user information variables include the following:

Table 80: System/User Information Variables

Name	Function
\$eth_add	Server Ethernet address - media access control (MAC) address of the first interface
\$lpw	Local password
\$portid	User port number executing the macro
\$prompt	Prompt strings, depending on user privilege
\$sess	Number of user sessions outstanding (on this port)
\$priv	Privilege of the user executing the macro
\$user	User privilege value is 0 (user)
\$lrm	User privilege value is 1

Name	Function
\$gnm	User privilege value is 1
\$nm	Network manager, 1
\$lcerrs	Number of failed commands since the last execution of LoadConfigs. Value of -1 if no LoadConfigs command executed. The \$lcerrs variable lists the number of failed commands for the boot.cfg file only. For information about failed commands in a specialized configuration file — such as the GGSN configuration file (XGSN.cfg) — view the config.log file or the SysconF command menu Boot Statistics display. For details about listing the config.log file, see Boot, Capture, and Log Files Contents Review on page 116). For details about the Boot Statistics display, see the hardware user guide for your EOS platform.

\$priv contains the current privilege level of the user. Its value is equal to \$user or \$nm depending on the privilege level of the user. \$priv tests the privilege level of the user within a macro. \$lpw contains the passwords for Network Manager privilege levels. These variables are used to compare passwords within a macro.

Numeric Variables and String

Numeric variables store decimal values between 32767 and -32768. String variables can store any numeric value (within the described limit) or any character sequence. String variables can perform all the functions of numeric variables. They can be compared with other strings or numeric values, incremented or decremented, or assigned to another numeric variable. Numeric values incremented beyond 32767 become negative.

Readable and Writable

All variables, except the output variable \$>, are readable by the macro, which means that the values they store can be interpreted and compared by the macro in any expression. The \$> variable generates output to the user's screen.

Some variables are writable by the macro, which means they can be reassigned new values within the macro. Only the following variables are writable:

- \$#
- \$1
- \$2
- \$3
- \$4
- \$5
- \$>
- \$error
- \$brk
- \$global

Assigning values to non-writable variables causes a syntax error and aborts the macro. For example, if you want to change the global password, \$lpw = *<password>* does not work because the variable \$lpw is not writable.

D.2.1.2

Variables Comparison and Reassignment

Six comparison operators are available for testing macro variables against each other or constant values. Comparison operators are used most often in the if-else-end control structure to compare values. The following table lists the available comparison operators.

Table 81: Comparison Operators

Operator	Comparison Performed
==	Values are equal.
!=	Values are not equal.
>=	Value on left is greater than or equal to value on right.
<=	Value on left is less than or equal to value on right.
<	Value on left is less than value on right.
>	Value on left is greater than value on right.

Numeric and string variables can be compared with each other. These rules apply:

- When both variables are numeric, which can mean that they are string variables containing numeric values, the comparison is based on value. For example:
`123 == 00123`
- If any one of the variables is a string value (containing a character other than 0–9, excluding space and tab), a string comparison is performed. The difference between uppercase and lowercase is ignored. Only the first character is compared. If the first characters are equal, the next character is compared until a decision is made.
For example:
`ABC is equal to AbC`
`ABC is not equal to 123`
- Variables can be reassigned with statements such as the following:
`<variable> = <value>`
- Only numeric values or string variables containing numeric values can be assigned to numeric variables. Otherwise, a syntax error is detected and the macro execution is aborted.
- Variables can be incremented or decremented with plus and minus statements. The plus statement is used most often within a loop structure to increment a counter, which can then be tested against a value.
`variable ++`
`variable --`

D.2.1.3

Variable Substitutions

Immediately before a line is executed, the line is scanned and all variables are replaced with their values. Substitution can be done only once. Variables can appear anywhere in the line and still be substituted. For example:

```
Echo "My arguments are $# $1 $2 $3 $4 and $5"
```

```
REMOte 192.9.200.$1
```

Two dollar signs (\$) allow you to escape variable substitutions. For example, if you enter:

Echo "argument \$\$1 is \$1"

the following display appears:

argument \$1 is *<substituted_value>*

D.3

Control Structures

Control structures are the tools that can alter the sequences of execution. The syntax is similar to a C program. Control structures must begin and end within the boundary of the macro. For example, in the if-else-end structure, all three parts of the conditional statement (if, else, and end) must be contained within the macro. If any part of the structure is missing, a syntax error is detected and the macro aborts.

Control structures are free to nest within one another. For example, within one loop structure you can have several if-end structures. There is no limit to the number of nested control structures allowed.

If-Else-End

The if-else-end structure is used to make two-way decisions. The syntax is as follows:

```
if <expression>
commands ...
else
commands ...
end
```

The else part is optional. The *<expression>* is evaluated; if it is TRUE, the macro executes the immediately following commands. If it is FALSE and there is an else statement, then the commands following the else statement are executed. If it is FALSE and there is no else statement, then the commands following the end statement are executed. There can be any number of commands between if-else-end, including none.

The syntax for *<expression>* is:

<variable> *<op>* *<value>*

<op> can be one of the six comparison operations ==, !=, >=, >, <=, and <. A single variable must be on the left side of the comparison operator. *<value>* can be any string of characters and digits, with variables intermixed, or it can be empty.

Both *<variable>* and *<value>* can contain numeric values or strings. Both can contain more than one word, but only the first word is compared.

Switch-Case-End

The switch-case-end structure is a multidirectional decision maker. It is usually used with \$< to make a comparison based on your input. The syntax is as follows:

```
switch <value>
case <value>
commands ...
case <value>
commands ...
case *
commands ...
end
```

The switch structure tests whether the *<value>* immediately following the switch matches one of the *<values>* after case. If the values match, the macro executes the immediately following commands. There can be any number of commands after each case, including none.

In situations where there is no match, but there is a case* (wildcard character) before the end, the command following case* is executed. The * can appear after any case between the switch and end, but the case* is not required in the switch-case-end structure. If the value immediately following the switch does not match any case within the switch, the macro continues to execute the commands after the end.

Loop-End

Any of the following commands can appear within the loop structure:

```
loop
commands ...
end
```

The loop structure comprises a set of instructions that can be executed repeatedly while certain conditions prevail. You define the conditions within the loop structure using the comparison operators. You can terminate the loop using the “break,” “return,” and “exit” keywords inside the loop-end structure.

D.4

Keywords

If a keyword is not the first word in a line, it is not recognized. You can use the following keywords:

- audit
- break
- case
- continue
- else
- end
- exit
- if
- loop
- return
- switch

audit

This keyword generates an audit trail record of the macro information (MI) type. You can provide a string of data following the keyword.

break

This keyword terminates the current loop structure, and the macro execution continues after the end keyword. The current loop structure is the structure that contains the break keyword. The break function does not apply to the switch-end structure. Break is only meaningful within a loop-end structure.

continue

This keyword directs macro execution to the beginning of the enclosing loop-end structure. It is only meaningful within a loop-end structure.

exit

This keyword stops execution of all macros. It frees all associated buffers and returns you to command mode.

return

This keyword stops execution of the current macro. It resumes the previous calling macro, if any.

D.5

Macro Caching and Shared Macros

As macros become larger, more complex, and more heavily nested, even a relatively simple macro can require a large number of nested macros. The dependence on macro file service from the image on the local server becomes more critical. A sudden failure of the local disk drive can create a serious service interruption.

The server supports the following features to handle failures:

- The server keeps track of all the macros currently executed in a macro cache.
When a macro is not executed, it is kept in the cache memory as long as there is space available. The next time you request a macro file, the cache is searched first. If the macro is found, it is automatically executed. This reduces the dependence on the server and speeds up macro execution.
- The server links several users to a single copy of a macro instead of distributing many copies of it, thereby sharing the macro.
Sharing macros relieves the memory overhead associated with keeping numerous similar macros. There is no limit to the number of users that can be linked to a macro.

As long as a macro is linked to a user, that macro stays cached. If there are no users linked to the macro, the server keeps the macro in the cache, based on the amount of space available in the cache.

For example, if the macro cache is between 80% and 100% full, a macro that is not linked to any users is stored for up to 10 minutes.

The following table lists the macro cache aging algorithm.

Table 82: Macro Cache Aging

Cache Usage Level	Aging
Below 50%	No aging
Below 80% but higher than 50%	Cached up to 8 hours
Below 100% but higher than 80%	Cached up to 10 minutes

If the cache overflows, the server rebuilds its cache memory, which frees all macros that are not linked.

Macro caching can cause a discrepancy between the `D0 <macro_name>` and `SHow MACros <macro_name>` commands. The `D0` command searches for the file, first in the cache and then in the local image or macro server. The `SHow` command reads the macro file directly from the local image or macro file server and never checks the cache. If the file stored in the cache is not the same as the one on the local image or file server, you get different results. If the network manager modifies the macro files, the cache aging algorithm may not pick up new macros until after the aging period. You can use the `FLush MACros` command to force the server to flush its cache.

D.6

Larger Macros

The size limit for each macro that is stored is 1023 bytes. In many applications, this is barely enough for a moderately sophisticated macro. To solve this problem, the network manager can create macros that are the necessary size using the plus sign (+).

To store a macro larger than 1023 bytes, you must split the macro into smaller macros such as `m1`, `m2`. The name of the large macro is `m1+m2` when it is cached in the memory. The macro cache stores up to eight bytes of the macro name. When the cache is searched for a macro, only the first 14 bytes of the macro name are considered significant. If a macro name exceeds the limit, it is truncated. For example, the following two macro names are considered the same in the cache:

```
macro1+macro2+macro3
```

```
macro1+macro2+XX
```

Spaces are not allowed around the plus sign (+). To execute the macro, use the `D0` command `m1+m2`, which informs the server that two macros must be read and concatenated into one large macro. The command `D0 m1 <space>+m2 <space>` means execute `+m2` as its argument.

You can break up control structures across two or more macros that are concatenated with the plus sign (+), because concatenated macros are considered one macro. For example, `macro1` could contain the if-else-end control structure, and `macro2` could contain the else and end parts of the structure.

There is no set limit on the number of plus signs (+) that you can use; however, the + operation requires a great deal of memory from the large server buffers and should not be overused.

D.7

Macro Nesting

A macro can call other macros (including itself), similar to a subroutine call in any computer language. These calls to other macros create macro nesting.

The limit to the number of nested macros is 10. Exceeding this limit causes an error and aborts the macro service. You can set up the variable `<$error>` to automatically capture the error and start a new macro service.

The following example shows macros using features such as variables and conditional statements. This example executes the `TraceRoute` command on the address specified in the first argument to the macro and loops the number of times specified in the second argument to the macro.

```
S6000 [1]# define tracen = (
if $# !=2
echo "USAGE: tracen <destination_IP_addr> <times_to_loop>"
exit
else
if $2 ==0
echo "TRACEN: iterations must be > 0"
exit
end
loop
echo " "
$>=#$2-
traceroute $1
```

```
$2--  
if $2 <= 0  
exit  
end  
end  
end)
```


Appendix E

Supported MIBs and Traps

This appendix lists all management information base (MIB) modules supported by EOS devices and the software packages that run on these systems. It also describes the traps generated by EOS and supported on EOS devices.

E.1

Supported Operations

All simple objects and tables support the Get and GetNext operations. The Set operation is supported only by objects that provide write access. Set operations take effect immediately, and changes are saved to the disk so new configurations are not lost after reboot.

The tables within the Motorola Solutions-defined MIBs support a subset of the functionality provided by the rowStatus textual convention. Row creation is allowed using only the *createAndGo* method. With the createAndGo method, the Status object of the table is set to createAndGo(4) within the same protocol data unit (PDU) that carries the other columnar values; the result is that the new row is immediately marked as active(1). Once active, with a few exceptions, the row cannot be modified. You can make changes to rows only by first deleting the row, then recreating it with the proper values.

E.2

Port Numbering Conventions in SNMP

Throughout this appendix, references to port-numbering assume the format used in the EOS software user interface. The current implementation of Simple Network Management Protocol (SNMP) uses a different port-numbering scheme.

The following table shows the relationship between the two port numbering schemes for MNR S2500.

Table 83: Mapping UI to SNMP Port Numbers (MNR S2500)

MNR S2500	
UI Port Label	SNMP Port Label
1	1
2	2
3	3
Reserved	4-73

The following table shows the relationship between the two port numbering schemes for MNR S6000 routers. You can calculate the mapping with the following formula:

SNMP port # = ((timeslot - 1) * 32) + slot# + ('port letter' * 2)

where 'port letter' = 0-11, with A=1

For example, for UI port 4K.24:

SNMP port # = ((24-1) * 32) + 4 + (10 * 2) = 736 + 4 + 20 = 760

Table 84: Mapping UI Port Numbers to SNMP Port Numbers (S6000 Routers)

UI Port Label	SNMP Port Label	UI Port Label	SNMP Port Label	UI Port Label	SNMP Port Label	UI Port Label	SNMP Port Label
1	1						
2	2						
3	3						
4A.1	4	4A.2	36	4A.3	68	.	4A.31 964
5A.1	5	5A.2	37	5A.3	69	.	5A.31 965
4B.1	6	4B.2	38	4B.3	70	.	4B.31 966
5B.1	7	5B.2	39	5B.3	71	.	5B.31 967
4C.1	8	4C.2	40	4C.3	72	.	4C.31 968
5C.1	9	5C.2	41	5C.3	73	.	5C.31 969
4D.1	10	4D.2	42	4D.3	74	.	4D.31 970
5D.1	11	5D.2	43	5D.3	75	.	5D.31 971
4E.1	12	4E.2	44	4E.3	76	.	4E.31 972
5E.1	13	5E.2	45	5E.3	77	.	5E.31 973
4F.1	14	4F.2	46	4F.3	78	.	4F.31 974
5F.1	15	5F.2	47	5F.3	79	.	5F.31 975
4G.1	16	4G.2	48	4G.3	80	.	4G.31 976
5G.1	17	5G.2	49	5G.3	81	.	5G.31 977
4H.1	18	4H.2	50	4H.3	82	.	4H.31 978
5H.1	19	5H.2	51	5H.3	83	.	5H.31 979
4I.1	20	4I.2	52	4I.3	84	.	4I.31 980
5I.1	21	5I.2	53	5I.3	85	.	5I.31 981
4J.1	22	4J.2	54	4J.3	86	.	4J.31 982
5J.1	23	4J.2	55	5J.3	87	.	5J.31 983
4K.1	24	4K.2	56	4K.3	88	.	4K.31 984
5K.1	25	5K.2	57	5K.3	89	.	5K.31 985
4L.1	26	4L.2	58	4L.3	90	.	4L.31 986
5L.1	27	5L.2	59	5L.3	91	.	5L.31 987
V1	1028	.	.	V101	1128	.	V2048 3075

E.3

MIBs Supported by EOS Devices

The EOS SNMP agent supports two enterprise trees:

- The **Motorola Solutions enterprise tree** includes the CCGW and MNR product lines and supports Motorola Solutions Network Management Alliance (NMA) standards. EOS devices support the following private objects and MIBs:

- MNR NMA WAN Slot Object
- MNR NMA LAN Path Object
- MNR NMA Device Status Object
- Ethernet WAN Link Object MIB
- MNR Service ID MIB
- MNR Mobile IP Link Fault Object
- MNR Mobile IP Service Fault Object
- MNR DBReplicator Fault Object
- MNR DHCP Service Fault Object
- MNR CCGW Private MIBs



NOTE: The MNR CCGW private MIBs are supported only on MNR S2500 and GGM 8000 gateways and MCG 8000.

In addition, EOS devices support the following generic Motorola Solutions NMA MIBs:

- NMA Time MIB
- Data Interface Version MIB
- The **3Com enterprise tree** has MNR legacy MIB implementations and does not support Motorola Solutions NMA standards. EOS devices support the following MNR legacy MIBs:
 - SNMPv1/v2c MIBs defined by the Internetworking Engineering Task Force (IETF), Motorola Solutions, IBM, and Novell.
 - MNR private MIBs

In addition, EOS devices support SNMPv3 MIBs, which are supported by the SNMPv3 Common Agent.

The MIB files are located within the EOS tar file, in the following directories:

18.2.x.xx/mibs/ccgwmibs (Motorola Solutions NMA CCGW private MIBs)

18.2.x.xx/mibs/ccgwmibs (MNR NMA private MIBs)

18.2.x.xx/mibs/ccgwmibs (MNR Legacy MIBs)

where 18.2.x.xx is the specific version of EOS software; for example, 18.2.0.12.



NOTE: Other than the SNMPv3 MIBs, which are accessible only via SNMPv3, the MIBs are accessible via SNMPv1, SNMPv2c, or SNMPv3.

E.3.1

MNR NMA WAN Slot Object

The MNR WAN slot object (`mnrWanSlotObj`) reports an aggregated status for the MNR S2500 or MNR S6000 WAN I/O module(s) on the T1/E1 interfaces on the GGM 8000 gateway.

- For the S2500 and S6000, the `mnrWanSlotObj` object exists only when at least one WAN module is installed in one of the I/O slots.
- For the GGM 8000, the `mnrWanSlotObj` object always exists because T1/E1 interfaces are included as part of the GGM 8000 base unit.
- For the MCG8000, the `mnrWanSlotObj` object does not exist.

The MNR WAN slot object is not created until a minute or more after bootup. This avoids the generation of multiple unnecessary traps during the bootup process when the device configuration is not complete and path connections are established.

The software starts a one-minute alarm during bootup and waits for the device configuration to finish. If the configuration did not finish in a minute, which could happen in large configurations, the alarm is reset for an additional minute. This alarm resetting process continues until the configuration is complete. At this point, the MNR WAN slot object is created, and the device sends a state change trap. If the device receives any queries to the MNR WAN slot object before the object is created, it returns an error.

The S2500 supports the following WAN modules:

- WAN/Telco (T1/E1)
- FlexWAN

The S6000 router supports the following WAN modules:

- UltraWAN and UltraWAN II (T1/E1)
- 12-port T1/E1 and 12-port T1/E1 II
- FlexWAN
- T3/E3
- HSSI

The GGM 8000 gateway supports the following WAN modules:

- T1/E1 (provided as part of base system)



NOTE: Each WAN module is associated with multiple physical paths. All paths should be configured correctly and administratively enabled or disabled based on their usage. If the paths are not configured correctly (for example, if paths that are not used in the device are not administratively disabled), the status returned by the MNR WAN slot object may be misleading. The MNR WAN slot object cannot identify device misconfigurations.

The following table describes how the states returned by the MNR WAN slot object are determined and lists the severity of each state as reported in the network management application.

The EOS device sends a state change trap whenever a state change occurs in the WAN paths that are associated with the MNR WAN Slot object. Events that could cause a state change trap include:

- Device bootup – After the device bootup is complete, the MNR WAN Slot object state is determined, and a trap is sent with the current state.
- WAN path down – A physical WAN path goes down.
- WAN path up – A physical WAN path comes up.
- Switchover – A CWR switchover occurs.

For a complete list of the MIBs and traps (and the corresponding OIDs) supported in the MNR WAN slot object, see the MIB file.

Table 85: MNR WAN Slot Object MIB States

mnrWanSlotObj State	Path Administrative Status	Path Operational Status	Severity
Disabled	All paths are administratively disabled or no paths exist.	Not applicable	Warning
UP	All paths are administratively enabled.	All corresponding paths are up.	Normal/Clear
Down	All paths are administratively enabled.	All corresponding paths are down and CWR control for this slot is disa-	Critical

mnrWanSlotObj State	Path Administrative Status	Path Operational Status	Severity
		bled or CWR control for this slot is enabled and its state is active.	
Dormant	All paths are administratively enabled.	All corresponding paths are down and CWR control for this slot is enabled and its state is standby.	Normal/Clear
Degraded	All paths are administratively enabled.	Some corresponding paths are down/disabled/not available/other.	Minor
Malfunction	All paths are administratively enabled.	All corresponding paths are disabled/not available/other.	Critical

E.3.2

MNR NMA LAN Path Object

The MNR LAN path object (`mnrLanPathObj`) reports the state of the physical LAN (Ethernet) port on the EOS device. This object exists in all supported hardware platforms at all times as there is always one LAN path on the EOS device.

The MNR LAN path object is not created until a minute or more after bootup. This avoids the generation of multiple unnecessary traps during the bootup process when the device configuration is not complete and path connections are established.

The software starts a one-minute alarm during the bootup and waits for the device configuration to finish. If the configuration is not finished in a minute, which could happen in large configurations, the alarm is reset for an additional minute. The alarm resetting process continues until the configuration is complete. At this point, the MNR LAN path object is created, and the device sends a state change trap. If the device receives any queries to the MNR LAN path object before the object is created, it returns an error.



NOTE: All paths should be configured correctly and administratively enabled or disabled based on their usage. If the paths are not configured correctly (for example, if paths that are not used in the device are not administratively disabled), the status returned by the MNR LAN path object may be misleading. The MNR LAN path object cannot identify device misconfigurations.

The following table describes how the states returned by the MNR LAN path object are determined and lists the severity of each state as reported in the network management application.

The EOS device sends a state change trap whenever a state change occurs in the LAN path associated with the MNR LAN Path object.

For a complete list of the MIBs and traps (and the corresponding OIDs) supported in the MNR WAN slot object, see the MIB file.

Table 86: MNR LAN Path Object MIB States

mnrLanPathObj State	Path Administrative Status	Path Operational Status	Severity
Malfunction	Enabled	Not available	Critical

mnrLanPathObj State	Path Administrative Status	Path Operational Status	Severity
Malfunction	Disabled	Other	Critical
Disabled (default)	Not applicable	Disabled	Warning
Down	Not applicable	Down	Critical
Up	Not applicable	Up	Normal/Clear

E.3.3

MNR NMA Device Status Object

EOS now supports the MNR Device Status object (`mnrDeviceStatusObj`). Fault managers can use this object to determine if a device is enabled.

Supported devices determine the status of the Device Status object and report it through the SNMP interface by following NMA fault object standards. The following table lists the states returned by the MNR Device Status object and the severity of each state as reported in the network management application.



NOTE: The MNR Device Status object (`mnrDeviceStatusObj`) is supported only on the GGM 8000, MCG 8000 and the MNR S6000 hardware platforms; it is not supported on the MNR S2500 hardware platform.

Table 87: Device Status Object State Definitions

mnrDeviceStatusState	Value	Description	Severity
Initialization	1	The MNR device booted and is in the process of initializing all services.	Warning
Enabled	2	The MNR device booted completely and all required services are enabled.	Clear

Supported devices send a state change trap when the Device Status changes.

Table 88: Device Status Object State Change Trap Definition

SNMP Trap OID	Packet Type	Description	Category Name
<code>mnrDeviceStatusReg.0.1</code>	Inform	This trap is sent when the device status changes: Variables: <code>mnrDeviceStatusState.mnrDeviceStatusInstanceID</code> <code>mnrDeviceStatusCause.mnrDeviceStatusInstanceID</code> <code>mnrDeviceStatusText.mnrDeviceStatusInstanceID</code> <code>nmaTimeInSec</code> <code>nmaTimeInMilliSec</code>	NMA State Cause Events

SNMP Trap OID	Packet Type	Description	Category Name
		<code>mnrDeviceStatusInstanceID</code> – A unique identifier for the Device Status object. <code>mnrDeviceStatusState</code> – The state of the MNR device: 1 – Initialization 2 – Enabled	
		<code>mnrDeviceStatusCause</code> – The cause for the most recent device status state change: 1 – Boot process started (The MNR device received the boot-start command, and the boot process started). 2 – Boot completed (The MNR device boot is complete).	
		<code>mnrDeviceStatusText</code> – Supplemental text that provides additional information about the Device Status state change trap. This variable is meaningful only in a trap PDU. If it is requested through Get or GetNext, a null string is returned. <code>nmaTimeInSec</code> – The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). <code>nmaTimeInMilliSec</code> – The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	

E.3.4

Ethernet WAN Link Object MIB

EOS devices determine the status of the Ethernet WAN links and report this status through the SNMP interface by following the NMA fault object standards. The devices send an Ethernet WAN link state change trap when the BFD status changes (for site links or GGSN-to-BR links) or the BGP peer status changes (for interzone links). The Ethernet WAN link object (`mnrEthWanIfFlt`) reports the state of the Ethernet WAN links (up, down, disabled, degraded, or dormant). The following table lists the states returned by the Ethernet WAN link object and the severity of each state as reported in the UEM.

Table 89: MNR Ethernet WAN Link Object MIB State Definitions

<code>mnrEthWanIfFlt</code> State	Value	Description	Severity
up	1	The Ethernet WAN link is connected and BFD determined that the link is up (site links or GGSN-to-BR links) or the BGP peers are up (interzone links). For GGSN-to-BR links, the VRRP status is Master.	Clear

mnEthWanIfFlt State	Value	Description	Severity
down	2	The Ethernet WAN link is down. BFD determined that the Ethernet WAN link is down (site links or GGSN-to-BR links) or a BGP peer is down (interzone links). For GGSN-to-BR links, the VRRP status is either Master or Disabled.	Critical
disabled	3	The Ethernet WAN link been administratively disabled.	Clear
degraded	4	The Ethernet WAN link is connected but degraded. BFD (site links) or BGP (interzone links) determined that the link or BGP peers are up and an Ethernet WAN link performance threshold was exceeded.	Major
dormant	5	The Ethernet WAN link is down, and the VRRP status is Backup. This state applies to GGSN-to-BR links only.	Clear

Ethernet WAN Link Object Traps and Alarms

EOS sends:

- An Ethernet WAN link state change trap when the BFD status changes (for site links or GGSN-to-BR links) or the BGP peer status changes (for interzone links). The following table describes the events that trigger a state change.

Table 90: MNR Ethernet WAN Interface Fault Object State Change Triggers

Event	Event Text	Description
Evt-up (1)	Link established	This event is generated when the Ethernet WAN link is established and BFD (site links) or BGP (interzone links) determined that the link is up.
Evt-down (2)	Link down	This event is generated when BFD determined that the link is down (site links) or when the BGP peer is down (interzone links).
Evt-disabled (3)	Link administratively disabled	This event is generated when the Ethernet WAN link was administratively disabled.
Evt-limit-exceeded (4)	Static measurement exceeded configured threshold value	This event is generated when the performance measurement on the Ethernet WAN link exceeds one of the following configured threshold values: Average round trip IPTD value (AvgRtXfer)

Event	Event Text	Description
		Average IPDV 99th percentile (Avg99Percent) IPLR loss rate (PktRate)
Evt-within-limit (5)	Static measurement within configured threshold value	This event is generated when the performance measurement on the Ethernet WAN link is within the configured threshold values for the following measurements: Average round trip IPTD value (AvgRtXfer) Average IPDV 99th percentile (Avg99Percent) IPLR loss rate (PktRate)
ggsn-br-evt-up (6)	GGSN BR Link established	This event is generated when the GGSN-to-BR link is established, BFD determined that the link is up, and the VRRP status is Master.
ggsn-br-evt-down (7)	GGSN BR Link Down	This event is generated when the GGSN-to-BR link is down, and the VRRP status is Master or Disabled.
ggsn-br-evt-disabled (8)	GGSN BR Link administratively disabled	This event is generated when the GGSN-to-BR link was administratively disabled.
ggsn-br-evt-standby (9)	GGSN BR Standby Link	This event is generated when BFD determined that the GGSN-to-BR link is down and the VRRP status is Backup.

For a complete list of the MIBs and traps (and corresponding OIDs) supported in the Ethernet WAN link object, see the MIB file.

- A transient trap and an alarm when the threshold is exceeded on the Ethernet WAN interface for the following IPDV (IP packet delay variation) values:
 - Average round-trip IPTD (IP packet transfer delay)
 - Average IPDV 99th percentile value
 - IPLR (IP packet loss rate)

For more information about these measurements, see [Delay and Jitter Performance Management Tool on page 1032](#).

The following table describes the alarms sent when an IPDV threshold value is exceeded on an Ethernet WAN interface.

Table 91: Ethernet WAN Interface IPDV Alarms

Alarm Code	Alarm Text	Criticality
1	Threshold value exceeded for average round-trip IPTD measurement	Major
2	Threshold value exceeded for average IPDV 99th percentile measurement	Major
3	Threshold value exceeded for IPLR measurement	Major

E.3.5**MNR Service ID MIB**

The MNR service ID MIB (*mnrSysServiceID*) identifies whether or not a device is a CCGW or a GGSN.

The possible values for the variable are:

1 – Undefined

This value is returned if the device was identified as neither a CCGW nor a GGSN or both a CCGW and a GGSN.

2 – CCGW

If the `SETDefault -CCGW CONTROL = Enabled` command is read from the `boot.cfg` file when the device boots, the device is identified as a CCGW, and it remains identified as such until the next reboot. The device remains identified as a CCGW even if the `SETDefault -CCGW CONTROL = Disabled` command is executed. To not be identified as a CCGW, the device must be rebooted with a `boot.cfg` file that includes the `SETDefault -CCGW CONTROL = Disabled` command or no `SETDefault -CCGW CONTROL` command (since the `CCGW CONTROL` parameter is set to `Disabled` by default).



NOTE: You can also use the absence or presence of the `mnrWanSlotObj` variable to determine if the device is a CCGW only or both a CCGW and a site router/gateway.

- If the *mnrWanSlotObj* variable does not exist, and the *mnrSysServiceID* variable returns a value of 2 (CCGW), the device is a CCGW only.
- If the *mnrWanSlotObj* variable exists, and the *mnrSysServiceID* variable returns a value of 2 (CCGW), the device is both a CCGW and a site router/gateway.

3 – GGSN

If the `SETDefault -GGSN CONTROL = Enabled` command is read from the `xgsn.cfg` file when the device boots, the device is identified as a GGSN, and it remains identified as such until the next reboot. The device remains identified as a GGSN even if the `SETDefault -GGSN CONTROL = Disabled` command is executed. To not be identified as a GGSN, the device must be rebooted with an `xgsn.cfg` file that includes the `SETDefault -GGSN CONTROL = Disabled` command or no `SETDefault -GGSN CONTROL` command (since the `GGSN CONTROL` parameter is set to `Disabled` by default).

E.3.6**MNR GGSN Fault Object**

The MNR GGSN fault object (*mnrGgsnFltObj*) reports the status of the GGSN configuration in the network. The MNR GGSN fault object is supported only on the MNR S6000 and GGM 8000 hardware platforms.



NOTE: The MNR GGSN fault object should not be removed from a device in a configured system.

Supported devices determine the status of the GGSN configuration and report it through the SNMP interface following NMA fault object standards. The following table lists the states returned by the MNR GGSN fault object and the severity of each state as reported in the network management application.

Table 92: MNR GGSN Fault Object State Definitions

mnrGgsnFltObj State	Value	Description	Severity
up	1	The GGSN configuration file was executed, and the GGSN service is enabled.	Clear
down	2	The GGSN configuration file may or may not have been executed, and the GGSN service is disabled.	Critical

GGSN Fault Object Traps

Supported devices send a state change trap when the GGSN configuration status changes and a transient trap when an error occurs during configuration file processing.

Table 93: MNR GGSN Fault Object Trap Definitions

SNMP Trap OID	Packet Type	Description	Category Name
mnrGgsnFltReg.0.1	Inform	This trap is sent when the GGSN configuration state changes. Variables: mnrGgsnFltState.mnrGgsnFltInstanceID mnrGgsnFltCause.mnrGgsnFltInstanceID nmaTimeInSec nmaTimeInMilliSec mnrGgsnFltState – The state of the GGSN configuration: 1 – The GGSN configuration is UP. 2 – The GGSN configuration is DOWN. mnrGgsnFltInstanceID – A unique identifier for the GGSN fault object. mnrGgsnFltCause – The cause for the most recent GGSN configuration state change: 1 – GGSN Ready. The GGSN configuration is complete (GGSN-Service-Enabled), and the GGSN is ready. 2 – GGSN not Ready. The GGSN configuration is incomplete (GGSN-Service-Disabled), and the GGSN is not ready.	NMA State Cause

SNMP Trap OID	Packet Type	Description	Category Name
		nmaTimeInSec – The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec – The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrGgsnFltReg.0.2	Trap	This trap is sent when an error is reported during configuration file execution. Variables: mnrGgsnFltAlarmCode.mnrGgsnFltInstanceID mnrGgsnFltText.mnrGgsnFltInstanceID nmaTimeInSec nmaTimeInMilliSec mnrGgsnFltAlarmCode – A code indicating the nature of the error reported during configuration file processing: 1 – boot.cfg configuration file execution error (Major alarm) 2 – xgsn.cfg configuration file execution error (Major alarm) 3 – boot.cfg and xgsn.cfg configuration file execution error (Major alarm) This variable is meaningful only in a trap PDU. If it is requested through Get or GetNext, a null string is returned. mnrGgsnFltInstanceID – A unique identifier for the GGSN fault object. mnrGgsnFltText – Supplemental text that provides additional information about the GGSN transient fault: Boot config error count – The number of commands in the boot.cfg file that failed.	NMA Transient Events

SNMP Trap OID	Packet Type	Description	Category Name
		GGSN config error count – The number of commands in the <code>xgsn.cfg</code> file that failed. Boot and GGSN error counts – The number of commands in both the <code>boot.cfg</code> and <code>xgsn.cfg</code> files that failed. This variable is meaningful only in a trap PDU. If it is requested through Get or GetNext, a null string is returned. <code>nmaTimeInSec</code> – The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). <code>nmaTimeInMilliSec</code> – The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	

E.3.7

MNR Mobile IP Link Fault Object

The MNR Mobile IP Link fault object (`mnrMobileIPLinkObj`) reports the state of the link between the Mobile IP agents: the foreign agent running on the GGSN and the home agent running on the border router. This link state is determined based on the multihop BFD status between the two Mobile IP agents. Both the foreign agent and the home agent reported this object.

Supported devices determine the status of the Mobile IP Link fault object and report it via the SNMP interface following NMA fault object standards. The following are the states returned by the MNR Mobile IP Link fault object and the severity of each state as reported in the network management application.

Table 94: Mobile IP Link Fault Object State Definitions

<code>mnrMobileIPLink-State</code>	Value	Description	Severity
Up	1	The BFD link between the home agent and the foreign agent comes up.	Clear
Down	2	The BFD link between the home agent and the foreign agent goes down.	Critical
Dormant	3	The BFD link between the home agent and the foreign agent is in standby state.	Clear

mnrMobileIPLink-State	Value	Description	Severity
Critical Malfunction	4	The foreign agent Care Of Address does not match the configured value on the home agent.	Major

Mobile IP Link Fault Object Traps

Supported devices send a state change trap when Mobile IP link status changes and a transient trap when there is an error reported after a Mobile IP Registration Request is sent to the home agent. The following table provides descriptions of these traps.

Table 95: Mobile IP Link Fault Object Trap Definitions

SNMP Trap OID	Packet Type	Description	Category Name
mnrMobileI-PLinkReg.0.1	Inform	This trap is sent when the state of the BFD link between the Mobile IP agents (the foreign agent running on the GGSN and the home agent running on the border router) changes: Variables: mnrMobileIPLinkAlias.mnrMobileIPLinkInstanceID mnrMobileIPLinkState.mnrMobileIPLinkInstanceID mnrMobileIPLinkCause.mnrMobileIPLinkInstanceID mnrMobileIPLinkText.mnrMobileIPLinkInstanceID nmaTimeInSec nmaTimeInMilliSec mnrMobileIPLinkAlias – A text string indicating the type of mobility agent: Home Agent or Foreign Agent. mnrMobileIPLinkInstanceID – A unique identifier for the Mobile IP Link fault object. mnrMobileIPLinkState – The state of the BFD link between the Mobile IP agents: 1 – Up 2 – Down 3 – Dormant 4 – Critical Malfunction mnrMobileIPLinkCause – The cause for the most recent Mobile IP link state change: 1 – Mobility Agent reachable (The BFD link between the foreign agent and the home agent went up).	NMA State Cause Events

SNMP Trap OID	Packet Type	Description	Category Name
		2 – Mobility Agent not reachable (The BFD link between the foreign agent and the home agent went down). 3 – Mobility Agent in StandBy (The mobility agent is acting as a standby to the active foreign agent or home agent). 4 – Care Of Address does not match configured value (The Care Of Address does not match the configured value on the home agent). mnrMobileIPLinkText – Supplemental text that provides additional information about the Mobile IP link state change trap. This variable is meaningful only in a trap PDU. If it is requested through Get or GetNext, a null string is returned. nmaTimeInSec – The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec – The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrMobileIPLinkReg0.2	Trap	This trap is sent when an error is reported after a Mobile IP Registration Request is sent to the home agent. Variables: mnrMobileIPLinkAlias.mnrMobileIPLinkInstanceID mnrMobileIPLinkAlarmCode.mnrMobileIPLinkInstanceID mnrMobileIPLinkText.mnrMobileIPLinkInstanceID nmaTimeInSec nmaTimeInMilliSec mnrMobileIPLinkAlias – A text string indicating the type of mobility agent: Home Agent or Foreign Agent. mnrMobileIPLinkInstanceID – A unique identifier for the Mobile IP link fault object. mnrMobileIPLinkAlarmCode – A code indicating the nature of the error reported after the Mobile IP Registration Request was sent to the home agent:	NMA Transient Events

SNMP Trap OID	Packet Type	Description	Category Name
		0 – "" default empty text; alarm code 0 is not considered a trap to the network manager (Info) 1 – No registration reply from home agent (Major alarm) 2 – Registration reply from home agent: Registration Request rejected (Major alarm) 3 – Registration rejected by home agent: Invalid home agent/home address/ Care Of Address (Major alarm) 4 – Deregistration rejected by Home Agent: Care Of Address mismatch with existing mobile node (Major alarm) This variable is meaningful only in a trap PDU. If it is requested through Get or GetNext, a null string is returned. mnrMobileIPLinkText – Supplemental text that provides additional information about the Mobile IP link transient fault. This variable is meaningful only in a trap PDU. If it is requested through Get or GetNext, a null string is returned. nmaTimeInSec – The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec – The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	

E.3.8

MNR Mobile IP Service Fault Object

The MNR Mobile IP Service fault object (`mnrMobileIPServiceObj`) reports the status of Mobile IP service on either the foreign agent or the home agent.

Supported devices determine the status of the Mobile IP Service fault object and report it via the SNMP interface following NMA fault object standards. The following are the states returned by the MNR Mobile IP Service fault object and the severity of each state as reported in the network management application.

Table 96: Mobile IP Service Fault Object State Definitions

mnrMobileIPServiceState	Value	Description	Severity
InService	1	The Mobile IP service is in service on this device because the VRRP	Clear

mnrmobileIPServiceState	Value	Description	Severity
		state is "VRRP Master" (this device is the master VRRP).	
OutOfService	2	The Mobile IP service is out of service on this device because the MOBILEIP service was disabled.	Critical
Dormant	3	The Mobile IP service is dormant on this device because the VRRP state is "VRRP Backup" (this device is the backup VRRP).	Clear

Mobile IP Service Fault Object Trap

Supported devices send a state change trap when the Mobile IP service status changes.

Table 97: Mobile IP Service Fault Object Trap Definition

SNMP Trap OID	Packet Type	Description	Category Name
mnrmobileIPServiceReg.0.1	Inform	This trap is sent when the state of the Mobile IP service changes: Variables: mnrmobileIPServiceAlias.mnrmobileIPServiceInstanceID mnrmobileIPServiceState.mnrmobileIPServiceInstanceID mnrmobileIPServiceCause.mnrmobileIPServiceInstanceID mnrmobileIPServiceText.mnrmobileIPServiceInstanceID nmaTimeInSec nmaTimeInMilliSec mnrmobileIPServiceAlias – A text string indicating the type of mobility agent: Home Agent or Foreign Agent. mnrmobileIPServiceInstanceID – A unique identifier for the Mobile IP Service fault object. mnrmobileIPServiceState – The state of the Mobile IP service on this mobility agent: 1 – InService 2 – OutOfService 3 – Dormant mnrmobileIPServiceCause – The cause for the most recent Mobile IP service state change:	NMA State Cause Events

SNMP Trap OID	Packet Type	Description	Category Name
		1 – Service up with Peer (This device becomes the master VRRP, and a redundancy peer is detected). 2 – Service up without Peer (This device becomes the master VRRP, and a redundancy peer is NOT detected). 3 – Service in Standby (This device becomes the backup VRRP). 4 – Service not Initialized (The MOBILEIP service is disabled on this device).	
		<code>mnrMobileIPServiceText</code> – Supplemental text that provides additional information about the Mobile IP service state change trap. This variable is meaningful only in a trap PDU. If it is requested through Get or GetNext, a null string is returned. <code>nmaTimeInSec</code> – The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). <code>nmaTimeInMilliSec</code> – The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	

E.3.9

MNR DBReplicator Fault Object

The MNR DBReplicator fault object (`mnrDBReplicatorObj`) reports the replication status for the associated replication database clients. Replication of the client database is triggered when a device is informed of a peer via VRRP.

- The backup VRRP router initiates a three-way state exchange control protocol (handshake) with its peer (master VRRP router).
- Once the handshake is complete, the devices transition from a "Down" or "Malfunction" state to a state of being inter-connected.
- The devices enter the "Synchronizing" state when the master VRRP router initiates replication of its database toward the backup VRRP router. This state is an intermediate state, and should last only a few seconds.
- The devices become "Synchronized" when the database replication is successfully completed.
- If connectivity between the peers is disrupted, both devices transition from "Synchronized" to "Malfunction".

Supported devices report the state of the MNR DBReplicator fault object through the SNMP interface following NMA fault object standards. The following are the states returned by the MNR DBReplicator fault object and the severity of each state as reported in the network management application.

Table 98: DBReplicator Fault Object State Definitions

mnrDBReplicatorState	Value	Description	Severity
Synchronized	1	Database replication completed successfully	Clear
Synchronizing	2	The replication protocol on this device completed a three-way handshake with its replication peer and is in the process of synchronizing with the peer.	Warning
Malfunction	3	The replication protocol on this device is in “Malfunction” state and is waiting for its peer to communicate.	Clear

MNR DBReplicator Fault Object Trap

Supported devices send a state change trap when the replication status changes.

Table 99: DBReplicator Fault Object Trap Definition

SNMP Trap OID	Packet Type	Description	Category Name
mnrDBReplicatorReg.0.1	Inform	This trap is sent when the state of the Mobile IP service changes: Variables: mnrDBReplicatorAlias.mnrDBReplicatorInstanceID mnrDBReplicatorState.mnrDBReplicatorInstanceID mnrDBReplicatorCause.mnrDBReplicatorInstanceID mnrDBReplicatorText.mnrDBReplicatorInstanceID nmaTimeInSec nmaTimeInMilliSec mnrDBReplicatorAlias – A text string indicating the type of mobility agent: Home Agent or Foreign Agent. mnrDBReplicatorInstanceID – A unique identifier for the DBReplicator fault object. mnrDBReplicatorState – The database replication status on this mobility agent: 1 – Synchronized 2 – Synchronizing 3 – Malfunction mnrDBReplicatorCause – The cause for the most recent database replication state change:	NMA State Cause Events

SNMP Trap OID	Packet Type	Description	Category Name
		1 – Replicator – Bulk update with Peer completed (Database replication completed successfully). 2 – Replicator – handshake with Peer succeeded (The replication protocol on this device exchanged a handshake with its peer). 3 – Replicator – Bulk update with Peer failed (Database replication did not complete, possibly due to unavailability of the peer). 4 – Replicator – Peer down (No peer communication took place).	
		mnrDBReplicatorText – Supplemental text that provides additional information about the replication database state change trap. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. nmaTimeInSec – The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec – The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	

E.3.10

MNR DHCP Service Fault Object

The MNR DHCP service fault object (**mnrDHCPServiceObj**) reports the state of the DHCP service running on the border router.

Supported devices determine the status of the DHCP service fault object and report it via the SNMP interface following NMA fault object standards. The following are the states returned by the MNR DHCP service fault object and the severity of each state as reported in the network management application.

Table 100: DHCP Service Fault Object State Definitions

mnrDHCPServiceState	Value	Description	Severity
InService	1	The DHCP service is in service because this device is the master VRRP.	Clear
OutOfService	2	The DHCP service is out of service either because there is no DHCP pool on the device or no enabled DHCP port is available.	Critical

mnrdHCPServiceState	Value	Description	Severity
Dormant	3	The DHCP service is dormant because this device is the backup VRRP.	Clear

DHCP Service Fault Object Traps

Supported devices send a state change trap when the DHCP service status changes and a transient trap when there is an error reported during configuration file execution.

Table 101: DHCP Service Fault Object Trap Definitions

SNMP Trap OID	Packet Type	Description	Category Name
mnrdHCPServiceReg.0.1	Inform	This trap is sent when the state of the DHCP service on the border router changes: Variables: mnrdHCPServiceAlias.mnrdHCPServiceInstanceID mnrdHCPServiceState.mnrdHCPServiceInstanceID mnrdHCPServiceCause.mnrdHCPServiceInstanceID mnrdHCPServiceText.mnrdHCPServiceInstanceID nmaTimeInSec nmaTimeInMilliSec mnrdHCPServiceAlias – A text string indicating whether the DHCP service is Active or Inactive on this device. mnrdHCPServiceInstanceID – A unique identifier for the DHCP Service fault object. mnrdHCPServiceState – The state of the DHCP service on this border router: 1 – In Service 2 – OutOfService 3 – Dormant mnrdHCPServiceCause – The cause for the most recent DHCP service state change: 1 – DHCP Service up with peer (This device becomes the master VRRP, and a redundancy peer is detected). 2 – DHCP Service Up without Peer (This device becomes the master VRRP, and a redundancy peer is NOT detected). 3 – DHCP Service is Standby (This device becomes the backup VRRP).	NMA State Cause Events

SNMP Trap OID	Packet Type	Description	Category Name
		4 – DHCP Service is not Initialized (No DHCP pool exists on the device or no DHCP port is enabled). <code>mnrDHCPServiceText</code> – Supplemental text that provides additional information about the DHCP service state change trap. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. <code>nmaTimeInSec</code> – The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). <code>nmaTimeInMilliSec</code> – The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>mnrDHCPServiceReg0.2</code>	Trap	This trap is sent when an error is reported during configuration file execution. Variables: <code>mnrDHCPServiceAlias.mnrDHCPServiceInstanceID</code> <code>mnrDHCPServiceAlarmCode.mnrDHCPServiceInstanceID</code> <code>mnrDHCPServiceText.mnrDHCPServiceInstanceID</code> <code>nmaTimeInSec</code> <code>nmaTimeInMilliSec</code> <code>mnrDHCPServiceAlias</code> – A text string indicating whether the DHCP service is Active or Inactive on this device. <code>mnrDHCPServiceInstanceID</code> – A unique identifier for this DHCP service fault object. <code>mnrDHCPServiceAlarmCode</code> – A code indicating the nature of the error reported during configuration file execution: 0 – "" default empty text; alarm code 0 is not considered a trap to the network manager (Info) 1 – DHCP runs out of dynamic IP addresses to allocate to a mobile subscriber (Major alarm) 2 – DHCP already assigned the requested IP address for renewal to a different mobile	NMA Transient Events

SNMP Trap OID	Packet Type	Description	Category Name
		subscriber (Major alarm)  NOTE: Alarm code 2 is supported in the EOS code but, per the current system design, is not sent by EOS devices. This variable is meaningful only in a trap PDU. If it is requested through Get or GetNext, a null string is returned. mnrDHCPServiceText – Supplemental text that provides additional information about the DHCP service transient fault. This variable is meaningful only in a trap PDU. If it is requested through Get or GetNext, a null string is returned. nmaTimeInSec – The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec – The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	

E.3.11

MNR CCGW Private MIBs

MNR CCGW private MIBs are MNR private MIBs supported by the MNR S2500s and the GGM 8000 gateways configured as conventional channel gateways (CCGWs) and the MCG 8000 platform.

These MIBs follow Motorola Solutions Network Management Alliance (NMA) standards. The following are the filenames, descriptions, and MIB names for each of the CCGW private MIBs.

Table 102: MNR CCGW Private MIBs

Filename	Description	MIB Name
cgcw-agent-reg-mib.mib	Generic agent MIB	CCGW-AGENT-REG-MIB
cgcw-ccgw-mib.mib	Product-specific MIB for the CCGW box	CONSOLE-CCGW-MIB
cgcw-ccgw-reg-mib.mib	Registration MIB for the CCGW box	CCGW-CCGW-REG-MIB
cgcw-channel-mib.mib	Product-specific MIB for the conventional channels	CONSOLE-CCGW-CHANNEL-MIB
cgcw-channel-reg-mib.mib	Registration MIB for the conventional channels	CCGW-CHANNEL-REG-MIB

Filename	Description	MIB Name
cgcw-zccp-mib.mib	Product-specific MIB for the control path from the CCGW to the zone controller	CONSOLE-SITE-ZC-CONTROL-PATH-MIB
cgcw-zccp-reg-mib.mib	Registration MIB for the control path from the CCGW to the zone controller	CCGW-ZCCP-REG-MIB
cgcw-rnglink-mib.mib	Product-specific MIB for the LAPD link between the CCGW and the RNG. This link supports IV&D applications.	CCGW-RNG-LINK-MIB
cgcw-rnglink-reg-mib.mib	Registration MIB for the LAPD link between the CCGW and the RNG. This link supports IV&D applications.	CCGW-RNG-LINK-REG-MIB
cgcw-rng-bcastlink-mib.mib	Product-specific MIB for the broadcast link between the CCGW and the RNG. This link supports IV&D broadcast applications.	CCGW-RNG-BCASTLINK-MIB
cgcw-rng-bcastlink-reg-mib.mib	Registration MIB for the broadcast link between the CCGW and the RNG. This link supports IV&D broadcast applications.	CCGW-RNG-BCASTLINK-REG-MIB

E.3.12

MNR NMA Private MIBs

MNR NMA private MIBs are MNR product MIBs that follow the Motorola Solutions Network Management Alliance (NMA) standards.

The following are the filenames, descriptions, and MIB names for each of the MNR NMA private MIBs.

Table 103: MNR NMA Private MIBs

Filename	Description	MIB Name
mnr-agent-reg-mib.mib	Generic agent MIB	MNR-AGENT-REG-MIB
mnr-ethwanflt-mib.mib	Product-specific MIB for the Ethernet WAN link fault object	MNR-ETHWANFLT-MIB
mnr-ethwanflt-reg-mib.mib	Registration MIB for the Ethernet WAN link fault object	MNR-ETHWANFLT-REG-MIB
mnr-ethwanperf-mib.mib	Product-specific MIB for the Ethernet WAN link performance object	MNR-ETHWANPERF-MIB
mnr-ethwanperf-reg-mib.mib	Registration MIB for the Ethernet WAN link performance object	MNR-ETHWANPERF-REG-MIB
mnr-lanpath-mib.mib	Product-specific MIB for the MNR LAN path object	MNR-LANPATH-MIB

Filename	Description	MIB Name
mnr-lanpath-reg-mib.mib	Registration MIB for the MNR LAN path object	MNR-LANPATH-REG-MIB
mnr-wanslot-mib.mib	Product-specific MIB for the MNR WAN slot object	MNR-WANSLOT-MIB
mnr-wanslot-reg-mib.mib	Registration MIB for the MNR WAN slot object	MNR_WANSLOT-REG-MIB

E.3.13

Standard MIBs

EOS devices support the following MIB modules defined by the Internetworking Engineering Task Force (IETF), Motorola Solutions, IBM, and Novell:

- RFC 1213 (MIB II with interface group obsoleted by RFC 1573)
 - RFC 1253 (OSPF MIB)
 - RFC 1271 (RMON Alarm and Event MIB)
 - RFC 1284 (Ethernet-like MIB)
 - RFC 1286 (MIB fragment to get MacAddress)
 - RFC 1304 (SIP MIB)
 - RFC 1315 (Frame Relay DTE MIB)
Except for frCircuitCommittedBurst and frCircuitExcessBurst
 - RFC 1354 (IP Forwarding MIB)
 - RFC 1493 (Bridge MIB)
 - RFC 1657 (BGP4 Standard MIB)
 - RFC 1659 (RS232 Hardware Devices)
 - RFC 2127 (ISDN MIB using SMIv2)
 - RFC 2863 (MIB II ifTable, ifStackTable, ifXTable)
-  **NOTE:** RFC 2863 replaces the obsolete RFC 1573.
- RFC IANAifType (Internet Assigned Numbers Authority ifType MIB)

E.3.14

SNMPv3 MIBs

The SNMPv3 Common Agent supports the following MIBs:

- USM
- VACM
- SNMP-FRAMEWORK-MIB
- snmpMPDStats
- snmp_mib

- MMR



NOTE: The SNMPv3 Common Agent uses NET-SNMP, which is a Motorola Solutions-approved software product.

E.3.15

MNR Private MIBs

MNR private MIB modules control various services.

The following table lists the filename, description, and MIB name for each of the Motorola Solutions private MIBs.

Table 104: Motorola Solutions Non-CCGW Private MIBs

Filename	Description	MIB Name
_moca_0004.mib	Generic MIB	AMotorolaCa0004-GENERIC
_moca_0019.mib	Remote Poll MIB	AMotorolaCa0019-RMON-RE-MOTE_POLL
_moca_0024.mib	Event Extension MIB	AMotorolaCa0024-EVENT-EXTENSION
_moca_0027.mib	RMON Extensions MIB	AMotorolaCa0027-RMON-EXTENSIONS
_moca_products.mib	Enterprise MIB	AMotorolaCa-products-MIB
audlMIB.mib	AuditLog MIB	AMotorolaCa-AUDL-r4-MIB
bgpMIB.mib	BGP MIB	MOCA-BGP-r1-MIB
brMIB.mib	Bridge Extension MIB	AMotorolaCa-Bridge-r5-MIB
cdrMIB.mib	Call Detailed Records MIB	AMotorolaCa-CDR-r2-MIB
MIB.mib	CRYPTO (encryption) MIB	AMotorolaCa-Crypto-TRAP-r1-MIB
cwr.mib	Cooperative WAN Routing MIB	AMotorolaCa-CWR-r1-MIB
dhcpMIB.mib	DHCP MIB	AMotorolaCa-DHCP-r1-MIB
dvmrpMIB.mib	DVMRP MIB	AMotorolaCa-Dvmrp-r1-MIB
envMon.mib	Environment Monitoring MIB	AMotorolaCa-ENVMON-r1-MIB
filtMIB.mib	Filtering MIB	AMotorolaCa-Filter-r7-MIB
frMIB.mib ¹	Frame Relay MIB	AMotorolaCa-FR-r1-MIB
ipMIB.mib	IP MIB	AMotorolaCa-IPextns-r7-MIB
ipripMIB.mib	RIP IP MIB	AMotorolaCa-RIP-IPextns-r5-MIB
ipsoMIB.mib	IP Security Options MIB	AMotorolaCa-IPSO-r1-MIB
mipMIB.mib	Multicast IP MIB	AMotorolaCa-Mip-r1-MIB
mnr_ggsn.mib	GGSN MIB	AMotorolaCa-GGSN-r1-MIB
mnr_ospf.mib	OSPF MIB	AMotorolaCa-OSPF-r1-MIB
mnrSysTools.mib	System (EOS) Tools MIB	AMotorolaCa-SysTools-r1-MIB

Filename	Description	MIB Name
mnrtraps.mib	MNR Router MIB 2 Traps	AMotorolaCa-MIB2TRAPS-MIB
mnr_WANQoS.mib	IPQoS MIB	AMotorolaCa-WANQos-r1-MIB
natMIB.mib	NAT MIB	AMotorolaCa-NAT-r1-MIB
perfMIB.mib	PerformanceMIB	AMotorolaCa-Perf-r3-MIB
pimMIB.mib	PIM MIB	AMotorolaCa-PIM-r1-MIB
ppMIB.mib	Port and Path MIB	AMotorolaCa-PortPath-r15-MIB
rasMIB.mib	Remote Access MIB	AMotorolaCa-RemoteAccess-r2-MIB
rdpMIB.mib	Router Discovery MIB	AMotorolaCa-Rdp-r1-MIB
rempextMIB.mib	Remote Poll Extension MIB	AMotorolaCa-RMON-REMOTE-POLL-EXTENSION-r2-MIB
sysMIB.mib	System MIB	AMotorolaCa-System-r15-MIB
t1e1.mib	DS1 (T1)/E1 MIB	AMotorolaCa-DS1E1-r1-MIB
vrrpMIB.mib	VRRP MIB	AMotorolaCa

¹ The frMIB.mib MIB includes a MIB table (AMotorolaCaFrLmiStatsEntry) to support LMI statistics. For more information about the supported LMI statistics, see [FR Service Statistics on page 971](#).

E.4

Traps Supported on EOS Devices

Traps generated by EOS and supported on EOS devices are divided into the following three categories:

EOS Non-Enterprise-Specific Traps

MIB-2 traps defined by standard MIBs and RFCs under the auspices of the Internet Engineering Task Force (IETF).

These traps **do not** contain the trap event timestamp and SysName fields. For details about these fields, see [SysName and Trap Event Timestamp in EOS Enterprise-Specific Traps on page 878](#).

You can enable or disable these traps by using the SNMP service Mib2TrapControl parameter.

For information on the non-enterprise-specific traps generated by EOS and supported on EOS devices, see [EOS Non-Enterprise-Specific Traps on page 872](#).

EOS Enterprise-Specific Traps

EOS enterprise-specific traps in the “aMotorolaCa” MIB group; customized to meet specific private enterprise information requirements.

These traps **do** contain the trap event timestamp and SysName fields. For details about these fields, see [SysName and Trap Event Timestamp in EOS Enterprise-Specific Traps on page 878](#).

Except for the GENERIC traps, described in [EOS Derived MIB-2 Enterprise-Specific Traps on page 876](#), these traps are always enabled.

For information on the enterprise-specific traps generated by EOS and supported on EOS devices, see [EOS Enterprise-Specific Traps on page 872](#).

EOS Derived MIB-2 Enterprise-Specific Traps

Enterprise-specific traps derived from the MIB-2 non-enterprise-specific traps listed in [Table 105: EOS Non-Enterprise-Specific Traps on page 872](#).

These traps **do** contain the trap event timestamp and SysName fields. For details about these fields, see [SysName and Trap Event Timestamp in EOS Enterprise-Specific Traps on page 878](#).

You can enable or disable these traps by using the EOS SNMP service `Mib2TrapControl` parameter.

For information on the derived MIB-2 enterprise-specific traps generated by EOS and supported on EOS devices, see [EOS Derived MIB-2 Enterprise-Specific Traps on page 876](#).

E.4.1

EOS Non-Enterprise-Specific Traps

The following table lists the EOS non-enterprise-specific traps supported on EOS devices. For further details about these traps, see [Trap Descriptions on page 888](#).



NOTE: The Trap Profiles listed in the third column of the following table refer to the predefined trap profiles supported by EOS. For details about how to configure trap profiles and map them to network management stations, see [Traps and Trap Profiles on page 664](#). For information about how to add and delete custom trap profiles, see the description of the `TrapProfile` parameter in “SNMP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

For trap descriptions, see [Trap Descriptions on page 888](#).

Table 105: EOS Non-Enterprise-Specific Traps

Trap Object ID (OID)	Trap Name	Trap Profiles	MIB Group	MIB File
1.3.6.1.2.1.15.7.0.2	BGP_PEER_DOWN	ALL	MIB-2	RFC1657.mib
1.3.6.1.2.1.15.7.0.1	BGP_PEER_UP	ALL	MIB-2	RFC1657.mib
1.3.6.1.2.1.17.0.1	newRoot	All, Dot1dBridge	MIB-2	RFC1493.mib
1.3.6.1.2.1.17.0.2	topologyChange	ALL, Dot1dBridge	MIB-2	RFC1493.mib
1.3.6.1.2.1.10.32.1	frDLCIStatusChange	ALL, FrameRelay	MIB-2	RFC1315.mib
1.3.6.1.2.1.10.20.2.0.1	isdMibCallInformation	ALL, ISDN	MIB-2	RFC2127.mib
1.3.6.1.2.1.16.0.2	fallingAlarm	ALL, RMon	MIB-2	RFC1271.mib
1.3.6.1.2.1.16.0.1	risingAlarm	ALL, RMon	MIB-2	RFC1271.mib

E.4.2

EOS Enterprise-Specific Traps

The following lists the EOS enterprise-specific traps supported on EOS devices. For further details about these traps, see [Trap Descriptions on page 888](#).



NOTE: The Trap Profiles listed in the third column of [Table 106: EOS Enterprise-Specific Traps on page 873](#) refer to the predefined trap profiles supported by EOS. For details about how to configure trap profiles and map them to network management stations, see [Traps and Trap Profiles on page 664](#). For information about how to add and delete custom trap profiles, see the description of the `TrapProfile` parameter in “SNMP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Table 106: EOS Enterprise-Specific Traps

Trap Object ID (OID)	Trap Name	Trap Profiles	MIB Group	MIB File
1.3.6.1.4.1.43.2.2 9.0.101	audlCfgChange	ALL, AUDitLog	aMotorolaCa	audlMIB.mib
1.3.6.1.4.1.43.2.2 9.0.102	audlLoginFail	ALL, AUDitLog	aMotorolaCa	audlMIB.mib
1.3.6.1.4.1.43.2.2 9.0.103	audlFileAccess	ALL, AUDitLog	aMotorolaCa	audlMIB.mib
1.3.6.1.4.1.43.2.2 9.0.104	audlSyslogMessage	ALL, AUDitLog	aMotorolaCa	audlMIB.mib
sysObjectID ¹	authenticationFailure	ALL, GEnr	aMotorolaCa	rfc1157 ²
sysObjectID ¹	coldStart	ALL, GEnr	aMotorolaCa	rfc1157 ²
sysObjectID ¹	egpNeighborLoss	ALL, GEnr	aMotorolaCa	rfc1157 ²
sysObjectID ¹	linkDown	ALL, GEnr	aMotorolaCa	rfc1157 ²
sysObjectID ¹	linkUp	ALL, GEnr	aMotorolaCa	rfc1157 ²
1.3.6.1.4.1.43.2.1 6.0.1	aMotorolaCaPathDown	ALL, PATH	aMotorolaCa	ppMIB.mib
1.3.6.1.4.1.43.2.1 6.0.2	aMotorolaCaPathUp	ALL, PATH	aMotorolaCa	ppMIB.mib
1.3.6.1.4.1.43.2.4 7.0.7	aMotorolaCafrAvcStatusChange	ALL, FrameRelay	aMotorolaCa	frMIB.mib
1.3.6.1.4.1.43.2.4 5.0.2	pimExcessiveRegistersDR	ALL, Multicast	aMotorolaCa	pimMIB.mib
1.3.6.1.4.1.43.2.4 5.0.1	pimExcessiveRegistersRP	ALL, Multicast	aMotorolaCa	pimMIB.mib
1.3.6.1.4.1.43.2.4 5.0.4	pimExcessiveRouteEntries	ALL, Multicast	aMotorolaCa	pimMIB.mib
1.3.6.1.4.1.43.2.4 5.0.5	pimNoReceivers	ALL, Multicast	aMotorolaCa	pimMIB.mib
1.3.6.1.4.1.43.2.4 5.0.6	pimRecdNullBSM	ALL, Multicast	aMotorolaCa	pimMIB.mib
1.3.6.1.4.1.43.2.4 5.0.7	pimBsrReElection	ALL, Multicast	aMotorolaCa	pimMIB.mib
1.3.6.1.4.1.43.2.1 7.0.1	aMotorolaCaPortDown	ALL, PORT	aMotorolaCa	ppMIB.mib
1.3.6.1.4.1.43.2.1 7.0.2	aMotorolaCaPortUp	ALL, PORT	aMotorolaCa	ppMIB.mib
1.3.6.1.4.1.43.2.4 0.0.4	aMotorolaCaRasUserConnectionAttemptFail	ALL, REMote	aMotorolaCa	rasMIB.mib

Trap Object ID (OID)	Trap Name	Trap Profiles	MIB Group	MIB File
1.3.6.1.4.1.43.2.4 0.0.1	aMotorolaCaRasU serLogin	ALL, REMote	aMotorolaCa	rasMIB.mib
1.3.6.1.4.1.43.2.4 0.0.3	aMotorolaCaRasU serLoginFail	ALL, REMote	aMotorolaCa	rasMIB.mib
1.3.6.1.4.1.43.2.4 0.0.2	aMotorolaCaRasU serLogout	ALL, REMote	aMotorolaCa	rasMIB.mib
1.3.6.1.4.1.43.2.4 2.0.1	remPollLogFileT rap	ALL	aMotorolaCa	rempextMIB.mib
1.3.6.1.4.1.43.0.7 5	remPollFailureT rap	ALL	aMotorolaCa	_moca_0019.mib
1.3.6.1.4.1.43.0.7 4	remPollSuccessT rap	ALL	aMotorolaCa	_moca_0019.mib
1.3.6.1.4.1.43.2.4 6.0.1	vrrpMasterDown	ALL, Multicast	aMotorolaCa	vrrpMIB.mib
1.3.6.1.4.1.43.2.4 6.0.2	vrrpMasterUp	ALL, Multicast	aMotorolaCa	vrrpMIB.mib
1.3.6.1.4.1.43.2.4 8.1.1.0.1	mnrGtpContextEx ceeded	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.1.0.2	mnrGtpCreateFai l	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.1.0.3	mnrGtpFlowMisma tch	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.1.0.5	mnrGtpPdgUnreac hable	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.1.0.6	mnrGtpRedundant StatusActive	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.1.0.7	mnrGtpRedundant StatusStandby	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.1.0.8	mnrGtpRedundant StatusSolitary	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.1.0.9	mnrGtpRedundant StatusNoStandby	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.2.0.1	mnrGgsnDHCPPrim aryFail	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.2.0.2	mnrGgsnDHCPSeCo ndaryFail	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.2.0.3	mnrGgsnDHCPPool Empty	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.2.0.4	mnrGgsnDHCPAssi gned	ALL	aMotorolaCa	mnr_ggsn.mib

Trap Object ID (OID)	Trap Name	Trap Profiles	MIB Group	MIB File
1.3.6.1.4.1.43.2.4 8.1.3.0.1	mnrGgsnRADIUSSw itchOver	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.3.0.2	mnrGgsnRADIUSAs signed	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.4.0.1	mnrGgsnBRLinkUp	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.4.0.2	mnrGgsnBRLinkDo wn	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.5.0.1	mnrGgsnHADSync0 bjUp	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.5.0.2	mnrGgsnHADSync0 bjDown	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.5.0.3	mnrGgsnHADSync0 bjInfo	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.5.0.4	mnrGgsnHADSync0 bjSingleUpdFail ed	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.5.0.6	mnrGgsnHADSync0 bjEarlyFailOver	ALL	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 9.1.0.1	mnrDS1E1LineCha ngeTrap	ALL	aMotorolaCa	t1e1.mib
1.3.6.1.4.1.43.2.5 0.0.1	mnrEnvMonFaultC ritical	ALL, ENVMON	aMotorolaCa	envMon.mib
1.3.6.1.4.1.43.2.5 0.0.2	mnrEnvMonFaultW arning	ALL, ENVMON	aMotorolaCa	envMon.mib
1.3.6.1.4.1.43.2.5 0.0.3	mnrEnvMonFaultN ormal	ALL, ENVMON	aMotorolaCa	envMon.mib
1.3.6.1.4.1.43.2.5 1.1.2	mnrDeviceFaultC ritical	ALL	aMotorolaCa	mnrSysTools.mib
1.3.6.1.4.1.43.2.5 2.0.1	mnrOverloadEnte r	ALL, SYS	aMotorolaCa	mnrSysTools.mib
1.3.6.1.4.1.43.2.5 2.0.2	mnrOverloadLeav e	ALL, SYS	aMotorolaCa	mnrSysTools.mib
1.3.6.1.4.1.43.2.3 4.0.1	mnrVpnTunnelUp	ALL	aMotorolaCa	vpnMIB.mib
1.3.6.1.4.1.43.2.5 3.1.0.1	mnrCWRStatusCha ngeTrap	ALL, CWR	aMotorolaCa	cwr.mib
1.3.6.1.4.1.43.2.3 4.0.7	cryptoMibSACrea ted	ALL, VPN	aMotorolaCa	cryptoMIB.mib
1.3.6.1.4.1.43.2.3 4.0.8	cryptoMibSADele ted	ALL, VPN	aMotorolaCa	cryptoMIB.mib

Trap Object ID (OID)	Trap Name	Trap Profiles	MIB Group	MIB File
1.3.6.1.4.1.43.2.3 4.0.9	cryptoMibSAException	ALL, VPN	aMotorolaCa	cryptoMIB.mib
1.3.6.1.4.1.43.2.3 4.0.10	mnrcryptoSelfTestFailed	ALL, VPN	aMotorolaCa	cryptoMIB.mib
1.3.6.1.4.1.43.2.5 4.1.1.0.1	mnrospfMD5KeyMismatch	ALL, OSPF	aMotorolaCa	mnr_ospf.mib
1.3.6.1.4.1.43.2.2 7.7.0.1	mnrmultiCastGroupsExceeded	ALL, SYS	aMotorolaCa	mipMIB.mib
1.3.6.1.4.1.43.2.4 8.1.4.0.1	mnrggsnCenLinkUp	ALL, SYS	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.4 8.1.4.0.2	mnrggsnCenLinkDown	ALL, SYS	aMotorolaCa	mnr_ggsn.mib
1.3.6.1.4.1.43.2.3 4.0.11	mnraproTOCOLAuthFailed	ALL, VPN	aMotorolaCa	cryptoMIB.mib
1.3.6.1.4.1.43.2.5 5.1.1	mnrwanIfQueueDepthCaution ³	ALL, QOS	aMotorolaCa	mnr_WAN-QoS.mib
1.3.6.1.4.1.43.2.5 5.1.2	mnrwanIfQueueDepthExceeded	ALL, QOS	aMotorolaCa	mnr_WAN-QoS.mib
1.3.6.1.4.1.43.2.5 5.1.3	mnrwanIfTotalQueueDepthExceeded ³	ALL, QOS	aMotorolaCa	mnr_WAN-QoS.mib

¹ sysObjectID = 1.3.6.1.4.1.43.1.4.[37|38] where:

1.3.6.1.4.1.43.1.4.37 = MNR S2500

1.3.6.1.4.1.43.1.4.38 = MNR S6000 router

1.3.6.1.4.1.43.1.4.39 = GGM 8000 gateway

1.3.6.1.4.1.43.1.4.41 = MCG 8000

² The form and content of the generic traps are based on RFC 1157 itself, not a MIB file.

³ In EOS version 16.0 and forward, the mnrcryptoSelfTestFailed, mnrospfMD5KeyMismatch, mnrmultiCastGroupsExceeded, and mnrggsnCenLinkUp traps replace the frQueueDepthCaution (OID 1.3.6.1.4.1.43.2.47.0.2), frQueueDepthExceeded (OID 1.3.6.1.4.1.43.2.47.0.3), and frTotalQueueDepthExceeded (OID 1.3.6.1.4.1.43.2.47.0.4) traps supported in earlier releases. The new traps support QOS for both the IP and the Frame Relay services.

E.4.3

EOS Derived MIB-2 Enterprise-Specific Traps

The following lists the EOS derived MIB-2 enterprise-specific traps supported on EOS devices. For further details about these traps, see [Trap Descriptions on page 888](#).



NOTE: The Trap Profile(s) listed in the third column of [Table 107: EOS Derived MIB-2 Enterprise-Specific Traps on page 877](#) refer to the predefined trap profiles supported by EOS. For details about how to configure trap profiles and map them to network management stations, see [Traps and Trap Profiles on page 664](#). For information about how to add and delete custom trap profiles, see the description of the TrapProfile parameter in “SNMP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Table 107: EOS Derived MIB-2 Enterprise-Specific Traps

Trap Object ID (OID)	Trap Name	Trap Profile(s)	MIB Group	MIB File
1.3.6.1.4.1.43.2.1 00.15.7.0.2	MNR_BGP_PEER_DOWN	ALL	MIB-2	mnrtraps.mib
1.3.6.1.4.1.43.2.1 00.15.7.0.1	MNR_BGP_PEER_UP	ALL	MIB-2	mnrtraps.mib
1.3.6.1.4.1.43.2.1 00.17.0.1	MNR_newRoot	All, Dot1dBridge	MIB-2	mnrtraps.mib
1.3.6.1.4.1.43.2.1 00.17.0.2	MNR_topology-Change	ALL, Dot1dBridge	MIB-2	mnrtraps.mib
1.3.6.1.4.1.43.2.4 7.0.1	MNR_frDLCIStatusChange	ALL, FrameRelay	MIB-2	frMIB.mib
1.3.6.1.4.1.43.2.1 00.10.20.2.0.1	MNR_isdnMibCallInformation	ALL, ISDN	MIB-2	mnrtraps.mib
1.3.6.1.4.1.43.2.1 00.16.0.2	MNR_fallingAlarm	ALL, RMon	MIB-2	mnrtraps.mib
1.3.6.1.4.1.43.2.1 00.16.0.1	MNR_risingAlarm	ALL, RMon	MIB-2	mnrtraps.mib

Standard and Private Generic Traps

The following generic traps are implemented in EOS as both standard and private generic traps:

- authenticationFailure
- coldStart
- egpNeighborLoss
- linkDown
- linkUp

The traps are standard generic traps because they are well-known traps, defined by RFC-1157. The traps are private generic traps because they are assigned to EOS, private, enterprise-specific OIDs (sysObject ID). In EOS version 11.7.3 and later, these generic traps are handled as follows:

- Both the standard and the private generic traps are controlled via the SNMP service `Mib2TrapControl` parameter.
- The standard generic traps **do not** contain the new event trap timestamp and SysName fields.
- The private generic traps **do** contain the new event trap timestamp and SysName fields.
- Both the standard generic trap OID and the private generic trap OID are assigned to the sysObject ID.
- The event format for the generic traps accommodates both the standard and the private generic traps, as it is set up to unconditionally describe and display the event trap timestamp and SysName fields, whether or not these fields are actually contained in the generic trap received.

E.4.4

Trap Delay Queueing and Event-Based Trap Suppression

EOS implements a trap delay queueing mechanism to reduce bursts and duplicates of certain traps, and configure the traps to be dequeued and processed at a regular interval. In addition, the affected traps are

suppressed during cooperative WAN routing (CWR) switchover and device bootup, thereby minimizing the performance issues and network congestion that may be caused by the large number of traps triggered by these events. The trap delay queueing and event-based trap suppression feature allows you to identify more important traps more easily in the Unified Event Manager (UEM).

E.4.5

SysName and Trap Event Timestamp in EOS Enterprise-Specific Traps

All EOS *enterprise-specific* traps include the following fields to aid in correlating and diagnosing network errors:

- The `SysName` field indicates the system name (`SysName`) of the device that generated the trap. If the `SysName` parameter was not configured or is blank, then the value for `SysName` is a blank character.
- The `trap event timestamp` field indicates the time at which the event that generated the trap occurred - in ASCII and integer format.



NOTE: To derive the intended event correlation benefit from the trap event timestamp, you should initialize the EOS system time via synchronization with an NTP server and set the `SYStem service TimeZone` parameter to the appropriate local time zone.

You can specify whether the ASCII-formatted timestamp is displayed in UTC time (Coordinated Universal Time, formerly Greenwich Mean Time) or in local time by using the `SNMP service TrapTimeScale` parameter.

The format of the trap timestamp is as follows:

- Integer whole seconds since the Epoch (00:00:00 UTC, January 1, 1970)
- Integer fractional seconds within the second. This field represents the partial seconds to the accuracy of the clock of the device generating the trap event.
- ASCII local or UTC time with the following format:

YYYYMMDD_HHMMSS.mmm aaaaaaaa

where:

YYYY = Year

MM = Month

DD = Day

HH = Hours since midnight

MM = Minutes after the hour

SS = Seconds after the minute

mmm = Milliseconds after the second

aaaaaaaa = Time zone label

If the `TrapTimeScale` parameter is set to `UTC` (the default), the time zone label is `UTC`. If the `TrapTimeScale` parameter is set to `LOCAL`, the time zone label is either the time zone value specified with the `SYStem service DSTime` parameter or (if no time zone value was specified with the `DSTime` parameter) a null string.

E.4.6

Support for SNMPv2 64-bit Counters

EOS supports SNMPv2 64-bit counters. This feature provides the following functions:

- SNMPv1 and SNMPv2 access to data in the Interface MIB table, ifXTable.
-  **NOTE:** EOS versions prior to 11.8 supported the IF-MIB as defined by RFC1573. EOS version 11.8 and higher supports the IF-MIB as defined by RFC2863.
- Support for the SNMPv2 data type Counter64, for 64-bit counters in the Interface MIB table, ifXTable.
- SNMPv2 GET operation for the Interface MIB table, ifXTable.
- SNMPv2 GETNEXT operation for the Interface MIB table, ifXTable.
- Physical layer MIB variables and a line status change trap for all T1 (DS1) and E1 interfaces.
The DS1 (T1)/E1 MIBs are 64-bit counters that provide a view of physical layer performance and help in interface debugging if an interface is dropping packets or is down.

 **NOTE:** For details about the DS1/E1 and ifXTable performance statistics, see [Performance Statistics on page 687](#).

The DS1/E1 trap is required to understand the physical layer faults of T1 (DS1) and E1 interfaces.
EOS does not support the SNMPv2 GETBULK and SET operations at this time.

E.4.7

SNMPv3 and SNMPv1 Trap Generation

In SNMPv3, traps are received through the managers configured with the `ADD -SNMP V3MANager` command. You are not required to delete these managers before changing to SNMPv1/v2c operation. However, when you switch from SNMPv3 to SNMPv1/v2c operation, you must add v1 managers with the `ADD -SNMP MANager` command in order to receive SNMPv1 traps.

 **NOTE:** EOS does not support SNMPv2c traps.

E.4.8

S6000 and GGM 8000 Supply System / Hardware Fault Monitoring Traps

MNR S6000 routers, MCG 8000 and GGM 8000 gateways support environmental monitoring traps that monitor the operational status of the power/cooling supply system.

MNR S6000 routers, MCG 8000 and GGM 8000 gateways support a supply system consisting of a power subsystem component and an cooling component. The device can monitor the operational health of its power and cooling supplies and report significant state changes and/or failure states to the network administrator using SNMP traps.

- EOS sends an `mnrEnvMonCritical` trap when the S6000, the MCG 8000 or GGM 8000 supply system transitioned to a critical state, indicating abnormal operating conditions that potentially pose a severe problem, with system failure likely. Immediate service is recommended. As long as the critical state persists, EOS sends a trap once every 24 hours.
- EOS sends a trap when the S6000, the MCG 8000 or GGM 8000 supply system transitioned to a warning state, indicating abnormal operating conditions that compromised system fault resilience. These problems can be addressed by scheduling a service maintenance visit. As long as the warning state persists, EOS sends a trap once every 24 hours.

- EOS sends a trap when the S6000, the MCG 8000 or GGM 8000 supply system transitioned to a normal state from a warning or critical state, indicating that the supply system is within normal operating conditions.



NOTE: MNR S6000 base units with model numbers CLN1780A and higher and GGM 8000 gateways and MCG 8000 support a single power/cooling supply system. MNR S6000 base units with model numbers lower than CLN1780A support a redundant power/cooling supply system.

E.4.9

CCGW Fault Traps

The CCGW fault traps listed in the following table are supported by MNR S2500s and GGM 8000 gateways configured as CCGWs and MCG 8000 platform. The MNR service ID MIB (`mnrSysServiceID`) identifies whether or not a device is a CCGW. For details, see [MNR Service ID MIB on page 854](#).

The CCGW agent reports these traps to the configured CCGW Network Manager.



NOTE: For the CCGW fault traps only, the device system IP address (SysIP) is sent as the trap source IP address.

Table 108: CCGW Fault Traps

Trap Name	Description and Variables ¹	Enterprise ²
<code>ccgwStateChangeTrap</code>	This trap is sent when the CCGW box changes states. Variables: <code>ccgwFltState</code> - The current operational state of the CCGW. <code>ccgwFltCause</code> - The cause for the most recent CCGW state change. If the cause of the last CCGW state change is unknown, the value for this variable defaults to 0: 1 - CCGW Configuration complete 2 - User requested CCGW Disable 10 - User disabled all channels 11 - User enabled all channels 19 - Between 100% and 50% (inclusive) channels available 20 - Less than 50% channels available <code>ccgwFltText</code> - Supplemental text that provides additional information about the CCGW state change. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. <code>nmaTimeInSec</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). <code>nmaTimeInMilliSec</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the	<code>ccgwCCGWReg</code>

Trap Name	Description and Variables ¹	Enterprise ²
	accuracy of the clock of the device generating the trap events.	
ccgwTransientTrap	This trap is sent when a CCGW transient alarm occurs. Variables: ccgwFltAlarm - The alarm code for transient alarms: (1) CCGW NTP is unconfigured. (2) CCGW NTP server is unavailable. (6) One DSP Down (7) Both DSPs Down (8) One DSP Up (9) Both DSPs Up This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. ccgwFltText - Supplemental text that provides additional information about the CCGW transient fault. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. nmaTimeInSec - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	ccgwCCGWReg
ccgwZCCPStateChange-Trap	This trap is sent when a control path between the CCGW and the zone controller changes states. When CCGW is disabled on a dynamic system resilience (DSR) system, EOS generates two critical state change traps (one trap for each CCGW-to-zone controller connection). Because the traps do not include the IP addresses of the zone controllers, the traps appear to be redundant. However this is not the case: one trap indicates a state change on the current (active) zone controller link, while the other indicates a state change on the backup zone controller link. Variables: ccgwZCCPFltState - The current operational state of the CCGW to zone controller control	ccgwZCCPReg

Trap Name	Description and Variables ¹	Enterprise ²
	path. See Table 110: CCGW-to-Zone Controller State Definitions on page 886 for control path state definitions. <code>ccgwZCCPF1tCause</code> - The cause for the most recent control path state change. If the cause of the last control path state change is unknown, the value for this variable defaults to 0: 1 - CCGW to zone controller connection established. 2 - CCGW to zone controller connection lost. <code>ccgwZCCPF1tTxt</code> - Supplemental text that provides additional information about the control path fault. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. <code>nmaTimeInSec</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). <code>nmaTimeInMilliSec</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>ccgwZCCPRedStateChangeTrap</code>	This trap is sent when the redundant control path between the CCGW and the zone controller changes states. Variables: <code>ccgwZCCPF1tRedObjState</code> - The operational state of the CCGW to zone controller redundant control path. <code>ccgwZCCPF1tRedObjCause</code> - The cause for the most recent redundant control path state change. If the cause of the last redundant control path state change is unknown, the value for this variable defaults to 0: 1 - CCGW- to-zone controller redundant control path becomes active. 2 - CCGW-to-zone controller redundant control path becomes inactive. <code>ccgwZCCPF1tText</code> - Supplemental text that provides additional information about the control path fault. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. <code>nmaTimeInSec</code> - The time at which the event that generated the trap occurred, in integer format,	<code>ccgwZCCPReg</code>

Trap Name	Description and Variables ¹	Enterprise ²
	represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
ccgwChannelStateChangeTrap	State change trap for the channel managed object. This trap is sent when the channel changes states. Variables: ccgwChannelFltState - The current operational state of the channel. ccgwChannelFltCause - The cause for the most recent channel state change. If the cause of the last channel state change is unknown, the value for this variable defaults to 0: 2 - User disabled channel 3 - Channel recovered 4 - Channel failed 5 - User unconfigured channel 6 - User configured channel 7 - Analog link monitor tone (ALMT) undetect from the DSP (mixed mode channels only) 8 - Failure of the brstat protocol between the CCGW and the base repeater / Call Manager (IP conventional channels only) 9 - Analog Link Failed 10 - Both Links Failed ccgwChannelFltText - Supplemental text that provides additional information about the channel state change. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. nmaTimeInSec - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	ccgwChannelReg

Trap Name	Description and Variables ¹	Enterprise ²
ccgwChannelTransient-Trap	This trap is sent when a conventional channel transient alarm occurs. Variables: ccgwChannelFltAlarmCode - The alarm code for transient alarms: (1) voiceBufferDiscarded - Audio discarded due to voice/data contention. (2) voiceCallDenied - New console call rejected because a current call is in progress. (3) configIgnored - Configuration from LDAP server ignored because channel type or interface type in CCGWDB file does not match LDAP server configuration. This can occur if the CCGWDB file is not removed prior to CCGW upgrade or migration. (4) invalidToneTable (5) unexpectedTalkGroup This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. ccgwChannelFltText - Supplemental text that provides additional information about the conventional channel transient fault. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. nmaTimeInSec - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	ccgwChannelReg
ccgwRNGLinkStateChangeTrap	This trap is sent when the link between the CCGW and the RNG changes states. Variables: ccgwRNGLinkFltState - The current operational state of the RNG link. ccgwRNGLinkFltCause - The cause for the most recent RNG link state change. If the cause of the last RNG link state change is unknown, the value for this variable defaults to 0:	ccgwRNGLinkReg

Trap Name	Description and Variables ¹	Enterprise ²
	1 - Connection established 2 - Connection lost 3 - Link deleted cogwRNGLinkFltText - Supplemental text that provides additional information about the RNG link state change. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. nmaTimeInSec - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
cogwRNGBCastLinkState-ChangeTrap	This trap is sent when the broadcast link between the CCGW and the RNG changes states. Variables: cogwRNGBCastLinkFltState - The current operational state of the broadcast RNG link. cogwRNGBCastLinkFltCause - The cause for the most recent broadcast RNG link state change. If the cause of the last broadcast RNG link state change is unknown, the value for this variable defaults to 0: 1 - Connection established 2 - Connection lost 3 - Link deleted 4 - Link disabled cogwRNGBCastLinkFltText - Supplemental text that provides additional information about the broadcast RNG link state change. This variable is meaningful only in a trap PDU. If it is requested via Get or GetNext, a null string is returned. nmaTimeInSec - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). nmaTimeInMilliSec - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	cogwRNGBCastLink-Reg

Trap Name	Description and Variables ¹	Enterprise ²
¹ The siteID, pathID, and channelID variables are defined as a table index and are attached to the state and cause OID.		
² The Enterprise OID is the managed object registration ID (ccgwReg, ccgwZCCPReg, ccgwChannelReg, ccgwRNGLinkReg, or CCGWRNGBcastLinkReg).		

Table 109: CCGW Box State Definitions

State Name	State Number	Severity	DAR	Description
Unconfigured	1	Minor	Unavailable	The CCGW is not configured.
Disabled	2	Warning	Unavailable	The CCGW configuration is complete, and the CCGW is disabled.
Up	4	Normal	Available	Netcom was established and one or two control paths are up.
Major Malfunction	6	Major	Available	(applies to analog conventional only) One of the two DSPs failed.
Critical Malfunction	7	Critical	Unavailable	All channels failed or the E&M or V.24 module failed.
Minor Malfunction	8	Minor	Available	More or equal to 50% channels available (but less than 100% channels available)

Table 110: CCGW-to-Zone Controller State Definitions

State Name	State Number	Severity	DAR	Description
LinkDown	1	Critical	Unavailable	The Netcom connection was lost or was not established
LinkUp	2	Normal	Available	The Netcom connection is up and established

Table 111: CCGW-to-Zone Controller Redundant Control Path State Definitions

State Name	State Number	Severity	Description
Non-Redundant	1	Normal	No redundancy.
Standby	2	Normal	The redundant control path is the standby control path.
Active	3	Normal	The redundant control path is the active control path.

Table 112: CCGW Conventional Channel State Definitions

State Name	State Number	Severity	DAR	Description
Up	1	Normal	Available	The conventional channel is up.
Down	2	Warning	Unavailable	The conventional channel is either disabled or it failed.
Unconfigured	3	Normal	Unavailable	The conventional channel is unconfigured.
Partial	4	Warning	Partially available	Degraded state: the V.24 port is operational, but the corresponding E&M port is either missing or not operational. This state definition applies only to mixed mode channels.

Table 113: CCGW RNG Link State Definitions

State Name	State Number	Severity	DAR	Description
Down	1	Warning	Unavailable	The link between the RNG and the CCGW is lost or was not established.
Up	2	Normal	Available	The link between the RNG and the CCGW is established.
Unconfigured	3	Normal	Unavailable	The link between the RNG and the CCGW is not configured or was deleted. This trap does not get sent to the network manager. The common agent converts it to an object deletion trap.

Table 114: CCGW RNG Broadcast Link State Definitions

State Name	State Number	Severity	DAR	Description
Down	1	Warning	Unavailable	The broadcast link between the RNG and the CCGW is lost or was not established.
Up	2	Normal	Available	The broadcast link between the RNG and the CCGW is established.
Unconfigured	3	Normal	Unavailable	The broadcast link between the RNG and the CCGW is not configured or was deleted. This trap does not get sent to the network manager. The common agent converts it to an object deletion trap.

State Name	State Number	Severity	DAR	Description
Disabled	4	Normal	Unavailable	The broadcast link between the RNG and the CCGW was administratively configured to be disabled.

E.4.10

Trap Descriptions

The following table lists the HP® OpenView® trap definition format, description, and severity for each of the traps generated by EOS and supported on EOS devices.

The variables \$1, \$2, \$3, and so forth (in the Trap Format column of the following table) map the SNMP trap variables to HP OpenView event descriptors. Traps are considered as events in HP OpenView, and trap variables are called Event Attributes or variables. These variables can be of any value. For example, a variable can be a port number or an IP address, or a trap time, and so forth. Each event has any number of variables or possibly none at all. The value of the variables depends on the particular trap.

The HP OpenView trap definition format applies only if you are using HP OpenView as your management application.

For more information on the trap variable for a particular EOS device trap, select the trap (alarm) in the **Alarms Browser** window and from the menu bar select **Actions** → **Alarm Details**. An **Alarm Details** dialog box opens, and if trap variable definitions exist, the **Event/Alarm description** field in the **Alarm Details** dialog box defines the trap variables.

Table 115: Details about EOS Traps Supported on EOS Devices

Trap Name	HP OpenView Trap Definition Format	Description	Severity
BGP_PEER_DOWN	BGP_PEER_DOWN BGP Peer status moves from established state to down state: BGP Peer Remote Address is \$1; BGP Last Error code and Subcode is \$2; BGP Peer Connection State is \$3	A BGP FSM moved from a higher numbered state to a lower numbered state. Variables: (1) bgpPeerRemoteAddr - The remote IP address of this entry's BGP peer. (2) bgpPeerLastError - The last error code and subcode seen by this peer on this connection. If no error occurred, this field is zero; otherwise, the first byte of this two-byte OCTET STRING contains the error code, and the second byte contains the subcode. (3) bgpPeerState - The BGP peer connection state.	Critical
BGP_PEER_UP	BGP_PEER_UP BGP Peer Status enters Up state: BGP Peer Remote Address is \$1; BGP Last Error code	A BGP FSM entered the ESTABLISHED state. Variables:	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	and Subcode is \$2; BGP Peer Connection State is \$3	(1) bgpPeerRemoteAddr - The remote IP address of this entry's BGP peer. (2) bgpPeerLastError - The last error code and subcode seen by this peer on this connection. If no error occurred, this field is zero; otherwise, the first byte of this two-byte OCTET STRING contains the error code, and the second byte contains the subcode. (3) bgpPeerState - The BGP peer connection state.	
newRoot	DBRIDGE_New_Root Bridge Spanning Tree New Root	The sending agent becomes the new root of the Spanning Tree. The trap is sent by a bridge soon after its election as the new root, for example upon expiration of the Topology Change Timer immediately subsequent to its election.	Warning
topologyChange	DBRIDGE_Tp_Change Bridge Topology Changed	A configured port on a bridge transitioned from the Learning state to the Forwarding state or from the Forwarding state to the Blocking state. This trap is not sent if a newRoot trap is sent for the same transition.	Warning
frDLCIStatusChange	FR_DLCI_Change Virtual Circuit Change: interface index = \$1; DLCI = \$2 new circuit state = \$3	The specified virtual circuit changed state. It was either created or invalidated, or it toggled between the active and inactive states.  NOTE: This trap is sent only if the Frame Relay MIB variable frTrapState was enabled. Variables: (1) frCircuitIfIndex - The ifIndex value of the ifEntry onto which this virtual circuit is layered. (2) frCircuitDlci - The Data Link Connection Identifier for this virtual circuit. (3) frCircuitState - Indicates whether the particular virtual circuit is operational.	Minor
isdnMibCallInformation	ISDN_Call_Info	An ISDN line received a call. Variables:	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	ISDN Call Information: ifIndex \$1; isdnBearerOperStatus \$2; isdnBearerPeerAddress \$3; isdnBearerPeerSubAddress \$4; isdnBearerCallSetupTime \$5; isdnBearerInfoType \$6; isdnBearerCallOrigin \$7	(1) ifIndex - A unique value for each interface, ranging between 1 and the value of ifNumber. The value for each interface must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization. (2) isdnBearerOperStatus - The current call control state for this port: idle (1) - The B channel is idle; no call or call attempt is occurring. connecting (2) - A connection attempt (outgoing call) is made on this interface. connected (3) - An incoming call is in the process of validation. active (4) - A call is active on this interface. (3) isdnBearerPeerAddress - The ISDN address to which the current or last call was connected. (4) isdnBearerPeerSubAddress - The ISDN subaddress to which the current or last call was connected. (5) isdnBearerCallSetupTime - The value of sysUpTime when the ISDN setup message for the current or last call was sent or received. If, since system startup, there was no call on this interface, this object has a value of zero. (6) isdnBearerInfoType - The Information Transfer Capability for the current or last call. (7) isdnBearerCallOrigin - The call origin for the current or last call. If, since system startup, there was no call on this interface, this object has a value of unknown (1).	
fallingAlarm	RMON_Falling_Alarm RMON Falling Alarm: \$2 fell below threshold \$5; value = \$4. (Sample type = \$3; alarm index = \$1)	An alarm entry crossed its falling threshold and generated an event that is configured for sending SNMP traps. (1) alarmIndex - An index that uniquely identifies an entry in the alarm table. Each such entry defines	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		a diagnostic sample at a particular interval for an object on the device. (2) <code>alarmVariable</code> - The object identifier of the particular variable to be sampled. Only variables that resolve to an ANS.1 primitive type of INTEGER (INTEGER, Counter, Gauge, or TimeTicks) may be sampled. (3) <code>alarmSampleType</code> - The method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is <code>absoluteValue</code> (1), the value of the selected variable is compared directly with the thresholds at the end of the sampling interval. If the value of this object is <code>deltaValue</code> (2), the value of the selected variable at the last sample is subtracted from the current value, and the difference compared with the thresholds. (4) <code>alarmValue</code> - The value of the statistic during the last sampling period. The value during the current sampling period is not made available until the period is completed. (5) <code>alarmFallingThreshold</code> - A threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single event is generated. After a falling event is generated, another such event is not be generated until the sampled value rises above this threshold and reaches the <code>alarmRisingThreshold</code>.	
risingAlarm	RMON_Rise_Alarm RMON Rising Alarm: \$2 exceeded threshold \$5; value = \$4. (Sample type = \$3; alarm index = \$1)	An alarm entry crossed its rising threshold and generated an event that is configured for sending SNMP traps. Variables: (1) <code>alarmIndex</code> - An index that uniquely identifies an entry in the alarm table. Each such entry defines a diagnostic sample at a particular interval for an object on the device.	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(2) <code>alarmVariable</code> - The object identifier of the particular variable to be sampled. Only variables that resolve to an ANS.1 primitive type of INTEGER (INTEGER, Counter, Gauge, or TimeTicks) may be sampled. (3) <code>alarmSampleType</code> - The method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is <code>absoluteValue</code> (1), the value of the selected variable is compared directly with the thresholds at the end of the sampling interval. If the value of this object is <code>deltaValue</code> (2), the value of the selected variable at the last sample is subtracted from the current value, and the difference compared with the thresholds. (4) <code>alarmValue</code> - The value of the statistic during the last sampling period. The value during the current sampling period is not made available until the period is completed. (5) <code>alarmRisingThreshold</code> - A threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single event is generated. After a rising event is generated, another such event is not generated until the sampled value falls below this threshold and reaches the <code>alarmFallingThreshold</code>.	
<code>audlCfgChange</code>	AUDL_Cfg_Change Audit Log Configuration Change. The last command or set operation was \$1. Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5	The device's configuration was changed as a result of either a configuration command or an SNMP set operation. Variables: (1) <code>aMotorolaCaAudlLastTrapCommand</code> - The last command string or SNMP set object ID (in ASCII) which caused one of the enterprise audit/change traps to be sent.	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
audlLoginFail	AUDL_Login_Fail Audit: User Login Failed. The last operation was \$1. Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5	A user login attempt failed. Variables: (1) aMotorolaCaAudlLastTrapCommand - The last command string or SNMP set object ID (in ASCII) which caused one of the enterprise audit/change traps to be sent. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Minor
audlFileAccess	AUDL_File_Access Audit: File Access. The last operation was \$1.	A file operation that creates, renames, or removes a file occurred. Variables:	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5	(1) aMotorolaCaAudlLastTrapCommand - The last command string or SNMP set object ID (in ASCII) which caused one of the enterprise audit/change traps to be sent. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
audlSyslogMessage	AUDL_Syslog_Message Audit: System Log Info. SysLogSeqNumber = \$1; SysLogId = \$2; SysLogTimestamp = \$3; SysLogSeverity = \$4; SysLogSource = \$5; SysLogText = \$6. Sysname = \$7, TimeStampAscii = \$8, TimeStampSecs = \$9, TimeStampMilliSecs = \$10	The information for a message written to the system log. Variables: (1) aMotorolaCaAudlSyslogSeqNumber - The sequence number of the syslog message. (2) aMotorolaCaAudlSyslogSyslogId - The message identifier of the syslog message. (3) aMotorolaCaAudlSyslogTimestamp - The timestamp that is added to a syslog message when the message is logged at the local device. (4) aMotorolaCaAudlSyslogSeverity - The priority of the syslog message: logEMERG (1) logALERT (2) logCRIT (3) logERR (4) logWARNING (5) logNOTICE (6) logINFO (7)	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		logDEBUG (8) (5) aMotorolaCaAudlSyslogSource - The entity or source of the syslog message. (6) aMotorolaCaAudlSyslogText - The text of the syslog message. (7) Sysname - The system name (Sysname) of the device that generated the trap. (8) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (9) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (10) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
authentication-Failure	AMotorolaCa_Auth-en_Failure Incorrect Community Name (authentication-FailureTrap) \$+1, \$+2, \$+3, \$+4	The sending protocol entity is the addressee of a protocol message that is not properly authenticated. While implementations of the SNMP must be capable of generating this trap, they must also be capable of suppressing the emission of such traps via an implementation-specific mechanism. This generic trap is implemented in EOS as both a standard and a private generic trap. The private generic trap contains the variables listed below; the standard generic trap does not. Variables: (1) Sysname - The system name (Sysname) of the device that generated the trap. (2) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (3) TimeStampSecs - The time at which the event that generated the	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (4) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
coldStart	AMotorola-Ca_Cold_Start Agent Up with Possible Changes (coldStart Trap) \$+1, \$+2, \$+3, \$+4	The sending protocol entity is re-initializing itself such that the agent's configuration or the protocol entity implementation may be altered. This generic trap is implemented in EOS as both a standard and a private generic trap. The private generic trap contains the variables listed below; the standard generic trap does not. Variables: (1) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (2) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (3) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (4) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Major
egpNeighborLoss	AMotorola-Ca_EGP_Down EGP Neighbor Down (egpNeighborLoss Trap) \$+2, \$+3, \$+4, \$+5, Neighbor = \$1	An EGP neighbor for which the sending protocol was an EGP peer is marked down and the peer relationships no longer applies. Variables: (1) <code>egpNeighAddr</code> - The IP address of this entry's EGP neighbor. This generic trap is implemented in EOS as both a standard and a private	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		generic trap. The private generic trap contains the additional variables listed below; the standard generic trap does not. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
linkDown	AMotorola- Ca_Link_Down Agent Interface Down (linkDown Trap) \$+2, \$+3, \$+4, \$+5 Interface = \$1	The sending protocol entity recognizes a failure in one of the communication links represented in the agent's configuration. Variables: (1) ifIndex — A unique value for each interface, ranging between 1 and the value of ifNumber. The value for each interface must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization. This generic trap is implemented in EOS as both a standard and a private generic trap. The private generic trap contains the additional variables listed below; the standard generic trap does not. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format.	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(4) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
linkUp	AMotorolaCa_Link_Up Agent Interface Up (linkUp Trap) \$+2, \$+3, \$+4, \$+5 Interface = \$1	The sending protocol entity recognizes that one of the communication links represented in the agent's configuration is up. Variables: (1) <code>ifIndex</code> - A unique value for each interface, ranging between 1 and the value of <code>ifNumber</code>. The value for each interface must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization. This generic trap is implemented in EOS as both a standard and a private generic trap. The private generic trap contains the additional variables listed below; the standard generic trap does not. (2) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (3) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (4) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		accuracy of the clock of the device generating the trap events.	
aMotorolaCaPath-Down	AMotorolaCa_Path-Down MotorolaCa Mrouter Path (\$3) Down on PathIndex \$1; Path was assigned to Port \$2. Sysname = \$6, TimeStampAscii = \$7, TimeStampSecs = \$8, TimeStampMillisecs = \$9	The physical link transitioned into a down state such that communication on that physical link is not successful. Variables: (1) aMotorolaCaPathIndex - A unique identifier for each path. The value for each path must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization. (2) aMotorolaCaPathPort - The Port table to which this path entry was assigned. Unassigned entries have a value of zero. this object allows adding and deleting paths from ports. In order for any modifications to take effect, the path(s) associated with this command must be enabled (re-enabled if they are already enabled). Paths that have been assigned to the dial pool (pool of dynamic paths) have port values of -1. For a read(get) operations, a value of -1 is returned if the path is not assigned to any port but belongs to the dial pool. A write operation with a value of -1 is not legal and generates an error return. (3) aMotorolaCaPathName - A descriptive name for the path. (4) aMotorolaCaPathAdminStatus - This object enables or disables a path. (5) aMotorolaCaPathOperStatus - The actual state of the path. (6) Sysname - The system name (Sysname) of the device that generated the trap. (7) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (8) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, rep-	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		resented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (9) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>aMotorolaCaPathUp</code>	<code>AMotorolaCa_PathUp MotorolaCa Mrouter Path (\$3) Up on PathIndex \$1; Path was assigned to Port \$2.</code> <code>Sysname = \$6, TimeStampAscii = \$7, TimeStampSecs = \$8, TimeStampMilliSecs = \$9</code>	The physical link transitioned into an up state such that communication on that physical link is successful. Variables: (1) <code>aMotorolaCaPathIndex</code> - A unique identifier for each path. The value for each path must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization. (2) <code>aMotorolaCaPathPort</code> - The index of the entry in the Port table to which this path entry was assigned. Unassigned entries have a value of zero. this object allows adding and deleting paths from ports. In order for any modifications to take effect, the path(s) associated with this command must be enabled (re-enabled if they are already enabled). Paths that were assigned to the dial pool (pool of dynamic paths) have port values of -1. For a read(get) operations, a value of -1 is returned if the path is not assigned to any port but belongs to the dial pool. A write operation with a value of -1 is not legal and generates an error return. (3) <code>aMotorolaCaPathName</code> - A descriptive name for the path. (4) <code>aMotorolaCaPathAdminStatus</code> - This object enables or disables a path. (5) <code>aMotorolaCaPathOperStatus</code> - The actual state of the path. (6) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap.	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(7) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (8) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (9) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
aMotorolaCafrAvc-StatusChange	FR_AVC_STATUS_CHANGE FR Avc Status Change: Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMilliSecs = \$6, New Status of AVC = \$2, AVC Port ID = \$1	A Multilink Frame Relay port changed its status. This trap is sent out either when the AVC comes UP or when it goes DOWN. Variables: (1) <code>aMotorolaCafrMfrAvcIfIndex</code> - The AVC port ID. (2) <code>aMotorolaCafrMfrAvcStatus</code> - The current status of the AVC: 1 (DOWN) or 2 (UP). (3) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (4) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (5) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Minor
pimExcessiveRegistersDR	PIM_EXCESSIVE_REGISTER_DR	The DR did not receive Register Stop for more than the configured number of encapsulated packets. Variables:	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	PIM Excessive Registers AT RP: Group IpAddress = \$1 Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5	(1) GroupIpAddress - The IP multicast group address with which the RP (rendezvous point) is associated. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
pimExcessiveRegistersRP	PIM_EXCESSIVE_REGISTER_RP PIM Excessive Registers AT DR: Group IpAddress = \$1 Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5	The number of registered messages (@RP) exceeded the configured value. Variables: (1) GroupIpAddress - The IP multicast group address with which the DR (designated router) is associated. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
pimExcessiveRouteEntries	PIM_EXCESSIVE_ROUTE_ENTRIES PIM Excessive Route Entries: RtEntries = \$1 Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5	The number of route entries is greater than a configured value. Variables: (1) RtEntries - The number of route entries at or above which a pimExcessiveRouteEntries trap is generated. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Warning
pimNoReceivers	PIM_NO_RECEIVERS PIM No Receivers: Group IpAddress = \$1 Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5	A PIM register was received, but there are no receivers. Variables: (1) GroupIpAddr - The IP multicast group address. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		accuracy of the clock of the device generating the trap events.	
pimRecdNullBSM	PIM_RCVD_NULL_BSM PIM null and valid BSM received: Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5, Group IpAddress = \$1.	A device that is a candidate bootstrap router (CBR) and has some entries in its RP-Set received a null bootstrap message. Variables: (1) GroupIpAddress - The IP address of the bootstrap router sending the null bootstrap message. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Critical
pimBsrReElection	PIM_BSR_RE_ELECTION PIM BSR ReElection: Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5, Group IpAddr = \$1.	A device that is a candidate bootstrap router aged out some other router as the elected bootstrap router (EBSR) and elected itself as the EBSR. Variables: (1) GroupIpAddress - The IP address of the old bootstrap router, which was aged out. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		the Epoch (00:00:00 UTC, January 1, 1970). (5) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
aMotorolaCaPort-Down	AMotorolaCa_Port-Down MotorolaCa Mrouter Port (\$3) Down on PortIndex \$1; Port was assigned to Path \$2; If Virtual Port, DLCI is \$-4 Sysname = \$7, TimeStampAscii = \$8, TimeStampSecs = \$9, TimeStampMilliSecs = \$10	The port transitioned into a down state such that communication on that port is not successful. Variables: (1) <code>aMotorolaCaPortIndex</code> - A unique value for each port. This corresponds to <code>ifIndex</code> for MIB II. The interface table from MIB II contains additional information for this port, including its type, its name, and its description. (2) <code>aMotorolaCaPortPath</code> - The path which is assigned to this port. In some cases, one or more paths are assigned to a single port. In such cases, this object returns the lowest number path. For those situations where more than one port is assigned to a single path, this object returns the <code>PathIndex</code> (or parent port). You must specify this object when adding a new virtual port. In the case where the dial pool (pool of dynamic paths) is assigned to a port, a value of -1 is used. When doing a <code>get(read)</code> of either a virtual or non-virtual dial port, a return value of -1 indicates that the dial pool is assigned to that port. A <code>write(set)</code> operations using a path number value of -1 is required only in the case of creating a virtual dynamic dial port; if you are creating a non-virtual dynamic dial port, you are not required to set this object. (3) <code>aMotorolaCaPortName</code> - The user-assigned name given to this port. This name cannot contain any blanks. (4) <code>aMotorolaCaPortCircuitID</code> - This object applies only to virtual	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		ports; it specifies the unique data link circuit or channel identifier assigned to this virtual port. For example, if this virtual port is attached to a Frame Relay network, this object would specify the DLCI associated with this virtual port. The parent of this virtual port is identified by aMotorolaCaPortPath (which also identifies the physical path with which this port is associated. You must specify this object when adding a new virtual port. (5) ifAdminStatus - The desired state of the interface. The testing(3) state indicates that no operational packets can be passed. (6) ifOperStatus - The current operational status of the interface. The testing(3) state indicates that no operational packets can be passed. (7) Sysname - The system name (Sysname) of the device that generated the trap. (8) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (9) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (10) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
aMotorolaCaPortUp	AMotorolaCa_PortUp MotorolaCa Mrouter Port (\$3) Up on PortIndex \$1; Port was assigned to Path \$2; If Virtual Port, DLCI is \$4 Sysname = \$7, TimeStampAscii = \$8, Time-	The port transitioned into an up state such that communication on that port is successful. (1) aMotorolaCaPortIndex - A unique value for each port. This corresponds to ifIndex for MIB II. The interface table from MIB II contains additional information for this port, including its type, its name, and its description.	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	StampSecs = \$9, TimeStampMilliSecs = \$10	(2) aMotorolaCaPortPath - The path which is assigned to this port. In some cases, one or more paths are assigned to a single port. In such cases, this object returns the lowest number path. For those situations where more than one port is assigned to a single path, this object returns the PathIndex (or parent port). You must specify this object when adding a new virtual port. In the case where the dial pool (pool of dynamic paths) is assigned to a port, a value of -1 is used. When doing a get(read) of either a virtual or non-virtual dial port, a return value of -1 indicates that the dial pool is assigned to that port. A write(set) operations using a path number value of -1 is required only in the case of creating a virtual dynamic dial port; if you are creating a non-virtual dynamic dial port, you are not required to set this object. (3) aMotorolaCaPortName - The user-assigned name given to this port. This name cannot contain any blanks. (4) aMotorolaCaPortCircuitID - This object applies only to virtual ports; it specifies the unique data link circuit or channel identifier assigned to this virtual port. For example, if this virtual port is attached to a Frame Relay network, this object would specify the DLCI associated with this virtual port. The parent of this virtual port is identified by aMotorolaCaPortPath (which also identifies the physical path with which this port is associated. You must specify this object when adding a new virtual port. (5) ifAdminStatus - The desired state of the interface. The testing(3) state indicates that no operational packets can be passed. (6) ifOperStatus - The current operational status of the interface. The	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		testing(3) state indicates that no operational packets can be passed. (7) Sysname - The system name (Sysname) of the device that generated the trap. (8) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (9) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (10) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
aMotorolaCaRasUserConnectionAttemptFail	AMotorolaCa_RasUser_ConnFail RAS User: Generic #G, Specific \$S Sysname = \$7, TimeStampAscii = \$8, TimeStampSecs = \$9, TimeStampMilliSecs = \$10	A Remote Access Server connection attempt failed before PPP authentication. Variables: (1) aMotorolaCaPathIndex - A unique identifier for each path. The value for each path must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization. (2) aMotorolaCaRasPathUserName - The name of the currently-logged-in user. (3) aMotorolaCaRasPathConnectionAttemptFailReason - The reason for the connection attempt failure: lcp-Fail (1) - The failure occurred in the Link Control Protocol of PPP.ncp-Fail (2) - The failure occurred in the Network Control Protocol of PPP. (4) aMotorolaCaPathName - A descriptive name for the path. (5) aMotorolaCaRasPathCurrentTime - The value of system time	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		when the session is referenced, in Universal Time Coordinates. (6) <code>aMotorolaCaPathLabel</code> - The descriptive label assigned to this path by the device. These labels are used by the User Interface and are useful when comparing information retrieved via SNMP with information retrieved via the User Interface. (7) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (8) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (9) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (10) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>aMotorolaCaRasUserLogin</code>	<code>AMotorolaCa_RasUser_Login</code> RAS User: Generic #<code>\$G</code>, Specific #<code>\$S</code> Sysname = <code>\$11</code>, TimeStampAscii = <code>\$12</code>, TimeStampSecs = <code>\$13</code>, TimeStampMilliSecs = <code>\$14</code>	A Remote Access Server user logged in. Variables: (1) <code>aMotorolaCaPortName</code> - The user-assigned name given to this port. No blanks are allowed in this name. (2) <code>ifIndex</code> - The <code>ifIndex</code> value of the interface associated with this communication port. (3) <code>aMotorolaCaRasPortUserName</code> - The name of the currently-logged-in user. (4) <code>aMotorolaCaRasPortUserPrivilege</code> - The privilege level of the currently-logged-in user; 'user' level is assigned to an ordinary user. (5) <code>aMotorolaCaRasPortCurrentTime</code> - The value of system time	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		when the session is referenced, in Universal Time Coordinates. (6) <code>aMotorolaCaPortLabel</code> - The descriptive label assigned to this port by the device. These labels are used by the User Interface and are useful when comparing information retrieved via SNMP with information retrieved via the User Interface. <code>aMotorolaCaRasPortCompType</code> - The type of compression used on the port/path: none (1) indicates that compression is disabled. <code>portHistory</code> (2) indicates that compression is enabled on a port basis and history is preserved across the frames. <code>portPerPacket</code> (3) indicates that compression is enabled on a port basis but history is not preserved across the frames. <code>pathTinyGram</code> (4) indicates that compression is enabled on a path basis and applies to zero-padded frames. (8) <code>aMotorolaCaRasPortAuthHashType</code> - The type of hash algorithm used on the port/path for message authentication: none (1) indicates that no hash algorithm is used; this generally means that no message authentication is used. other (2) indicates that a hash algorithm other than those specifically enumerated is used for message authentication. hmacMD5 (3) indicates that the hash algorithm Message Digest 5 is used for message authentication. hmacSha1 (4) indicates that the hash algorithm Secure Hash Algorithm 1 is used for message authentication. (9) <code>aMotorolaCaRasPortEncrypType</code> - The type of encryption used on the port/path:	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		none (1) indicates that no encryption method is used for encrypting data on this port; this generally means that encryption is not performed. other (2) indicates that an encryption method other than those specifically enumerated is used for encrypting data on this port. des (3) indicates that the U.S. Data Encryption Standard (FIPS Pub. 46-2) is used for encrypting data on this port; it uses a single 56-bit key. des3 (4) indicates that a special form of the U.S. Data Encryption Standard (FIPS Pub. 46-2) is used for encrypting data on this port; this form specifies that DES is applied three times using three different keys. des3-2key (5) indicates that a special form of the U.S. Data Encryption Standard (FIPS Pub. 46-2) is used for encrypting data on this port; this form specifies that DES is applied three times using two different keys. rc4Mppe (6) indicates that Rivest's Cipher 4 with Microsoft Peer-to-peer encryption is used for encrypting data on this port. rc5 (7) indicates the Rivest's Cipher 5 encryption is used for encrypting data on this port. (10) aMotorolaCaPathLabel - The descriptive label assigned to this path by the device. These labels are used by the User Interface and are useful when comparing information retrieved via SNMP with information retrieved via the User Interface. (11) Sysname - The system name (Sysname) of the device that generated the trap. (12) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format.	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(13) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (14) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
aMotorolaCaRasUserLoginFail	AMotorolaCa_RasUser_LoginFail RAS User: Generic #G, Specific \$S Sysname = \$7, TimeStampAscii = \$8, TimeStampSecs = \$9, TimeStampMilliSecs = \$10	A Remote Access Server user login failed. Variables: (1) aMotorolaCaPathName - A descriptive name for the path. (2) aMotorolaCaPathIndex - A unique identifier for each path. The value for each path must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization. (3) aMotorolaCaRasPathUserName - The name of the currently-logged-in user. (4) aMotorolaCaRasPathLogin-FailReason - The reason for the user login failure: authenticate_fail (1) - None of the other failure reasons apply. user_ID_too_long (2) - The user's ID exceeds the maximum number of characters allowed. invalid_user_ID (3) - The user ID entered by the user does not match any of those configured. password_too_long (4) - The user's password exceeds the maximum number of characters allowed.	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		password_mismatch (5) - The password entered by the user does not match the password configured for the user ID. challenge_mismatch (6) - The session's MD5-Challenge authentication protocol failed. ras_not_enabled (7) - The Remote Access feature is disabled on the device. insufficient_resources (8) - The device does not have enough ports and/or paths to support the login. (5) aMotorolaCaRasPathCurrentTime — The value of system time when the session is referenced, in Universal Time Coordinates. (6) aMotorolaCaPathLabel — The descriptive label assigned to this path by the device. These labels are used by the User Interface and are useful when comparing information retrieved via SNMP with information retrieved via the User Interface. (7) Sysname — The system name (Sysname) of the device that generated the trap. (8) TimeStampAscii — The time at which the event that generated the trap occurred, in ASCII format. (9) TimeStampSecs — The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (10) TimeStampMilliSecs — The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
aMotorolaCaRasUserLogout	AMotorolaCa_RasUser_Logout RAS User: Generic #G, Specific #S	A remote access connection was dropped. Variables:	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	Sysname = \$15, TimeStampAscii = \$16, TimeStampSecs = \$17, TimeStampMilliSecs = \$18	(1) aMotorolaCaPortName - The user-assigned name given to this port. No blanks are allowed in this name. (2) ifIndex - The ifIndex value of the interface associated with this communication port. (3) aMotorolaCaRasPortUserName - The name of the currently-logged-in user. (4) aMotorolaCaRasPortLoginTime - The value of system time when the session is first initiated, in Universal Time Coordinates. (5) aMotorolaCaRasPortCurrentTime - The value of system time when the session is referenced, in Universal Time Coordinates. (6) aMotorolaCaPortInPackets - The number of correctly received packets in this session. (7) aMotorolaCaRasPortOutPackets - The number of correctly transmitted packets in this session. (8) aMotorolaCaRasPortInBytes - The number of correctly received bytes in this session. (9) aMotorolaCaRasPortOutBytes - The number of correctly transmitted bytes in this session. (10) aMotorolaCaRasPortConnectionDropReason - The reason for the connection drop: normal-User-Logout (1) - Normal logout procedure used; equivalent to none (1) other (2) - Used when no other values are applicable. idle-Timeout (3) - Session terminated due to inactivity. maximum-Connection-Time (4) - Maximum allowed time for connection reached.	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		compress-Fail (5) - Session denied as a result of unsupported or misconfigured compression algorithm. encrypt-Fail (6) - Session denied as a result of unsupported or misconfigured cryptography algorithm. admin-Logout (7) - Logout forced to occur on the device by action of administrator. carrier-Loss (8) - Session terminated due to loss of carrier signal. (11) aMotorolaCaPortLabel - The descriptive label assigned to this port by the device. These labels are used by the User Interface and are useful when comparing information retrieved via SNMP with information retrieved via the User Interface. (12) aMotorolaCaRasPortCompType - The type of compression used on the port/path: none (1) indicates that compression is disabled. portHistory (2) indicates that compression is enabled on a port basis and history is preserved across the frames. portPerPacket (3) indicates that compression is enabled on a port basis but history is not preserved across the frames. pathTinyGram (4) indicates that compression is enabled on a path basis and applies to zero-padded frames. (13) aMotorolaCaRasPortEncryptType - The type of encryption used on the port/path: none (1) indicates that no encryption method is used for encrypting data on this port; this	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		generally means that encryption is not performed. other (2) indicates that an encryption method other than those specifically enumerated is used for encrypting data on this port. des (3) indicates that the U.S. Data Encryption Standard (FIPS Pub. 46-2) is used for encrypting data on this port; it uses a single 56-bit key. des3 (4) indicates that a special form of the U.S. Data Encryption Standard (FIPS Pub. 46-2) is used for encrypting data on this port; this form specifies that DES is applied three times using three different keys. des3-2key (5) indicates that a special form of the U.S. Data Encryption Standard (FIPS Pub. 46-2) is used for encrypting data on this port; this form specifies that DES is applied three times using two different keys. rc4Mppe (6) indicates that Rivest's Cipher 4 with Microsoft Peer-to-peer encryption is used for encrypting data on this port. rc5 (7) indicates the Rivest's Cipher 5 encryption is used for encrypting data on this port. (14) <code>aMotorolaCaPathLabel</code> - The descriptive label assigned to this path by the device. These labels are used by the User Interface and are useful when comparing information retrieved via SNMP with information retrieved via the User Interface. (15) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (16) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (17) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		the Epoch (00:00:00 UTC, January 1, 1970). (18) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
remPollLogFile-Trap	Remote_Poll_LogFile Remote Poll log file threshold \$1 is reached Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5	The remote polling log file size exceeded the threshold with a write operation. (1) <i>remPollDefaultLogThreshold</i> - The percentage of the size of the log file that must be exceeded to generate this trap. (2) <i>Sysname</i> - The system name (Sysname) of the device that generated the trap. (3) <i>TimeStampAscii</i> - The time at which the event that generated the trap occurred, in ASCII format. (4) <i>TimeStampSecs</i> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Major
remPollFailure-Trap	Remote_Poll_Failure Remote Poll failed for Address \$1, Protocol \$2 Information \$3 Sysname = \$4, TimeStampAscii = \$5, TimeStampSecs = \$6, TimeStampMilliSecs = \$7	The remPollTable failed to receive a reply to a poll. This trap is generated by the EventTable (if the eventEntry is configured to send traps). Variables: (1) <i>remPollAddress</i> - The network address of the device to monitor. Typically this is an IP, MAC, or IPX address. (2) <i>remPollProtocol</i> - The protocol with which to monitor the device. (3) <i>remPollInformation</i> - After each successful poll, this field is updated with some protocol-specific	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		information that was gleaned from the packet exchange. For the IP protocol, this string contains the value of Time-To-Live from the IP packet header. (4) Sysname - The system name (Sysname) of the device that generated the trap. (5) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (6) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (7) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
remPollSuccessTrap	Remote_Poll_Success Remote Poll success for Address \$1, Protocol \$2 Information \$3 Sysname = \$4, TimeStampAscii = \$5, TimeStampSecs = \$6, TimeStampMilliSecs = \$7	The remPollTable received a reply to a poll after a sequence of unsuccessful polls. This trap is generated by the EventTable (if the eventEntry is configured to send traps). Variables: (1) remPollAddress - The network address of the device to monitor. Typically this is an IP, MAC, or IPX address. (2) remPollProtocol - The protocol with which to monitor the device. (3) remPollInformation - After each successful poll, this field is updated with some protocol-specific information that was gleaned from the packet exchange. For the IP protocol, this string contains the value of Time-To-Live from the IP packet header. (4) Sysname - The system name (Sysname) of the device that generated the trap.	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(5) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (6) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (7) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>vrrpMasterDown</code>	VRRP_MASTER_DOWN VRRP Master Router Down: VRRP master router sysname = \$1, VRID = \$2, Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMilliSecs = \$6.	The VRRP master router is down. If the VRRP master router is configured with a SysName, the trap sends the SysName; otherwise, it sends the IP address of the VRRP master router. Variables: (1) <code>vrrpMasterSysName</code> - The SysName of the VRRP master router. (2) <code>VRID</code> - The virtual router identifier (VRID) for a port on the owner router. (3) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (4) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (5) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Warning
<code>vrrpMasterUp</code>	VRRP_MASTER_UP VRRP Master Router Up: VRRP master router	The VRRP master router is up. If the VRRP master router is configured with a SysName, the trap sends the	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	sysname = \$1, VRID = \$2, Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMilliSecs = \$6.	SysName; otherwise, it sends the IP address of the VRRP master router. Variables: (1) vrrpMasterSysName - The SysName of the VRRP master router. (2) VRID - The virtual router identifier (VRID) for a port on the owner router. (3) Sysname - The system name (Sysname) of the device that generated the trap. (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (5) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
MNR_BGP_PEER_DOWN	MNR_BGP_PEER_DOWN BGP Peer status moves from established state to down state: BGP Peer Remote Address is \$1; BGP Last Error code and Subcode is \$2; BGP Peer Connection State is \$3 Sysname = \$4, TimeStampAscii = \$5, TimeStampSecs = \$6, TimeStampMilliSecs = \$7	A BGP FSM moved from a higher numbered state to a lower numbered state. This is an EOS derived MIB-2 enterprise-specific trap. Variables: (1) bgpPeerRemoteAddr - The remote IP address of this entry's BGP peer. (2) bgpPeerLastError - The last error code and subcode seen by this peer on this connection. If no error occurred, this field is zero; otherwise, the first byte of this two-byte OCTET STRING contains the error code, and the second byte contains the subcode. (3) bgpPeerState - The BGP peer connection state. (4) Sysname - The system name (Sysname) of the device that generated the trap.	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(5) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (6) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (7) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
MNR_BGP_PEER_UP	MNR_BGP_PEER_UP BGP Peer Status enters Up state: BGP Peer Remote Address is \$1; BGP Last Error code and Subcode is \$2; BGP Peer Connection State is \$3 Sysname = \$4, TimeStampAscii = \$5, TimeStampSecs = \$6, TimeStampMilliSecs = \$7	A BGP FSM entered the ESTABLISHED state. This is an EOS derived MIB-2 enterprise-specific trap. Variables: (1) <code>bgpPeerRemoteAddr</code> - The remote IP address of this entry's BGP peer. (2) <code>bgpPeerLastError</code> - The last error code and subcode seen by this peer on this connection. If no error occurred, this field is zero; otherwise, the first byte of this two-byte OCTET STRING contains the error code, and the second byte contains the subcode. (3) <code>bgpPeerState</code> - The BGP peer connection state. (4) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (5) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (6) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (7) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
MNR_newRoot	MNR_DBRIDGE_New_Root Bridge Spanning Tree New Root Sysname = \$1, TimeStampAscii = \$2, TimeStampSecs = \$3, TimeStampMilliSecs = \$4	The sending agent becomes the new root of the Spanning Tree. This trap is sent by a bridge soon after its election as the new root, for example upon expiration of the Topology Change Timer immediately subsequent to its election. This is an EOS derived MIB-2 enterprise-specific trap. Variables: (1) Sysname - The system name (Sysname) of the device that generated the trap. (2) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (3) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (4) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Warning
MNR_topology-Change	MNR_DBRIDGE_Tp_Change Bridge Topology Changed Sysname = \$1, TimeStampAscii = \$2, TimeStampSecs = \$3, TimeStampMilliSecs = \$4	A configured port on a bridge transitioned from the Learning state to the Forwarding state or from the Forwarding state to the Blocking state. The trap is not sent if a newRoot trap is sent for the same transition. This is an EOS derived MIB-2 enterprise-specific trap. (1) Sysname - The system name (Sysname) of the device that generated the trap. (2) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (3) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, rep-	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		resented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (4) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
MNR_frDLCIStatusChange	MNR_FR_DLCI_Change Virtual Circuit Change: interface index = \$1; DLCI = \$2 , new circuit state = \$3 Sysname = \$4, TimeStampAscii = \$5, TimeStampSecs = \$6, TimeStampMilliSecs = \$7	The specified virtual circuit changed state. It was either created or invalidated, or it toggled between the active and inactive states. This is an EOS derived MIB-2 enterprise-specific trap. This trap is sent only if the Frame Relay MIB variable <i>frTrapState</i> was enabled. Variables: (1) <i>frCircuitIfIndex</i> - The <i>ifIndex</i> value of the <i>ifEntry</i> onto which this virtual circuit is layered. (2) <i>frCircuitDlci</i> - The Data Link Connection Identifier for this virtual circuit. (3) <i>frCircuitState</i> - Indicates whether the particular virtual circuit is operational. (4) <i>Sysname</i> - The system name (<i>Sysname</i>) of the device that generated the trap. (5) <i>TimeStampAscii</i> - The time at which the event that generated the trap occurred, in ASCII format. (6) <i>TimeStampSecs</i> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (7) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
MNR_isdnMibCallInformation	MNR_ISDN_Call_Info ISDN Call Information: ifIndex \$1; isdnBearerOperStatus \$2; isdnBearerPeerAddress \$3; isdnBearerPeerSubAddress \$4; isdnBearerCallSetupTime \$5; isdnBearerInfoType \$6; isdnBearerCallOrigin \$7 Sysname = \$8, TimeStampAscii = \$9, TimeStampSecs = \$10, TimeStampMillisecs = \$11	An ISDN line received a call. This is an EOS derived MIB-2 enterprise-specific trap. Variables: (1) ifIndex - A unique value for each interface, ranging between 1 and the value of ifNumber. The value for each interface must remain constant at least from one re-initialization of the entity's network management system to the next re-initialization. (2) isdnBearerOperStatus - The current call control state for this port: idle (1) - The B channel is idle; no call or call attempt is occurring. connecting (2) - A connection attempt (outgoing call) is made on this interface. connected (3) - An incoming call is in the process of validation. active (4) - A call is active on this interface. (3) isdnBearerPeerAddress - The ISDN address to which the current or last call was connected. (4) isdnBearerPeerSubAddress - The ISDN subaddress to which the current or last call was connected. (5) isdnBearerCallSetupTime - The value of sysUpTime when the ISDN setup message for the current or last call was sent or received. If, since system startup, there was no call on this interface, this object has a value of zero. (6) isdnBearerInfoType - The Information Transfer Capability for the current or last call. (7) isdnBearerCallOrigin - The call origin for the current or last call. If, since system startup, there was	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		no call on this interface, this object has a value of unknown (1). (8) Sysname - The system name (Sysname) of the device that generated the trap. (9) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (10) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (11) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
MNR_fallingAlarm	MNR_RMON_Falling_Alarm RMON Falling Alarm: \$2 fell below threshold \$5; value = \$4. (Sample type = \$3; alarm index = \$1) Sysname = \$6, TimeStampAscii = \$7, TimeStampSecs = \$8, TimeStampMilliSecs = \$9	An alarm entry crossed its falling threshold and generated an event that is configured for sending SNMP traps. This is an EOS derived MIB-2 enterprise-specific trap. (1) alarmIndex - An index that uniquely identifies an entry in the alarm table. Each such entry defines a diagnostic sample at a particular interval for an object on the device. (2) alarmVariable - The object identifier of the particular variable to be sampled. Only variables that resolve to an ANS.1 primitive type of INTEGER (INTEGER, Counter, Gauge, or TimeTicks) may be sampled. (3) alarmSampleType - The method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is absoluteValue (1), the value of the selected variable is compared directly with the thresholds at the end of the sampling interval. If the value of this object is deltaValue (2), the value of the selected variable at the last sample is subtracted from the current value,	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		and the difference compared with the thresholds. (4) <code>alarmValue</code> - The value of the statistic during the last sampling period. The value during the current sampling period is not made available until the period is completed. (5) <code>alarmFallingThreshold</code> - A threshold for the sampled statistic. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval was greater than this threshold, a single event is generated. After a falling event is generated, another such event is not generated until the sampled value rises above this threshold and reaches the <code>alarmRisingThreshold</code>. (6) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (7) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (8) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (9) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
MNR_risingAlarm	MNR_RMON_Rise_Alarm RMON Rising Alarm: \$2 exceeded threshold \$5; value = \$4. (Sample type = \$3; alarm index = \$1) Sysname = \$6, TimeStampAscii = \$7, TimeStampSecs = \$8,	An alarm entry crossed its rising threshold and generated an event that is configured for sending SNMP traps. This is an EOS derived MIB-2 enterprise-specific trap. Variables: (1) <code>alarmIndex</code> - An index that uniquely identifies an entry in the alarm table. Each such entry defines a diagnostic sample at a particular interval for an object on the device.	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	TimeStampMilliSecs = \$9	(2) <code>alarmVariable</code> - The object identifier of the particular variable to be sampled. Only variables that resolve to an ANS.1 primitive type of INTEGER (INTEGER, Counter, Gauge, or TimeTicks) may be sampled. (3) <code>alarmSampleType</code> - The method of sampling the selected variable and calculating the value to be compared against the thresholds. If the value of this object is <code>absoluteValue</code> (1), the value of the selected variable is compared directly with the thresholds at the end of the sampling interval. If the value of this object is <code>deltaValue</code> (2), the value of the selected variable at the last sample is subtracted from the current value, and the difference compared with the thresholds. (4) <code>alarmValue</code> - The value of the statistic during the last sampling period. The value during the current sampling period is not made available until the period is completed. (5) <code>alarmRisingThreshold</code> - A threshold for the sampled statistic. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval was less than this threshold, a single event is generated. After a rising event is generated, another such event is not generated until the sampled value falls below this threshold and reaches the <code>alarmFallingThreshold</code>. (6) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (7) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (8) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970).	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(9) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>mnrGtpContextExceeded</code>	MNR_GTP_CONTEXT_EXCEEDED GTP maximum number of subscribers per box was reached: Sysname = \$1, TimestampAscii = \$2, TimeStampSecs = \$3, TimeStampMillisecs = \$4.	The maximum number of GGSN subscribers per Packet Data Gateway (PDG) was reached. Variables: (1) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (2) <code>TimestampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (3) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (4) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Major
<code>mnrGtpCreateFail</code>	MNR_GTP_CONTEXT_CREATE_FAILED GTP failed to create new PDP Context: Sysname = \$4, TimestampAscii = \$5, TimeStampSecs = \$6, TimeStampMillisecs = \$7, PDGAddr = \$1, ImsiId = \$2, CreateFailReason = \$3.	GTP failed to create a new PDP context. Variables: (1) <code>PDGAddr</code> - The IP address of the PDG. (2) <code>ImsiId</code> - The International Mobile Subscriber Identity (IMSI) number of the mobile subscriber. (3) <code>CreateFailReason</code> - The reason for the PDP context creation failure. (4) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (5) <code>TimestampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (6) <code>TimeStampSecs</code> - The time at which the event that generated the	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (7) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrGtpFlowMismatch	MNR_GTP_FLOW_MISMATCH Mismatch between Flow Label and an active PDP Context: Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6, PDGAddr = \$1, FlowLabel = \$2.	The radio subscriber ID does not match an active PDP context. Variables: (1) <i>PDGAddr</i> - The IP address of the PDG. (2) <i>FlowLabel</i> - Identifies unambiguously a GTP flow (radio subscriber). (3) <i>Sysname</i> - The system name (Sysname) of the device that generated the trap. (4) <i>TimeStampAscii</i> - The time at which the event that generated the trap occurred, in ASCII format. (5) <i>TimeStampSecs</i> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Warning
mnrGtpPdgUnreachable	MNR_GTP_PDG_UNREACHABLE PDG is unreachable: Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMillisecs = \$5, PDGAddr = \$1.	The PDG is unreachable. Variables: (1) <i>PDGAddr</i> - The IP address of the PDG. (2) <i>Sysname</i> - The system name (Sysname) of the device that generated the trap. (3) <i>TimeStampAscii</i> - The time at which the event that generated the trap occurred, in ASCII format.	Critical

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(4) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>mnrGtpRedundant-StatusActive</code>	MNR_REDUNDANT_STATUS_ACTIVE The device becomes the active GGSN. This may be accompanied by a trap from the device which just becomes standby. Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMillisecs = \$5 , PDGAddr = \$1	This trap indicates that the device becomes the active GGSN. Variables: (1) <code>PDGAddr</code> - The IP address of the PDG. (2) <code>SysName</code> - The SysName of the GGSN that generated the trap. (3) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (4) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Major
<code>mnrGtpRedundant-StatusStandby</code>	MNR_REDUNDANT_STATUS_STANDBY The device becomes the standby GGSN. This may be accompanied by a trap from the device which just becomes active. Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMillisecs =	This trap indicates that the device becomes the standby GGSN. Variables: (1) <code>PDGAddr</code> - The IP address of the PDG. (2) <code>SysName</code> - The SysName of the GGSN that generated the trap. (3) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (4) <code>TimeStampSecs</code> - The time at which the event that generated the	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	\$5 , PDGAddr = \$1	trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrGtpRedundant- StatusSolitary	MNR_REDUN- DANT_STATUS_SOLI- TARY The device is not configured for GGSN redundancy. Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMillisecs = \$5 , PDGAddr = \$1	This trap indicates that the device is not configured for GGSN redundancy. Variables: (1) <i>PDGAddr</i> - The IP address of the PDG. (2) <i>SysName</i> - The SysName of the GGSN that generated the trap. (3) <i>TimeStampAscii</i> - The time at which the event that generated the trap occurred, in ASCII format. (4) <i>TimeStampSecs</i> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Minor
mnrGtpRedundant- StatusNoStandby	MNR_REDUN- DANT_STA- TUS_NO_STANDBY The active GGSN lost communication with the standby GGSN. Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMillisecs = \$5 , PDGAddr = \$1	This trap is sent by the active GGSN when it lost communication with the standby GGSN. Variables: (1) <i>PDGAddr</i> - The IP address of the PDG. (2) <i>SysName</i> - The SysName of the GGSN that generated the trap. (3) <i>TimeStampAscii</i> - The time at which the event that generated the trap occurred, in ASCII format. (4) <i>TimeStampSecs</i> - The time at which the event that generated the	Critical

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrGgsnDHCPPri- maryFail	MNR_GGSN_PRI- MARY_DHCP_FAILED GGSN DHCP switched over from Primary to Secondary DHCP Server: Sysname = \$3, TimestampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6, CurrDHCPServer = \$1, APNName = \$2.	GGSN switched from the Primary to the Secondary DHCP server because there was no response from the Primary DHCP server. Variables: (1) <i>CurrDHCPServe</i> - The current DHCP server IP address used by this APN. (2) <i>APNName</i> - The access point name (APN); the name that identifies the CEN access point. (3) <i>Sysname</i> - The system name (Sysname) of the device that generated the trap. (4) <i>TimeStampAscii</i> - The time at which the event that generated the trap occurred, in ASCII format. (5) <i>TimeStampSecs</i> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Warning
mnrGgsnDHCPSecon- daryFail	MNR_GGSN_SECON- DARY_DHCP_FAILED GGSN DHCP switched over from Secondary to Primary DHCP Server: Sysname = \$3, TimestampAscii = \$4, TimeStampSecs = \$5,	GGSN switched from the Secondary to the Primary DHCP server because there was no response from the Secondary DHCP server. Variables: (1) <i>CurrDHCPServe</i> - The current DHCP server IP address used by this APN.	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	TimeStampMillisecs = \$6, CurrDHCPsServer = \$1, APNName = \$2.	(2) APNName - The access point name (APN); the name that identifies the CEN access point. (3) Sysname - The system name (Sysname) of the device that generated the trap. (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (5) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrGgsnDHCPPoolEmpty	MNR_GGSN_DHCP_POOL_EMPTY GGSN DHCP internal address pool is empty: Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMillisecs = \$5, APNName = \$1.	The GGSN, functioning as an "internal" DHCP server, runs out of internal dynamic IP addresses to allocate to a mobile subscriber. Variables: (1) APNName - The access point name (APN); the name that identifies the CEN access point. (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
mnrGgsnDHCPAs-signed	MNR_GGSN_DHCP_AS-SIGNED GGSN DHCP allocated IP Address was previously assigned to an MS: Sysname = \$3, TimestampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6, UserIpAddress = \$1, APNName = \$2.	The IP address allocated to the mobile subscriber by the GGSN, functioning as an "internal" DHCP server, was previously allocated to a mobile subscriber. (1) UserIpAddress - The host IP address of the GGSN mobile user. (2) APNName - The access point name (APN); the name that identifies the CEN access point. (3) Sysname - The system name (Sysname) of the device that generated the trap. (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (5) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) TimeStampMillisecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Warning
mnrGgsnRADIUS-switchOver	MNR_GGSN_RADIUS_SWITCH_OVER GGSN RADIUS switch-over: Sysname = \$3, TimestampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6, APNName = \$1, CurrRADIUSServer = \$2	GGSN switched from the primary to the secondary RADIUS server (or vice versa) because there was no response from the primary (or secondary) server. If only one RADIUS server is configured, this trap is generated when that server becomes unreachable. (1) APNName - The access point name (APN); the name that identifies the CEN access point. (2) CurrRADIUSServer - The IP address of the failed RADIUS server. (3) Sysname - The system name (Sysname) of the device that generated the trap. (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format.	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(5) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>mnrGgsnRADIUSassigned</code>	<code>MNR_GGSN_RADIUS_ASSIGNED</code> GGSN RADIUS allocated IP address was previously assigned to a mobile subscriber: <code>Sysname</code> = \$3, <code>TimeStampAscii</code> = \$4, <code>TimeStampSecs</code> = \$5, <code>TimeStampMilliSecs</code> = \$6, <code>APNName</code> = \$1, <code>UserIPAddress</code> = \$2	RADIUS server assigned an IP address belonging to the same CEN that was previously assigned to a mobile subscriber. (1) <code>APNName</code> - The access point name (APN); the name that identifies the CEN access point. (2) <code>UserIPAddress</code> - The host IP address of the GGSN mobile user. (3) <code>Sysname</code> - The system name (<code>Sysname</code>) of the device that generated the trap. (4) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (5) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Warning
<code>mnrGgsnBRLinkUp</code>	<code>MNR_BR_LINK_UP</code> The active GGSN established a connection to the BR. <code>Sysname</code> = \$4, <code>TimeStampAscii</code> = \$5, <code>TimeStampSecs</code> = \$6, <code>TimeStampMilliSecs</code> = \$7	This trap is sent by the active GGSN when it established a connection to the border router (BR) in the CEN. Variables: (1) <code>APNName</code> - The access point name (APN); the name that identifies the CEN access point.	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	, APNName = \$1 , Priority = \$2 , BRIpAddress = \$3	(2) <i>Priority</i> - The session priority for the IPIP tunnel connecting the GGSN and the BR. This priority is specified when the virtual port over which the IPIP tunnel runs is configured. (3) <i>BRIpAddress</i> - The IP address of the border router (BR). (4) <i>SysName</i> - The SysName of the GGSN that generated the trap. (5) <i>TimeStampAscii</i> - The time at which the event that generated the trap occurred, in ASCII format. (6) <i>TimeStampSecs</i> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (7) <i>TimeStampMilliSecs</i> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrGgsnBRLinkDown	MNR_BR_LINK_DOWN The connection to the BR went down. Sysname = \$4, TimeStampAscii = \$5, TimeStampSecs = \$6, TimeStampMillisecs = \$7 APNName = \$1 Priority = \$2 BRIpAddress = \$3	This trap is sent by the active GGSN when the connection to the border router (BR) in the CEN went down. Variables: (1) <i>APNName</i> - The access point name (APN); the name that identifies the CEN access point. (2) <i>Priority</i> - The session priority for the IPIP tunnel connecting the GGSN and the BR. This priority is specified when the virtual port over which the IPIP tunnel runs is configured. (3) <i>BRIpAddress</i> - The IP address of the border router (BR). (4) <i>SysName</i> - The SysName of the GGSN that generated the trap. (5) <i>TimeStampAscii</i> - The time at which the event that generated the trap occurred, in ASCII format. (6) <i>TimeStampSecs</i> - The time at which the event that generated the	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (7) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>mnrGgsnHADSyncObjUp</code>	MNR_HAD_SYNC_OBJ_UP Context database synchronization with the HA GGSN peer was established. Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6 , HADSyncObjPeerName = \$1 , FltTxt = \$2	This trap is sent by the active GGSN when context database synchronization with the HA GGSN peer was established. Variables: (1) <code>HADSyncObjPeerName</code> - SysName of the HA GGSN peer. If no SysName was specified for the peer when it was configured (using the GGSN service <code>HAStandBy</code> parameter), this variable lists the IP address of the HA GGSN peer. (2) <code>FltText</code> - Supplemental text that provides further detail about the trap. (3) <code>SysName</code> - The SysName of the GGSN that generated the trap. (4) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (5) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Normal
<code>mnrGgsnHADSyncObjDown</code>	MNR_HAD_SYNC_OBJ_DOWN Context database synchronization with the HA GGSN peer failed.	This trap is sent by the active GGSN when context database synchronization with the HA GGSN peer failed. Variables:	Critical

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6 , HADSyncObjPeerName = \$1 , FltTxt = \$2	(1) HADSyncObjPeerName - SysName of the HA GGSN peer. If no SysName was specified for the peer when it was configured (using the GGSN service HAStandBy parameter), this variable lists the IP address of the HA GGSN peer. (2) FltText - Supplemental text that provides further detail about the trap. (3) SysName - The SysName of the GGSN that generated the trap. (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (5) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrGgsnHADSyncObjInfo	MNR_HAD_SYNC_OBJ_INFO Context database synchronization status update. Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6 , HADSyncObjPeerName = \$1 , FltTxt = \$2	This informational trap reports a context database synchronization status update. Variables: (1) HADSyncObjPeerName - SysName of the HA GGSN peer. If no SysName was specified for the peer when it was configured (using the GGSN service HAStandBy parameter), this variable lists the IP address of the HA GGSN peer. (2) FltText - Supplemental text that provides further detail about the trap. (3) SysName - The SysName of the GGSN that generated the trap. (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format.	Info

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(5) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>mnrGgsnHADSyncObjSingleUpdFailed</code>	<code>MNR_HAD_SYNC_OBJ_SINGLE_UPD_FAILED</code> Single context update failure. <code>Sysname</code> = \$3, <code>TimeStampAscii</code> = \$4, <code>TimeStampSecs</code> = \$5, <code>TimeStampMilliSecs</code> = \$6 , <code>HADSyncObjPeerName</code> = \$1 , <code>FltTxt</code> = \$2	This trap indicates that a context database synchronization update for a single context failed. Variables: (1) <code>HADSyncObjPeerName</code> - <code>SysName</code> of the HA GGSN peer. If no <code>SysName</code> was specified for the peer when it was configured (using the GGSN service <code>HAStandBy</code> parameter), this variable lists the IP address of the HA GGSN peer. (2) <code>FltText</code> - Supplemental text that provides further detail about the trap. (3) <code>SysName</code> - The <code>SysName</code> of the GGSN that generated the trap. (4) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (5) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Minor
<code>mnrGgsnHADSyncObjEarlyFailOver</code>	<code>MNR_HAD_SYNC_OBJ_EARLY_FAIL_OVER</code> Premature context database resynchronization.	This trap indicates that a premature context database resynchronization occurred (the standby GGSN becomes active before it is fully	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6 , HADSyncObjPeerName = \$1 , FltTxt = \$2	synchronized with the old active GGSN). Variables: (1) HADSyncObjPeerName - SysName of the HA GGSN peer. If no SysName was specified for the peer when it was configured (using the GGSN service HASTandBy parameter), this variable lists the IP address of the HA GGSN peer. (2) FltText - Supplemental text that provides further detail about the trap. (3) SysName - The SysName of the GGSN that generated the trap. (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (5) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) TimeStampMillisecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrDS1E1LineChangeTrap	MNR_DS1E1_LINE_CHANGED DS1/E1 interface line status changed: Sysname = \$5, TimeStampAscii = \$6, TimeStampSecs = \$7, TimeStampMillisecs = \$8, PathLabel = \$1, PathType = \$2, LineStatus = \$3, LoopbackType = \$4.	The physical line status of the T1(DS1)/E1 interface changed. (1) PathLabel - The path label identifying the T1/E1 path; for example: DS1E1_PA_2, DS1E1_PA_3A, DS1E1_PA5. (2) PathType - The interface type associated with the path: DS1 or E1. (3) LineStatus - The line status of the interface, displayed as a bit map so multiple failures and loopback states can be represented simultaneously. (4) LoopbackType - The type of loopback on the interface This information is provided only if the dsx1Loop-	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		backState bit is set in the LineStatus. (5) Sysname - The system name (Sysname) of the device that generated the trap. (6) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (7) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (8) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
mnrPidMonFault-Critical  NOTE: This trap is implemented only on the MNR S6000, GGM 8000 and MCG 8000 platforms. In addition, MNR S6000 base units with model numbers CLN1780A and higher and GGM 8000 gateways support a single power/cooling supply system. Supplemental trap messages pertaining to the redundant supply system do not apply to and are not reported for these units.	MNR_ENV-MON_FAULT_CRITICAL The supply system transitioned to a Critical state. Sysname = \$3, TimestampAscii = \$4, TimestampSecs = \$5, TimestampMillisecs = \$6, mnrPidState = \$1, mnrPidText = \$2.	This trap reports the critical-level faults provided by the operational status of the supply system. Variables: (1) mnrPidState - The operational state of the supply system. Critical - Immediate subsystem service warning (2) mnrPidText - Supplemental text included in critical-level fault monitoring traps. This text varies by hardware platform; examples include but are not limited to: Power Supply fans failed - One of the following events occurred: - Any two redundant supply system fans are rotating too slowly. - Both auxiliary cooling component fans are rotating too slowly. - Both redundant power subsystem component fans are rotating too slowly. Input voltage too low - Redundant power subsystem voltage on both AC inputs is too low. Output voltage out of range - Redundant power subsystem voltage is outside the required operating range by more than 5% (<4.75V or >5.25V). Output current too high - Redundant power subsystem total output current is too high (> 12 amps). Power Supply temp too high - Redundant power subsystem internal temperature is too high. (3) Sysname - The system name (Sysname) of the device that generated the trap. (4) TimestampAscii - The time at which the event that generated the trap occurred, in ASCII format. (5) TimestampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since	Critical

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		the Epoch (00:00:00 UTC, January 1, 1970). (6) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>mnrEnvMonFault-Warning</code>  NOTE: This trap is implemented only on the MNR S6000, GGM 8000 and MCG 8000 platforms. In addition, MNR S6000 base units with model numbers CLN1780A and higher and GGM 8000 gateways and MCG 8000 support a single power/cooling supply system. Supplemental trap messages pertaining to the redundant supply system do not apply to and are not reported for these units.	<code>MNR_ENV-MON_FAULT_WARNING</code> The supply system transitioned to a Warning state. <code>Sysname = \$3,</code> <code>TimeStampAscii = \$4,</code> <code>TimeStampSecs = \$5,</code> <code>TimeStampMilliSecs = \$6,</code> <code>mnrObjState = \$1,</code> <code>mnrObjText = \$2.</code>	This trap reports the warning-level faults provided by the operational status of the supply system. Variables: (1) <code>mnrObjState</code> - The operational state of the supply system. Warning - Schedule subsystem service warning (2) <code>mnrObjText</code> - Supplemental text included in warning-level fault monitoring traps. This text varies by hardware platform; examples include but are not limited to: CPU Fan failed - The auxiliary cooling component CPU fan is failing. Aux Fan 1 failed - Auxiliary cooling component fan #1 is failing. Aux Fan 2 failed - Auxiliary cooling component fan #2 is failing. AC1 input too low - Redundant power subsystem voltage on AC input #1 is too low (or unplugged). AC2 input too low - Redundant power subsystem voltage on AC input #2 is too low (or unplugged). One Power Supply fan failed - One redundant power subsystem fan is rotating too slowly. Output voltage out of range - Redundant power subsystem voltage is outside the required operating range by plus or minus 2% (<4.9V or >5.1V).	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		One power supply failed - One of the following problems was detected on an individual power supply: - Shorted output - Open output - Under-voltage - Over-voltage Load imbalance - Redundant power subsystem current load imbalance. The individual power supplies are not sharing the load equally (>20% difference). (3) Sysname - The system name (Sysname) of the device that generated the trap. (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (5) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
mnrEnvMonFault- Normal  NOTE: This trap is implemented only on the MNR S6000, GGM 8000 and MCG 8000 platforms. In addition, MNR S6000 base units with model numbers CLN1780A and higher and GGM 8000 gateways and MCG 8000 support a single power/cooling supply system. Supplemental trap messages pertaining to the redundant supply system do not apply to and are not reported for these units.	MNR_ENV- MON_FAULT_NORMAL The supply system transitioned to a Normal state. Sysname = \$3, TimestampAscii = \$4, TimestampSecs = \$5, TimestampMillisecs = \$6, mnObjState = \$1, mnObjText = \$2.	This trap indicates that the supply system transitioned to a Normal state from a Critical state or the indicated Warning state. Variables: (1) mnObjState - The operational state of the supply system. (2) mnObjText - Supplemental text included in warning-level fault monitoring traps. This text varies by hardware platform; examples include but are not limited to: CPU Fan failed - The auxiliary cooling component CPU fan is failing. Aux Fan 1 failed - Auxiliary cooling component fan #1 is failing. Aux Fan 2 failed - Auxiliary cooling component fan #2 is failing. AC1 input too low - Redundant power subsystem voltage on AC input #1 is too low (or unplugged). AC2 input too low - Redundant power subsystem voltage on AC input #2 is too low (or unplugged). One Power Supply fan failed - One redundant power subsystem fan is rotating too slowly. Output voltage out of range - Redundant power subsystem voltage is outside the required operating range by plus or minus 2% (<4.9V or >5.1V). One power supply failed - One of the following problems was detected on an individual power supply: - Shorted output - Open output - Under-voltage - Over-voltage Load imbalance - Redundant power subsystem current load imbalance. The individual power	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		supplies are not sharing the load equally (>20% difference). (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (5) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) TimeStampMillisecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrDeviceFault-Critical	MNR_DE- VICE_FAULT_CRITICAL The device experienced a critical failure. Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6 , Reason = \$1 , Supp Text =\$2	This trap indicates that the device experienced a critical failure. Variables: (1) Reason - A string that states the reason for the device failure. One of the following strings is reported: Clock Pin Error - Replace Router - This text is an alarm warning that there is a hardware issue which warrants router replacement. The trap and alarm continue to be reported until the router is replaced. (2) SuppText - An optional string that provides additional diagnostic detail about the device failure. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMillisecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Critical

Trap Name	HP OpenView Trap Definition Format	Description	Severity
mnrOverloadEnter	MNR_OVERLOAD_ENTER The EOS device entered Overload state: Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMillisecs = \$5, OverLoadCause = \$1	This trap indicates that the device entered Overload state, with the cause being either CPU utilization, memory usage, or queue drops. Variables: (1) OverLoadCause - The cause of the device entering Overload state: 1 - CPU utilization 2 - Memory usage 3 - Queue drops (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMillisecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Major
mnrOverloadLeave	MNR_OVERLOAD_LEAVE The EOS device exited from Overload state: Sysname = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMillisecs = \$5, OverLoadCause = \$1	This trap indicates that the device exited from Overload state, with the cause being either CPU utilization, memory usage, or queue drops. Variables: (1) OverLoadCause - The cause of the device exiting from Overload state: 1 - CPU utilization 2 - Memory usage 3 - Queue drops (2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format.	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(4) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
<code>mnrVpnTunnelUp</code>	<code>MNR_VPN_TUNNEL_UP</code> A VPN tunnel was established. <code>Sysname</code> = \$14, <code>TimeStampAscii</code> = \$15, <code>TimeStampSecs</code> = \$16, <code>TimeStampMilliSecs</code> = \$17, <code>PortName</code> = \$12, <code>PortId</code> = \$4, <code>Tunnel_type</code> = \$13, <code>LocalAddress</code> = \$2, <code>RemoteAddress</code> = \$3.	This trap indicates that a VPN tunnel was established. (2) <code>LocalAddress</code> - The IP address of the local side of the VPN tunnel. (3) <code>RemoteAddress</code> - The IP address of the remote side of the VPN tunnel. (4) <code>PortId</code> - The port number on which the VPN tunnel is created. (13) <code>Tunnel_type</code> - The protocol used to establish the tunnel. (14) <code>Sysname</code> - The system name (<code>Sysname</code>) of the device that generated the trap. (15) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (16) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (17) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	Normal
<code>mnrCWRStatusChangeTrap</code>	<code>MNR_CWR_STATUS_CHANGE</code> A CWR-enabled 12-port T1/E1 module changed its CWR state from ac-	This trap indicates that a CWR-enabled 12-port T1/E1 module changed its CWR state. Variables: (1) <code>ModuleNumber</code> - The slot number in which the CWR-enabled 12-port T1/E1 module is installed: 4 or 5.	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	tive to inactive or vice versa: Sysname = \$3, TimeStampAscii = \$4, TimeStampSecs = \$5, TimeStampMillisecs = \$6, ModuleNumber = \$1 , CWRStatus = \$2	(2) CWRStatus - The new state of the CWR-enabled module: 1 - Active 2 - Inactive (3) Sysname - The system name (Sysname) of the device that generated the trap. (4) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (5) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (6) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
cryptoMibSACreated	MNR_CRYPT-SA_CREATED Crypto Security Association Created. SAId = \$1, SPI = \$2, PortId = \$3, SourceAddress = \$4, DestinationAddress = \$5, Mode = \$6, KeyManagementMode = \$7, Protocol = \$8, AuthAlgorithm = \$9, EncryptAlgorithm = \$10, PolicyId = \$11, CurrentTime = \$12, PortLabel = \$13, PortName = \$14, PeerAddress = \$15, LifetimeTime = \$16, LifetimeData = \$17, IkeClientId = \$18, SysName = \$19,	This trap indicates that a security association (SA) was established. For dynamic keys, this trap is sent after phase 2 negotiation completes and the first security association of a bundle is created. For manual keys (statically configured), this trap is sent when the security association is configured. For further detail about encryption-specific variables, see Table 118: Encryption-Specific Trap Variable Descriptions on page 963. Variables: (1) SAId (2) SPI (3) PortId (4) SourceAddress (5) DestinationAddress (6) Mode (7) KeyManagementMode (8) Protocol (9) AuthAlgorithm	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	TimeStampAscii = \$20, TimeStampSecs = \$21, TimeStampMilliSecs = \$22.	(10) EncryptAlgorithm (11) PolicyId (12) CurrentTime (13) PortLabel (14) PortName (15) PeerAddress (16) LifetimeTime (17) LifetimeData (18) IkeClientId (19) Sysname - The system name (Sysname) of the device that generated the trap. (20) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (21) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (22) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
cryptoMibSADeleted	MNR_CRYPTOSADELETED Crypto Security Association Deleted. SAId = \$1, SPI = \$2, PortId = \$3, SourceAddress = \$4, DestinationAddress = \$5, Mode = \$6, KeyManagementMode = \$7, Protocol = \$8, AuthAlgorithm = \$9, EncryptAlgorithm = \$10, PolicyId = \$11, StartTime = \$12,	This trap indicates that a security association was deleted. For dynamic keys, this trap is sent after the last security association of a bundle is deleted. For manual keys (statically configured), this trap is sent when the security association is deleted. Variables: (1) SAId (2) SPI (3) PortId (4) SourceAddress (5) DestinationAddress (6) Mode (7) KeyManagementMode (8) Protocol	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	CurrentTime = \$13, TerminateReason = \$14, PortLabel = \$15, PortName = \$16, PeerAddress = \$17, LifetimeTime = \$18, LifetimeData = \$19, SysName = \$20, TimeStampAscii = \$21, TimeStampSecs = \$22, TimeStampMilliSecs = \$23.	(9) AuthAlgorithm (10) EncryptAlgorithm (11) PolicyId (12) StartTime (13) CurrentTime (14) TerminateReason (15) PortLabel (16) PortName (17) PeerAddress (18) LifetimeTime (19) LifetimeData (20) Sysname - The system name (Sysname) of the device that generated the trap. (21) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (22) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (23) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
cryptoMibSAException	MNR_CRYPTOS_SA_EXCEPTION Crypto Security Association Exception. SAId = \$1, SPI = \$2, PortId = \$3, SourceAddress = \$4, DestinationAddress = \$5, Mode = \$6, KeyManagementMode = \$7, Protocol = \$8, AuthAlgorithm = \$9, EncryptAlgorithm =	This trap indicates that a security association exception was encountered. The Exception variable indicates the reason for the exception. Variables: (1) SAId (2) SPI (3) PortId (4) SourceAddress (5) DestinationAddress (6) Mode (7) KeyManagementMode (8) Protocol	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	\$10, PolicyId = \$11, StartTime = \$12, CurrentTime = \$13, Exception = \$14, PortLabel = \$15, PortName = \$16, PeerAddress = \$17, LifetimeTime = \$18, LifetimeData = \$19, IkeClientId = \$20, SysName = \$21, TimeStampAscii = \$22, TimeStampSecs = \$23, TimeStampMilliSecs = \$24.	(9) AuthAlgorithm (10) EncryptAlgorithm (11) PolicyId (12) StartTime (13) CurrentTime (14) Exception (15) PortLabel (16) PortName (17) PeerAddress (18) LifetimeTime (19) LifetimeData (20) IkeClientId (21) Sysname - The system name (Sysname) of the device that generated the trap. (22) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (23) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (24) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrCryptoSelf-TestFailed	MNR_CRYPTO_SELF_TEST_FAILED FIPS self test failure. Reason= \$1, SysName = \$2, TimeStampAscii = \$3, TimeStampSecs = \$4, TimeStampMilliSecs = \$5.	This trap indicates that one of the following self tests failed: Cryptographic power-on self-test - These tests cover the range of cryptographic operations that are available on the device. Each test executes a cryptographic operation on known input data where the output data is also known and expected. A test is said to pass if the output data compares exactly with the expected output. Continuous random number generator test - Whenever a random number is needed, for example to	Critical

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		generate a cryptographic key, the random number generated is compared to the previously generated random number to ensure they are different. Conditional bypass test - Whenever the user changes the state of the routed data from clear to encrypted, the conditional bypass test is executed to insure that the affected cyryptographic operations perform as desired before the data is subjected to these operations. Package authentication test - Whenever a software package is downloaded to or uploaded from a secure device, the package is authenticated to ensure that the executable code contained in the package was created by a trusted source. Secure feature authentication - One of the following failure cases occurred: • The boot image does not contain the public key cryptographic payload required to enable the encryption engine. • The feature configuration file required to enable the encryption engine does not exist. • The contents of the feature configuration file required to enable the encryption engine are not intact. • The secure feature authentication process failed. • The enabled encryption engine failed a self test.  NOTE: Secure feature authentication is supported only on the GGM 8000 platform. Variables: (1) Reason - A string stating which self test failed and indicating the reason for the failure.	

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.  NOTE: When hardware encryption is enabled in the system, the router performs a reboot after the failed self test and the trap is not sent.	
mnrOspfMD5KeyMismatch	MNR_OSPF_MD5_KEY_MISMATCH MNR OSPF MD5 Key Mismatch. Neighbor IP = \$1, Key Id = \$2, OSPF Port Label = \$3, Reason = \$4, SysName = \$5, TimeStampAscii = \$6, TimeStampSecs = \$7, TimeStampMilliSecs = \$8.	This trap indicates that an OSPF MD5 key mismatch occurred and an OSPF packet was dropped due to authentication failure. The device generating the trap cannot authenticate the OSPF packet from the neighbor specified by the Neighbor IP variable.  NOTE: This trap is generated only by the OSPF designated router (DR) or backup DR. Variables: (1) Neighbor IP - The IP address of the neighbor router from which the packet that cannot be authenticated was received. (2) Key Id - The key ID for the OSPF MD5 key (0-255). (3) OSPF Port Label - The port on which the MD5 key is configured. (4) Reason - A string indicating the reason for the key mismatch:	Critical

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		● Key used is bad - The key in the received packet is not a valid OSPF MD5 key. ● Key not configured - The key in the received packet is not configured on the device generating the trap. ● Key is Out of Service - The key in the received packet expired or is not yet active. ● Secret doesn't match for the two keys (typo) - The key in the received packet does not match the OSPF MD5 key configured on the device generating the trap. (5) Sysname - The system name (Sysname) of the device that generated the trap. (6) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (7) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (8) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrMultiCastGroupsExceeded	MNR_MULTI-CAST_GROUP_EXCEEDED The multicast group limit was exceeded: Sysname = \$6, TimeStampAscii = \$7, TimeStampSecs = \$8, TimeStampMilliSecs = \$9, Multicast Group IP Address = \$1, Msg Source IP Address	This trap indicates that the configured multicast group limit was met or exceeded and groups are rejected. Variables: (1) Multicast Group IP Address - The IP address of the multicast group (2) Msg Source IP Address - The source IP address of the packet that was dropped because the multicast group limit was met or exceeded.	Critical

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	= \$2, Port Label = \$3, Number of groups rejected = \$4, Reason = \$5.	(3) Port Label - The port associated with the IGMP join message. If the trap is caused by one of the PIM messages, then this variable reports N/A. (4) Number of groups rejected - The number of multicast groups rejected because the multicast group limit was met or exceeded. (5) Reason - The reason (type of message) that caused the configured multicast group limit to be met or exceeded: 1 - IGMP Joins dropped by PIM 2 - PIM Joins/Prunes 3 - PIM Registers 4 - PIM Multicast Data Packets (6) Sysname - The system name (Sysname) of the device that generated the trap. (7) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (8) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (9) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrGgsnCenLinkUp	MNR_CEN_LINK_UP A GGSN connection to the CEN was established: Sysname = \$4, TimeStampAscii = \$5, TimeStampSecs = \$6, TimeStampMillisecs = \$7, APN Name= \$1, Priority = \$2, BR IPaddress = \$3.	This trap indicates that a connection between the GGSN and the CEN was established. Variables: (1) APN Name - The access point name (APN). This is the name that identifies the CEN access point. (2) Priority - The GGSN priority. This is the priority value that is assigned when the virtual port is bound to the APN using the GGSN service AccessPointName parameter.	Normal

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(3) <code>BR_IPAddress</code> - The IP address of the border router. (4) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (5) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (6) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (7) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrGgsnCenLink-Down	MNR_CEN_LINK_DOWN A GGSN connection to the CEN is down: Sysname = \$4, TimeStampAscii = \$5, TimeStampSecs = \$6, TimeStampMillisecs = \$7, APN Name = \$1, Priority = \$2, BR_IPAddress = \$3.	This trap indicates that a connection between the GGSN and the CEN went down. Variables: (1) <code>APN Name</code> - The access point name (APN). This is the name that identifies the CEN access point. (2) <code>Priority</code> - The GGSN priority. This is the priority value that is assigned when the virtual port is bound to the APN using the GGSN service <code>AccessPointName</code> parameter. (3) <code>BR_IPAddress</code> - The IP address of the border router. (4) <code>Sysname</code> - The system name (Sysname) of the device that generated the trap. (5) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (6) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970).	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(7) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrProtocolAuth-Failed	MNR_PROTOCOL_AUTHENTICATION_FAILED MNR Protocol authentication failed: Sysname = \$9, TimeStampAscii = \$10, TimeStampSecs = \$11, TimeStampMillisecs = \$12, cryptoSPI = \$1, cryptoSourceAddress = \$2, cryptoDestinationAddress = \$3, cryptoProtocol = \$4, aMotorolaCaPortName = \$5, cryptoPeerAddress = \$6, mnrAuthFailedProtocol = \$7, mnrAuthFailureReason = \$8.	This trap indicates that a protocol packet was discarded because of an authentication failure. The device generating the trap cannot authenticate the packet from a peer router. Variables: (1) cryptoSPI - The security parameters index (SPI) for the SA. This value may be 0 when a failure to create an SA occurs. (2) cryptoSourceAddress - The IP address of the source of the SA. This variable is reported as a string rather than an IP address to support IPv6 as well as IPv4 addresses. (3) cryptoDestinationAddress - The IP address of the destination of the SA. This variable is reported as a string rather than an IP address to support IPv6 as well as IPv4 addresses. (4) cryptoProtocol - The security mechanism used by the SA: AH or ESP (5) aMotorolaCaPortName - The user-assigned name given to this port. The name cannot contain any blanks. (6) cryptoPeerAddress - The IP address of the SA's peer. This variable is reported as a string rather than an IP address to support IPv6 as well as IPv4 addresses. (7) mnrAuthFailedProtocol - The protocol for which authentication failed. For example: PIM. (8) mnrAuthFailureReason - A string stating the reason for the authentication failure. (9) Sysname - The system name of the device that generated the trap.	Critical

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		(10) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (11) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (12) <code>TimeStampMillisecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrWANIfQueue-DepthCaution	MNR_WAN_IF_QUEUE_DEPTH_CAUTION WAN Link priority queue depth caution: Sysname = \$5, TimeStampAscii = \$6, TimeStampSecs = \$7, TimeStampMillisecs = \$98, IndexLabel = \$1, QoSPriority = \$2, QoSCautionPercent = \$3, QosErrTime = \$4.	This trap indicates that the WAN link exceeded the user-specified caution threshold (a percentage of the queue limit specified with the <code>CautionTrap PErcent</code> action attribute in the IP-QoS service <code>POLicy</code> parameter for Ethernet links and with the Frame Relay service <code>CautionTrap</code> parameter for Frame Relay links). This trap is sent as a warning concerning the number of packets in the queue. Variables: (1) <code>IndexLabel</code> - The virtual port number on which the WAN link is configured. (2) <code>QoSPriority</code> - The priority number of the queue for either the virtual port (for Ethernet links) or the DLCI associated with this virtual circuit (for Frame Relay links). Possible values are 0-4 for Ethernet links: 0 = Urgent 1 = High 2 = Medium 3 = Normal 4 = Low and 3-7 for Frame Relay links, with 7 being the highest priority queue. The EOS code puts priority 0, 1, and 2 packets in the priority 3 queue. (3) <code>QoSCautionPercent</code> - Sets the queue depth at which a caution trap	Warning

Trap Name	HP OpenView Trap Definition Format	Description	Severity
		is generated, as a percentage of the queue limit. For Ethernet links, the queue limit is fixed at 300 packets for all five queues (0-4). For Frame Relay links, the queue limit is calculated based on the bandwidth allocated to the link and can range from 300 to 1200 packets. For example, setting a <code>CautionLimit</code> of 50 for queue 1 on an Ethernet link would mean that the trap is generated when the queue exceeds 150 packets (and any packets above the 150-packet (50%) threshold would be dropped). (4) <code>QoSErrTime</code> - The value of the <code>SysUpTime</code> variable when the queue depth caution limit was exceeded. (5) <code>Sysname</code> - The system name (<code>Sysname</code>) of the device that generated the trap. (6) <code>TimeStampAscii</code> - The time at which the event that generated the trap occurred, in ASCII format. (7) <code>TimeStampSecs</code> - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (8) <code>TimeStampMilliSecs</code> - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrWANIfQueueDepthExceeded	MNR_WAN_IF_QUEUE_DEPTH_EXCEEDED WAN Link priority queue depth exceeded: Sysname = \$5, TimeStampAscii = \$6, TimeStampSecs = \$8, TimeStampMilliSecs = \$9, IndexLabel = \$1, QoSPriority = \$2,	This trap indicates that WAN link exceeded the queue limit. The queue was full, and the packet could not be added. Variables: (1) <code>IndexLabel</code> - The virtual port number on which the WAN link is configured. (2) <code>QoSPriority</code> - The priority number of the queue for either the virtual port (for Ethernet links) or the DLCI associated with this virtual circuit	Minor

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	QoSTotalPacketsDropped = \$3, QoSErrTime = \$4.	(for Frame Relay links). Possible values are 0-4 for Ethernet links: 0 = Urgent 1 = High 2 = Medium 3 = Normal 4 = Low and 3-7 for Frame Relay links, with 7 being the highest priority queue. The EOS code puts priority 0, 1, and 2 packets in the priority 3 queue. (3) QoSTotalPacketsDropped - The number of packets dropped because the queue was full for this priority queue. (4) QoSErrTime - The value of the SysUpTime variable when the queue limit was exceeded. (5) Sysname - The system name (Sysname) of the device that generated the trap. (6) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (7) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (8) TimeStampMillisecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	
mnrWANIfTotalQueueDepthExceeded	MNR_WAN_IF_TOTAL_QUEUE_DEPTH_EXCEEDED WAN Total priority queue depth exceeded: Sysname = \$2, TimeStampAscii = \$3 , TimeStampSecs = \$4 , TimeStampMillisecs =	This trap indicates that the total number of queued packets in all priority queues for the WAN link exceeded the sum of the queue limits for all priority queues. Variables: (1) QoSErrTime - The value of the SysUpTime variable when the total queue limit was exceeded.	Major

Trap Name	HP OpenView Trap Definition Format	Description	Severity
	\$5, QoSErrTime = \$1.	(2) Sysname - The system name (Sysname) of the device that generated the trap. (3) TimeStampAscii - The time at which the event that generated the trap occurred, in ASCII format. (4) TimeStampSecs - The time at which the event that generated the trap occurred, in integer format, represented as whole seconds since the Epoch (00:00:00 UTC, January 1, 1970). (5) TimeStampMilliSecs - The partial seconds portion of the integer trap event timestamp. This variable represents the partial seconds to the accuracy of the clock of the device generating the trap events.	

Table 116: T1/E1 Line Statuses Represented by LineStatus Bit Positions

Variable Value	Definition	Meaning
1	dsx1NoAlarm	No alarm present
2	dsx1RcvFarEndLOF	Far end LOF (yellow alarm)
4	dsx1XmtFarEndLOF	Near end sending LOF indication
8	dsx1RcvAIS	Far end sending AIS
16	dsx1XmtAIS	Near end sending AIS
32	dsx1LossOfFrame	Near end LOF (red alarm)
64	dsx1LossOfSignal	Near end Loss of Signal
128	dsx1LoopbackState	Near end is looped
8192	dsx1UnavailSisState	Near end in unavailable signal state
16384	dsx1NetEquipOOS	Carrier equipment out of service

Table 117: Loopbacks Indicated by LoopbackType Variable

Variable Value	Definition
For S6000 and GGM 8000 Platforms	
(0)	No Loopback
(1)	Local Simulated Remote Line Loopback
(2)	Local Simulated Remote Payload Loopback
(3)	Remote ANSI BOP Payload Loopback
(4)	Remote ANSI BOP Line Loopback

Variable Value	Definition
(5)	Remote ATT Inband Line Loopback
(6)	Remote ATT FDL Payload Loopback
For S2500 Platforms	
(0)	No Loopback
(1)	Local Simulated Remote Line Loopback
(2)	Local Simulated Remote Payload Loopback
(3)	Local V54 Loopback
(4)	Local Line Loopback
(5)	Local Payload Loopback
(6)	Remote ANSI BOP Payload Loopback
(7)	Remote ANSI BOP Line Loopback
(8)	Remote ATT Inband Line Loopback
(9)	Remote ATT FDL Payload Loopback
(10)	TDM Echo Loopback
(11)	TDM Interrupt Line Loopback

Table 118: Encryption-Specific Trap Variable Descriptions

Variable	Description
SAId	The security association (SA) ID number. This number uniquely identifies the SA in the crypto table. You must compare the SAId variable logged in the encryption traps against the results from the <code>EOS SHOW -CRYPTO IKESecAssoc</code> command in order to find the right session.
SPI	The security parameters index (SPI) for the created SA. This variable may be zero when a failure to create the SA occurs. You must compare the SPI variable logged in the encryption traps against the results from the <code>EOS SHOW -CRYPTO IKESecAssoc</code> command in order to find the right session.
PortId	The network interface on which these SA services are applied.
SourceAddress	The source IP address of the SA. This variable is reported as a string rather than an IP address to support IPv6 as well as IPv4 addresses.
DestinationAddress	The destination IP address of the SA. This variable is reported as a string rather than an IP address to support IPv6 as well as IPv4 addresses.
Mode	The mode in which the SA is run: (1) <code>transport</code> - A PPTP or L2TP session, or FRF17. All traps generated on FRF.17 circuits are logged as mode="transport". (2) <code>tunnel</code> - An IPIP tunnel session.
KeyManagementMode	The key management scheme used by this SA: (1) <code>manual</code> - manual key management is used.

Variable	Description
	(2) <code>dynamic</code> - dynamic key management (IKE) is used.
Protocol	The security mechanism that was applied to this SA: (1) <code>none</code> - Neither authentication nor encryption is performed. (2) <code>ah</code> - (not supported for this release) Authentication is performed by an authentication header. This mechanism is used only when the <code>Mode</code> is <code>transport</code>. (3) <code>esp</code> - Encryption is performed by an encapsulating security payload. This mechanism is used only when the <code>Mode</code> is <code>transport</code>. (4) <code>ahEsp</code> - (not supported for this release) Authentication is performed by an authentication header and encryption is performed by an encapsulating security payload. This mechanism is used only when the <code>Mode</code> is <code>transport</code>. (5) <code>espAuth</code> - Encryption is performed by an encapsulating security payload. This mechanism is used only when the <code>Mode</code> is <code>tunnel</code>.
AuthAlgorithm	The algorithm that is used for authentication by this security association: (1) <code>none</code> - Authentication is not performed. (2) <code>other</code> - Authentication is performed by an unlisted algorithm. (3) <code>hmacMD5</code> - Authentication is performed by the Message Digest 5 hash function. (4) <code>hmacSHA1</code> - Authentication is performed by the revised Secure Hash Algorithm hash function.
EncryptAlgorithm	The algorithm that is used for encryption by this security association: (1) <code>none</code> - Encryption is not performed. (2) <code>other</code> - Encryption is performed by an unlisted algorithm. (3) <code>desCbc</code> - Encryption is performed by the Data Encryption Standard algorithm with Cipher Block Chaining. (4) <code>tripleDesCbc</code> - Encryption is performed by the Triple Data Encryption Standard algorithm with Cipher Block Chaining. (5) <code>rc5</code> - (not supported for this release) Encryption is performed by the Rivest Cipher 5 algorithm. (6) <code>tripleDes2key</code> - Encryption is performed by the Triple Data Encryption Standard algorithm with Cipher Block Chaining, using two keys. (7) <code>aes128</code> - Encryption is performed by 128-bit Advanced Encryption Standard algorithm. (8) <code>aes192</code> - Encryption is performed by 192-bit Advanced Encryption Standard algorithm. (9) <code>aes256</code> - Encryption is performed by 256-bit Advanced Encryption Standard algorithm.
PolicyId	A name identifying the policy applied to this SA. This is the policy name specified with the CRYPTO service <code>DynamicPOLicy</code> or <code>manualPOLicy</code> parameter.
StartTime	The system time when the SA was first established.

Variable	Description
	This variable applies only to the <code>CryptoMibSADeleted</code> and <code>CryptoMibSAException</code> traps.
<code>CurrentTime</code>	The system time when the SA was referenced.
<code>TerminateReason</code>	The reason that the SA was terminated: (1) <code>other</code> - The SA was deleted due to a reason other than those described below. (2) <code>normal</code> - A manual or dynamic rekey occurred. (3) <code>administration</code> - The CRYPTO service <code>CONTROL</code> parameter was set to <code>Disabled</code> (<code>SETDefault -CRYPTO CONTROL = Disabled</code>). (4) <code>noResource</code> - The system is out of memory due to a memory leak caused by a software failure or oversubscription. This variable applies to only the <code>CryptoMibSADeleted</code> trap.
<code>Exception</code>	The reason that the SA exception occurred: (1) <code>other</code> – Encryption is not performed. (2) <code>ahHashFailure</code> – (N/A because AH is not supported for this release) (3) <code>ahDuplicateSequence</code> – (N/A because AH is not supported for this release) (4) <code>ahOutOfRange</code> – (N/A because AH is not supported for this release) (5) <code>ahUnknownAlgorithm</code> – (N/A because AH is not supported for this release) (6) <code>espCryptoFailure</code> – An error occurred in encryption/decryption. For example a packet may be corrupted, or the header padding or alignment may be incorrect. (7) <code>espHashFailure</code> – The packet was decrypted but the hash comparison failed. (8) <code>espDuplicateSequence</code> – A duplicate IPSec sequence number was received. (9) <code>espOutOfRange</code> – The IPSec sequence number is out of range. (10) <code>espUnknownAlgorithm</code> – An unknown encryption algorithm was encountered. (11) <code>policyInvalidSpi</code> – An invalid security parameter index (SPI) value was received. (12) <code>policySAExpired</code> – The security association lifetime expired. (13) <code>policyClearText</code> – A clear text packet was received. (14) <code>isakmpFailure</code> – The IKE negotiations to generate a security association failed. This variable applies to only the <code>CryptoMibSAException</code> trap.
<code>PortLabel</code>	The descriptive label assigned to this port by the EOS device. These labels are used by the User Interface and are useful when comparing information retrieved via SNMP with information retrieved via the User Interface.

Variable	Description
PortName	The user-assigned name given to this port. No blanks are allowed in this name.
PeerAddress	The IP address of this SA's peer. This address is the same as either the <code>SourceAddress</code> or the <code>DestinationAddress</code> . This variable is reported as a string rather than an IP address to support IPv6 as well as IPv4 addresses.
LifetimeTime	The SA's lifetime, in seconds. A value of zero indicates no time restriction on the SA's lifetime; however, the SA may have a data throughput restriction on its lifetime, as indicated by the <code>LifetimeData</code> variable.
LifetimeData	The SA's lifetime, in kilobytes of data throughput. A value of zero indicates no data throughput restriction on the SA's lifetime; however, the SA may have a time restriction on its lifetime, as indicated by the <code>LifetimeTime</code> variable.
IkeClientId	The Identity Key Exchange (IKE) client ID. The IKE client ID identifies the local identity or remote identity attribute used during Internet Security Association Key Management protocol negotiation. Its value, in text form, can be an IPv4 node address, an IPv4 subnet, an IPv4 address range, a fully-qualified domain name, a fully-qualified user name, an encoded X.500 distinguished or general name, or a vendor-specific key identity.

Appendix F

Statistics Displays

To display statistics, use the `SHoW -SYS STATistics -<service> <option>` syntax. The statistics displayed are based on the time period in which they were gathered. For example, during the busiest minute, specified by the `SETDefault -SYS SampleOption` command, and on the time interval in which you want to see the statistics, specified by the `SETDefault -SYS SampleTime` command. A statistical display showing a blank in any field indicates that the service was not configured for the specified port.

You can use the `FLush -SYS STATistics -service` command to clear the current statistics. After you enter the `FLush` command, it may take several seconds before statistics sampling is restarted. For more information on the `FLush` command, see “SYS Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

When a 64-bit statistic value reaches the maximum value, it is not incremented further unless you flush it by issuing the `FLush -SYS STATistics` command. For details about the 64-bit counters supported by EOS see [Performance Statistics on page 687](#).

For more information on FR statistics, see “FR Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

EOS also supports retrieval of various statistics, including physical layer performance statistics, via SNMP. For details, see [Remote Polling Protocol Configuration on page 708](#).

F.1

ARP Service Statistics

The display generated by the `SHoW -SYS STATistics -ARP` command includes the following statistics:

Data Pkts Discarded

`Aged` – Number of packets discarded on a port as a result of timeouts waiting for a response.

`Queue Full` – Number of packets discarded on a port as a result of a full routing queue. The maximum allowable number of packets waiting for Internet address resolution is 10.

`Addr List Full` – Number of packets discarded on a port as a result of a full address list. The maximum allowable number of addresses waiting for resolution is 10.

`Data Pkts in Queue` – Number of packets on a port still in the queue waiting for address resolution.

Requests Received

`All Requests Rcvd` – Number of Internet address resolution requests received on a port.

`All Requests Rspd` – Number of responses to Internet address resolution requests on a port.

`Proxy Requests Rcvd` – Number of proxy requests received on a port for an Internet address not on the network where the request originated.

`Proxy Requests Rspd` – Number of responses to proxy requests responded to on a port.

`Agent Requests Rcvd` – Number of agent requests received on a port. An agent is a designated router that can respond to Address Resolution Protocol (ARP) requests for a PC; for example, when the ARP Service is not implemented.

`Agent Requests Rspd` – Number of responses to agent requests on a port.

Requests Sent

`All Requests Sent` – Number of Internet address resolution requests sent on a port since the last boot or since dynamic information in routing tables was last flushed.

Repeat Req. Sent – Number of Internet address resolution repeat requests sent on a port since the last boot or since dynamic information in routing tables was last flushed.

Refresh Req. Sent – Number of Internet address resolution refresh requests sent on a port since the last boot or since dynamic information in routing tables was last flushed.

InARP Statistics

InARP Request (out) – Number of InARP requests sent on a port.

InARP Response (in) – Number of InARP responses received on a port.

InARP Request (in) – Number of InARP requests received on a port.

InARP Response (out) – Number of InARP responses sent on a port.

Pkts Discarded – Number of packets discarded on a port.

RARP Statistics

Outgoing Requests – Number of outgoing Reverse Address Resolution Protocol (RARP) requests transmitted on a port by a RARP client.

Incoming Responses – Number of incoming RARP responses received on a port by a RARP client.

Incoming Requests – Number of incoming RARP requests received on a port by the RARP server.

Outgoing Responses – Number of outgoing RARP responses transmitted on a port by the RARP server.

F.2

BGP Service Statistics

The display generated by the `SHoW -SYS STATistics -BGP` command includes the following statistics:

BGP Statistics for All Peers – The following statistics apply to the entire EOS device:

Bytes in
 Bytes out
 Updates In
 Updates Out
 Notification Messages
 Notifications Out
 Keepalives in
 Keepalives Out
 Networks
 AS paths
 Network Policy discards
 AS policy discards

Per-peer Statistics – The following statistics apply to a particular peer:

InUpdates
 OutUpdates
 InMessages
 OutMessages

F.3

BRidge Service Statistics

The display generated by the `SHoW -SYS STATistics -BRidge` command includes the following statistics:

InFrames – Number of good incoming frames.
 InDiscards – Number of incoming discarded frames.

OutFrames – Number of outgoing good frames.
 OutDiscards – Number of outgoing discarded frames.
 MtuDiscards – Number of packets discarded as a result of exceeding maximum packet size. Mixed media configurations only.
 BCLDiscards – Number of packets discarded by the broadcast limit mechanism.
 BCLInvoked – Number of timer intervals in which the broadcast limit threshold was exceeded.
 IPFragmented – Total number of IP packets fragmented.
 IPFragments – Total number of IP fragments generated.

While InFrames and InDiscards counters are accurate, OutFrames and OutDiscards counters report the number of packets the bridge tried to forward. If the packets are dropped for other reasons after the bridge tried to forward it, these packets do not show up in the OutDiscards counters, but in the port statistics. InDiscards are a normal condition (that is, the destination media access control (MAC) address is on the same network segment as the bridge from which it was received).

If the display indicates that there are very few InDiscards compared to OutFrames on a LAN port, the LAN segments connected by your device bridge may not be evenly distributed. The bridge routinely forwards all or most of the frames received on that port. If the display indicates that there is an excessively large number of OutDiscards on a LAN port, the port may be highly saturated, the output port may not be forwarding, or the -Bridge BroadCastLimit parameter may not be set appropriately.

F.4 CRYPTO Service Statistics

The display generated by the `SHOW STATISTICS -CRYPTO` command includes the following statistics:

CRYPTO Statistics

Encrypted – The number of encrypted packets received on the port.
 Decrypted – The number of decrypted packets received on the port.
 Authenticated Out – The number of authenticated packets transmitted on the port.
 Authenticated In – The number of authenticated packets received on the port.
 Encrypt_AuthOut – The number of encrypted and authenticated packets transmitted on the port.
 Decrypt_AuthIn – The number of decrypted and authenticated packets received on the port.
 No Operation Out – The number of packets transmitted on the port on which no encryption or authentication was performed.
 No Operation In – The number of packets received on the port on which no encryption or authentication was performed.
 New Key Late – The number of packets for which the rekey process was started but not completed.
 Compressed – The number of compressed packets received on the port.
 Decompressed – The number of decompressed packets received on the port.
 Compression Ratio (%) – The compression ratio used on the port.

Cause of Pkt Discard:

Xmit Policy – The number of packets discarded due to nonconformance to the configured transmit policy.
 Rcv Policy – The number of packets discarded due to nonconformance to the configured receive policy.
 Anti Replay – The number of discarded packets detected by anti replay.
 Dest Unreachable – The number of packets discarded due to an unreachable destination IP address.
 ICV Failed – Not used.
 Bad Pad Length – The number of packets discarded due to an invalid pad length.
 Encrypt Failed – The number of packets discarded due to failed encryption.

Decrypt Failed – The number of packets discarded due to failed decryption.
 AuthOut Failed – The number of outgoing packets discarded due to failed authentication.
 AuthIn Failed – The number of incoming packets discarded due to failed authentication.
 Comp Failed – The number of packets discarded due to failed compression.
 Decompress Failed – The number of packets discarded due to failed decompression.
 Out of Memory – The number of packets discarded due to insufficient device memory.
 Seq Number Cycled – The number of packets discarded due to a cycled sequence number.
 Unsupported IP Option – The number of packets because they were not supported by AH.
 Bad Pkt In – The number of packets discarded due to a bad ESP, AH, or IP header.
 CE Queue Overflow – The number of packets discarded due to a full transmit queue.
 Other Xmit Errors – The number of packets discarded to unspecified transmission errors.
 DCS Queue Overflow – The number of packets that were not compressed due to a queueing failure.
 Other Recv Errors – The number of packets discarded to unspecified reception errors.

Cause of NonCompress:

Threshold – The number of packets that were not compressed because their length was less than the threshold (90).
 Expanded – Not used.

IKE Stats:

Recv Err-Unknown Peer – The total number of packets received on this interface with an unknown peer.
 Recv Err-Unknown SPI – The total number of packets received on this interface with an unknown SPI.
 Xmit Err-Unk Peer/SPI – The total number of packets transmitted on this interface with an unknown peer and/or SPI.
 IKE up/FRF17 DIS – The total number of packets received on this interface with IKE up and FRF.17 disabled. This state is encountered when FRF.17 changed state from enabled to disabled, and IKE is in the process of changing state from up to down.
 IKE up/FRF 17, EN, NT, OP – The total number of packets received on this interface with IKE up and FRF.17 enabled but not operational. This state is encountered when IKE changed state from down to up, and FRF.17 is in the process of changing state from not operational to operational.

F.5

DVMRP Service Statistics

The display generated by the `SHoW -SYS STATistics -DVMRP` command includes the following statistics:

DVMRP Statistics

Rcvd from MOSPF – Total number of packets received from Multicast Open Shortest Path First Protocol (MOSPF) domains.
 Sent to MOSPF – Total number of packets transmitted to MOSPF domains.
 Pruned by MOSPF – Total number of packets pruned by MOSPF domains for no listener in the MOSPF domains.
 Graft from MOSPF – Number of packets grafted by MOSPF domains for listeners in the MOSPF domain.
 Sent by local client – Number of packets sent by the local client-a directly-attached (normally via LAN) multicast station.
 Transmit error – Number of packets that could not be transmitted due to transmit errors.

Pkts Received – Total number of DVMRP packets received.

Reports – Number of route update packets received.
 Prunes – Number of Prune packets received from a downstream neighbor router.
 Grafts – Number of Graft packets received from a downstream neighbor router.

Graft Acks – Number of Graft Acknowledge packets received from the upstream parent router.

Pkts Transmitted – Total number of DVMRP packets transmitted.

Reports – Number of route update packets transmitted.

Prunes – Number of Prune packets transmitted to the upstream parent router.

Grafts – Number of Graft packets transmitted to the upstream parent router.

Graft Acks – Number of Graft Acknowledge packets transmitted to any downstream neighbor router.

Pkts Forwarded – Total number of packets forwarded.

Pkts Discarded

NoRoute – Number of packets received when there is no route to the source.

WrongPort – Number of packets received on a port that is not used to forward to the source.

Unknown Type – Number of packets received that are of unknown type.

MiscErr – Number of packets received when the system is out of resources.

IP over IP Statistics

Pkts Received – Total number of IP-over-IP packets received.

Pkts Discarded – Total number of IP-over-IP packets discarded.

F.6

FR Service Statistics

The display generated by the `SHOW -SYS STATistics -FR` command includes the following statistics:

Frame Relay Statistics

DTE LMI Frames Xmit – The number of DTE status enquiry packets sent on this interface.

DCE LMI Frames Xmit – The number of DCE status enquiry packets sent on this interface.

DTE LMI Frames Recv – The number of DTE status enquiry packets received on this interface.

DCE LMI Frames Recv – The number of DCE status enquiry packets received on this interface.

DTE LMI LIV Time Outs – The number of times the status enquiry message was not received on this interface within the T391 keepalive timer, configured with the `KATime` parameter.

DTE LMI SEQ Errs – The total number of times a loss of sequencing in received LMI messages was detected by the DTE on this interface.

DCE LMI POLL Time Outs – The number of times the status enquiry message was not received on this interface within the T392 timer, configured with the `PVTime` parameter.

DCE LMI SEQ Errs – The total number of times a loss of sequencing in received LMI messages was detected by the DCE on this interface.

DCE LMI Protocol Errs – The total number of times protocol errors in received LMI messages were detected by the DCE on this interface.

Invalid DLCI frames – The number of invalid DLCI frames.

Small Frames Recv – The number of small frames successfully received.

Port Status Change – The number of times the underlying port was brought *down* after LMI exhausts its retry attempts. This statistic does not increment when the port comes back up. The Port Status Change statistic does not increment when you *administratively* bring the port down.

SVC Frames Xmit – The number of Frame Relay frames transmitted for Frame Relay SVCs.

SVC Frames Recv – The number of Frame Relay frames received for Frame Relay SVCs.

F.7

GGSN Service Statistics

You can display the following GGSN performance statistics:

- **GTP statistics for the GGSN service as a whole.**
To view statistics for the GGSN service as a whole, you must use:
`SHoW -GGSN GTPSTATistics`
- **Performance statistics for a specific APN.**
To view performance statistics for a particular APN (CEN), you must use:
`SHoW !<APN> -GGSN CENSTATistics`
where *<APN>* is the access point name.
To view performance statistics for all APNs, you must use:
`SHoW -GGSN CENSTATistics`
- **RADIUS statistics**
To view GGSN RADIUS statistics, you must use:
`SHoW -GGSN AAASSTATistics`
If you want to view RADIUS statistics for a particular APN, you must use:
`SHoW !<APN> -GGSN AAASSTATistics`
where *<APN>* is the name of the APN for which you want to view statistics.
To flush GGSN RADIUS statistics, you must use:
`FLush -GGSN AAASSTATistics`
- **Charging statistics**
To view GGSN charging statistics, you must use:
`SHoW -GGSN CHGSTATistics`

F.7.1

GTP Statistics for the GGSN Service as a Whole

The system returns the following global GTP statistics:

GTP Peer IP Address – The IP address of the GTP tunnel peer.

State – The operational status of the GTP tunnel (UP or DOWN).

Number of PDP Contexts – The currently-configured maximum number of PDP contexts supported on the GGSN.

Rcvd Control Packets – The total number of control packets received from all PDGs (Packet Data Gateways).

Sent Control Packets – The total number of control packets sent to any PDG.

Rcvd Data Packets – The total number of data packets received from all PDGs.

Rcvd Data Bytes – The total number of data bytes received from all PDGs.

Sent Data Packets – The total number of data packets sent to any PDG.

Sent Data Bytes – The total number of data bytes sent to any PDG.

F.7.2

APN Statistics

The system returns the following statistics for each APN:

User Data

Sent Packets – The total number of packets sent to the CEN through this APN.

Rcvd Packets – The total number of packets received from the CEN through this APN.

Sent Bytes – The total number of bytes sent to the CEN through this APN.

Rcvd Bytes – The total number of bytes received from the CEN through this APN.

DDNS

Sent DDNS Update Req – The number of DDNS update request messages sent to the DNS server in the CEN.

Rcvd DDNS Update Rsp – The number of DDNS update response messages received from the DNS server in the CEN.

Rcvd DDNS Update Err – The number of received DDN update responses which indicate error in the DNS update.

ReXmt DDNS Update Req – The number of retransmissions for DDNS update request messages.

DDNS Update Timeout – The number of DDNS update request messages that were retransmitted and eventually timed out.

DHCP

Sent DHCP Discover – The number of DHCP discover messages sent to the DHCP server in the CEN.

Rcvd DHCP Offer – The number of DHCP offer messages received from the DHCP server in the CEN.

Sent DHCP Req – The number of DHCP request messages sent to the DHCP server in the CEN.

Rcvd DHCP Ack – The number of DHCP Ack messages received from the DHCP server in the CEN.

Sent DHCP Release – The number of DHCP release messages sent to the DHCP server in the CEN.

Rcvd DHCP Rsp Err – The number of received DHCP response messages which indicate error.

DHCP Discover Timeout – The number of DHCP discover messages that were retransmitted and eventually timed out.

DHCP Req Timeout – The number of DHCP request messages that were retransmitted and eventually timed out.

DHCP Switchover – The number of switchovers from primary to secondary DHCP server or vice versa.

DHCP – The IP address of the current external DHCP server.

RADIUS Auth. – The IP address of the current RADIUS authentication server.

RADIUS Acct. – The IP address of the current RADIUS accounting server.

F.7.3

RADIUS Statistics

Authentication

Switchovers – The number of switchovers from primary to secondary RADIUS authentication server, or vice versa.

Accepts – The number of authentication accept messages received from the RADIUS server.

Rejects – The number of authentication reject messages received from the RADIUS server.

Timeout – The number of context creation failures caused by an authentication timeout.

Retries – The number of authentication message retransmissions.

Max RoundTrip – The maximum response time from the first authentication request message to the time when the response is received. This may include retransmissions.

Accounting

Switchovers – The number of switchovers from primary to secondary RADIUS accounting server, or vice versa.

Success – The number of accounting start response messages received from the RADIUS server.

Failure – The number of failures on the accounting start response messages received from the RADIUS server.

Timeout – The number of timeouts on accounting start request messages.

Duplicated IP – The number of RADIUS-allocated IP addresses which are duplicated in the GGSN.

Max RoundTrip – The maximum response time from the first accounting start request message to the time when the start response is received. This may include retransmissions.

F.7.4

GGSN Charging Statistics

The system returns the following charging statistics:

Current Charging Gateway – The IP address of the charging gateway that is currently in operation.

rcvd Node Alive Request – The number of node alive request messages received from the charging gateway.

sent Node Alive Response – The number of node alive response messages sent to the charging gateway.

rcvd Redirection Request – The number of redirection request messages received from the charging gateway.

sent Redirection Response – The number of redirection response messages sent to the charging gateway.

sent Data Record Transfer Request – The number of data record transfer (DRT) request messages sent to the charging gateway.

rcvd Data Record Transfer Response – The number of data record transfer (DRT) response messages received from the charging gateway.

sent Echo Request (to Charging Gateway) – The number of GTP' echo request messages sent to the charging gateway.

rcvd Echo Request (from Charging Gateway) – The number of GTP' echo request messages received from the charging gateway.

sent Echo Response (to Charging Gateway) – The number of GTP' echo response messages sent to the charging gateway.

rcvd Echo Response (from Charging Gateway) – The number of GTP' echo response messages received from the charging gateway.

rcvd Unknown Charging Messages – The number of unknown GTP' messages received from the charging gateway.

F.8

IP Service Statistics

The display generated by the `SHoW -SYS STATistics -IP` command includes the following statistics:

Datagram Rates (pkts/s)

Rcvd from Net – The total number of input datagrams received per second on a port including those received in error.

Txmit to Net – The total number of IP-datagrams transmitted per second to the port.

IP Datagrams

NORMAL PATH

Totals

Rcvd from Net – The total number of input datagrams received on a port including those received in error.

Rcvd and Fwd – The number of input datagrams for which this port was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination.

Rcvd by Local Client – The number of input datagrams passed on to the local client.

Txmit to Net – The number of datagrams transmitted on the network through this port. The statistics collected here are the sum of the statistics collected in "Txmit by Forwarding" and "Txmit by Local Client" as described below.

Txmit by Forwarding – The total number of IP-datagrams transmitted to the port. Note that this statistic is different from the "Txmit by Local Client" statistic described below in the sense that it only includes those datagrams which were transmitted to the EOS device by various other routers in the network because this device happens to be the device in the path to the final destination. This counter does not include any datagrams maintained in the "Originated by system" statistics.

Txmit by Local Client – The total number of IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission.

Unicasts

Rcvd from Net – The total number of unicast datagrams received on a port including those received in error.

Rcvd and Fwd – The number of unicast datagrams for which this port was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination.

Txmit to Net – The number of unicast datagrams transmitted on the network through this port. The statistics collected here are the sum of the statistics collected in "Txmit by Forwarding" and "Txmit by Local Client" as described below.

Txmit by Forwarding – The total number of unicast IP-datagrams transmitted to the port. Note that this statistic is different from the "Txmit by Local Client" statistic described below in the sense that it only includes those datagrams which were transmitted to the EOS device by various other routers in the network because this device happens to be the device in the path to the final destination. This counter does not include any datagrams maintained in the "Originated by system" statistics.

Txmit by Local Client – The total number of unicast IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission.

Rcv Tunnelled – Number of unicast IP datagrams received using IP-IP tunneling.

Multicasts

Rcvd from Net – The total number of multicast datagrams received on a port including those received in error.

Rcvd and Fwd – The number of multicast datagrams for which this port was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination.

Txmit to Net – The number of multicast datagrams transmitted on the network through this port. The statistics collected here are the sum of the statistics collected in "Txmit by Forwarding" and "Txmit by Local Client" as described below.

Txmit by Forwarding – The total number of multicast IP-datagrams transmitted to the port. Note that this statistic is different from the "Txmit by Local Client" statistic described below in the sense that it only includes those datagrams which were transmitted to the EOS device by various other routers in the network because this device happens to be the device in the path to the final destination. This counter does not include any datagrams maintained in the "Originated by system" statistics.

Txmit by Local Client – The total number of multicast IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission.

Broadcasts

Rcvd from Net – The total number of broadcast datagrams received on a port including those received in error.

Rcvd and Fwd – The number of broadcast datagrams for which this port was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination.

Txmit to Net – The number of broadcast datagrams transmitted on the network through this port. The statistics collected here are the sum of the statistics collected in "Txmit by Forwarding" and "Txmit by Local Client" as described below.

Txmit by Forwarding – The total number of broadcast IP-datagrams transmitted to the port. Note that this statistic is different from the "Txmit by Local Client" statistic described below in the sense that it only includes those datagrams which were transmitted to the EOS device by various other routers in the network because this device happens to be the device in the path to the final destination. This counter does not include any datagrams maintained in the "Originated by system" statistics.

Txmit by Local Client – The total number of broadcast IP datagrams which local IP user-protocols (including ICMP) supplied to IP in requests for transmission.

OPTIMIZED PATH

Totals

Rcvd and Fwd – The total number of datagrams received through fast path on a port including those received in error.

Txmit by Forwarding – The total number of multicast IP-datagrams transmitted to the port through fast path.

Rcv Tunnelled – The total number of IP packets received using the optimized path.

Xmit Tunnelled – The total number of IP packets transmitted using IP-IP tunneling and the optimized path.

Multicasts

Rcvd from Net – The total number of multicast datagrams received through fast path on a port including those received in error.

Rcvd and Fwd – The number of multicast datagrams receive through the fast path for which this port was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination.

Txmit by Forwarding – The total number of multicast IP-datagrams transmitted to the port through fast path.

IP Fragmentation

Datagrams Fragmented – The number of IP datagrams that were successfully fragmented at this port.

Fragments Generated – The number of IP datagram fragments that were generated as a result of fragmentation at this port.

Fragmentation Failure – The number of IP datagrams that were discarded because they needed to be fragmented at this port but could not be; for example, because their Don't Fragment flag was set.

Fragments Received – The number of IP fragments received which needed to be reassembled at this entity.

Datagrams Assembled – The number of IP datagrams successfully reassembled.

Reassembly Failures – The number of failures detected by the IP reassembly algorithm (for whatever reason: timed out, errors, and so on).

Errors

Filtering Discards – The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded due to the filter rules; for example, Firewall denying certain IP packets depending on the filter rules configured by the user.

Rcvd Bad Header – The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, and errors discovered in processing their IP options

Rcvd Bad IP Addr – The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this port. This count includes invalid addresses (e.g., 0.0.0.0) and addresses of unsupported Classes (e.g., Class E).

Rcvd Unknown Proto - The number of locally-addressed datagrams discarded because of an unknown or unsupported protocol.

Other Receive Errs - The number of input datagrams discarded because of the errors which do not belong to any of the above categories. This counter is incremented under any of the following conditions:

Received fragmented UDP broadcast datagrams addressed to the EOS device.

Directed broadcast packet is received by the EOS device and **NoFwdSubnetBcast** set in the **IP CONTRol** parameter.

Received an unknown type of ICMP packet.

Incoming datagram could not be delivered to the local client because of mailbox error (for example, the mailbox was full).

EOS device received a MAC broadcast packet and the destination IP address in the IP header was not a directed broadcast address.

Route Lookup Failed - The number of IP datagrams discarded because no route could be found to transmit them to their destination.

Invalid IP option - The number of IP datagrams which could not be transmitted because of invalid options in the packet; for example, the options were either not supported or error occurred while processing the packet.

TTL expired - The number of IP datagrams which could not be transmitted because of invalid (0) value in the TTL field.

Buffer Error - The number of IP datagrams which could not be transmitted because of the buffer unavailability.

Other Transmit Errs - The number of output IP datagrams which could not be transmitted on the network due to any of the following:

Bad IP header in the transmit packet.

TTL field is 0

Invalid option specified

Lower layer did not transmit for whatever reasons.

ICMP Messages

Totals

Received Messages - The total number of ICMP packets received on the port.

Messages Discarded - The number of ICMP messages which the port received but determined as having ICMP specific errors (bad ICMP checksums, bad length, and so on).

Transmit Messages - The total number of ICMP messages transmitted on the port.

Transmit Failures - The number of ICMP messages which this port did not send due to problems discovered within ICMP such as the lack of buffers.

ICMP Queries

Echo Txmit - The number of ICMP echo (request) messages transmitted.

Echo Rcvd - The number of ICMP echo (request) messages received.

Addr Mask Txmit - The number of ICMP Address Mask Request messages transmitted.

Addr Mask Rcvd - The number of ICMP Address Mask Request messages received.

RDP Solicits Txmit - The number of RDP solicitation messages transmitted on the port.

RDP Solicits Rcvd - The number of RDP solicitation messages received on the port.

ICMP Responses

Echo Reply Txmit - The number of ICMP echo reply messages transmitted.

Echo Reply Rcvd - The number of ICMP echo reply messages received.

Addr Mask Reply Txmit - The number of ICMP Address Mask Reply messages transmitted on the port.

Addr Mask Reply Rcvd - The number of ICMP Address Mask Reply messages received on the port.

RDP Advertise Txmit - The number of RDP Advertisement messages transmitted on the port.

RDP Advertise Rcvd – The number of RDP Advertisement messages received on the port.

ICMP Redirects

Redirects Txmit – The number of ICMP Redirect messages transmitted on the port.

Redirects Rcvd – The number of ICMP Redirect messages received on the port.

ICMP Dest Unreach

Des Unreach Txmit – The number of ICMP Destination Unreachable transmitted on the network.

Dest Unreach Rcvd – The number of ICMP Destination Unreachable received

ICMP TTL Msgs

Time Exceed Txmit – The number of ICMP time exceeded messages transmitted.

Time Exceed Rcvd – The number of ICMP time exceeded messages received.

ICMP TimeStamp Msgs

Time Stamp Txmit – The number of ICMP TimeStamp messages transmitted.

Time Stamp Rcvd – The number of ICMP TimeStamp messages received.

BFD Msgs

BFD Control Txmit – The number of BFD control messages transmitted.

BFD Control Rcvd – The number of BFD control messages received.

BFD Reply Rcvd – The number of BFD reply messages received.

BFD Reply Txmit – The number of BFD reply messages transmitted.

BFD Control Drops – The number of BFD control messages dropped.

F.9

IPv6 Service Statistics

The display generated by the `SHoW -SYS STATistics -IPv6` command includes the following IPv6 statistics:

IPv6 Data Messages

Received from Net

Forwarded

Passed to Clients

Sent to Clients

Passed to ICMPv6

ICMPv6 Inf Messages

Echo Request Sent

Echo Request Rcv

Echo Reply Sent

Echo Reply Rcv

ICMPv6 Error Messages

DestUnreach Sent – The number of ICMPv6 destination unreachable error messages sent.

DestUnreach Rcv – The number of ICMPv6 destination unreachable error messages received.

PktTooBig Sent – The number of ICMPv6 packet too big error messages sent.

PktTooBig Rcv – The number of ICMPv6 packet too big error messages received.

TimeExceeded Sent – The number of ICMPv6 time exceeded error messages sent.

TimeExceeded Rcv – The number of ICMPv6 time exceeded error messages received.

ParamProb Sent – The number of ICMPv6 parameter problem error messages sent.

ParamProb Rcv – The number of ICMPv6 parameter problem error messages received.

RateLimit Sent drops – The number of packets that were dropped because the ICMP rate limit was exceeded.

NbrDiscovery Messages

NbrSol/MC Sent
NbrSol/UC Sent
NbrSol Rcv
NbrAdv Sent
NbrAdv Rcv
RtrAdv Sent
RtrAdv Rcv
RtrSol Sent
RtrSol Rcv
Redirect Sent
Redirect Rcv

Messages Discarded

Bad Header
Invalid Dest Addr
Invalid Src Addr
Link Local MC/Src
Hop Limit Exceeded
Unknown Next Hdr
Unknown Option
No Router
NbrSol Queue Full
Routing Disabled
Checksum Errors
Receive Errors
Forwarding Errors
Transmit Errors
Bad Tunnel Pkts
Bad Ipv6 Tunnel Pkts – The number of invalid IPv6 tunnel packets received.
Fragmentation Failur – The number of IPv6 fragmentation failures.
Fragmentation OK – The number of IPv6 packets that were successfully fragmented.
Reassemble Req Pkts – The number of IPv6 fragments received that required reassembly.
Reassemble Req Fail – The number of IPv6 fragments received that could not be reassembled.
Reassemble Req OK – The number of IPv6 packets that were reassembled successfully.

Rcv IPv6 Tunnelled – The number of valid IPv6 tunnel packets received.

F.10

MIP Service Statistics

The display generated by the `SHoW -SYS STATistics -MIP` command includes the following statistics:



NOTE: IGMP is not a service in the user interface; however, you can still obtain IGMP statistics through the MIP Service.

Multicast IP Datagram (Total) – The total number of multicast IP datagrams received.

Rcvd from Net – The number of multicast IP datagrams received from the network on each port.

Rcvd and Fwd – The number of multicast IP datagrams received and forwarded on each port.

Rcvd and Discard – The number of multicast IP datagrams received and discarded on each port.

Rcv: Pkt/Sec – The number of multicast IP datagrams received per second on each port.

Transmit to Net – The number of multicast IP datagrams transmitted to the network on each port.

Xmit: Pkt/Sec – The number of multicast IP datagrams transmitted per second on each port.

Discard by TTL expiry – The number of multicast data packets discarded by the EOS device on each port because the TimeToLive value expired for the packet. This is a possible indication of a loop in the network.

Txmit by Local Client – The number of multicast IP datagrams transmitted by the local client (normally a client directly attached to the EOS device, for example via the LAN) on each port.

IGMP Pkts Rcvd (Total) – The total number of Internet Group Management Protocol (IGMP) packets received, including DVMRP packets.

Queries – The number of Host Query packets received on each port.

Reports – The number of Host Report packets received on each port.

Leaves – The number of Host Leave Group packets received on each port.

IGMP Pkts Txmit – The total number of IGMP packets transmitted, including DVMRP packets, on each port.

IGMP Pkts Discard (Total) – The total number of IGMP packets discarded.

TooShort – The number of packets received with the data length too short on each port.

Version Err – The number of packets received with a bad version number on each port.

Chksum Err – The number of packets received with bad checksums on each port.

Unknown Type – The number of packets received with unknown type on each port.

Rejected by Limit – The number of packets rejected on each port because the configured multicast group limit was met or exceeded.

F.11

MOBILEIP Service Statistics

System-level statistics are maintained for the MOBILEIP service. To view these statistics, enter the following command from either the foreign agent (GGSN) or the home agent (border router):

```
SHoW -SYS STATistics -MOBILEIP
```

The display generated by the SHoW -SYS STATistics -MOBILEIP command on the foreign agent includes the following statistics:

Mobile IP Encapsulated Pkts:

Total Xmit Pkts – the total number of Mobile IP encapsulated packets transmitted from the home agent

Total Xmit SimulBind Pkts – the total number of simulcast binding packets transmitted from the home agent

Total Rcvd Pkts – the total number of Mobile IP encapsulated packets received on the foreign agent

Mobile IP Registration Pkts:

Total Xmit Reg-Request – the total number of mobile node registration requests transmitted from the foreign agent

Total Rcvd Reg-Request – the total number of mobile node registration requests received on the home agent

Total Rcvd Valid Regs – the total number of valid registrations: mobile node registration requests received on the home agent and notifications of a node being added sent from the active replicator to its replication peer

Total Rcvd Reg Delete – the total number of received registration deletes: mobile node de-registration requests received on the home agent; registration requests rejected on the foreign agent, and accepted mobile node de-registration requests

Total Xmit Reg-Reply – the total number of replies to mobile node registration requests transmitted from the home agent

Total Rcvd Reg-Reply – the total number of replies to mobile node registration requests received on the foreign agent

Total Rcvd Accepted Reg-Reply – the total number of accepted mobile node registration replies received on the foreign agent

Total Rcvd Rejected Reg-Reply – the total number of rejected mobile node registration replies received on the foreign agent

Total Half Lifetime ReReg – the total number of registration requests sent out for renewal

Total Expired Reg – the total number of mobile registrations that expired (the configured MblNodeLifetime expired)

Mobile IP Registration Errors:

Total Poorly Formed Reg – the total number of poorly formed registration requests received on the home agent.

Identification Mismatches – the total number of registration requests received on the home agent with mismatched IDs

Requests from UnKnown CoA – the total number of registration requests received on the home agent with unknown Care of Addresses (CoAs)

Total Receive Errors – the total number of registration request packets received with errors on the home agent

Total Reg Timeouts – the total number of mobile node registration errors on the foreign agent that were attributed to registration timeouts

Total DeReg Timeouts – the total number of mobile node registration errors on the foreign agent that were attributed to de-registration timeouts

Total Reg Dropped – the total number of mobile node registrations that were dropped on the foreign or home agent

Total DeReg Dropped – the total number of mobile node de-registrations that were dropped on the foreign or home agent

Total Dereg COA Mismatch – the total number of de-registration requests received on the home agent with a Care of Address (CoA) that did not match any in the home agent's memory

Total No Memory Errors – the total number of mobile node registration errors on the foreign or home agent that were attributed to memory issues

Mobile IP Replication:

Total Master with Peer – the total number of times that the local device became the VRRP master with a replication peer

Total Master w/o Peer – the total number of times the local device became the VRRP master without a replication peer

Total Standby – the total number of times the local device became the VRRP backup

Total Bulk – the total number of times the local device (either the home agent or the foreign agent) received a start of bulk update packets to the home agent mobility database or foreign agent visitor table

Total Single Add – the total number of times the local device (either the home agent or the foreign agent) sent or received a single addition to the home agent mobility database or foreign agent visitor table

Total Single Add Fail – the total number of failed attempts by the local device (either the home agent or the foreign agent) to send or receive a single addition to the home agent mobility database or foreign agent visitor table

Total Single Delete – the total number of times the local device either the home agent or the foreign agent) sent or received a single deletion to the home agent mobility database or foreign agent visitor table

Total Single Delete Fail – the total number of failed attempts by the local device (either the home agent or the foreign agent) to send or receive a single deletion to the home agent mobility database or foreign agent visitor table

Total Flush – the total number of times the local device (either the home agent or the foreign agent) received a replication Flush message for the home agent mobility database or foreign agent visitor table

Total Full – the total number of times the local device (either the home agent or the foreign agent) received a replication database Full message for the home agent mobility database or foreign agent visitor table

Total Reset Sent – the total number of times the local device (either the home agent or the foreign agent) sent a replication reset message

Mobile IP BFD Events:

Total Restart Sent – the total number of restart messages sent from the Mobile IP client on the home agent to the BFD service

Total Restart Rcvd – the total number of restart messages received from the BFD service on the Mobile IP client on the foreign agent

Mobile IP Client Reqs – the total number of requests sent by Mobile IP, as a client of GTP, to the foreign agent to create a mobile node registration request

Total Add Request – the total number of requests received from GTP on to add a mobile node or new mobile node registration requests received on the foreign agent

Total Add Fail – the total number of times that add requests on the foreign agent failed

Total Add Response – the total number of times that a registration request received a reply that was accepted by the home agent; this statistic is reported on the foreign agent only

Total Del Request – the total number requests received from GTP on the foreign agent to delete a mobile node or send a Mobile IP de-registration request to the home agent

Total Del Fail – the total number of failed Mobile IP de-registration requests sent from the foreign agent to the home agent

Total Del Response – the total number of times that a de-registration request sent from the foreign agent to the home agent passed and a notification was sent to GTP

Total Bad Request – the total number of invalid request types received from the Mobile IP client (GGSN) and requests that contain an invalid home agent /home address.

Total Bad Config – the total number of times that the foreign agent received a registration request request when no Care of Address (CoA) or BFD session to the home agent was configured

Total Invalid COA – the total number of invalid deregistration errors sent to the mobile IP client (foreign agent on the GGSN)

Total Flush Node Req – the total number of times the foreign agent received a request from GTP that caused a mobile node to be de-registered from the home agent mobility database or foreign agent visitor table

Total Flush Fail – the total number of failed requests from GTP to de-register a mobile node from the home agent mobility database or foreign agent visitor table; this statistic is reported on the foreign agent only

Total Flush All – the total number of times the foreign agent received a message from GTP that caused all mobile nodes to be de-registered from the home agent mobility database or foreign agent visitor table

Total Standby Rcv Err – The total number of times that the standby foreign agent received a registration request from the client (GTP)

The display generated by the `SHoW -SYS STATistics -MOBileIP` command on the home agent includes the following statistics:

Mobile IP Encapsulated Pkts:

Total Xmit Pkts – the total number of Mobile IP encapsulated packets transmitted from the home agent

Total Xmit SimulBind Pkts – the total number of simulcast binding packets transmitted from the home agent

Total Rcvd Pkts – the total number of Mobile IP encapsulated packets received on the home agent

Mobile IP Registration Pkts:

Total Xmit Reg-Request – The total number of mobile node registration requests transmitted from the home agent.

Total Rcvd Reg-Request – The total number of mobile node registration requests received on the home agent.

Total Rcvd Valid Regs – The total number of valid mobile node registration requests received on the home agent.

Total Rcvd Reg Delete – The total number of mobile node de-registration requests received on the home agent.

Total Xmit Reg-Reply – The total number of replies to mobile node registration requests transmitted from the home agent.

Total Rcvd Reg-Reply – The total number of replies to mobile node registration requests received on the home agent

Total Rcvd Accepted Reg-Reply – The total number of accepted replies to mobile node registration requests received on the home agent.

Total Rcvd Rejected Reg-Reply – The total number of rejected replies to mobile node registration requests received on the home agent.

Total Half Lifetime ReReg – the total number of mobile node registrations that were renewed because half of the configured lifetime (MblNodeLifetime) expired.

Total Expired Reg – the total number of mobile registrations that expired (the configured MblNodeLifetime expired).

Mobile IP Registration Errors:

Total Poorly Formed Reg – the total number of mobile node registration errors on the home agent that were attributed to improperly formatted registration requests

Identification Mismatches – the total number of mobile node registration errors on the home agent that were attributed to home agent address mismatches

Requests from UnKnown CoA – the total number of mobile node registration errors on the home agent that were attributed to an unknown COA

Total Receive Errors – the total number of mobile node registration errors on the home agent that were attributed to receive errors on the home agent

Total Reg Timeouts – the total number of mobile node registration errors on the home agent that were attributed to registration timeouts

Total DeReg Timeouts – the total number of mobile node registration errors on the home agent that were attributed to de-registration timeouts

Total Reg Dropped – the total number of mobile node registrations that were dropped on the home agent

Total DeReg Dropped – the total number of mobile node de-registrations that were dropped on the home agent

Total Dereg COA Mismatch – the total number of mobile node registrations that were de-registered on the home agent due to a COA mismatch

Total No Memory Errors – the total number of mobile node registration errors on the home agent that were attributed to memory issues

Mobile IP Replication:

Total Master with Peer – the total number of times that the local device became the VRRP master with a replication peer

Total Master w/o Peer – the total number of times the local device became the VRRP master without a replication peer

Total Standby – the total number of times the local device became the VRRP backup

Total Bulk – the total number of times the local device sent a bulk home agent mobility table update to its replication peer

Total Single Add – the total number of times the local device sent a single addition to the home agent mobility table to its replication peer

Total Single Add Fail – the total number of failed attempts by the local device to send a single addition to the home agent mobility table to its replication peer

Total Single Delete – the total number of times the local device sent a single deletion to the home agent mobility table to its replication peer

Total Single Delete Fail – the total number of failed attempts by the local device to send a single deletion to the home agent mobility table to its replication peer

Total Flush – the total number of times the local device sent a replication Flush message for the home agent mobility table

Total Full – the total number of times the local device sent a replication database Full message for the home agent mobility table

Total Reset Sent – the total number of times the local device sent a replication reset message

F.12

MOSPF Service Statistics

The display generated by the `SHoW -SYS STATistics -MOSPF` command includes the following statistics:

MOSPF Statistics

Cache flushed – Number of times forwarding cache was flushed, possibly due to network topology changes.

SPF calculations – Number of route computations, using Dijkstra's algorithm.

Rcvd from DVMRP – Total number of packets received from DVMRP domains.

Sent to DVMRP – Total number of packets transmitted to DVMRP domains.

Prune back DVMRP – Number of Prune messages sent back to DVMRP because received multicast packets have no listener in the MOSPF domain.

Resource error – Number of failures in obtaining memory resources.

Internal Client Transmit

Internal Client Receive

Internal Client Errors

Internal GrLSA Dropped Errors

Receive

Good – Number of good receive packets.

No Route – Number of packets discarded due to no recipient.

Bad MAC Address – Number of packets for which a unicast or multicast MAC address expected, but the incoming packets were not using the correct MAC address format.

Wrong Source – Number of packets that arrived from an upstream neighbor other than the correct upstream neighbor.

TTL exceeded – Number of packets for which the time-to-live value was exceeded on the receive packet.

Recv Disabled – Number of packets received for which MOSPF is disabled on the interface.

GroupLSA dropped – Number of Group LSA (link state advertisement) packets dropped for any reason.

Transmit

Good – Number of good transmit packets.

TTL too small – TTL too small to reach any of the downstream destinations, or TTL smaller than the Threshold (in the MIP service) parameter.

Xmit Disabled – MOSPF is disabled on the interface.

F.13

NAT Service Statistics

NAT replication statistics are maintained on both a per-device and per-port basis. To view these statistics, enter the following command:

```
SHoW -SYS STATistics -NAT
```

The display generated by the `SHoW -SYS STATistics -NAT` command includes the following port-level replication statistics:

Send update – the number of times the local device sent a TCP connection creation request or state update on this port.

Send update fail – the number of failed attempts by the local device to send a TCP connection request or state update on this port.

Send delete – the number of times that the local device sent a TCP connection closure request on this port.

Send delete fail – the number of failed attempts by the local device to send a TCP connection closure request on this port.

Rcv update – The number of times the local device received a TCP connection creation request or state update on this port.

Rcv update fail – the number of failed attempts by the local device to receive a TCP connection request or state update on this port.

Rcv delete – the number of times that the local device received a TCP connection closure request on this port.

Rcv delete fail – the number of failed attempts by the local device to receive a TCP connection closure request on this port.

The display generated by the `SHoW -SYS STATistics -NAT` command includes the following device-level replication statistics:

Master with peer – the number of times that the local device became the VRRP master with a replication.

Master without peer – the number of times the local device became the VRRP master without a replication peer.

Standby – the number of times the local device became the VRRP backup.

Send Bulk – the number of times the local device sent a bulk NAT sessions table update.

Send Bulk Fail – the number of failed attempts by the local device to send a bulk NAT sessions table update.

Send Bulk End – the number of times the local device sent a bulk NAT sessions table update End message.

Send Bulk End Fail – the number of failed attempts by the local device to send a bulk NAT sessions table update End message.

Send Reset – the number of times the local device sent a replication reset message.

Send Reset Fail – the number of failed attempts by the local device to send a replication reset message.

Flush – the number of times the local device sent a replication Flush message for the NAT sessions table.

Full – the number of times the local device sent a replication database Full message for the NAT sessions table.

F.14

OSPF Service Statistics

The display generated by the `SHoW -SYS STATistics -OSPF` command includes the following statistics:

OSPF Statistics

- `SPF calculations` – Number of times the EOS device performed the SPF calculation.
- `Resource error` – Number of times OSPF failed to obtain buffers for packet transmission or LSA storage.
- `Hello Rcvd` – Number of Hello messages received by the EOS device.
- `Hello Xmit` – Number of Hello messages sent by the EOS device.
- `DD Rcvd` – Number of data description packets received by the EOS device.
- `DD Xmit` – Number of data description packets transmitted by the EOS device.
- `LSR Rcvd` – Number of link state information request packets received by the EOS device.
- `LSR Xmit` – Number of link state information request packets sent by the EOS device.
- `LSA Rcvd` – Number of link state acknowledgment packets received by the EOS device.
- `LSA Xmit` – Number of link state acknowledgment packets sent by the EOS device.
- `LSU Rcvd` – Number of link state update packets received by the EOS device.
- `LSU Xmit` – Number of link state update packets sent.
- `Number of DR Election` – Number of times designated router election was performed.
- `Adjacency UP Events` – Number of times that an adjacency went from Down to Up state.
- `Adjacency DOWN Events` – Number of times that an adjacency went from Up to Down state.

Errors

- `Xmit Fail` – Number of transmission congestions experienced while transmitting OSPF packets. Congestion happens when the transmit queue overflows, and the OSPF packets are dropped.
- `Rcv bad packet header` – Number of errors received by the EOS device because of faulty packet headers.
- `Mismatch HelloTime` – Number of Hello packets received with mismatched HelloTimes. In order for two OSPF systems to establish an adjacency, both must have identical HelloTime values.
- `Mismatch RouterDeadTime` – Number of Hello packets received with mismatched RouterDeadTimes. In order for two OSPF systems to establish an adjacency, both must have identical RouterDeadTime values.
- `Mismatch subnet/mask` – Number of Hello packets received with mismatched subnet or mask. In order for two OSPF systems to establish an adjacency, both must have identical IP subnets and masks.
- `Mismatch area ID` – Number of adjacency rejections due to mismatched area ID. In order for two neighbors to become adjacent, they must both be configured with identical OSPF area IDs.
- `Unknown packet type` – Number of OSPF packets received that are not one of the known (Hello, DD, LSR, LSA and LSU) packet types.
- `Authentication Error` – Number of packets received that fail authentication.
- `Packet Checksum Error` – Number of packets received with checksum errors.
- `LSA Checksum Error` – Number of link state advertisements with checksum errors.

F.15

PATh Service Statistics

The display generated by the `SHoW -SYS STATistics -PATh` command includes the following statistics:

`Rcvd Packets` – Number of good packets received on the interface.

`Bytes` – Number of good bytes received on the specified interface. Includes headers but not cyclic redundancy check (CRC) bytes.

`Error:`

CRC – Number of frames that were received but failed the cyclic redundancy check.

Framing – Number of frames that were received but were not on a 16-bit Too Long boundary.

Too Long – Number of frames discarded because the packet length was longer than the maximum packet length allowed.

Lost – Number of receptions aborted because the CPU could not provide the controller chip with memory quickly enough.

Parity – Number of asynchronous parity receive errors.

Break – Number of asynchronous break receive errors.

Xmit Packets – Number of good packets transmitted.



NOTE: Failure to terminate your Ethernet network results in the false detection of transmission on the non-terminated BNC connector. This is due to Ethernet module sensitivity to RF transmissions from nearby boards. If this occurs, the Xmit Packets number displayed in the PATH Service statistics are incorrect.

Bytes – Number of bytes transmitted from a port.; includes headers but not CRC bytes.

Error:

Deferred – Number of frames that could not be transmitted because of existing traffic on the link. Transmission would have resulted in a collision. A later attempt was made to transmit the frame.

Collision – Number of frames that experienced a collision during the first attempt to transmit.

Late Collisions – Number of frames that received a collision outside of the preamble.

Xcess Collision – Number of frames not discarded after 16 consecutive collisions.

Carrier Loss – Number of frames that experienced a loss of the carrier signal during transmission.

Underrun – Number of transmissions aborted because the CPU could not provide the controller chip with data fast enough.

Discard:

Buf Overrun – Number of good frames lost because of memory overrun. This occurs when the system does not have enough memory to transfer the packet internally for further processing.

Congestion – Number of frames that could not be transmitted because of transmit queue overflow.

Utilization – The percentage of time the carrier sense signal was active during the specified interval.

The percentage of utilization displayed for HSS ports is based on a full-duplex link. For example, a 64 Kbps circuit can transmit and receive 64 Kbps simultaneously. If this link were transmitting at 64 Kbps and receiving nothing, the percentage of utilization would be 50%.

Rcv Good:

Pkt/Sec – Number of good packets received per second.

Byte/Sec – Number of good bytes received per second.

Xmit Good:

Pkt/Sec – Number of good packets transmitted per second.

Byte/Sec – Number of good bytes transmitted per second.

F.16

PIM-SM Service Statistics

The display generated by the `SHOW -SYS STATISTICS -PIM` command lists both per-port and global statistics, with the per-port statistics listed first

PIM Packets Received

Hellos – The number of PIM hello packets received on this port from any other PIM router on that port.

Join-Prunes – The number of PIM join/prune packets received on this port. This can be (*, G), (S, G), or (S, G) RPT join/prune packets. Aggregated join/prune packets, containing more than one join/prune entry, are counted as one packet for the purposes of this statistic.

Assert – The number of PIM assert packets received on this port, which can be an upstream/downstream port for an existing MRT entry. If there is no corresponding MRT entry, the statistic is incremented even though the assert packet is ignored.

Bootstrap – The number of multicast or unicast PIM bootstrap messages received on this port.

PIM Packets Transmitted

Hellos – The number of PIM hello packets transmitted on this port.

Join-Prunes – The number of PIM join/prune packets transmitted on this port. This can be (*, G), (S, G), or (S, G) RPT join/prune packets. Aggregated join/prune packets, containing more than one join/prune entry, are counted as one packet for the purposes of this statistic.

Assert – The number of PIM assert packets transmitted on this port.

Bootstrap – The number of multicast or unicast PIM bootstrap messages transmitted/forwarded on this port.

Events

Cache Miss – Not used.

Wrong Interface – Not used.

IGMP Join – The number of IGMP membership reports received on this PIM port.

IGMP Leave – The number of IGMP leaves received on this PIM port.

PIM Errors

Hellos – The number of invalid hello packets received on this port.

Hello Opt-Unknown – The number of hello packets received on this port with an unknown hello option value.

Joins-Prunes – The number of invalid join/prune packets received on this port. This statistic is incremented for each error encountered within a packet, so this count may be greater than the Join-Prunes received count. The errors can be any of the following:

Join/prune packet size too small.

Invalid encoded multicast address/ masklen / source address in the packet.

Assert – The number of invalid assert packets received on this port.

Asserts Rate Limit – The number of assert messages not sent because of rate-limiting asserts. A non-zero asserts rate limit count indicates that either there is a loop, or the network is unstable and devices do not agree on unicast routing.

Bootstrap – The number of invalid bootstrap packets received on this port.

Packet Too Small – The number of PIM control packets received on this port that are invalid because the length of the packet is smaller than the PIM header.

Version Error – The number of PIM packets received on this port with an incorrect version number (not equal to 2).

Checksum Error – The number of PIM packets received on this port with a bad checksum. This statistic excludes PIM register packets.

Ignored Bootstraps – The number of bootstrap messages ignored on this port. This is not necessarily an error condition. Bootstrap messages may be ignored in any of the following scenarios:

When an elected bootstrap router receives a multicast bootstrap message.

When an EOS device boots up, waits for unicast routing to converge, and receives a multicast bootstrap message.

When there is no valid RPF route for the bootstrap router.

When the source of the bootstrap router packet is not a valid PIM neighbor on the port.

When the received bootstrap message contains an entry from a device that lost the bootstrap router election.

When a unicast bootstrap message is received after the first unicast bootstrap message was received.

When the fragment tag in the bootstrap message is incorrect.

Unknown Packet Typ – Not used.

Unknown Neighbor – The number of dropped packets from unknown neighbors on this port. If a hello packet is not received from a device, other link local control packets from that device are dropped.

Other Transmit Err – The number of PIM control packets that could not be transmitted due to lower layer issues.

The following PIM-SM statistics allow you to monitor the effects of event-triggered RPF updates. These statistics are global rather than port-specific:

RPF Lookups – Total number of RPF lookups.

RPF Changes – Number of times a route changed during RPF lookup.

Unexpected Changes – Number of route changes detected with no trigger. A nonzero value for this statistic indicates a routing change that was not communicated to PIM-SM.

Triggers – Number of times PIM-SM received a trigger message.

False Triggers – Number of triggers received with no changes detected.

Ignored Triggers – Number of trigger messages for which the count did not match the current value of the global counter tracking the number of route updates and thus the message was ignored and not processed.

PT Insertions – The number of insertions in the PIM-SM multicast routing table (MRT) Patricia Tree.

PIM Multicast Groups Rejected by Limit – The number of PIM multicast routing table entries that were denied because the configured multicast group limit was met or exceeded.

IGMP Joins – Number of IGMP joins that were discarded by PIM because the configured multicast group limit was met or exceeded.

Joins/Prunes – Number of PIM joins and prunes that were discarded because the configured multicast group limit was met or exceeded.

Registers – Number of PIM registers that were discarded because the configured multicast group limit was met or exceeded.

Multicast Data Packets – Number of multicast data packets that were discarded because the configured multicast group limit was met or exceeded.

F.17

Port Service Statistics

The display generated by the `SHoW -SYS STATistics -POrt` command includes the following statistics:

Rcvd

Packets – Number of good packets received on a specified port.

Bytes – Number of good bytes received on a specified port.

Multicast – Number of multicast packets received. Multicast packets are sent to more than one station on the network.

Broadcast – Number of broadcast packets received. Broadcast packets are sent to the entire network.

Xmit

Packets – Number of good packets sent by ports.

Bytes – Number of good bytes sent by ports.

Multicast – Number of multicast packets sent. Multicast packets are sent to more than one station on the network.

Broadcast – Number of broadcast packets sent. Broadcast packets are sent to the entire network.

Filter

Custom – Number of packets that matched the custom filters and were therefore discarded.

Discard

Buffer Overrun – Number of packets discarded because of buffer overrun. This occurs when the system does not have enough memory to transfer the packet internally for further processing.

Congestion – Number of packets discarded because of congestion. This occurs when a packet cannot be transmitted within a specified amount of time.

DialOnDemand Mode

DodCallsMade – Number of outgoing calls successfully initiated by the port operating in dial-on-demand mode.

DodCallsFail – Number of outgoing calls unsuccessfully initiated by the port operating in dial-on-demand mode.

DodUpTime – Length of time in seconds that the primary path of a port is up while operating in dial-on-demand mode.

DodPktsOut – Number of user-data packets sent out by the port operating in dial-on-demand mode.

F.18

PPP Service Statistics

The display generated by the `SHoW -SYS STATistics -PPP` command includes the following statistics:

LCP path statistics

Rcvd

Conf-Request – Number of received LCP packets of code 1 for configure request.

Conf-Ack – Number of received LCP packets of code 2 for configure acknowledgment.

Conf-Nak – Number of received LCP packets of code 3 for configure negative acknowledgment.

Conf-Reject – Number of received LCP packets of code 4 for configure rejection.

Term-Request – Number of received LCP packets of code 5 for terminate request.

Term-Ack – Number of received LCP packets of code 6 for terminate acknowledgment.

Code-Reject – Number of received LCP packets of code 7 for code rejection.

Protocol-Reject – Number of received LCP packets of code 8 for protocol rejection.

Echo-Request – Number of received LCP packets of code 9 for echo request.

Echo-Reply – Number of received LCP packets of code 10 for echo reply.

Discard-Req – Number of received LCP packets of code 11 for discard request.

Identification – Number of received LCP identification packets.

Unknown Code – Number of received LCP packets of unknown code.

Xmit

Conf-Request – Number of transmitted LCP packets of code 1 for configure request.

Conf-Ack – Number of transmitted LCP packets of code 2 for configure acknowledgment.

Conf-Nak – Number of transmitted LCP packets of code 3 for configure negative acknowledgment.

Conf-Reject – Number of transmitted LCP packets of code 4 for configure rejection.

Term-Request – Number of transmitted LCP packets of code 5 for terminate request.

Term-Ack – Number of transmitted LCP packets of code 6 for terminate acknowledgment.

Code-Reject – Number of transmitted LCP packets of code 7 for code rejection.

Protocol-Reject – Number of transmitted LCP packets of code 8 for protocol rejection.

Echo-Request – Number of transmitted LCP packets of code 9 for echo request.

Echo-Reply – Number of transmitted LCP packets of code 10 for echo reply.

Discard-Req – Number of transmitted LCP packets of code 11 for discard request.

Identification – Number of transmitted LCP identification packets.

Unknown Code – Number of transmitted LCP packets of unknown code.

F.18.1

PPP Over Ethernet Statistics

The display generated by the `SHoW -SYS STATistics -PPPOE` command includes the following statistics:

- PADI Packets
- PADO Packets
- PADR Packets
- PADS Packets
- PADT Packets
- Session Packets

F.19

RIPIP Service Statistics

The display generated by the `SHoW -SYS STATistics -RIPIP` command includes the following statistics:

Incoming Packets

Request Updates – Number of incoming RIP request packets on a port. RIP request packets request routing information.

Response Updates – Number of incoming RIP response packets on a port. RIP response packets are sent to convey routing information.

Discarded Updates – Number of update packets discarded on a port because they originated from an unconfigured neighbor.

V2 Packets – Number of incoming RIPv2 packets on port. This count is divided into the following categories:

- Request Updates

- Response Updates

- Error - RIPv2 on RIPv1 interface

- Error - RIPv1 on RIPv2 interface

- Error - Authenticate

Outgoing Pkts

Request Updates – Number of outgoing request update packets sent on a port. Request updates are sent either when the EOS device is booted or when RIP was configured.

Regular Responses – Number of outgoing regular update packets sent on a port. Regular updates are router information packets that are sent out by the EOS device at regular intervals.

Triggered Responses – Number of triggered response packets sent on a port. Triggered responses are router information packets sent out immediately when a network becomes unreachable.

V2 Packets – Number of outgoing RIPv2 packets on port. This count is divided into the following categories:

- Total

- Multicast

- Error - RIPv2 on RIPv1 interface

- Error - RIPv1 on RIPv2 interface

F.20

RSVP Service Statistics

The display generated by the `SHoW -SYS STATistics -RSVP` command includes the following statistics:

RSVP Port Statistics

- No. of PATH msgs rcvd – The number of PATH messages received.
- RESV msgs rcvd – The number of RESV messages received.
- PATH ERR msgs rcvd – The number of PATH ERROR messages received.
- RESV ERR msgs rcvd – The number of RESV ERROR messages received.
- PATH TEAR msgs rcvd – The number of PATH TEAR messages received.
- RESV TEAR msgs rcvd – The number of RESV TEAR messages received.
- CONFIRM msgs rcvd – The number of CONFIRM messages received.
- REPORT msgs rcvd – The number of REPORT messages received.
- DREQ msgs rcvd – The number of DREQ messages received.
- DREP – The number of DREP messages received.
- No. of PATH msgs sent – The number of PATH messages sent.
- RESV msgs sent – The number of RESV messages sent.
- PATH ERR msgs sent – The number of PATH ERROR messages sent.
- RESV ERR msgs sent – The number of RESV ERROR messages sent.
- PATH TEAR msgs sent – The number of PATH TEAR messages sent.
- RESV TEAR msgs sent – The number of RESV TEAR messages sent.
- CONFIRM msgs sent – The number of CONFIRM messages sent.
- REPORT msgs sent – The number of REPORT messages sent.
- DREQ msgs sent – The number of DREQ messages sent.
- DREP – The number of DREP messages sent.
- admission failures – The number of admission failures.
- other resv failures – The number of other reservation failures.
- Resv sent due to scope – The number RESV messages sent due to scope.
- WF resv w/o scope – The number of wild card RESV messages without scope.
- blockade events – The number of blockade events.
- resv timeouts – The number RESV timeouts.
- path timeouts – The number of PATH time outs.
- path/ptear rcv ttl0 – The number of PATH and PATH TEAR messages received with TimeToLive specification of 0.
- path/ptear snt ttl0 – The number of PATH and PATH TEAR messages sent with TimeToLive specification of 0.
- Remaining bw (Bps) – The amount of remaining bandwidth.
- Configured bw (Bps) – The amount of configured bandwidth.
- RSVP data pkts sent – The number of RSVP data packets sent.
- RSVP bytes sent – The number of RSVP bytes sent.
- Non-RSVP data pkts – The number of non-RSVP data packets.

F.21

SNMP Service Statistics

The display generated by the `SHoW -SYS STATistics -SNMP` command includes the following statistics:

- Incoming SNMP PDUs – Number of PDUs delivered to SNMP.

Get Requests – Number of Get-Request PDUs processed by SNMP.

Get-Next Requests – Number of Get-Next Request PDUs processed by SNMP.

Set Requests – Number of Set-Request PDUs processed by SNMP.

Bad PDUs – Number of PDUs delivered to but not processed by SNMP, including unsupported version, unknown community name, not allowed operation by the named community, ASN.1 parsing errors, and unknown PDU type.

Get Responses – Number of responses to SNMP Get requests. This count is divided into the following categories:

Too Big Errors

No Such Name Errors

Bad Value Errors

Read Only Errors

General Errors

Outgoing SNMP PDUs – Number of PDUs generated by SNMP.

Get Responses – Number of Get-Response PDUs generated by SNMP.

Error Responses – Number of valid SNMP PDUs generated by SNMP and for which the value of the error-status is not noError (0), including tooBig (1), noSuchName (2), badValue (3), readOnly (4), and genErr (5).

Traps – Number of Trap PDUs generated by SNMP.

Get Requests – Number of Get-Request PDUs generated by SNMP.

Get-Next Requests – Number of Get-Next Request PDUs generated by SNMP.

Set Requests – Number of Set-Request PDUs generated by SNMP.

Incoming SNMPv3 PDUs – Number of PDUs delivered to SNMPv3.

Get Requests – Number of Get-Request PDUs processed by SNMPv3.

Get-Next Requests – Number of Get-Next Request PDUs processed by SNMPv3.

Set Requests – Number of Set-Request PDUs processed by SNMPv3.

Get Responses – Number of responses to SNMPv3 Get requests.

Incoming Traps – Number of Trap PDUs processed by SNMPv3.

Bad PDUs – Number of PDUs delivered to but not processed by SNMPv3 due to the following errors:

Too Big Errors

No Such Name Errors

Bad Value Errors

Read Only Errors

General Errors

Outgoing SNMPv3 PDUs – Number of PDUs generated by SNMPv3.

Get Requests – Number of Get-Request PDUs generated by SNMPv3.

Get-Next Requests – Number of Get-Next Request PDUs generated by SNMPv3.

Get Responses – Number of Get-Response PDUs generated by SNMPv3.

Set Requests – Number of Set-Request PDUs generated by SNMPv3.

Error Responses – Number of error response PDUs generated by SNMPv3, including the following:

Too Big Errors

No Such Name Errors

Bad Value Errors

Outgoing Traps – Number of SNMPv3 traps sent.

USM Unsupported Sec Levels – Number of SNMPv3 messages with unsupported USM security levels.

USM Unknown User Names – Number of SNMPv3 messages with unknown USM user names.

USM Unknown Engine IDs – Number of SNMPv3 messages with unknown USM engine IDs.

USM Wrong Digests – Number of SNMPv3 messages with an incorrect authentication protocol.

USM Decryption Errors – Number of SNMPv3 messages that failed decryption.

F.22

STP Service Statistics

The display generated by the `SHoW -SYS STATistics -STP` command includes the following statistics:

Forwarding \State Count – Number of times the specified port was put in forwarding state.

Blocking \State Count – Number of times the specified port was put in blocking state because of a possible loop.

Bad Config BPDU Receiv – Number of bridge protocol data units (BPDUs) received with information indicating that the transmitting bridge does not recognize another high-priority bridge on the network (possibly because the transmitting bridge just booted up).

Looped Config \BPDU Re – Number of BPDUs the bridge received its own.

Message Age\Timeouts – Number of times a neighboring high-priority bridge went out of service.

F.23

TCP Service Statistics

The display generated by the `SHoW -SYS STATistics -TCP` command includes the following statistics:

TCP Packets

Transmitted – Number of TCP packets transmitted within a specified interval.

Received – Number of TCP packets received within a specified interval.

Retransmitted – Number of TCP packets retransmitted within a specified interval. Packets are retransmitted when the previous attempt fails or when they are timed out.

TCP Connections

Initiated – Number of TCP connections attempted.

Accepted – Number of successful TCP connection attempts.

Failed – Number of failed TCP connection attempts.

Reset – Number of TCP connections reset. Connections are reset if they were aborted because of error.

F.24

UDP Service Statistics

The display generated by the `SHoW -SYS STATistics -UDP` command includes the following statistics:



NOTE: The UDP Service does not appear in the user interface. However, you can still obtain UDP statistics.

Datagrams transmitted – Number of packets sent to other networks.

Good datagrams received – Number of error-free packets received.

Datagrams with errors – Number of packets containing errors, such as checksum errors unable to give packets to UDP user.

Datagrams with unknown port – Number of packets received for which UDP does not have matching port number in its port table.

ICMP Datagrams received – Number of Internet Control Management Protocol packets received.

Out of buffers – Number of packets that could not be transmitted because no system buffers were available.

F.25

UDPHelp Service Statistics

The display generated by the `SHoW -SYS STATistics -UDPhelp` command includes the following statistics:



NOTE: The system displays the UDPHELP statistics on a system-wide basis, rather than on a port-by-port basis.

BOOTP/UDP/IP Broadcast Helper Statistics

Packets Received – Number of packets received.

Packets Forwarded – Number of packets forwarded to servers.

Packets ReBroadcasted – Number of packets forwarded to all interfaces.

Packets Discarded – Number of packets discarded because the packet did not broadcast a destination address.

Miscellaneous Error – Number of times a packet could not be forwarded because of a combination of heavy traffic on the network and memory constraint.

Appendix G

Static Tables

This appendix lists the number of entries you can have in statically configured routing tables – by supported hardware platform.

The following table lists the different EOS features and device types, and the maximum number of routing table entries you can have per feature for each hardware platform.

Table 119: Number of Entries Allowed in Static Tables

EOS Feature	MNR S6000 / GGM 8000	MNR S2500
Bridge		
Bridge Table	No limit	No limit
IP		
Static Routing Table	No limit	No limit
Static Address Table	20480	20k
Secondary IP Addresses	No limit	No limit
IP-OSPF		
ExteriorPolicy	No limit	No limit
InteriorPolicy	No limit	No limit
StaticPolicy	No limit	No limit
ReceivePolicy	No limit	No limit
Neighbors	No limit	No limit
Virtual Link	No limit	No limit

Appendix H

Syslog Messages

This appendix lists the EOS syslog messages. The following table describes syslog messages identified by record type and status codes.

By default, all syslog messages include a timestamp in the following format:

<day><month><date><hour:minutes:seconds><year>.

To configure EOS to modify the syslog message header so the timestamp format is compliant with RFC 5424, you can set the AuditLog service SysLogTimFormat paramter to Rfc. RFC 5424 specifies the following timestamp format:

<year>-<month>-<date>T<hour:minutes:seconds>Z

RFC 5424 compliance is required for systems that use the optional Centralized Event Logging feature.



NOTE: Italicized items in the log string are substituted with actual values.

Table 120: Syslog Messages

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
AuditLog	100	Auditlog file creation failure	"AuditLog log file not created - <i><reason></i> "	LogERror	AuditLog service control & config
	101		<i><reason></i> can be:		
	102		- bad pdmp size from DUMP		
	103		- fail memory allocation for DUMP		
	104		- file write error		
	105		- media not present		
	106		- media write protected		
	107		- directory is full		
	108		- unsupported media		
AuditLog	109	Auditlog file read failure	- file system error	LogERror	AuditLog service control & config
			- missing signature		
AuditLog	109	Auditlog file read failure	"AuditLog Read logfile or Read auditlog.dmp error"	LogERror	AuditLog service control & config
AuditLog	110	Auditlog file access status	"AuditLog no audlog or no log file error"	LogERror	AuditLog service control & config
AuditLog	111	Auditlog file created successful	"AuditLog log file %s/%s created"	LogNOtice	AuditLog service

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
					control & config
AuditLog	112	Auditlog file logging disabled	"AuditLog file logging disabled - MaxNumLogFile = 0"	LogNotice	AuditLog service control & config
AuditLog	160	SNMPv3 user created	"USM user <username> with engineID <engineID> created"  NOTE: If the length of the log string exceeds the LocalLog buffer size, the actual value of the engine ID is replaced with "...".	LogNotice	AuditLog service control and config
AuditLog	161	SNMPv3 user deleted	"USM user <username> with engineID <engineID> deleted" If the length of the log string exceeds the LocalLog buffer size, the actual value of the engine ID is replaced with "...".	LogNotice	AuditLog service control and config
AuditLog	162	SNMPv3 user activated	"USM user <username> with engineID <engineID> activated" If the length of the log string exceeds the LocalLog buffer size, the actual value of the engine ID is replaced with "...".	LogNotice	AuditLog service control and config
AuditLog	164	SNMPv3 user authentication or privacy (encryption) passphrase changed	"USM user <username> with engineID <engineID> [Authentication Privacy] passphrase changed" If the length of the log string exceeds the LocalLog buffer size, the actual value of the engine ID is replaced with "...".	LogNotice	AuditLog service control and config
AuditLog	165	SNMPv3 user security level changed	"USM user <username> with engineID <engineID> security lev-	LogNotice	AuditLog service con-

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			el changed to [AuthNo-Priv NoAuthNoPriv]" If the length of the log string exceeds the Local-Log buffer size, the actual value of the engine ID is replaced with "...".		trol and config
DHCP	300	DHCP address pool out of address	"dhcp_discover out of addresses"	LogAlert	AuditLog service control & config
DHCP	301	DHCP address offering	"OFFERING IP(<ip_addr>) to client(<client_id>)"	LogAlert	AuditLog service control & config
DHCP	302	DHCP address released	"receives DHCPRELEASE and release IP (<ip_addr>) from client (<client_id>)"	LogAlert	AuditLog service control & config
DHCP	303	DHCP RAS address allocation	"dhcp_ras_getip - Got IP address (<ip_addr>)"	LogAlert	AuditLog service control & config
System	400	Bandwidth status	"BW TARGET MET FOR PORT <port_num>: NORMAL BAND WIDTH IS ZERO"	LogInfo	AuditLog service control & config
System	401	Bandwidth status	"BW TARGET NOT MET FOR PORT <port_num>: <reason>" <reason> can be: DIAL POOL EMPTY/ USED DIAL NO. LIST EMPTY NO MATCH FOUND NO MORE MEMORY PORT NOT DIALABLE NO MORE PATH/DIAL NO. PORT MIS-CONFIGURED REASON UNKNOWN	LogInfo	AuditLog service control & config
System	402	Dial-on-Demand port status	"DOD PORT <port_num> DOWN - NO PATH(s) AVAILABLE"	LogInfo	AuditLog service

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
					control & config
System	403	Outgoing call failure	"RETRY COUNT ON PORT <i><port_num></i> EXCEEDED"	LogInfo	AuditLog service control & config
System	404	Call Termination status	"INITIATING HANGUP ON PATH <i><path_num></i> "	LogInfo	AuditLog service control & config
System	405	Outgoing call failure	"CALL ON PATH <i><path_num></i> REJECTED, NO CARRIER"	LogInfo	AuditLog service control & config
System	406	Outgoing call failure	"CALL ON PATH <i><path_num></i> REJECTED, PATH DID NOT COME UP"	LogInfo	AuditLog service control & config
System	407	Incoming call status	"INCOMING CALL ON DIALPOOL PATH <i><path_num></i> "	LogInfo	AuditLog service control & config
System	408	Incoming call failure	"INCOMING CALL ON DIALPOOL PATH <i><path_num></i> REJECTED"	LogInfo	AuditLog service control & config
System	409	Incoming call failure	"INCOMING CALL ON PATH <i><path_num></i> REJECTED. NO PORT."	LogInfo	AuditLog service control & config
System	410	Incoming call success	"INCOMING CALL ON DIALPOOL PATH <i><path_num></i> CONNECTED"	LogInfo	AuditLog service control & config
System	411	Incoming call success	"INCOMING CALL ON PATH <i><path_num></i> , PORT <i><port_num></i> CONNECTED"	LogInfo	AuditLog service control & config
System	412	Call termination status	"DISCONNECT ON PATH <i><path_num></i> "	LogInfo	AuditLog service control & config
System	413	Incoming call failure	"CALLER ID MISSING ON DIAL PATH <i><path_num></i> ."	LogInfo	AuditLog service control & config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
System	414	Call termination status	<i><hangup_initiator></i> INITIATING HANGUP ON PATH <i><path_num></i>, PORT <i><port_num></i> <i><hangup_initiator></i> can be: USER, BOD, DOD, PATH NOT COMING UP, PATH IS IDLE	LogInfo	AuditLog service control & config
System	415	Call termination status	<i><hangup_initiator></i> INITIATING HANGUP ON PATH <i><hangup_initiator></i> can be: USER, BOD, DOD, PATH NOT COMING UP, PATH IS IDLE	LogInfo	AuditLog service control & config
System	416	Incoming call failure	"CALLER ID <i><caller_id></i> ON PATH <i><path_num></i> NOT FOUND."	LogInfo	AuditLog service control & config
System	417	Outgoing call status	"INITIATING CALL ON PATH <i><path_num></i> , PORT <i><port_num></i> "	LogInfo	AuditLog service control & config
System	418	Call termination status	"INITIATING HANGUP ON PATH <i><path_num></i> , PORT <i><port_num></i> "	LogInfo	AuditLog service control & config
System	419	Outgoing call failure	"CALL ON PATH <i><path_num></i> , PORT <i><port_num></i> REJECTED, NO CARRIER"	LogInfo	AuditLog service control & config
System	420	Call failure	"CALL ON PATH <i><path_num></i> , PORT <i><port_num></i> REJECTED, PATH DID NOT COME UP"	LogInfo	AuditLog service control & config
System	421	Incoming call status	"INCOMING CALL ON PATH <i><path_num></i> , PORT <i><port_num></i> "	LogInfo	AuditLog service control & config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
System	422	Incoming call failure	"INCOMING CALL ON PATH <path_num>, PORT <port_num> REJECTED"	LogInfo	AuditLog service control & config
System	423	Outgoing call failure	"CALL ON PATH <path_num>, PORT <port_num> REJECTED, PREVIOUS HANGUP NOT YET COMPLETED"	LogInfo	AuditLog service control & config
System	424	Outgoing call failure	"CALL ON PATH <path_num>, PORT <port_num> REJECTED"	LogInfo	AuditLog service control & config
System	425	Outgoing call status	"OUTGOING CALL ON PATH <path_num>, PORT <port_num> CONNECT- ED"	LogInfo	AuditLog service control & config
System	426	Incoming call status	"INCOMING CALL ON PATH <path_num>, PORT <port_num> CONNECT- ED"	LogInfo	AuditLog service control & config
System	427	Call termina- tion status	"DISCONNECT ON PATH <path_num>, PORT <port_num>"	LogInfo	AuditLog service control & config
System	428	Incoming call failure	"INCOMING CALL ON PATH <path_num> REJECTED; PORT <port_num> DISA- BLED"	LogInfo	AuditLog service control & config
System	429	Incoming call failure	"INCOMING CALL ON PATH <path_num> REJECTED; PORT <port_num> OWN- ER IS NOT PPP"	LogInfo	AuditLog service control & config
System	430	Incoming call failure	"INCOMING CALL ON PATH <path_num> RE- JECTED BY PORT <port_num>"	LogInfo	AuditLog service control & config
System	431	Incoming call failure	"INCOMING CALL ON PATH <path_num> RE- JECTED; PATH NOT IN PORT <port_num> POOL PREF LIST"	LogInfo	AuditLog service control & config
System	432	Incoming call status	"INCOMING CALL ON PATH <path_num>"	LogInfo	AuditLog service

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			BOUND TO PORT <port_num>"		control & config
System	433	Leased path status	"DYNAMIC LEASED PATH <path_num> BOUND TO PORT <port_num>"	LogInfo	AuditLog service control & config
System	434	Outgoing call failure	"ISDN CALL ON PATH %path_num%, PORT %port_num% REJECT- ED, CAUSE CODE = %code_num%"	LogInfo	AuditLog service control & config
System	435	Outgoing call failure	"ISDN CALL ON PATH <path_num>, PORT <port_num> REJECTED, CAUSE CODE = <code_num> (<code_string>)" <code_string> can be: "NORMAL CALL CLEAR- ING" "USER BUSY" "NO USER RESPOND- ING" "CALL REJECTED" "DESTINATION OUT OF ORDER" "NO CIRCUIT AVAILA- BLE" "NETWORK OUT OF OR- DER" "TEMPORARY FAILURE" "NETWORK CONGES- TION" "INVALID INFO. ELE- MENT CONTENTS" "MESSAGE INCOMPATI- BLE WITH CALL STATE" "PREVIOUS HANGUP NOT YET COMPLETED"	LogInfo	AuditLog service control & config
System	436	Outgoing call failure	"CALL ON PATH <path_num>, PORT <port_num> REJECTED, CODE = <reject_code>"	LogInfo	AuditLog service control & config
System	437	Outgoing call status	<initiator> INITIAT- ING CALL ON PATH	LogInfo	AuditLog service

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			<path_num>, PORT <port_num>" <initiator> can be: "PATH NOT COMING UP," "PATH IS IDLE," "USER" "DOD" "BOD"		control & config
System	438	Incoming call failure (Caller Line Identification)	"CLI SECURITY VIOLATION ON PATH <path_num>, PORT <port_num>"	LogInfo	AuditLog service control & config
System	439	Incoming call failure (Caller Line Identification)	"CLI IS MISSING IN INCOMING CALL ON PATH <path_num>, PORT <port_num>"	LogInfo	AuditLog service control & config
System	440	Incoming call failure (Caller Line Identification)	"CLI <cli_code> SECURITY VIOLATION ON PATH <path_num>, PORT <port_num>"	LogInfo	AuditLog service control & config
PORT	441		EXPECTING CALL BACK ON PORT %port_num%	LogInfo	AuditLog service control & config
PORT	442		DOD RESET FAILED ON PORT port_num PATH %path_num%: %reason% reason is: DIAL NO. LIST EMPTY NO MATCH DIAL NO. NO MORE MEMORY	LogInfo	AuditLog service control & config
PORT	443		ISDN CALL ON PATH path_num, PORT %port_num% REJECTED, DUE TO CALL BACK	LogInfo	AuditLog service control & config
PORT	444		CALL BACK MATCH CLI cli_code ON PATH %path_num% , PORT %port_num%	LogInfo	AuditLog service control & config
PORT	445		CALL BACK TIME OUT EXPIRED ON PORT %port_num%	LogInfo	AuditLog service control & config
PORT	446		CALL BACK TIME OUT CANCELLED ON PORT %port_num%	LogInfo	AuditLog service

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
					control & config
PORT	447		bm_bring_down_bw: BM MEMORY ALLOCATION ERROR	LogInfo	AuditLog service control & config
PORT	448		BM PANIC -- %text%	LogInfo	AuditLog service control & config
FireWall	501	Firewall action log (protocol = TCP)	Tx/Rx !<port_label>src_ip_addr (<src_port_name>->dest_ip_addr (<dest_port_name>) TCP, Permit/Deny, <info> <info> can be: TinyFragment, SrcSpoof, IPTunnel, InFilter, DefActionIn, SrcRoute, RecordRoute, TimeStamp, OutFilter, DefActionOut	User configured via firewall setup, default is LogInfo	firewall log set to SysLog
FireWall	502	Firewall action log (protocol == ICMP)	Tx/Rx !<port_label>src_ip_addr >><dest_ip_addr> ICMP (<icmp_type>), Permit/Deny, <info> <icmp_type> can be: DestUnreach SrcQuench Redirect EchoReq EchoRsp TimeExcd ParamProblem DestAdmUnreach info can be:	User configured via firewall setup, default is LogInfo	firewall log set to SysLog

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			TinyFragment, SrcSpoof, IPTunnel, InFilter, DefActionIn, SrcRoute, RecordRoute, TimeStamp, OutFilter, DefActionOut		
FireWall	503	Firewall action log (protocol is not TCP) and (protocol is not ICMP)	Tx/Rx !<port_label src_ip_addr> ><dest_ip_addr protocol>, Permit/Deny, info <protocol> can be: UDP, OSPF, ESP, AH, GRE, RESERVED <info> can be: TinyFragment, SrcSpoof, IPTunnel, InFilter, DefActionIn, SrcRoute, RecordRoute, TimeStamp, OutFilter, DefActionOut	User configured via firewall setup, default is LogNfo	firewall log set to SysLog
IPsec	600	Crypto/IPsecLog	"<number_of_msgs> log messages suppressed"	LogWarning	Crypto LogLevel and LogDest
IKE PI /IKE PII/ IPsec	601	Crypto/IPsecLog	"ISAKMP (Phase <<number>>): SA Negotiation with <peer address>failed: <reason><detail>" "Cannot negotiate Phase 1 SA: Local certificate not installed"	LogError	Crypto LogLevel and LogDest

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			"Cannot negotiate Phase 1 SA: Local certificate identity is too long" "Cannot negotiate Phase 1 SA: Invalid local certificate identity type"		
Crypto	650	Crypto	"New Key Encryption Key Generated"	LogWarning	Crypto LogLevel and LogDest
Crypto	651	Crypto	"Key Encryption Key was Zeroized"	LogWarning	Crypto LogLevel and LogDest
NAT	800	NAT log overflow (when log rate is over 10 per second)	<num_msg> previously unreported log messages, due to overflow	LogInfo	NAT log set to Syslog
NAT	801	NAT Session Successful for TCP/UDP	Inbound/Outbound ! <port_label> TCP/UDP (<src_ip>,<src_tuport>)->(<dest_ip>,<dest_tuport>), translated to <pkt_direction> (<xlate_ip>,<xlate_port>) <pkt_direction> can be: "be destined to", "originate from"	configured by user using the TcpUdpPort-Mapparameter under the NAT service (default is LogInfo)	NAT log set to Syslog
NAT	802	NAT Session Failure for TCP/UDP	Inbound/Outbound ! <port_label> TCP/UDP (<src_ip>,<src_tuport>)->(<dest_ip>,<dest_tuport>), translation failed - <failure_reason> <failure_reason> can be: "No matching MAP found.", "TUport not found in free tuport list.", "Packet direction is invalid.",	configured by user using the TcpUdpPort-Mapparameter under the NAT service (default is LogInfo)	NAT log set to Syslog

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			"IP protocol not permitted.", "Ran out of NAT addresses.", "Ran out of TCP/UDP ports.", "Ran out of Load share hosts.", "No matching NAT address.", "Could be due to premature session time-out.", "Cannot create any more sessions.", "Too many retries without a response.", "Packet data length is too small."		
NAT	803	NAT Session successful for protocol = ICMP	Inbound/Outbound ! <i><port_label></i> ICMP(<i><icmp_type></i>), <i><src_ip></i> > <i><dest_ip></i> translated to <i><pkt_direction></i> <i>xlate_ip></i> <i><icmp_type></i> can be: DestUnreach SrcQuench Redirect EchoReq EchoRsp TimeExcd ParamProblem DestAdmUnreach <i>pkt_direction</i> can be: "be destined to", "originate from"	configured by user using the TcpUdpPort-Map parameter under the NAT service (default is Log-Info)	NAT log set to Syslog
NAT	804	NAT Session Failure for protocol = ICMP	Inbound/Outbound ! <i><port_label></i> ICMP(<i><icmp_type></i>), <i><src_ip></i> > <i><dest_ip></i> translation failed: <i><fail_reason></i>	configured by user using the TcpUdpPort-Map parameter under the NAT service	NAT log set to Syslog

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			<icmp_type> can be: DestUnreach SrcQuench Redirect EchoReq EchoRsp TimeExcd ParamProblem DestAdmUnreach fail_reason can be: "No matching MAP found.", "TUport not found in free tuport list.", "IP protocol not permitted.", "Packet direction is invalid.", "Ran out of NAT addresses.", "Ran out of TCP/UDP ports.", "Ran out of Load share hosts.", "Ran out of fragmem.", "No matching NAT address.", "Could be due to premature session time-out.", "Cannot create any more sessions.", "Too many retries without a response.", "Cannot get internal buffer.", "Packet data length is too small."	(default is Log-Info)	
NAT	805	NAT Session Successful for protocol is not one of the following:	Inbound/Outbound ! <port_label protocol_num src_ip>><dest_ip> translated to	configured by user using the TcpUdpPort-Map parameter under the NAT service	NAT log set to Syslog

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
		(TCP UDP ICMP)	<i><pkt_direction></i> <i><late_ip></i> <i><pkt_direction></i> can be: "be destined to", "originate from"	(default is Log-Info)	
NAT	806	NAT Session Failure for protocol != (TCP UDP ICMP)	Inbound/Outbound ! <i><port_label></i> <i><protocol></i> <i><src_ip></i> <i><dest_ip></i> translation failed: <i><fail_reason></i> <i><fail_reason></i> can be: "No matching MAP found.", "Tuport not found in free tuport list.", "IP protocol not permitted.", "Packet direction is invalid.", "Ran out of NAT addresses.", "Ran out of TCP/UDP ports.", "Ran out of Load share hosts.", "Ran out of fragmem.", "No matching NAT address.", "Could be due to premature session time-out.", "Cannot create any more sessions.", "Too many retries without a response.", "Cannot get internal buffer.", "Packet data length is too small."	configured by user using the "TcpUdpPort Map" parameter under the NAT service (default is Log-Info)	NAT log set to Syslog
RAS	900	RAS LogInSuccess	User <i><username></i> Logged In On Port	LogAlert	RAS Log parameter set to SysLog

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
		(with no tunnel IP address)	<port_num>, Path <path_num>		and auditLog server ip
RAS	901	RAS LogInSuccess (with tunnel IP address)	User <username> Logged In On Port <port_num>, Path <path_num> - Tunnel Remote IP Addr <IP_address>	LogAlert	RAS Log parameter set to SysLog and auditLog server ip
RAS	902	RAS LogOut	User <username> Idled/Logged Out of Port <port_num>, path <path_num> - IP address <IP_address> - Tunnel Remote IP Addr <IP_address> - Session start time <time_string> - Duration <duration_time> - Pkts In <num_pkts> Out <num_pkts> - Bytes In <num_bytes> Out <num_bytes>	LogInfo	RAS Log parameter set to SysLog and auditLog server ip
RAS	903	RAS LogInFailure	<username> Login Failed on Path <path_num>: <reason> <reason> can be: Authorization Failure UserID Exceeds Max Length UserID Invalid Password Exceeds Max Length Password Invalid Challenge Response Invalid RAS Disabled No Port Available	LogAlert	RAS Log parameter set to SysLog and auditLog server ip
RAS	904	RAS ConnFailure	<LCP NCP> Failure on Path <path_num>, User <username> Note: either LCP or NCP is specified	LogNotice	RAS Log parameter set to SysLog and auditLog server ip
RAS	905	RAS PathLogInSuccess	User <username> Logged In On TunnelSwitch	LogAlert	RAS Log parameter set

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
		(with no tunnel IP address)	Port <port_num>, Path <path_num>		to SysLog and auditLog server ip
RAS	906	RAS PathLogInSuccess (with tunnel IP address)	User <username> Logged In On TunnelSwitch Port <port_num>, Path <path_num> - Tunnel Remote IP Address <IP_address>	LogAlert	RAS Log parameter set to SysLog and auditLog server ip
RAS	907	RAS PathLogOut	User <username> Idled/Logged Out Of TunnelSwitch Port <port_num>, Path <path_num> - Pkts In <num_pkts> Out <num_pkts> - Bytes In <num_bytes> Out <num_bytes>	LogInfo	RAS Log parameter set to SysLog and auditLog server ip
REMP	1000 to 1014	RemPolling poll result	Encoded remote polling results, check REMPPolling doc	LogInfo	AuditLog service control & config
SNMP	1200	SNMP set requests	"SNMP <snmp_oid value>" for snmp SET requests Example: SNMP 129.213.144.69 "1.3.6.1.2.1.1.5.0 119.121.99.95.100.112.101"	LogWarning	AuditLog service control & config
SNMP	1201		"%snmp_oid_value%" for SNMP SET requests with invalid community string	LogAlert	AuditLog service control & config
SNMP	1210	SNMPv3 persistent data file write error The presence of this log message indicates that changes to the SNMPv3 data may not be preserved over the reboot.	SNMPv3 data file write error —writing to temp file failed prior to reboot When the device receives a reboot command, the SNMPv3 persistent data file is written to a temporary file prior to the reboot. If the write operation succeeds, the temporary file is verified, the original SNMPv3 persistent data file is deleted, the temporary file is re-	LogWarning	AuditLog service control and config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			named as the original file, and the temporary file is deleted. If the write operation fails, the device sends the "SNMPv3 data file write error" syslog message.		
SNMP	1211	SNMPv3 persistent data file boot error The presence of this log message indicates that changes to the SNMPv3 data may not be preserved over the reboot.	SNMPv3 data file boot error— temp file verification failed during reboot When the device receives a reboot command, the SNMPv3 persistent data file is written to a temporary file. When the reboot occurs, EOS checks for the existence of a temporary file. If a temporary file exists, the file is verified. If the temporary file is valid, the original SNMPv3 persistent data file is deleted, the temporary file is renamed as the original file, and the temporary file is deleted. If the temporary file is not valid, the device sends the "SNMPv3 data file boot error" syslog message, and deletes the temporary file.	LogWarning	AuditLog service control and config
System	1300	Console login failure	"Login Failed - Invalid Credentials"	LogAlert	AuditLog service control & config
System	1301	Console login success	"Login Successful"	LogAlert	AuditLog service control & config
System	1302	Console session termination	"UI session terminated"	LogAlert	AuditLog service control & config
System	1303	Console listen command	"Listen executed"	LogAlert	AuditLog service

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
					control & config
System	1304	Console privilege change failure	"Set Privilege Failed"	LogAlert	AuditLog service control & config
System	1305	Reboot command executed	"monitor reboot"	LogNotice	AuditLog service control & config
System	1306	System start-up	"System Booted"	LogNotice	AuditLog service control & config
System	1307	Loop detected	"Loop Detected"	LogNotice	AuditLog service control & config
System	1308	System initialized	"System Initialized and Running"	LogNotice	AuditLog service control & config
System	1309	Path up	"Path <path_num> Up"	LogNotice	AuditLog service control & config
System	1310	Path down	"Path <path_num> Down"	LogNotice	AuditLog service control & config
System	1311	Path faulty	"Path <path_num> Faulty"	LogNotice	AuditLog service control & config
System	1312	path loopback	"Path <path_num> Loopback"	LogNotice	AuditLog service control & config
System	1313		System restart due to: %text%	LogAlert	AuditLog service control & config
System	1314		System restart without partial dump	LogAlert	AuditLog service control & config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
System	1315		Invalid Command: %cmd_text%	LogWarning	AuditLog service control & config
System	1316		Write to disk of New Password for %user% failed	LogAlert	AuditLog service control & config
System	1317		Failed adding user : '<un- known user name>'	LogAlert	AuditLog service control & config
System	1318		Failed adding '%user%'	LogAlert	AuditLog service control & config
System	1319		User '%user%' is added	LogAlert	AuditLog service control & config
System	1320		Failed adding user '%user%' : name duplicat- ed	LogAlert	AuditLog service control & config
System	1321		Failed adding user '%user%'	LogAlert	AuditLog service control & config
System	1322		Failed deleting user : <unknown user name>	LogAlert	AuditLog service control & config
System	1323		Failed deleting user '%user%'	LogAlert	AuditLog service control & config
System	1324		User '%user%' is deleted	LogAlert	AuditLog service control & config
System	1325		Failed changing pass- word : <unknown user name>	LogAlert	AuditLog service control & config
System	1326		Failed changing pass- word : '%user%'	LogAlert	AuditLog service

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
					control & config
System	1327		Password of '%user%' was changed by user	LogAlert	AuditLog service control & config
System	1328		Failed changing password for '%user%'	LogAlert	AuditLog service control & config
System	1329		Failed expiring user : <unknown user name>	LogAlert	AuditLog service control & config
System	1330		Failed expiring '%user%'	LogAlert	AuditLog service control & config
System	1331		User '%user%' expired	LogAlert	AuditLog service control & config
System	1332	Insufficient privilege to execute a command	"Insufficient privilege to execute: <command>"	LogWarning	AuditLog service control and config
System	1350	General system messages	<component> "<message_string>" <component> is the component in the EOS device that generated the message. Possible components are: GBF ROOT PPM2 SNMP LC LDAP AUtoStartUp INIT DHCP SCHELR	LogInfo	AuditLog service control & config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			<i><message_string></i> is the message logged. Example: "ISDN SwitchType changed; line is being initialized..." "DO flush_dmac" "Wed Feb 3 04:43:11 2001 Circuit UP LMAC 0A804440860A LSAP 04 RMAC 400040000002 RSAP 04 IP 10.5.146.250 "		
System	1351	General error system message	component "message_string" component is the component in the EOS device that generated the message. Possible components are: GBF ROOT PPM2 SNMP LC LDAP AUtoStartUp INIT DHCP SCHELR <i><message_string></i> is the message logged. Example: "ISDN SwitchType changed; line is being initialized..." "DO flush_dmac" "Wed Feb 3 04:43:11 2001 Circuit UP LMAC 0A804440860A LSAP 04 RMAC 400040000002 RSAP 04 IP 10.5.146.250 "	LogERror	AuditLog service control and config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
System	1361	Firmware authentication success	"Authentication of firm-ware succeeded"	LogNotice	AuditLog service control and config
System	1362	Firmware authentication failure	"Authentication of firm-ware failed"	LogERor	AuditLog service control and config
System	1398	General system messages	Status of messages: FAILED "Unsuccessful attempt to change date by user" "Unsuccessful attempt to change time zone by user" " Unsuccessful attempt to update time by NTP. NTP server: ..."	LogNotice	AuditLog service control and config
System	1399	General system messages	Status of messages: SUCCESS "Changed LNM password (encrypted): *****" "Time zone changed. Old date: ... " "Time updated by NTP. Old date: ... " ... LoadConfigs, Boot.cfg, init messages	LogNotice	AuditLog service control and config
System	1500	Telnet login failure	"Login Failed - Invalid Credentials"	LogAlert	AuditLog service control & config
System	1501	Telnet login success	"Login Successful"	LogAlert	AuditLog service control & config
System	1502	Telnet session termination	"UI session terminated"	LogAlert	AuditLog service control & config
System	1503	Telnet listen command	"Listen executed"	LogAlert	AuditLog service control & config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
System	1504	Telnet privilege change failure	"Set Privilege Failed"	LogAlert	AuditLog service control & config
System	1600	SSH login failure	"SSH Session from <IP Address>: Login Failed" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session. For details on how MNR routers and GGM 8000 gateways log unsuccessful SSH login attempts, see SSH "Login Failed" Messages on page 1024.	LOG_AUTH LogAlert	AuditLog service control & config
System	1601	SSH login successful	"SSH Session from IP Address: Login Successful" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session.	LOG_AUTH LogAlert	AuditLog service control & config
System	1602	SSH user-name change not allowed	"SSH Session from <IP Address>: Change of username or service not allowed" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session.	LOG_AUTH LogAlert	AuditLog service control & config
System	1603	SSH insufficient privilege	"SSH Session from <IP Address>: Insufficient Privilege" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session.	LOG_AUTH LogAlert	AuditLog service control & config
System	1604	SSH login timeout	"SSH Session from <IP Address>: Exceeded maximum graceful login time"	LOG_AUTH LogAlert	AuditLog service control & config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			In the actual message, <i><IP Address></i> is replaced by the IP address of the client that opened the SSH session.		
System	1605	SSH service disabled	"SSH Session from <i><IP Address></i>: SSH Service is disabled" In the actual message, <i><IP Address></i> is replaced by the IP address of the client that opened the SSH session.	LOG_AUTH LogAlert	AuditLog service control & config
System	1606	SSH public key exchange failure	"SSH Session from <i><IP Address></i>: User Public Key Authentication cannot decode key" In the actual message, <i><IP Address></i> is replaced by the IP address of the client that opened the SSH session.	LOG_AUTH LogAlert	AuditLog service control & config
System	1607	SSH public key exchange failure	"SSH Session from <i><IP Address></i>: User Public Key Authentication type mismatch" In the actual message, <i><IP Address></i> is replaced by the IP address of the client that opened the SSH session.	LOG_AUTH LogAlert	AuditLog service control & config
System	1608	SSH public key exchange failure	"SSH Session from <i><IP Address></i>: User Public Key Authentication failed" In the actual message, <i><IP Address></i> is replaced by the IP address of the client that opened the SSH session.	LOG_AUTH LogAlert	AuditLog service control & config
System	1609	SSH no client or server	"SSH Session from <i><IP Address></i>: Client or Server is NULL" In the actual message, <i><IP Address></i> is replaced by the IP address	LOG_AUTH LogAlert	AuditLog service control & config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			of the client that opened the SSH session.		
System	1610	SSH Cipher Mismatch	"SSH Session from <IP Address>: Cipher Mismatch: client <client name> server <server name>" In the actual message, <IP Address> and <client name> are replaced by the IP address and name of the client that opened the SSH session and <server name> is replaced by the name of the SSH server.	LOG_AUTH LogAlert	AuditLog service control & config
System	1611	SSH Unsupported Cipher	"SSH Session from <IP Address>: Unsupported Cipher <cipher name>" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session and <cipher name> is replaced by the name of the cipher algorithm.	LOG_AUTH LogAlert	AuditLog service control & config
System	1613	SSH Unsupported Mac	"SSH Session from <IP Address>: Unsupported mac <mac name>" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session and <mac name> is replaced by the name of the unsupported message authentication code (MAC).	LOG_AUTH LogAlert	AuditLog service control & config
System	1614	SSH No Key Exchange Algorithm	"SSH Session from <IP Address>: No kex algorithm" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session.	LOG_AUTH LogAlert	AuditLog service control & config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
System	1615	SSH Bad Key Exchange Algorithm	"SSH Session from <IP Address>: Bad Kex algorithm <algorithm name>" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session and <algorithm name> is replaced by the name of the bad key exchange algorithm.	LOG_AUTH LogAlert	AuditLog service control & config
System	1616	SSH No Hostkey Algorithm	"SSH Session from <IP Address>: No hostkey algorithm" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session.	LOG_AUTH LogAlert	AuditLog service control & config
System	1617	SSH Bad Hostkey Algorithm	"SSH Session from <IP Address>: Bad hostkey algorithm <algorithm name>" In the actual message, <IP Address> is replaced by the IP address of the client that opened the SSH session and <algorithm name> is replaced by the name of the bad hostkey algorithm.	LOG_AUTH LogAlert	AuditLog service control & config
System	1618	SSH Successful rekeying	"SSH Session from <IP Address>: SSH Session rekeying successful"	LOG_AUTH LogAlert	AuditLog service control and config
System	1619	SSH Compression algorithm mismatch	"SSH Session from <IP Address>: Compression algorithm mismatch: client <client algorithms> server <server algorithms>"	LOG_AUTH LogAlert	AuditLog service control and config
System	1620	SSH unsupported Compression algorithm	"SSH Session from <IP Address>: Unsupported compression algorithm <algorithm name>"	LOG_AUTH LogAlert	AuditLog service control and config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
MIP	1700	MIP multicast group limitation	"MIP MC Grp Lim, Port ! <port>, Rej IP = <<IP Address>, Src IP = <IP Address>" where <port> is the port number, <Rej IP address> is the multicast IP address of the rejected multicast group entry, and <Src IP address> is the source IP address of the rejected multicast group entry.	LogWarning	AuditLog service control and config
PIM	1701	PIM multicast group limitation	"PIM MC Grp Lim, <type>, Rej IP = <Rej IP address>, Src IP = <Src IP address>" where <type> is one of the following: IGMP Joins dropped by PIM PIM Joins/Prunes PIM Registers PIM Mcast Data Packets PIM Unknown Reject Type <Rej IP address> is the multicast IP address of the rejected multicast group entry, and <Src IP address> is the source IP address of the rejected multicast group entry	LogWarning	AuditLog service control and config
PKI	1801	Unsuccessful attempt to validate a certificate	Invalid certificate for subject: <subject> Reason: number (INVALID: <reason>) <reason> can be: "There is no trusted root certificate" "The certificate named '<profile_name>' expired" "The certificate named '<profile_name>' was revoked"	LOG_AUTH LogWarning	AuditLog service control and config

Service	Msg ID	Message name/type	Logged STRING	Priority	CONTROL Parameters
			"Some certificate had an invalid extension" "Some certificate had an invalid version number" "Some certificate had an invalid format" "Broken input certificate chain" "Some certificate had an invalid time range" "Some certificate had an invalid signature"		
System	9999	Console command	any configuration change command that was executed successfully Example: UI cmd "setd !4 -PATH CONTrol = e" "me -PATH" executed	LogWarning	AuditLog service control & config

H.1

SSH "Login Failed" Messages

Traps are not generated for SSH events; instead, MNR routers and GGM 8000 gateways send an SSH "Login Failed" log message for each failed login attempt (as opposed to each unsuccessful SSH session). In addition, a client that attempts to log in via an authentication method that is not supported by EOS also triggers a "Login Failed" message.

- If a session consists of one or two failed password login attempts followed by a successful login, the EOS software sends a "Login Failed" log message for each failed password login attempt, followed by a "Login Successful" message for the successful login.
- If a session consists of three failed password login attempts, the EOS software returns a "Login Failed" message for each failed attempt and disconnects the SSH session.
- If a session consists of a single successful password login attempt, the router or gateway may generate up to two "Login Failed" messages before the "Login Successful" message. This is because most clients try to log in with other authentication methods before attempting password authentication. The login attempts that use authentication methods not supported by EOS generate "Login Failed" messages prior to the "Login Successful" message that is sent upon the successful password authentication login.

Appendix I

Regular Expressions

This appendix describes the regular expressions used for creating and displaying ASpath- based filters (AsPolicyAll, AsPolicyExt, AsPolicyInt, AsPolicyPeer, and DisplayFilter parameters) in the BGP Service and for altering SHow and SHowDefault displays with the GREP command.

You can use regular expressions to specify a general string which can then be used for pattern matching.

A regular expression is a formula for generating a set of strings. If a particular string can be generated by a given regular expression, that string and regular expression match.

I.1

Regular Expression Components

A regular expression consists of different components described in the following table, each of which is used to build the string-generating program.

Table 121: Regular Expression Components

Regular Expression	Function
c	Use an ordinary ASCII character (excluding the special characters) to match that character. For example, c matches a lowercase c.
.	Use a period (.) to match any character except NEWLINE.
*	Use a regular expression followed by an asterisk (*) to match zero or more occurrences of the one-character regular expression. If there is any choice, the longest leftmost string that permits a match is chosen.
+	Use a regular expression followed by plus sign (+) to match one or more occurrences of the one-character regular expression. If there is any choice, the longest leftmost string that permits a match is chosen.
?	Use a regular expression followed by a question mark (?) to match zero or one occurrences of the one-character regular expression. If there is any choice, the longest leftmost string that permits a match is chosen.
 	Use two regular expressions separated by vertical bar () or NEWLINE to match either the first or the second (logical OR operation).
()	Use a regular expression enclosed in parentheses to match the regular expression.
\. *	Use a backslash (\) followed by any special character (period, asterisk, left square bracket, backslash) to match the special character.
\[\] \\	These characters are special except when they are enclosed within square brackets ([]). For example, * matches an asterisk (*).
[string]	Use a non-empty string of characters enclosed in square brackets to match any one character in that string. If the first character of the string is a caret (^), the one-character regular expression matches any character except NEWLINE and the remaining characters in the string. The caret has this special meaning only if it occurs first in the string.

Regular Expression	Function
	The hyphen (-) indicates a range of consecutive ASCII characters; for example, [0-9] is equivalent to [0123456789]. The hyphen loses this special meaning if it occurs first (after an initial ^, if any) or last in the string. The right square bracket (]) does not terminate a string when it is the first character within it (after an initial ^, if any); for example, []a-f] matches either a right square bracket or one of the letters a through f inclusive. The period, asterisk, left square bracket, and backslash represent themselves within such a string of characters.
Concatenation: The remaining regular expressions are for concatenation, a regular expression that matches the concatenation of the strings matched by each component of the regular expression.	
\<	Use the sequence \< in a regular expression to constrain the one-character regular expression immediately following it only to match something at the beginning of a "word;" for example, either at the beginning of a line, or just before a letter, digit, or underline and after a character not one of these.
\>	Use the sequence \> in a regular expression to constrain the one-character regular expression immediately following it only to match something at the end of a "word;" for example, either at the end of a line, or just before a character, which is neither a letter, digit, nor underline.
\(and \)	Use a regular expression enclosed between the character sequences \ (and \) to match whatever the unadorned regular expression matches.
\{m\}	Use a regular expression followed by \{m\}, \{m,\}, or \{m,n\} to match a range of occurrences of the regular expression. The values of m and n must be non-negative integers less than 256. \{m\} matches exactly m occurrences. \{m,\} matches at least m occurrences. \{m,n\} matches any number of occurrences between m and n inclusive. Whenever a choice exists, the regular expression matches as many occurrences as possible.
\{m,\}	
\{m,n\}	
\n	Use the expression \n to match the same string of characters that was matched by an expression enclosed between \ (and \) earlier in the same regular expression. n is a digit; the subexpression specified begins with the nth occurrence of \ (counting from the left). For example, the expression ^\(.*)\1\$ matches a line consisting of two repeated appearances of the same string.
^	Use the caret (^) at the beginning of a regular expression to constrain the regular expression to match an initial segment of a line.
\$	Use the dollar sign (\$) at the end of a regular expression to constrain the regular expression to match a final segment of a line. For example, ^entire regular expression \$ constrains the regular expression to match the entire line.

I.2

AS Filter Examples

You can use regular expressions in autonomous system (AS) filters by using the following syntax:

```
ADD -BGP ASFilter <AsfilterID> "<regular_expression>"
```

Blank spaces are represented here as underscores (_). When two spaces are shown together, a space is inserted between the underscores, for example __. You must enter a blank space for each underscore shown in the examples.

AS Filter Example 1

To create filter 1 that identifies an AS-path attribute containing AS25, you must use:

```
ADD -BGP ASFilter 1 "_25_"
```

AS Filter Example 2

To create filter 2 that identifies an AS-path attribute containing AS35 and AS50 (in this order), you must use:

```
ADD -BGP ASFilter 2 "_35_.*_50_"
```

The ".*" indicates a single character followed by any number of unspecified characters.

AS Filter Example 3

To create filter 3 that identifies an AS-path attribute containing AS35 and AS50 (in any order), you must use:

```
ADD -BGP ASFilter 3 "_35_.*_50_|_50_.*_35_"
```

The "|" indicates a logical OR operation.

AS Filter Example 4

To create filter 4 that identifies an AS-path attribute containing the AS sequence <AS5, AS46, AS32>, you must use:

```
ADD -BGP ASFilter 4 "<_5_ _46_ _32_>"
```

AS Filter Example 5

To create filter 5 that identifies an AS-path attribute containing the AS set [AS5, AS46, AS32], enter:

```
ADD -BGP ASFilter 5 "[_5_ _32_ _46_]"
```

AS sets are always sorted from lowest AS to highest AS.

I.3

GREP Command Examples

You can use regular expressions with the GREP filter by using the following syntax:

```
<COMMAND> <parameters> [<options>] | GREP [-v] [-i] <grep_pattern>
```

You can use the GREP command with the SHow and SHowDefault commands to alter the output to display the information you specify in the GREP pattern.

You cannot apply multiple GREP commands to a single UI command. In other words, the following command to show all IP routes that have the number 192 and 128 is not supported:

```
SHoW -IP AllRoutes | GREP 192 | GREP 128
```

Example 1

To display all IP routes, you normally enter the `SHoW` command. To show only those routes that have the number 152 in them, you can pipe the `SHoW` or `SHoWDefault` output to the `GREP` command by using:

```
SHoW -IP AllRoutes | GREP 152
```

A display similar to the following appears:

```
129.213.152.0      255.255.252.0      129.213.200.109    1    UP    RIP
```

The number 152 in regular expression form is a string of ASCII characters that are matched, generating the information you specified.

Example 2

To display all IP routes that do not have the number 152 in them, pipe the output to the `GREP` command by entering:

```
SHoW -IP AllRoutes | GREP -V 152
```

A display similar to the following appears:

----- IP Routing Table -----						
Total Routes = 12, Total Direct Networks = 1						
Destina- tion	Mask	Gateway	Metric	Status	TTL	Source
0.0.0.0	0.0.0.0	129.213.20 0.109	2	UP	170	RIP
		129.213.20 0.103	3	UP	150	RIP
129.213.16 .0	255.255.25 2.0	129.213.20 0.109	1	UP	170	RIP
		129.213.20 0.108	1	UP	170	RIP
129.213.32 .0	255.255.25 2.0	129.213.20 0.109	1	UP	170	RIP
129.213.48 .0	255.255.25 2.0	129.213.20 0.109	1	UP	170	RIP
		129.213.20 0.102	1	UP	170	RIP
129.213.72 .0	255.255.25 2.0	129.213.20 0.109	1	UP		RIP
129.213.96 .0	255.255.25 2.0	129.213.20 0.103	1	UP		RIP
129.213.20 0.0	255.255.25 2.0	129.213.20 3.16	0	UP	--	Connected
129.213.24 0.0	255.255.25 2.0	129.213.20 0.109	1	UP		RIP

The `-v` option for `GREP` performs the NOT (or invert) operation on the display information and lists all the IP routes that do not have the number 152 in them. For descriptions of the `GREP` options, refer to the `GREP` command description in “Global Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

Example 3

To display all IP routes containing the number 72 or 96, pipe the output to the `GREP` command by entering:

```
How -IP AllRoutes | GREP 72 | 96
```

A display similar to the following appears:

```
129.213.152.0      255.255.252.0      129.213.200.109    1    UP    RIP
-----
```

The first “|” represents the pipe to the `GREP` command. The “|” between the numbers is the regular expression for logical OR. In this example, the output matches any string that contains the numbers 72 or 96.

Appendix J

Wide Area Network Setup Information

This appendix provides information to help you set up your wide area serial ports.

J.1

Serial Port Utilization Percentage

The percentage of utilization displayed for FlexWAN and HSSI ports is based on a full-duplex link. For example, a 64-kbps circuit can transmit and receive 64 kbps simultaneously. If this link were transmitting at 64 kbps and receiving nothing, the percentage of utilization would be 50 percent.

To display utilization information, you must use:

```
SHow -SYS STATistics -PATH
```

J.2

Serial Line Connectivity

External Device Connections

When configuring your EOS device for remote communications with external clocking devices, Motorola Solutions recommends that you set the `-PATH CONNECTor` and `CLock` parameters to match the external devices before physically connecting the two devices.

External Device Cable Length

At serial line data rates greater than 56 kbps on RS-232 interfaces, Motorola Solutions recommends using cables less than 25 feet long.

Serial Line Clocking

When using an EOS device with a modem or channel service unit/data service unit (CSU/DSU), you must determine which piece of equipment provides the clocking signal.

- For an EOS device using an external device over dial-up lines, the external device must provide clocking.
- For an EOS device using an external device over a leased line, either device can provide clocking.

If you configure an external device to provide clocking to an EOS device, the external device must use the return clock provided by the EOS device. Neither the external device nor the cable used to connect it to the EOS device should loop back the transmit clock to the receive clock.

Serial Line Supported Data Rates

EOS software supports serial line data rates above 2,048 kbps for external clocking devices only. If the selected data rate is above 2,048 kbps with internal clocking (Test Mode) specified, the actual data rate used by the EOS device is 2,048 kbps.

When using external clocking devices, set the data rate of the serial line on the EOS device as close as possible to the external device. Packet processing is optimized to this value. Statistical reporting of line utilization is based on the data rate you configure.

J.3

Network Clock Synchronization

The Network Time Protocol (NTP) provides a standard mechanism to synchronize the computer clock in a distributed network. The NTP timescale is based on Universal Time Coordinated (UTC). The EOS device receives UTC time from the NTP server and adjusts it based on the configuration of the `TimeZone` and `DSTime` parameters in the EOS SYStem service before setting the system clock.

NTP support in EOS devices is based on RFC 1305. The EOS devices are clients in the NTP client /server model, and the EOS device NTP client does not use open source or third party software. EOS devices also support a broadcasting server configuration. If the broadcasting NTP server is configured, the EOS device NTP client uses that server to update the local system clock across zones if the NTP service is not available across the zones.

As NTP clients, EOS devices periodically poll the designated NTP server for time of day and process the reply to determine the local clock update. You identify the primary NTP server using the NTP service `PrimarySrvr` parameter and set the minimum polling interval using the NTP service `MinPollIntrvl` parameter.

If the primary NTP server fails, the client device switches over to the designated backup server. You identify the backup NTP server using the NTP service `BackupSrvr` parameter.

If either the primary or the secondary NTP server is configured at startup, the EOS device polls the server at the interval specified by the `MinPollIntrvl` parameter (the default is 64 seconds). If the server does not reply after 8 polls, the EOS device sends an “NTP server unreachable” trap. NTP continues to poll the server after sending the “NTP server unreachable” trap. When the server is reachable again, the EOS device then synchronizes with the server. If neither the primary nor the secondary NTP server is configured at startup, the device sends an “NTP server unconfigured” trap.

If the NTP service is not available (either because no server is configured or the configured server is unreachable), the EOS device runs on local clocks.

For details about NTP service parameters, see “NTP Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*.

J.3.1

Source Quality Clocking

An important consideration regarding NTP server switchover is the quality of the servers’ clocking sources. Clock precision is indicated by Stratum number—the lower the Stratum number, the higher the clock precision. The EOS device’s Stratum number is based on the timing source of the NTP server with which the device is synchronized:

$\text{EOS Device Stratum} = \text{NTP Server Timing Source Stratum} + 1$

You can display the EOS device’s Stratum number by issuing the following command:

```
SHoW -NTP STATUS
```

If the primary NTP server fails, the EOS device does not switch over to a backup NTP server with an inferior clock. Thus the Stratum number of the backup NTP server’s timing source must be equal to or less than the EOS device’s Stratum number in order for a switchover to occur.

For example, consider a scenario in which the EOS device is synchronized with a primary NTP server with Stratum number 1. The EOS device’s Stratum number is then $1+1=2$. In this case, the Stratum number of the backup NTP server’s timing source must be 1 or 2 in order for a switchover to occur. If the Stratum number of the backup NTP server’s timing source is 3, 4, or 5 the EOS device does not switch over to the backup server if the primary server fails.

Appendix K

Delay and Jitter Performance Management Tool

The EOS delay and jitter performance management tool allows you to determine if the link latency (delay), packet loss, and jitter cross tolerable limits on Ethernet WAN links.

EOS allows you to configure Ethernet WAN link performance thresholds (for link latency (delay), jitter, and packet loss) on a per link basis. A single EOS device may have many Ethernet WAN links configured, each with different link characteristics. You can customize the Ethernet WAN link performance thresholds for each link using the parameters in the EOS IPDV service. For details, see "IPDV Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*.

K.1

Delay and Jitter Statistics View

To view delay and jitter statistics, you can use:

```
SHoW !<port> -IPDV STATIStics
```

where <port> is the Ethernet port for which you want to view statistics. For example:

```
SHoW !v2 -IPDV STATIStics
```

Calculated statistics are displayed at the end of the measurement interval. The measurement interval is configured with the `MeasureInt` parameter. For a detailed description of the `SHoW -IPDV STATIStics` display, see the description of the `STATIStics` parameter in "IPDV Service Parameters" of the *Enterprise OS Software Version 18.2 Reference Guide*.

K.2

Delay and Jitter Performance Management Tool Functionality

The delay and jitter performance management tool uses streams of probes, sent as encapsulated UDP packets to port number 49402. Timers are implemented so that all measurements start at the beginning of the hour and minute and are accurate over a 60-millisecond period.

- For each Ethernet WAN link, the EOS device sends a stream of probes every five minutes to its peer.
- Within each five-minute interval, the EOS device sends a jitter stream and a latency (delay) stream. The packets include special bits that indicate whether they are part of a jitter stream or a latency stream. Only packets that are tagged for latency calculations are sent back to the originator.
 - The jitter stream consists of a total of 201 packets, separated by 60-millisecond intervals. The jitter stream is one way and is also used to calculate packet loss.
 - The latency stream consists of a total of five packets, separated by one-minute intervals. The latency stream is round trip.

EOS devices use BFD to detect link liveness on Ethernet site links and BGP keepalives to detect link liveness on Ethernet interzone links. When the BFD or BGP peer status changes, the EOS device sends an Ethernet WAN link object state change trap and notifies the delay and jitter performance management tool.

EOS also checks the link state at the beginning and end of each jitter packet stream and uses this information to determine whether or not an IPDV 99 percentile calculation should be performed for that stream. If the link

is down or not all jitter packets are received in a stream, that jitter measurement is skipped and not used for the IPDV 99 percentile calculation.



NOTE: The purpose of the delay and jitter performance management tool is to evaluate link quality as opposed to device performance; therefore, packets dropped due to priority queue limits being exceeded do not affect the delay and jitter tool measurements.

K.2.1

Latency (IPTD) Calculations

Latency, or IPTD (IP transfer delay) is the response time between link endpoints, in milliseconds. If the originator sends a probe stream at time T1; the receiver receives the probe stream at time T2; the receiver sends the probe stream back to the originator at time T3; and the originator receives the probe stream at time T4; latency is calculated as follows:

$$(T4 - T1) - (T3 - T2) = (\text{Total time for the response}) - (\text{time spent at the receiver})$$

EOS devices perform the following latency calculations each measurement interval:

- Number of successful measurements
- Sum of the measurements
This is a running total for each successful measurement.
- Average value of the measurements
- Sum of squares of all the measurements
- Minimum measurement
- Maximum measurement

You can set a threshold for the average IPTD value using the IPDV service `AvgRtXfer` parameter. For details see the description of this parameter in “IPDV Service Parameters” in the *Enterprise OS Software Version 18.2 Reference Guide*. When the configured threshold is exceeded, EOS sends a trap and alarm. For details, see [Ethernet WAN Link Object MIB on page 851](#).

K.2.2

Jitter (IPDV) Calculations

Jitter, or IPDV (IP delay variation), is the variation in latency between consecutive packets within a jitter stream sent from an originator to a receiver, in milliseconds. When the packets are received, the receiver stores the following timestamps:

- T1 — The timestamp sent with probe packet 1 by the originator.
- T2 — The receiver’s current timestamp when probe packet 1 is received.
- C1 — The timestamp sent with probe packet 2 by the originator.
- C2 — The receiver’s current timestamp when probe packet 2 is received.

Delay for probe packet one (D1) = T2 - T1

Delay for probe packet two (D2) = C2 - C1

Therefore jitter (J1) is calculated as:

$$J1 = D1 - D2 = (C2 - C1) - (T2 - T1) =$$

(Delay between when probe packet 2 was sent and received) – (Delay between when probe packet 1 was sent and received)

J2 calculates the jitter for probe packets 2 and 3. Each calculation of J1, J2, and so on is called the “nominal Delta.” For a jitter stream of 201 probe packets, there are 200 nominal Delta calculations. The “running sum”

is the sum of the current and all previous samples' nominal Deltas. The "minimum sum" is the minimum running sum recorded for the jitter stream. For each sample, IPDV is calculated as:

$IPDV = \text{runningSum} - \text{minimumSum}$

EOS devices perform the following IPDV calculations each measurement interval:

- Number of successful measurements
(the number of successful IPDV calculations)
- Sum of the measurements
- Average value of the measurements
- Sum of squares of the measurements

EOS devices also perform 99 percentile calculations for IPDV. The IPDV 99 percentile value is the smallest number that is greater than 99% of the IPDV values in the set.

EOS devices perform the following IPDV 99 percentile calculations each measurement interval:

- Number of successful measurements
The number of jitter streams successfully processed per configured measurement interval. If the interval is 15 minutes (configured with the `MeasureInt` parameter), then this number should be 3 (one stream per five minutes).
- Sum of the IPDV 99 percentile values
Each IPDV 99 percentile calculation returns a number; this is the sum of those numbers.
- Average IPDV 99 percentile value
- Minimum IPDV 99 percentile value
- Maximum IPDV 99 percentile value

You can set a threshold for the average IPDV 99th percentile value using the IPDV service `Avg99Percent` parameter. For details see the description of this parameter in "IPDV Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*. When the configured threshold is exceeded, EOS sends a trap and alarm. For details, see [Ethernet WAN Link Object MIB on page 851](#).

K.2.3

Packet Loss Rate (IPLR) Calculations

Packet loss is the number of packets transmitted by the sender but not received by the receiver. Packet loss calculations are performed using the same probe stream as that used for jitter calculations.

EOS devices perform the following packet loss calculations each measurement interval:

- Number of packets transmitted
- Number of packets not received by the receiver
- Packet loss rate (packets lost / packets transmitted), reported in units of one hundredth percent.

You can set a threshold for the packet loss rate value using the IPDV service `PktRate` parameter. For details see the description of this parameter in "IPDV Service Parameters" in the *Enterprise OS Software Version 18.2 Reference Guide*. When the configured threshold is exceeded, EOS sends a trap and alarm. For details, see [Ethernet WAN Link Object MIB on page 851](#).

Appendix L

Upgrading EOS Software

This appendix provides a recommended procedure for upgrading EOS software in a way that avoids scenarios in which the EOS device might be rebooted without a software image (`boot.ppc` file) during the upgrade process.



WARNING: EOS devices cannot boot without a `boot.ppc` file. For this reason, prior to upgrading EOS software, you should always verify that a `boot.ppc` file exists on the EOS device and that the device is configured to boot from the directory (primary or secondary) in which the `boot.ppc` file is located.

Procedure:

1. Enter:

```
SHow -SYS VERsion
```

2. Verify that the EOS device is configured to boot from the primary directory and perform one of the following actions:

If...	Then...
If the device is configured to boot from the primary directory,	go to step 3 .
If the device is configured to boot from the secondary directory,	perform the following actions: a. To open the Boot Sources menu, enter: <code>SysconF 7</code> b. To configure the device to boot from the primary directory, enter: <code>1</code> c. Return to the EOS device prompt.

3. Verify that the software in the primary directory matches the software that runs on the device by performing the following actions:
 - a. To display the current firmware version, enter:
`SHow -SYS VERsion`
 - b. To display the software package and version number stored in the device's primary directory, enter:
`SHow -SYS SOfwareInfo a:/primary/boot.ppc`
 - c. If the software versions in steps [step 3a](#) and [step 3b](#) do not match, stop the upgrade process, investigate how this mismatch occurred, and take actions as necessary to bring the device to a state where the software in the primary directory matches the software that is running on the device.

If a rollback of the EOS version is necessary, it is important to start from a known point.
4. Copy the software and configuration files from the primary to the secondary directory by entering:
`copy a:/primary/* a:/secondar`
5. To display the software package and version number stored in the device's secondary directory and verify that the copy operation was performed successfully, enter:

```
SHoW -SYS SOftwareInfo a:/secondar/boot.ppc
```

You now have a backup version of the current software in the device's secondary directory that can be used to boot the device if an error occurs when you are transferring the new software version to the device in the next step.

6. Load the new version of EOS software and the corresponding configuration files to the device's primary directory.



IMPORTANT: If the installed EOS version is 17.6.0.21 or newer you can upgrade it to any more recent EOS version keeping existing configuration files.

7. To display the software package and version number stored in the device's primary directory and verify that the copy operation was performed successfully, enter:

```
SHoW -SYS SOftwareInfo a:/primary/boot.ppc
```



NOTE: Starting with the 18.2.0.32 firmware, EOS stores passwords encrypted with the PBKDF2 algorithm. Passwords for the 'root' and 'admin' accounts are automatically converted to the new format during the upgrade. All accounts other than 'root' or 'admin (if defined)' are deleted and have to be recreated after the upgrade.

8. Configure the device to boot from the primary directory first and if that fails to boot from the secondary directory by performing the following actions:

- a. To open the **Boot Sources** menu, enter:

```
SysconF 7
```

- b. To configure the device to boot from the primary and secondary directories, enter: 2

- c. Return to the EOS device prompt.

9. Reboot the device by entering:

```
ReBoot
```

10. Verify that the software in the primary directory matches the software that is running on the device by performing the following actions:

- a. To display the current firmware version, enter:

```
SHoW -SYS VERSion
```

- b. To display the software package and version number stored in the device's primary directory, enter:

```
SHoW -SYS SOftwareInfo a:/primary/boot.ppc
```

If the software versions do not match, you must use

```
SHoW -SYS SOftwareInfo a:/secondar/boot.ppc
```

to display the software package and version number stored in the device's secondary directory. If that version matches the version displayed with the `SHoW -SYS VERSion` command, you know that the reboot from the primary directory failed, and the device booted from the old version of software that you copied to the secondary directory in [step 6](#). You must return to [step 7](#) and download the new version of EOS software to the device's primary directory again.

If the software versions in [step 10a](#) and [step 10b](#) match, the upgrade is complete.

11. If the software versions do not match, perform the following actions:

- a. To display the software package and version number stored in the secondary directory of the device, enter:

```
SHoW -SYS SOftwareInfo a:/secondar/boot.ppc
```


- If that version matches the version displayed with the `SHoW -SYS VERSiOn` command in [step 10a](#), the reboot from the primary directory failed, and the device booted from the old version of software that you copied to the secondary directory in [step 4](#).
- b.** Download the new version of EOS software to the device's primary directory again by repeating [step 6](#) to [step 10](#).

Appendix M

Downgrading EOS Software

Starting with 16.8.0.19 firmware, EOS is signed with the **SHA2withRSA2048** signature algorithm. Once the newer firmware is loaded, the previous firmware cannot be loaded again without changing the algorithm to **SHA1withRSA1024**. The downgrade is possible by using a command to both the EOS and the Boot Monitor that allows bypassing the requirement of the signature algorithm.

Starting with the 18.2.0.32 firmware, EOS stores user passwords encrypted with the PBKDF2 algorithm. Downgrade can be performed by issuing a command in EOS and converting user passwords through logging in.



WARNING:

By downgrading to the firmware releases that use the **SHA1withRSA1024** signature algorithm, you render the device non-compliant with Federal Information Processing Standards (FIPS).

By downgrading to the firmware releases that use the **MD5** signature algorithm, you render the device non-compliant with Federal Information Processing Standards (FIPS).

Procedure:

1. To safely downgrade EOS software (starting with 16.8.0.19), perform the following actions:

- a. To check the version of firmware on EOS device, enter:

```
sh ver
```

- b. Depending on the software version, perform one of the following actions:

- For version 16.8.0.18 or lower, use the **SHA1withRSA1024** signature algorithm.
- For version 16.8.0.19 or higher, go to [step 1c](#).



NOTE: Starting with version 16.8.0.20, the **SHA2withRSA2048** signature algorithm is set by default at the start of MNR devices.

- c. To change the **SWSignAlgorithm** parameter, enter:

```
setd -sys SwSignAlgorithm = SHA1withRSA1024
```



NOTE: If this command is not performed before loading the older firmware, an error message appears. After bootup, the MNR device may not have firmware to run if the MNR device is configured with only one boot source (a:/primary or a:/secondary) as the firmware location. The MNR device can be recovered from the Boot Monitor.

2. To safely downgrade EOS software (starting with 18.2.0.32), perform the following actions:

- a. To check the version of firmware on EOS device, enter:

```
sh ver
```

- b. Depending on the software version, perform one of the following actions:

- For version 18.2.0.31 or lower, no further action is required.
- For version higher than 18.2.0.31, go to [step 2c](#).

- c. Change the **PasswordKDF** parameter by logging on as **root** and entering:

```
setd -sys PWKDF = MD5
```

- d. Set the **RESolutionOrder** parameter to **Local** by entering:

```
setd -ac reso = Local
```



NOTE: This command disables the RADIUS authentication. If configured in `boot.cfg`, it is restored on reboot.

- e. Save the passwords in the old format by logging in as **root**, **admin** and other users (if defined).

The logging operations require locally set passwords. They may differ from passwords used by RADIUS.

All unconverted local users are shown in the **UserManager** menu, under the **List all not downgraded accounts** option.

If this command is not performed for **root** or **admin** users before booting the older firmware, a warning message appears after executing the `reboot` command. After the bootup, the MNR device has its **root** and **admin** passwords reset to the default values and any existing local users have to be recreated.

Appendix N

Glossary

This glossary provides definitions of acronyms used in this manual:

3DES

Triple DES (encryption algorithm)

4W

Four-wire interface (transmit and receive on different copper pairs)

AAA

Authentication, Authorization, and Accounting

ABR

Area Border Router

AC

Access Control

AC-SIG

Analog Conventional Signaling and Control

ACIM

ASTRO Console Interface Module

AES

Advanced Encryption Standard (encryption algorithm)

AIS

ASTRO® 25 infrastructure signaling

ALMT

Analog Link Monitor Tone

AM

Audio Manager

APN

Access Point Name

ARP

Address Resolution Protocol

AS

Autonomous System

ASBR

Autonomous System Boundary Router

ASN

Autonomous System Number

AT

Audio Transport

AVC

Aggregated Virtual Circuit

BDR

Backup Designated Router

BFD

Bi-Directional Forwarding Detection

BGP	Border Gateway Protocol
BOD	Bandwidth-on-Demand
BOOTPC	BootP Client
BPDU	Bridge Protocol Data Unit
BR	Border Router
BRI	Basic Rate Interface
BSR	Bootstrap Router
CBPDU	Configuration Bridge Protocol Data Unit
CBQ	Class-Based Queuing
CCC	Conventional Call Controller (a Zone Controller or Conventional Site Controller depending on the mode of the conventional site)
CCGW	Conventional Channel Gateway
CCITT	Consultative Committee for International Telegraph and Telephone
CCS	Common Channel Signaling (ISDN); Compact Configuration Store (CCS files in the <code>boot.cfg</code> file, for example)
CD	Carrier Detect (signal); compact disk; collision detection
CDR	Call Detail Record
CEN	Customer Enterprise Network
CG	Charging Gateway
CGF	Charging Gateway Functionality
CH	Conventional channel; a channel on which conventional audio via 4W analog interface is present. A conventional channel may be a physical conventional base station, or it could be a 4W analog signal attached to a comparator that manages several 4W conventional interfaces from conventional stations and presents it as a single 4W analog signal to the system.
CHAP	Challenge Handshake Authentication Protocol

CIR

Committed Information Rate

CLI

Command Line Interface

CM

Call Manager

CNI

Customer Network Interface

CoA

Care of Address (in conventional data redundancy (HA and DSR) topology)

COR

Carrier Operated Relay; another name for the base station's E-lead.

COS

Class of Service

CR

Carriage Return

CRC

Cyclic Redundancy Check

C RTP

Compressed Real-time Transport Protocol

CS

Communications Server

CSC

Conventional Site Controller

CSU

Channel Service Unit

CTS

Clear To Send

CU

Control Unit

CUG

Closed User Group

CVC

Constituent Virtual Circuit

CWR

Cooperative WAN Routing

DCD

Data Carrier Detected

DCE

Data Communications Equipment (EIA expansion); data circuit-terminating equipment (CCITT);
Distributed Computing Environment (OSF)

DEB

Destination Explicit Blocking

DEF

Destination Explicit Forwarding

DES

Data Encryption Standard (encryption algorithm)

DHCP

Dynamic Host Configuration Protocol

DIU

Digital Interface Unit

DLCI

Data Link Connection Identifier

DLT

Data Link Test

DN

Distinguished Name; Directory Number

DNS

Domain Name Service

DOD

Dial-on-demand

DR

Designated Router

DRT

Data Record Transfer

DS0

A single time slot of information in a TI/EI system

DSP

Digital Signal Processor

DSPD

DSP Driver

DSR

Dynamic System Resilience; also, Data Set Ready

DSU

Digital Service Unit

DTE

Data Terminal Equipment

DTR

Data Terminal Ready

DVMRP

Distance Vector Multicast Routing Protocol

E1

High capacity 2.048 Mbps transmission system used for voice and data primarily in Europe.

E-Lead

Ear lead; the CCGW's input closure, which is activated by the base station when the base station is receiving a radio call.

E & M

Ear and mouth

ECN

Explicit Congestion Notification

EMAD

E&M Analog Driver

ESP

Encapsulation Security Payload

FA

Foreign Agent (in conventional data redundancy (HA and DSR) topology)

FQDN

Fully Qualified Domain Name

FR

Frame Relay

G.728 START

An audio plane message defined in the XIS protocol indicating that audio packets are about to start for a transmission on a particular station.

G.728 STOP

An audio plane message defined in the XIS protocol indicating that audio packets are about to start for a transmission on a particular station.

Ga Interface

Charging data collection interface between a CDR-transmitting unit (GGSN or PDR) and a CDR-receiving unit (CGF).

G-CDR

GGSN Call Detail Record

GGSN

Gateway GPRS Support Node

Gi Interface

Reference point between GPRS and an external packet data network

GPRS

General Packet Radio Service

GSM

Global System for Mobile communication

GTP

GPRS Tunneling Protocol

GTP'

GTP Charging Protocol

GUI

Graphical User Interface

HA

Home Agent (in conventional data redundancy (HA and DSR) topology)

HA and DSR

Conventional data redundancy (HA and DSR) feature

HLR

Home Location Register

HSSI

High-Speed Serial Interface

ICMP

Internet Control Message Protocol

ICP	Inbound Control Packet
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IISIS	Integrated Intermediate System-to-Intermediate System
IKE	Internet Key Exchange
IMSI	International Mobile Subscriber Identity
IP	Internet Protocol
IPDV	IP Packet Delay Variation
IPIP	IP encapsulation within IP
IPLC	IP Link Converter
IPLR	IP Packet Loss Rate
IPTD	IP Packet Transfer Delay
IPQOS	IP Quality of Service
IPSec	Internet Security Protocol
IS	Intermediate System
ISDN	Integrated Services Digital Network
ISO	International Organization for Standardization
ITU-TSS	International Telecommunications Union–Telecommunications Standards Sector
IV&D	Integrated Voice and Data
KEK	Key Encryption Key
L2TP	Layer 2 Tunneling Protocol
LAN	Local Area Network

LAPB

Link Access Procedure, Balanced

LAPD

Link Access Procedure for D channel. A data transmission procedure used in ISDN systems on the D channel

LCN

Logical Channel Number

LCP

Link Control Protocol

LDAP

Lightweight Directory Access Protocol

LF

Line Feed

LMF

Line Management Function

LMI

Local Management Interface

LOBL

Line Operated Busy Light; describes the Vortex 2 system's detection of transmission of non-Vortex 2parallel console.

LS

Line Server

LSA

Link State Advertisement

LSR

Link State information Request

LSU

Link State Update

M-Lead

Mouth lead; the CCGW's output closure which is connected to the base station. When the conventional channel gateway activates its M-lead closure, it causes the base station to begin transmitting.

MAC

Media Access Control

MED

Multi-exit discriminator

MFR

Multilink Frame Relay

MIB

Management Information Base

MIP

Multicast Internet Protocol

MLP

Multilink Protocol

MMR

Multiple Manager Request

MNR	Motorola Solutions Network Router
MOSPF	Multicast Open Shortest Path First
MSK	Mobile Synchronization Key
MTU	Maximum Transmission Unit
NAT	Network Address Translation
NBMA	Non-Broadcast Multi-Access interfaces
NCP	Network Control Protocol
NFS	Network File System
NLPID	Network Layer Protocol Identifier
NMA	Network Management Alliance
NMS	Network Management System
NSA	National Security Agency
NSSA	Not So Stubby Area (OSPF)
NT1	Network Termination 1
NTP	Network Time Protocol
OCP	Outbound Control Packet
OSPF	Open Shortest Path First
PAP	Password Authentication Protocol
PCM	Physical Connection Management; Pulse Code Modulation (ISDN)
PDG	Packet Data Gateway; comprised of the PDR and the RNG
PDN	Packet Data Network
PDP	Packet Data Protocol

PDP Context

The information sets held in mobile subscribers and GPRS support nodes for a PDP address, identifies the user session within the GTP tunnel

PDR

Packet Data Router

PDU

Protocol Data Unit

PFS

Perfect Forward Secrecy

PIM-SM

Protocol Independent Multicasting - Sparse Mode

PPP

Point-to-Point Protocol

PPTP

Point-to-Point Tunneling Protocol

PRI

Primary Rate Interface

PSK

Pre-shared key

PSDN

Packet Switching Data Network

PSNP

Partial Sequence Number PDU

PSTN

Publish Switched Telephone Network; commercial land-based telecommunications

PVC

Permanent Virtual Circuit

QOS

Quality of Service

RADIUS

Remote Authentication Dial-in User Service

RARP

Reverse Address Resolution Protocol

RAS

Remote Authentication Service

RD

Route Designator; Received Data (signal)

RDP

Router Discovery Protocol

REMP

REMOte Polling

RF

Radio Frequency

RIB

Routing Information Base

RIP	Routing Information Protocol
RIIP	Routing Information Protocol for IP
RIPNG	Routing Information Protocol Next Generation
RFC	Request for Comments
RMON	Remote Monitoring
RNG	Radio Network Gateway
RNI	Radio Network Interface
RP	Rendezvous Point
RPB	Reverse Path Broadcasting
RPF	Reverse Path Forwarding
RPM	Reverse Path Multicasting
RIIP	Routing Information Protocol / Internet Protocol
RSVP	Resource Reservation Protocol
RTP	Routing Table Protocol
RTS	Request To Send
SA	Security Association
SCM	Site Conventional Manager
SCP	Secure Copy
SEB	Source Explicit Blocking
SEF	Source Explicit Forwarding
SGSN	Serving GPRS Support Node
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol

SPF

Shortest Path First

SPI

Security Parameters Index

SPID

Service Profile Identifiers

SPT

Shortest Path Tree

SSAP

Source Service Access Point

SSC

Solution Support Center

SSH

Secure Shell

STA

Spanning Tree Algorithm

STP

Spanning Tree Protocol

SVC

Switched Virtual Circuit

T1

High-capacity 1.544 Mbps transmission system used for voice and data primarily in North America

TA

Terminal Adapter

TACACS

Terminal Access Controller Access Control System

TCP

Transmission Control Protocol

TCP/IP

Transmission Control Protocol / Internet Protocol

TDV

Traffic Data Volume

TFTP

Trivial File Transfer Protocol

TGID

Talkgroup Identifier

TID

Tunnel ID

TOS

Type of Service

TRC

Tone Remote Control

TRSC

Tunnel Route Short Cut

TRPB

Truncated Reverse Path Broadcasting

TTL	Time-To-Live
UCS	User Configuration Server
UDP	User Datagram Protocol
UDPHelp	User Datagram Protocol Broadcast Helper
URID	Universal Resource Identifier
USM	User Security Model
VACM	View-based Access Model
VC	Virtual Connection (Frame Relay)
VCC	Virtual Channel Connection
VCI	Virtual Channel Identifier
VCID	Virtual Circuit Identifier
VLAN	Virtual LAN
VPI	Virtual Path Identifier
VOX	Voice Activity
VPN	Virtual Private Network
WAN	Wide Area Network
XIS	X-Zone Infrastructure Signaling
ZC	Zone Controller
ZNM	Zone Network Management