




Command Reference Guide

CoreBuilder® 3500
CoreBuilder 9000
CoreBuilder 9400
SuperStack® II Switch 3900
SuperStack II Switch 9300

<http://www.3com.com/>

Part No. 10013505
Published November 1999



3Com Corporation
5400 Bayfront Plaza
Santa Clara, California
95052-8145

Copyright © 1999, 3Com Corporation. All rights reserved. No part of this documentation may be reproduced in any form or by any means or used to make any derivative work (such as translation, transformation, or adaptation) without written permission from 3Com Corporation.

3Com Corporation reserves the right to revise this documentation and to make changes in content from time to time without obligation on the part of 3Com Corporation to provide notification of such revision or change.

3Com Corporation provides this documentation without warranty, term, or condition of any kind, either implied or expressed, including, but not limited to, the implied warranties, terms, or conditions of merchantability, satisfactory quality, and fitness for a particular purpose. 3Com may make improvements or changes in the product(s) and/or the program(s) described in this documentation at any time.

If there is any software on removable media described in this documentation, it is furnished under a license agreement included with the product as a separate document, in the hard copy documentation, or on the removable media in a directory file named LICENSE.TXT or !LICENSE.TXT. If you are unable to locate a copy, please contact 3Com and a copy will be provided to you.

UNITED STATES GOVERNMENT LEGEND

If you are a United States government agency, then this documentation and the software described herein are provided to you subject to the following:

All technical data and computer software are commercial in nature and developed solely at private expense. Software is delivered as "Commercial Computer Software" as defined in DFARS 252.227-7014 (June 1995) or as a "commercial item" as defined in FAR 2.101(a) and as such is provided with only such rights as are provided in 3Com's standard commercial license for the Software. Technical data is provided with limited rights only as provided in DFAR 252.227-7015 (Nov 1995) or FAR 52.227-14 (June 1987), whichever is applicable. You agree not to remove or deface any portion of any legend provided on any licensed program or documentation contained in, or delivered to you in conjunction with, this User Guide.

Unless otherwise indicated, 3Com registered trademarks are registered in the United States and may or may not be registered in other countries.

3Com, the 3Com logo, CoreBuilder, DynamicAccess, NETBuilder II, PACE, SmartAgent, SuperStack, and Transcend are registered trademarks of 3Com Corporation. 3Com Facts is a service mark of 3Com Corporation.

PostScript is a registered trademark of Adobe Systems, Inc. AppleTalk is a registered trademark of Apple Computer, Incorporated. Banyan and VINES are registered trademarks of Banyan Worldwide. DEC, DECnet, and PATHWORKS are registered trademarks of Compaq Computer Corporation. OpenView is a registered trademark of Hewlett-Packard Company. AIX, IBM, and NetView are registered trademarks and NetBIOS is a trademark of International Business Machines Corporation. Internet Explorer, Microsoft, MS-DOS, Windows, and Windows NT are registered trademarks of Microsoft Corporation. Netscape, Netscape Navigator, and the Netscape N and Ship's Wheel logos are registered trademarks of Netscape Communications Corporation in the United States and other countries. IPX, Novell, and NetWare are registered trademarks of Novell, Inc. Sun and SunNet Manager are trademarks of Sun Microsystems, Inc. Xerox and XNS are trademarks of Xerox Corporation. UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

All other company and product names may be trademarks of the respective companies with which they are associated.

CONTENTS

ABOUT THIS GUIDE

Using This Book	20
Finding Specific Information in This Guide	20
Command Information	22
Recommendations for Entering Commands	23
Conventions	23
Documentation Comments	25
Year 2000 Compliance	25

PART I GETTING STARTED

1 ADMINISTRATION OVERVIEW

Administration Console Overview	29
CoreBuilder 9000 System Management Overview	30
Management and Data Channels	31
CoreBuilder 9000 Management Features	33
EME Overview	33
Configuration Tasks	34
Accessing the Administration Console	35
Password Access Levels	35
Accessing Your System	36
Access Examples	37
Using Menus to Perform Tasks	39
Selecting Menu Options	40
Entering Values	41
Navigating Through the Menus	42

2 COMMAND SUMMARY

PART II SYSTEM-LEVEL FUNCTIONS

3 SYSTEM ENVIRONMENT

Menu Structure	68
system display	69
system fileTransfer	70
system console webHelpConfig	71
system console webAccess	72
system console consoleAccess	73
system console ctlKeys	74
system console password	75
system console screenHeight	76
system console security display	77
system console security define	78
system console security remove	80
system console security access	81
system console security message	82
system console timeout timeOut	83
system console timeout interval	84
system snapshot summary	85
system snapshot detail	86
system snapshot save	87
system softwareUpdate	89
system baseline display	90
system baseline set	91
system baseline requestedState	92
system serialPort terminalSpeed	93
system serialPort modemSpeed	95
system serialPort baudRate	96
system serialPort serialPortMode	98
system serialPort configModem	99
system serialPort enableModem	100
system name	101
system time	102
system time datetime	103
system time timezone	104
system time dst	106
system nvData save	107
system nvData restore	110

- system nvData examine 112
- system nvData reset 113
- system clearDiagBlock 114
- system diagErrLog 115
- system sntp display 116
- system sntp define 117
- system sntp modify 118
- system sntp remove 119
- system sntp state 120
- system sntp pollInterval 121
- system sntp tolerance 122
- system reboot 123
- script 124
- logout 126

4 MODULE ENVIRONMENT

- Menu Structure 128
 - module display 129
 - module snapshot summary 130
 - module snapshot detail 131
 - module baseline display 132
 - module baseline set 133
 - module baseline requestedState 134
 - module redundancy 135
 - module name 136
 - module time 137
 - module screenHeight 138
 - module nvData reset 139
 - module nvData emergencyDownload 140
 - module nvData displayDownload 141
 - module nvData staging 142
 - module clearDiagBlock 143
 - module diagErrLog 144
 - module reboot 145
 - disconnect 146

PART III ESTABLISHING MANAGEMENT ACCESS

5 OUT-OF-BAND MANAGEMENT

Menu Structure	150
management summary	151
management detail	153
management ip interface summary	156
management ip interface define	157
management ip interface modify	158
management ip interface remove	159
management ip route display	160
management ip route static	162
management ip route remove	163
management ip route flush	164
management ip route default	165
management ip route noDefault	166
management ip route findRoute	167
management ip arp display	168
management ip arp static	169
management ip arp remove	170
management ip arp flushAll	171
management ip arp flushDynamic	172
management ip rip display	173
management ip rip mode	174
management ip rip statistics	176
management ip ping	177
management ip advancedPing	179
management ip traceRoute	182
management ip advancedTraceRoute	184
management ip statistics	186

6 SIMPLE NETWORK MANAGEMENT PROTOCOL (SNMP)

Menu Structure	190
snmp display	191
snmp community	192
snmp trap display	193
snmp trap addModify	194
snmp trap remove	196
snmp trap flush	197
snmp trap smtProxyTraps	198
snmp rmonConfiguration	199
snmp writeDisable	200

PART IV PHYSICAL PORT PARAMETERS

7 ETHERNET PORTS

Menu Structure	203
ethernet summary	204
ethernet detail	207
ethernet autoNegotiation	212
ethernet portMode	213
ethernet flowControl	215
ethernet paceAccess	217
ethernet pacerInteractiveAccess	218
ethernet label	219
ethernet portState	220
ethernet monitoring summary	221
ethernet monitoring mode	222

8 FIBER DISTRIBUTED DATA INTERFACE (FDDI)

Menu Structure	223
fddi station display	224
fddi station connectPolicy	225
fddi station tNotify	227
fddi station statusReporting	228
fddi path display	229
fddi path tvxLowerBound	230
fddi path tmaxLowerBound	231
fddi path maxTreq	232
fddi mac summary	233
fddi mac detail	234
fddi mac frameErrorThreshold	237
fddi mac notCopiedThreshold	238
fddi mac llcService	239
fddi mac path	240
fddi port display	241
fddi port lerAlarm	242
fddi port lerCutoff	243
fddi port label	244
fddi port path	245
fddi stationMode display	246
fddi stationMode modify	247

PART V BRIDGING PARAMETERS

9 BRIDGE-WIDE PARAMETERS

Menu Structure	251
bridge display	252
bridge ipFragmentation	255
bridge ipxSnapTranslation	256
bridge addressThreshold	257
bridge agingTime	258
bridge spanningTree stpState	259
bridge spanningTree stpPriority	261
bridge spanningTree stpMaxAge	262
bridge spanningTree stpHelloTime	263
bridge spanningTree stpForwardDelay	264
bridge spanningTree stpGroupAddress	265
bridge gvrpState	266
bridge cos enable	267
bridge cos summary	268
bridge cos modify	269
bridge multicast igmp summary	270
bridge multicast igmp snoopMode	271
bridge multicast igmp queryMode	272
bridge multicast igmp queryIpAddress	273
bridge multicast igmp vlans	274
bridge multicast igmp groups	275
bridge multicast igmp desQuerier	276
bridge multicast igmp rPorts	277
bridge multicast igmp qPort	278

10 BRIDGE PORT PARAMETERS

Menu Structure	279
bridge port summary	280
bridge port detail	283
bridge port multicastLimit	288
bridge port stpState	289
bridge port stpCost	290
bridge port stpPriority	291
bridge port gvrpState	292
bridge port address list	293
bridge port address add	294
bridge port address remove	295

bridge port address find 296
bridge port address flushAll 297
bridge port address flushDynamic 298

11 TRUNKS

Menu Structure 300
bridge trunk autoMap summary 301
bridge trunk autoMap enable/disable 302
bridge trunk autoMap test 303
bridge trunk summary 304
bridge trunk detail 305
bridge trunk define 307
bridge trunk modify 312
bridge trunk remove 318

12 MULTIPoint LINK Aggregation (MPLA)

Menu Structure 321
bridge mpla summary 322
bridge mpla detail 323
bridge mpla mode 324
bridge mpla peerMacAddress 326

13 RESILIENT LINKS

Menu Structure 327
bridge link summary 328
bridge link detail 329
bridge link define 330
bridge link linkState 332
bridge link activePort 333
bridge link modify 334
bridge link remove 336

14 VIRTUAL LANs (VLANs)

Menu Structure 337
bridge vlan summary 338
bridge vlan detail 341
bridge vlan define (3500/9000 Layer 3) 345
bridge vlan define (3900/9300/9400/ 9000 Layer 2) 352

- bridge vlan modify (3500/9000 Layer 3) 355
- bridge vlan modify (3900/9300/9400/ 9000 Layer 2) 360
- bridge vlan remove 363
- bridge vlan mode 364
- bridge vlan stpMode 365
- bridge vlan vlanAwareMode 366

15 PACKET FILTERS

- Menu Structure 370
 - bridge packetFilter list 371
 - bridge packetFilter display 372
 - bridge packetFilter create portGroup 373
 - bridge packetFilter create custom 374
 - bridge packetFilter delete 376
 - bridge packetFilter edit 377
 - bridge packetFilter load 379
 - bridge packetFilter assign 382
 - bridge packetFilter unassign 384
 - bridge packetFilter portGroup list 386
 - bridge packetFilter portGroup display 387
 - bridge packetFilter portGroup create 388
 - bridge packetFilter portGroup delete 390
 - bridge packetFilter portGroup addPort 391
 - bridge packetFilter portGroup removePort 392

PART VI ROUTING PROTOCOLS

16 INTERNET PROTOCOL (IP)

- Menu Structure 396
 - ip interface summary 398
 - ip interface detail 400
 - ip interface define (3500/9000 Layer 3) 403
 - ip interface define (3900/9300/9400/ 9000 Layer 2) 406
 - ip interface modify 407
 - ip interface remove 408
 - ip interface arpProxy 409
 - ip interface broadcastAddress 411
 - ip interface directedBroadcast 412
 - ip interface icmpRedirect 413
 - ip interface icmpRouterDiscovery 415

- ip interface statistics 418
- ip route display 420
- ip route static 422
- ip route remove 423
- ip route flush 424
- ip route default 425
- ip route noDefault 426
- ip route findRoute 427
- ip arp display 428
- ip arp static 429
- ip arp remove 430
- ip arp flushAll 431
- ip arp flushDynamic 432
- ip arp age 433
- ip arp statistics 434
- ip dns display 436
- ip dns domainName 437
- ip dns define 438
- ip dns modify 439
- ip dns remove 440
- ip dns nslookup 441
- ip udpHelper display 442
- ip udpHelper define 443
- ip udpHelper remove 444
- ip udpHelper hopCountLimit 445
- ip udpHelper threshold 446
- ip udpHelper interface first 447
- ip udpHelper interface even 448
- ip udpHelper interface sequential 449
- ip routing 450
- ip rip display 451
- ip rip mode 453
- ip rip compatibilityMode 455
- ip rip cost 456
- ip rip poisonReverse 457
- ip rip routeAggregation Mode 458
- ip rip password 459
- ip rip addAdvertisement 460
- ip rip remove Advertisement 462
- ip rip policy summary 463
- ip rip policy detail 464
- ip rip policy define 465
- ip rip policy modify 469
- ip rip policy remove 471

- ip rip statistics 472
- ip ping 473
- ip advancedPing 475
- ip traceRoute 478
- ip advancedTraceRoute 480
- ip statistics 482

17 VIRTUAL ROUTER REDUNDANCY (VRRP)

- Menu Structure 485
 - ip vrrp summary 486
 - ip vrrp detail 488
 - ip vrrp define 492
 - ip vrrp modify 495
 - ip vrrp remove 498
 - ip vrrp mode 499
 - ip vrrp neighbor 500
 - ip vrrp statistics 501

18 IP MULTICAST

- Menu Structure 504
 - ip multicast dvmrp interface summary 505
 - ip multicast dvmrp interface detail 506
 - ip multicast dvmrp interface mode 507
 - ip multicast dvmrp interface metric 508
 - ip multicast dvmrp tunnels summary 509
 - ip multicast dvmrp tunnels define 511
 - ip multicast dvmrp tunnels remove 513
 - ip multicast dvmrp tunnels address 514
 - ip multicast dvmrp tunnels threshold 515
 - ip multicast dvmrp tunnels metric 516
 - ip multicast dvmrp routeDisplay 517
 - ip multicast dvmrp cacheDisplay 518
 - ip multicast dvmrp default 520
 - ip multicast igmp interface summary 521
 - ip multicast igmp interface detail 522
 - ip multicast igmp interface TTL 523
 - ip multicast igmp snooping 524
 - ip multicast igmp querying 525
 - ip multicast cache 526
 - ip multicast traceRoute 528

19 OPEN SHORTEST PATH FIRST (OSPF)

Menu Structure	530
ip ospf areas display	531
ip ospf areas defineArea	532
ip ospf areas modifyArea	533
ip ospf areas removeArea	534
ip ospf areas addRange	535
ip ospf areas modifyRange	536
ip ospf areas removeRange	537
ip ospf defaultRouteMetric display	538
ip ospf defaultRouteMetric define	539
ip ospf defaultRouteMetric remove	540
ip ospf interface summary	541
ip ospf interface detail	542
ip ospf interface statistics	544
ip ospf interface mode	548
ip ospf interface priority	549
ip ospf interface areaID	550
ip ospf interface cost	551
ip ospf interface delay	552
ip ospf interface hello	553
ip ospf interface retransmit	554
ip ospf interface dead	555
ip ospf interface password	556
ip ospf linkStateData databaseSummary	557
ip ospf linkStateData router	558
ip ospf linkStateData network	560
ip ospf linkStateData summary	561
ip ospf linkStateData external	563
ip ospf neighbors display	564
ip ospf neighbors add	565
ip ospf neighbors remove	566
ip ospf routerID	567
ip ospf partition display	569
ip ospf partition modify	570
ip ospf stubDefaultMetric display	571
ip ospf stubDefaultMetric define	572
ip ospf stubDefaultMetric remove	573
ip ospf virtualLinks summary	574
ip ospf virtualLinks detail	575
ip ospf virtualLinks statistics	577
ip ospf virtualLinks define	581
ip ospf virtualLinks remove	582

- ip ospf virtualLinks areaID 583
- ip ospf virtualLinks router 584
- ip ospf virtualLinks delay 585
- ip ospf virtualLinks hello 586
- ip ospf virtualLinks retransmit 587
- ip ospf virtualLinks dead 588
- ip ospf virtualLinks password 589
- ip ospf policy summary 590
- ip ospf policy detail 591
- ip ospf policy define 593
- ip ospf policy modify 598
- ip ospf policy remove 602
- ip ospf statistics 603

20 IPX

- Menu Structure 606
 - ipx interface display 607
 - ipx interface define 608
 - ipx interface modify 610
 - ipx interface remove 612
 - ipx interface SAPadvertising 613
 - ipx interface RIPadvertising 614
 - ipx route display 615
 - ipx route secondary 617
 - ipx route static 618
 - ipx route remove 620
 - ipx route flush 621
 - ipx server display 622
 - ipx server static 624
 - ipx server remove 626
 - ipx server flush 627
 - ipx server secondary 628
 - ipx forwarding 629
 - ipx rip mode 630
 - ipx rip triggered 631
 - ipx rip policy summary 632
 - ipx rip policy define 633
 - ipx rip policy modify 635
 - ipx rip policy remove 637
 - ipx sap mode 638
 - ipx sap triggered 639
 - ipx sap policy summary 640

- ipx sap policy detail 641
- ipx sap policy define 642
- ipx sap policy modify 645
- ipx sap policy remove 648
- ipx output-delay 649
- ipx statistics summary 650
- ipx statistics rip 651
- ipx statistics sap 652
- ipx statistics forwarding 653
- ipx statistics interface 655
- ipx oddLengthPadding 657
- ipx NetBIOS 658
- ipx secondary 659

21 APPLETALK

- Menu Structure 662
 - appletalk interface summary 663
 - appletalk interface detail 664
 - appletalk interface define 665
 - appletalk interface modify 667
 - appletalk interface remove 669
 - appletalk interface statistics 670
 - appletalk route display 672
 - appletalk route flush 673
 - appletalk aarp display 674
 - appletalk aarp remove 675
 - appletalk aarp flush 676
 - appletalk zone display network 677
 - appletalk zone display zone 678
 - appletalk forwarding 679
 - appletalk checksum 680
 - appletalk sourceSocket 681
 - appletalk ping 682
 - appletalk statistics ddp 683
 - appletalk statistics rtmp 684
 - appletalk statistics zip 685
 - appletalk statistics nbp 686

PART VII TRAFFIC POLICY

22 QUALITY OF SERVICE (QoS) AND RSVP

Menu Structure	690
qos classifier summary	691
qos classifier detail	692
qos classifier define	694
qos classifier modify	701
qos classifier remove	706
qos control summary	707
qos control detail	708
qos control define	710
qos control modify	718
qos control remove	724
qos ldap display	725
qos ldap enable	726
qos ldap disable	727
qos rsvp summary	728
qos rsvp detail	729
qos rsvp enable	730
qos rsvp disable	732
qos bandwidth display	733
qos bandwidth modify	734
qos excessTagging display	735
qos excessTagging enable	736
qos excessTagging disable	737
qos statistics interval	738
qos statistics receive	739
qos statistics transmit	741

PART VIII MONITORING

23 EVENT LOG

Menu Structure	748
log display	749
log devices	750
log services	752

24 ROVING ANALYSIS

Menu Structure	756
analyzer display	757
analyzer add	758
analyzer remove	760
analyzer start	761
analyzer stop	763

PART IX REFERENCE

A TECHNICAL SUPPORT

Online Technical Services	767
World Wide Web Site	767
3Com Knowledgebase Web Services	767
3Com FTP Site	768
3Com Bulletin Board Service	768
3Com Facts Automated Fax Service	769
Support from Your Network Supplier	769
Support from 3Com	769
Returning Products for Repair	771

INDEX

ABOUT THIS GUIDE

This *Command Reference Guide* provides information about the commands that you use to configure and manage your system or module after you install it. Before you use this guide, you should have already consulted documents such as your system *Getting Started Guide* or module *Quick Start Guide* and physically installed your system or module.

Several CoreBuilder® and SuperStack® II platforms are documented in this book. Table 1 lists the specific platforms and the current software release level of that platform as it relates to the information contained in this book:

Table 1 Platforms Covered in This Document

Platform	Release
CoreBuilder® 3500	3.0
SuperStack® II Switch 3900	3.0
CoreBuilder 9000	3.0
SuperStack II Switch 9300	3.0
CoreBuilder 9400	3.0

This guide is intended for the system or network administrator who is responsible for configuring, using, and managing the system. It assumes a working knowledge of local area network (LAN) operations and familiarity with communications protocols that are used on interconnected LANs.



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

Using This Book

This guide contains information for every command for the platforms listed at the beginning of this chapter. It includes specific information about command syntax, field descriptions, default values, and the possible range of values. Some command descriptions include a section called “Important Considerations” that contains additional information to be aware of when using the command. Where appropriate, examples help you to understand the commands.

This guide does not contain troubleshooting information or instructional material about why or when to use a particular command. For information about troubleshooting and networking tasks, see the *Implementation Guide* that is shipped with your system on a CD-ROM.

Finding Specific Information in This Guide

Use this chart to help you find information about specific tasks:

If you are looking for information about	Turn to
System administration and configuration tasks	Part I: Getting Started
Using command abbreviations	
Summary of commands for all platforms	
Displaying the system or module configuration	Part II: System-Level Functions
Using the snapshot feature	
Baselining statistics	
Configuring system parameters, such as name, date/time, and passwords	
Configuring system security	
Establishing system access through a Web browser	
Saving, restoring, and resetting nonvolatile data	
Running scripts of Console tasks	
Setting up the system for out-of-band management access through serial ports or using IP and setting up SNMP	Part III: Establishing Management Access
Administering the IP management interface	
Configuring SNMP community strings	
Administering trap reporting	
Administering Ethernet ports	Part IV: Physical Port Parameters
Displaying statistics for and labelling Ethernet ports	
Administering Fiber Distributed Data Interface (FDDI) ports	

If you are looking for information about	Turn to
Configuring bridge parameters such as bridge display, agingTime, stpState, and Class of Service Managing trunks Configuring bridge port parameters such as listing addresses, setting the port priority, and controlling the Spanning Tree Protocol (STP) on a bridge Displaying MultiPoint Link Aggregation (MPLA) parameters Configuring resilient links Configuring virtual LANs (VLANs) Configuring packet filters	Part V: Bridging Parameters
Configuring IP interfaces and IP protocol parameters Configuring Virtual Router Redundancy Protocol (VRRP) parameters Configuring IP multicast routing and filtering Configuring Open Shortest Path First (OSPF) routing Configuring IPX routing Configuring AppleTalk routing	Part VI: Routing Protocols
Configuring Quality of Service (QoS) classifiers, controls, Resource Reservation Protocol (RSVP), bandwidth, and excess tagging Viewing statistics	Part VII: Traffic Policy
Administering the event log Administering roving analysis	Part VIII: Monitoring
Technical support	Part IX: Reference
Quickly locating information on tasks and topics	Index

Command Information

Each software command has its own description in this guide. Each command description begins at the top of a page. A command description begins with these items:

- The full command name
 - Platforms on which this command is valid
- Under the command name is a list of 3Com switch platforms. The command is valid on every platform that has a check mark (✓) next to it.

Sample platform list

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

- A short description of the purpose of the command



*Some command descriptions begin with a sentence similar to this one: “**For CoreBuilder 9000: Applies to Layer n switching modules only.**” where n is either 2 or 3. Because the CoreBuilder 9000 system can house both Layer 2 modules and Layer 3 modules, this sentence alerts you to the fact that this particular command is valid only on Layer 2 modules or Layer 3 modules.*

The command description continues with one or more of the following sections:

- **Valid Minimum Abbreviation** — This section lists the shortest number of characters that you can type to issue the command.
- **Important Considerations** — These usage notes identify potential problems before you use the command.
- **Options** — If the command begins a configuration process or other procedure, this section presents each prompt that you see, its description, the possible values that you can enter, and the default value.
- **Fields** — If the command prompts the system to display information, this section lists the display parameters and their definitions.

- **Procedure** — Numbered steps walk you through complex commands.
- **Example** — Examples show the interactive display when the system provides additional useful information.

Recommendations for Entering Commands

Before you enter any command, 3Com recommends that you:

- Examine the system menu carefully for the full command string:
- Consult the documentation for the valid minimum abbreviation for the command string.



If you are unfamiliar with a particular system, always enter the entire command, even though the system accepts abbreviated commands. If you abbreviate commands, you may make errors or omissions that have undesirable consequences.

For example, on the CoreBuilder 9000, to list all addresses for a port, you use the `bridge port address list all` command. If you mistakenly enter `bridge port address all`, the system interprets it as an abbreviated version of the `bridge port address flushAll` command, which flushes the entire address table for the port and does not request that you confirm the command.

Conventions

Table 2 and Table 3 list icon and text conventions that are used throughout this guide.

Table 2 Notice Icons

Icon	Notice Type	Description
	Information note	Information that describes important features or instructions
	Caution	Information that alerts you to potential loss of data or potential damage to an application, system, or device
	Warning	Information that alerts you to potential personal injury

Table 3 Text Conventions

Convention	Description
Screen displays	This typeface represents information as it appears on the screen.
Command	<p>The word “command” means that you enter the command exactly as shown in the text and then press Return or Enter. Commands appear in bold. Example:</p> <p>To set flow control, enter the following command:</p> <p>ethernet flowControl</p> <div><div><i>i</i></div><div><p><i>This guide always gives the full form of a command. However, you can abbreviate commands by entering just enough characters to distinguish one command from another similar command, as shown in “Valid Minimum Abbreviations” under each command description. Commands are not case sensitive.</i></p></div></div>
The words “enter” and “type”	When you see the word “enter” in this guide, you must type something, and then press the Return or Enter key. Do not press Return or Enter when an instruction simply says “type.”
Keyboard key names	<p>If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example:</p> <p>Press Ctrl+C.</p>
Words in <i>italics</i>	<p>Italics are used to:</p> <ul style="list-style-type: none">■ Emphasize a point■ Denote a new term when it is defined in text

Documentation Comments

Your suggestions are very important to us. They help us to make our documentation more useful to you.

Please send e-mail comments about this guide to:

`sdtechpubs_comments@ne.3com.com`

Include the following information when commenting:

- Document title
- Document part number (found on the front or back page of each document)
- Page number (if appropriate)

Example:

Command Reference Guide


Part Number 10013505

Page 347

Year 2000 Compliance

For information on Year 2000 compliance and 3Com products, visit the 3Com Year 2000 Web page:


`http://www.3com.com/products/yr2000.html`




GETTING STARTED

Chapter 1 Administration Overview

Chapter 2 Command Summary





ADMINISTRATION OVERVIEW

This chapter introduces the Administration Console software that is supplied with your system, the types of commands that you use to perform network tasks, the valid syntax for command abbreviations, and some shortcuts to help you navigate through the menus. It also provides an overview of the management software that is specific to the CoreBuilder® 9000 Enterprise Switch. It introduces the EME (Enterprise Management Engine) Management Console for the CoreBuilder 9000 and describes its relationship to the Administration Console software.

The following topics are covered in this chapter:

- Administration Console Overview
- CoreBuilder 9000 System Management Overview
- CoreBuilder 9000 Management Features
- Configuration Tasks
- Accessing the Administration Console
- Using Menus to Perform Tasks

Administration Console Overview

The Administration Console software is installed at the factory in flash memory on the system processor. Because this software boots automatically from flash memory when you power on your system, the system is immediately ready for use in your network. However, you may need to:

- Configure certain parameters before the system can operate effectively in your networking environment.
- Connect to the Administration Console, if you have a CoreBuilder 9000.
- View important MAC, port, bridge, virtual LAN (VLAN), and IP statistics while you manage your system.

You use the Administration Console software to configure your system parameters (or, on the CoreBuilder 9000, to configure your module parameters) and display statistics and counters.



For more complete network management, you can use an external application, such as 3Com's Transcend® Network Control Services tool suite.



On the CoreBuilder 3500, CoreBuilder 9000, and CoreBuilder 9400, and on the SuperStack® II Switch 3900 and Switch 9300, you can also configure parameters and view statistics using the Web Management suite of HTML-based applications. See the Web Management User Guide for your system for additional information.

CoreBuilder 9000 System Management Overview

The CoreBuilder 9000 comes in a 7-slot, 8-slot, or 16-slot chassis in which you install switch fabric modules and interface modules. Before you begin to manage your CoreBuilder 9000 system, review the management-related information in this section.

The CoreBuilder 9000 system supports the following management interfaces:

- **EME Management Console**

Use the EME Management Console to manage EME and Enterprise Management Controller (EMC) functions, such as login table management, IP connectivity, event and trap logs, and software downloads to all modules. The EME Management Console also manages chassis functions, such as system inventory and power management features.

- **Administration Console**

Use the Administration Console to manage the CoreBuilder 9000 switch fabric modules and intelligent interface modules. These modules contain an on-board network management agent that allows this direct management.

- **ATM Local Management Application (LMA)**

Use the ATM LMA to manage the ATM Enterprise Switch, ATM Switch Fabric Module, and ATM interface modules. These modules contain an on-board network management agent to allow this direct management.

ATM LMA management of ATM switch fabric modules and ATM interface modules is outside of the scope of this guide. To learn about managing the ATM Enterprise Switch and ATM modules using the ATM LMA, see the *CoreBuilder 9000 ATM Enterprise Switch Management Guide*.

You cannot manage the EME using the ATM LMA, and you cannot manage ATM Switch Fabric Modules or ATM interface modules using the EME Management Console.

- Web Management

The Web Management suite of applications are an embedded part of the CoreBuilder 9000 system software image. They include the WebConsole and DeviceView applications. Additional installable applications include online Help. After you have set up your IP address for your system, you can access the Web Management applications directly from your Web browser by entering its IP address.

See the *Web Management User Guide for the CoreBuilder 9000 Enterprise Switch* for additional information about Web Management.


You manage the EME from a command line interface using EME management commands. You manage the switch fabric modules or interface modules through the Administration Console using module management commands or through the Web Management interface.

Management and Data Channels

The CoreBuilder 9000 system uses separate channels for network traffic and management traffic:

- The private management LAN (MLAN) handles management traffic. Management traffic travels to and from the EME, which acts as the single point of contact for all management traffic in the chassis.
- Switch fabric module backplane channels handle network traffic. Each interface module has one or two backplane ports that connect to the switch fabric module backplane, which allows network traffic to pass through the CoreBuilder 9000 system.

Figure 1 illustrates the MLAN channel and the switch fabric module backplane channels in the CoreBuilder 9000 16-slot chassis.

Figure 1 System Data Channels in the 16-slot Chassis

CoreBuilder 9000 Management Features

You can manage the CoreBuilder 9000 system through a terminal interface, through the Simple Network Management Protocol (SNMP), and through the 3Com Transcend[®] Network Control Services. The EME is the primary communication mechanism into the chassis and modules. You manage other intelligent modules within the chassis through the EME.

EME Overview

The EME is an SNMP-based network management module that manages and controls the 3Com CoreBuilder 9000 chassis and its modules. The EME has the following features:

- **Chassis Management Architecture** — Provides a cost-efficient management architecture that:
 - Provides a central point of contact for chassis management
 - Provides all controller functions, as well as EME functions
- **Intelligent Power Management** — Manages power use in the chassis by:
 - Preventing newly installed modules from receiving power when there is not enough power available
 - Allowing you to prioritize the order in which modules power off (if there is insufficient power available)
 - Allowing you to implement fault-tolerant power, which allows the chassis to reserve some of its power capacity to protect against a power supply failure

In the chassis:

- The EME exchanges information with all modules through the MLAN.
- Interface modules pass data through the switch fabric module.
- On modules that include their own agent, the EME “connects” to that module and then you can use the Administration Console management interface to manage that module.

Configuration Tasks

To help you configure your system, the top-level menu of the Administration Console groups the commands into types for certain tasks, as listed in Table 4.



Not all menus and tasks are available on all systems.

Table 4 Types of Commands Associated with Configuration Tasks

Type of Command	Top-Level Menus	Tasks
General	system	Set system or module parameters, handle nonvolatile (NV) data, set security, reboot
	module	
	script	Run scripts
Management setup	logout, disconnect	Leave the Administration Console
	management	Set up the out-of-band management interface
	snmp	Set up the system for SNMP and trap reporting
Port-based management	ethernet	Manage Ethernet ports
	fddi	Manage Fiber Distributed Data Interface (FDDI) ports
Bridging	bridge	Set bridge parameters for the entire system, including for Spanning Tree Protocol (STP) and Class of Service (CoS)
		Manage trunking of bridge ports
		Set and display MultiPoint Link Aggregation (MPLA) parameters
		Manage resilient links
		Set bridge parameters for specific bridge ports
		Manage virtual LANs (VLANs)
		Manage packet filtering for port groups
Routing	ip	Set up IP, IP multicast, and IP Open Shortest Path First (OSPF) routing
	ipx	Set up IPX routing
	appletalk	Set up AppleTalk routing

Table 4 Types of Commands Associated with Configuration Tasks (continued)

Type of Command	Top-Level Menus	Tasks
Quality of Service management	qos	Set up classifiers and controls for traffic-policy-based services
Monitoring	log	Set severity levels and services for event logging
	analyzer	Monitor the network using a network analyzer

Accessing the Administration Console

Depending on which system you are managing, you access the Administration Console in either two steps (for the CoreBuilder 9000) or one step (for all other systems). See “Accessing Your System” later in this section for details.

For all systems, the Administration Console supports three password levels, allowing you to provide different levels of access for a range of users.

Password Access Levels

Your access level determines which types of menu commands you can use, as described in Table 5.

Table 5 Password Access Levels

Access Level	For users who need to	Allows users to
Administer	Perform system or module setup and management tasks (usually a single network administrator)	Perform system-level or module-level administration (such as resetting the module or changing passwords)
Write	Perform active network management	Configure network parameters (such as setting the aging time for a bridge)
Read	Only view system or module parameters	Access only “display” menu items (like display, summary, and detail)

Accessing Your System

You access the Administration Console for your system in one of two ways:

- **For all systems except the CoreBuilder 9000** — Access the Administration Console for the first time at the *Administer* level and press Return at the password prompt (the initial password is null).

Subsequently, every time that you access the Administration Console, it prompts you for an access level and password, as shown here:

```
Select access level (read, write, administer):
```

```
Password:
```

The passwords are stored in nonvolatile (NV) memory. You must enter the password correctly before you can continue.

- **For the CoreBuilder 9000** — On this system, the Enterprise Management Engine (EME) controls passwords and access levels to manage the chassis and its installed modules.

To access a module in a CoreBuilder 9000 system, follow these steps:

- 1 Log in to the EME.
- 2 Access the module that you want to manage using the EME `connect` command.

Example: To connect to a module in slot 10, subslot 1, enter:

```
connect 10.1
```



All modules use subslot 1.

The system displays a prompt similar to the following:

```
CB9000@slot10.1 [20-E/FEN-TX-L2]
```

When you have connected to a module, you manage the module from the Administration Console with the same level of access that you have on the EME. For example, if you have logged in to the EME with *administer* privileges, you also have *administer* privileges for the module to which you are connected.



For additional information about the EME, see the CoreBuilder 9000 Enterprise Switch Getting Started Guide and the CoreBuilder 9000 Enterprise Management Engine User Guide.

Access Examples

The examples in this section show how the top-level menu structure of the Administration Console changes. The menus that you see in the Administration Console vary depending on:

- Which 3Com system you are viewing (as described in “Accessing Your System” earlier in this chapter).
- Your level of access.
- The optional interface modules, switch fabric modules, and other hardware options that you configure into your system. For example, you see the `fddi` menu only if you have installed the FDDI module on your CoreBuilder 3500 system.



These examples show the CoreBuilder 3500 menus. Menus for other platforms may differ. See the Command Quick Reference for your system for the list of commands on your system.

Administer Access Example

When you enter the Administration Console with *Administer* access, each menu contains all of the options for the system. Here is an example of a system menu for users with *Administer* access on the CoreBuilder 3500:

Select menu option: **system**

```
Menu options (CoreBuilder-2B4200): -----
display          - Display the system configuration
console          - Administer console-level functions
fileTransfer      - Set the file transfer protocol
snapshot         - Display all configuration and status information
softwareUpdate   - Load a new revision of system software
baseline         - Administer a statistics baseline
serialPort       - Administer the terminal and modem serial ports
name             - Set the system name
time             - Set the date and time
nvData           - Save, restore, or reset nonvolatile data
clearDiagBlock   - Clear the diagnostic block
diagErrLog       - Display Diagnose Error Log
snmp             - Administer the Simple Network Time Protocol
reboot          - Reboot the system
```

Type "q" to return to the previous menu or ? for help.

```
-----
Select menu option (system):
```

Write Access Example

When you enter the Administration Console with *write* access, the system menu contains a subset of the complete menu, focusing on the network, as shown in this example on the CoreBuilder 3500:

Select menu option: **system**

```
Menu options (CoreBuilder-2B4200): -----
display          - Display the system configuration
console          - Administer console-level functions
fileTransfer      - Set the file transfer protocol
snapshot         - Display all configuration and status information
baseline         - Administer a statistics baseline
serialPort       - Administer the terminal and modem serial ports
name             - Set the system name
diagErrLog       - Display Diagnose Error Log
sntp             - Administer the Simple Network Time Protocol
```

Type "q" to return to the previous menu or ? for help.

Read Access Example

When you enter the Administration Console with *read* access, the system menu contains the fewest options, as shown in this example on the CoreBuilder 3500:

Select menu option: **system**

```
Menu options (CoreBuilder-293300): -----
display          - Display the system configuration
snapshot         - Display all configuration and status information
baseline         - Administer a statistics baseline
diagErrLog       - Display Diagnose Error Log
sntp             - Administer the Simple Network Time Protocol
```

Type "q" to return to the previous menu or ? for help.

Using Menus to Perform Tasks

When you access the Administration Console, the top-level menu appears. You perform administrative tasks by selecting options from this menu and its submenus. A brief description accompanies each option in the display. The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The “Menu Structure” diagrams at the beginning of most chapters show the complete list of commands for all systems. See the checklist at the beginning of each command description in each chapter for whether your system supports the command.

The following example shows the CoreBuilder 9000 top-level menu when a Layer 2 switching module is installed:

Menu options:

```
-----  
module          - Administer module-level functions  
ethernet        - Administer Ethernet ports  
bridge          - Administer bridging/VLANs  
snmp            - Administer SNMP  
disconnect      - Disconnect and return to the Management Console
```

Type ? for help.

Select a menu option:



These examples show the CoreBuilder 9000 menu options for a Layer 2 switching interface module. Menus on other platforms may differ. See the Command Quick Reference for the list of commands on your system.

Selecting Menu Options

To select a menu option, at the prompt enter the menu option or enough of the name to uniquely identify it within the particular menu. Example: to access the `module` menu from the top level of the Administration Console on a module in the CoreBuilder 9000, enter:

```
Select a menu option: module
```



Menu options are not case sensitive.

When you enter a menu option or command correctly, either you move to the next menu in the hierarchy, or the Administration Console displays information (a prompt or a screen display) for the option that you entered.

If you enter the menu option incorrectly, a message indicates that your entry is not valid or is ambiguous. Reenter the option from the point at which it became incorrect or expand a truncated command until it becomes unambiguous.

When a new menu appears, the selection prompt (with its choices in parentheses) changes to reflect your progression through the menus.

Example: If you enter **bridge** at the top-level menu and then **agingTime** at the `module` prompt, the prompt changes to reflect the current level:

```
Select a menu option (bridge/agingTime):
```

Entering a Command String

After you become familiar with the menu structure, you can enter a string of menu options or commands to move immediately to a task. Example: The command string for setting the path cost for a port on a module looks like this:

```
Select a menu option: bridge port stpCost
```


Entering Abbreviated Commands

You can abbreviate command strings by typing only as much of the command as is necessary to make it unique:

Select a menu option: **b po stpc**

When you correctly enter either a full or an abbreviated command string, you move to the last menu level or option that is specified in the string. Information that is relevant to that option appears as a menu, a prompt, or a display.

If you enter a command string incorrectly, the Administration Console displays a message indicating that your entry was not valid or was ambiguous. Reenter the command from the point at which it became incorrect, or expand a truncated command until it becomes unambiguous.

Entering Values

When you reach the level at which you perform a task, the Administration Console prompts you for a value. The prompt usually shows all valid values (if applicable) and typically suggests a default value. The default may be the factory default value or the current value that you have defined for that parameter.

The Administration Console displays the valid values in parentheses and the default or current value in brackets. For example:

Enter a new value (disabled,enabled) [enabled]:

To accept the default or current value, press Enter.

Entering Values in Command Strings

A command string can also contain the value of a command parameter. If you enter a value at the end of a command string, the Administration Console executes the task and the previous menu appears on the screen.

For example, to set the path cost to the root through a port, from the top level of the Administration Console, enter:

bridge port stpcost 20


or

b po stpc 20

Navigating Through the Menus

The Administration Console provides several shortcuts:

- **Press Esc (the Escape key)** — To move quickly to the top-level menu without backtracking through each intermediate menu. The top-level menu immediately appears.
- **Enter ↵** —
 - To move up through the hierarchy, that is, to move to the menu that is one level higher in the hierarchy
 - To cancel an operation that is currently in progress. The previous menu appears.



COMMAND SUMMARY

Table 6 gives an overview of all the commands in this book.

Table 6 Command Summary

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
PART II: System-Level Functions						
Chapter 3 System Environment						
system display	✓			✓	✓	✓
system fileTransfer	✓					
system console webHelpConfig	✓			✓	✓	✓
system console webAccess	✓			✓	✓	✓
system console consoleAccess	✓			✓	✓	✓
system console ctlKeys	✓			✓	✓	✓
system console password	✓			✓	✓	✓
system console screenHeight	✓			✓	✓	✓
system console security display	✓			✓	✓	✓
system console security define	✓			✓	✓	✓
system console security remove	✓			✓	✓	✓
system console security access	✓			✓	✓	✓
system console security message	✓			✓	✓	✓
system console timeout timeOut	✓			✓	✓	✓
system console timeout interval	✓			✓	✓	✓

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
system snapshot summary	✓			✓	✓	✓
system snapshot detail	✓			✓	✓	✓
system snapshot save	✓			✓	✓	✓
system softwareUpdate	✓			✓	✓	✓
system baseline display	✓			✓	✓	✓
system baseline set	✓			✓	✓	✓
system baseline requestedState	✓			✓	✓	✓
system serialPort terminalSpeed	✓					
system serialPort modemSpeed	✓					
system serialPort baudRate				✓	✓	✓
system serialPort serialPortMode				✓	✓	✓
system serialPort configModem	✓			✓	✓	✓
system serialPort enableModem	✓			✓	✓	✓
system name	✓			✓	✓	✓
system time dateTime	✓			✓	✓	✓
system time timezone	✓			✓	✓	✓
system time dst	✓			✓	✓	✓
system nvData save	✓			✓	✓	✓
system nvData restore	✓			✓	✓	✓
system nvData examine	✓			✓	✓	✓
system clearDiagBlock	✓			✓	✓	✓
system diagErrLog	✓					
system sntp display	✓			✓	✓	✓
system sntp define	✓			✓	✓	✓

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
system sntp modify	✓			✓	✓	✓
system sntp remove	✓			✓	✓	✓
system sntp state	✓			✓	✓	✓
system sntp pollInterval	✓			✓	✓	✓
system sntp tolerance	✓			✓	✓	✓
system reboot	✓			✓	✓	✓
script	✓			✓	✓	✓
logout	✓			✓	✓	✓
Ch 4 Module Environment						
module display	✓	✓				
module snapshot summary	✓	✓				
module snapshot detail	✓	✓				
module baseline display	✓	✓				
module baseline set	✓	✓				
module baseline requestedState	✓	✓				
module redundancy display	✓					
module redundancy reset NonRedundant	✓					
module name	✓	✓				
module time	✓	✓				
module screenHeight	✓	✓				
module nvData reset	✓	✓				
module nvData emergencyDownload	✓	✓				
module nvData displayDownload	✓	✓				
module nvData staging	✓	✓				

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
module clearDiagBlock		✓	✓			
module diagErrLog			✓			
module reboot		✓	✓			
disconnect		✓	✓			
PART III: Establishing Management Access						
Ch 5 Out-of-Band Management						
management summary	✓			✓		✓
management detail	✓			✓		✓
management ip	✓					
management ip interface summary	✓					
management ip interface define	✓					
management ip interface modify	✓					
management ip interface remove	✓					
management ip route display	✓					
management ip route static	✓					
management ip route remove	✓					
management ip route flush	✓					
management ip route default	✓					
management ip route noDefault	✓					
management ip route findRoute	✓					
management ip arp display	✓					
management ip arp static	✓					
management ip arp remove	✓					
management ip arp flushAll	✓					

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
management ip arp flushDynamic	✓					
management ip rip display	✓					
management ip rip mode	✓					
management ip rip statistics	✓					
management ip ping	✓					
management ip advancedPing	✓					
management ip traceRoute	✓					
management ip advancedTraceRoute	✓					
management ip statistics	✓					
Ch 6 SNMP						
snmp display	✓	✓	✓	✓	✓	✓
snmp community	✓			✓	✓	✓
snmp trap display	✓	✓	✓	✓	✓	✓
snmp trap addModify	✓	✓	✓	✓	✓	✓
snmp trap remove	✓	✓	✓	✓	✓	✓
snmp trap flush	✓	✓	✓	✓	✓	✓
snmp trap smtProxyTraps	✓		✓			
snmp rmonConfiguration	✓		✓			
snmp writeDisable	✓		✓	✓	✓	✓
Part IV Physical Port Parameters						
Ch 7 Ethernet Ports						
ethernet summary	✓	✓	✓	✓	✓	✓
ethernet detail	✓	✓	✓	✓	✓	✓
ethernet autoNegotiation	✓	✓	✓	✓	✓	✓

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
ethernet portMode	✓	✓	✓		✓	
ethernet flowControl	✓	✓	✓	✓	✓	✓
ethernet paceInteractiveAccess	✓		✓			
ethernet paceAccess		✓			✓	
ethernet label	✓	✓	✓	✓	✓	✓
ethernet portState	✓	✓	✓	✓	✓	✓
ethernet monitoring summary		✓			✓	
ethernet monitoring mode		✓			✓	
Ch 8 FDDI						
fddi station display	✓		✓			
fddi station connectPolicy	✓		✓			
fddi station tNotify	✓		✓			
fddi station statusReporting	✓		✓			
fddi path display	✓		✓			
fddi path tvxLowerBound	✓		✓			
fddi path tmaxLowerBound	✓		✓			
fddi path maxTreq	✓		✓			
fddi mac summary	✓		✓			
fddi mac detail	✓		✓			
fddi mac frameErrorThreshold	✓		✓			
fddi mac notCopiedThreshold	✓		✓			
fddi mac llcService	✓		✓			
fddi mac path	✓		✓			
fddi port display	✓		✓			

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
fddi port lerAlarm	✓		✓			
fddi port lerCutoff	✓		✓			
fddi port label	✓		✓			
fddi port path	✓		✓			
fddi stationMode display	✓		✓			
fddi stationMode modify	✓		✓			
Part V Bridging Parameters						
Ch 9 Bridge-wide Parameters						
bridge display	✓	✓	✓	✓	✓	✓
bridge ipFragmentation	✓		✓			
bridge ipxSnapTranslation	✓		✓			
bridge addressThreshold	✓		✓			
bridge agingTime	✓	✓	✓	✓	✓	✓
bridge spanningTree stpState	✓	✓	✓	✓	✓	✓
bridge spanningTree stpPriority	✓	✓	✓	✓	✓	✓
bridge spanningTree stpMaxAge	✓	✓	✓	✓	✓	✓
bridge spanningTree stpHelloTime	✓	✓	✓	✓	✓	✓
bridge spanningTree stpForwardDelay	✓	✓	✓	✓	✓	✓
bridge spanningTree stpGroupAddress	✓	✓	✓	✓	✓	✓
bridge gvrpState	✓		✓			
bridge cos enable		✓		✓	✓	✓
bridge cos summary		✓		✓	✓	✓
bridge cos modify		✓		✓	✓	✓
bridge multicast igmp summary		✓		✓	✓	✓

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
bridge multicast igmp snoopMode	✓			✓	✓	✓
bridge multicast igmp queryMode	✓			✓	✓	✓
bridge multicast igmp queryIpAddress	✓			✓	✓	✓
bridge multicast igmp vlans	✓			✓	✓	✓
bridge multicast igmp groups	✓			✓	✓	✓
bridge multicast igmp desQuerier	✓			✓	✓	✓
bridge multicast igmp rPorts	✓			✓	✓	✓
bridge multicast igmp qPort	✓			✓	✓	✓
Ch 10 Bridge Ports						
bridge port summary	✓	✓	✓	✓	✓	✓
bridge port detail	✓	✓	✓	✓	✓	✓
bridge port multicastLimit	✓	✓	✓	✓	✓	✓
bridge port stpState	✓	✓	✓	✓	✓	✓
bridge port stpCost	✓	✓	✓	✓	✓	✓
bridge port stpPriority	✓	✓	✓	✓	✓	✓
bridge port gvrpState	✓		✓			
bridge port address list	✓	✓	✓	✓	✓	✓
bridge port address add	✓	✓	✓	✓	✓	✓
bridge port address remove	✓	✓	✓	✓	✓	✓
bridge port address find	✓	✓	✓	✓	✓	✓
bridge port address flushAll	✓	✓	✓	✓	✓	✓
bridge port address flushDynamic	✓	✓	✓	✓	✓	✓

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
Ch 11 Trunks						
bridge trunk autoMap summary		✓				
bridge trunk autoMap enable		✓				
bridge trunk autoMap disable		✓				
bridge trunk autoMap test		✓				
bridge trunk summary	✓	✓	✓	✓	✓	✓
bridge trunk detail	✓	✓	✓	✓	✓	✓
bridge trunk define	✓	✓	✓	✓	✓	✓
bridge trunk modify	✓	✓	✓	✓	✓	✓
bridge trunk remove	✓	✓	✓	✓	✓	✓
Ch 12 MultiPoint Link Aggregation						
bridge mpla summary				✓		
bridge mpla detail				✓		
bridge mpla mode				✓		
bridge mpla peerMacAddress				✓		
Ch 13 Resilient Links						
bridge link summary	✓			✓	✓	✓
bridge link detail	✓			✓	✓	✓
bridge link define	✓			✓	✓	✓
bridge link linkState	✓			✓	✓	✓
bridge link activePort	✓			✓	✓	✓
bridge link modify	✓			✓	✓	✓
bridge link remove	✓			✓	✓	✓

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
Ch 14 Virtual LANs (VLANs)						
bridge vlan summary	✓	✓	✓	✓	✓	✓
bridge vlan detail	✓	✓	✓	✓	✓	✓
bridge vlan define	✓	✓	✓	✓	✓	✓
bridge vlan modify	✓	✓	✓	✓	✓	✓
bridge vlan remove	✓	✓	✓	✓	✓	✓
bridge vlan mode	✓	✓	✓	✓	✓	✓
bridge vlan stpMode	✓		✓			
bridge vlan vlanAwareMode	✓		✓			
Ch 15 Packet Filters						
bridge packetFilter list	✓		✓			
bridge packetFilter display	✓		✓			
bridge packetFilter create	✓		✓			
bridge packetFilter delete	✓		✓			
bridge packetFilter edit	✓		✓			
bridge packetFilter load	✓		✓			
bridge packetFilter assign	✓		✓			
bridge packetFilter unassign	✓		✓			
bridge packetFilter portGroup list	✓		✓			
bridge packetFilter portGroup display	✓		✓			
bridge packetFilter portGroup create	✓		✓			
bridge packetFilter portGroup delete	✓		✓			
bridge packetFilter portGroup addPort	✓		✓			
bridge packetFilter portGroup removePort	✓		✓			

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
Part VI Routing Protocols						
Ch 16 IP						
ip interface summary	✓		✓	✓	✓	✓
ip interface detail	✓		✓			
ip interface define	✓		✓	✓	✓	✓
ip interface modify	✓		✓	✓	✓	✓
ip interface remove	✓		✓	✓	✓	✓
ip interface arpProxy	✓		✓			
ip interface broadcastAddress	✓		✓			
ip interface directedBroadcast	✓		✓			
ip interface icmpRedirect	✓		✓			
ip interface icmpRouterDiscovery	✓		✓			
ip interface statistics	✓		✓			
ip route display	✓		✓	✓	✓	✓
ip route static	✓		✓	✓	✓	✓
ip route remove	✓		✓	✓	✓	✓
ip route flush	✓		✓	✓	✓	✓
ip route default	✓		✓	✓	✓	✓
ip route noDefault	✓		✓	✓	✓	✓
ip route findRoute	✓		✓	✓	✓	✓
ip arp display	✓		✓	✓	✓	✓
ip arp static	✓		✓	✓	✓	✓
ip arp remove	✓		✓	✓	✓	✓
ip arp flushAll	✓		✓	✓	✓	✓

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
ip arp flushDynamic	✓		✓	✓	✓	✓
ip arp age	✓		✓			
ip arp statistics	✓		✓			
ip dns display	✓		✓	✓	✓	✓
ip dns domainName	✓		✓	✓	✓	✓
ip dns define	✓		✓	✓	✓	✓
ip dns modify	✓		✓	✓	✓	✓
ip dns remove	✓		✓	✓	✓	✓
ip dns nslookup	✓		✓	✓	✓	✓
ip udpHelper display	✓		✓			
ip udpHelper define	✓		✓			
ip udpHelper remove	✓		✓			
ip udpHelper hopCountLimit	✓		✓			
ip udpHelper threshold	✓		✓			
ip udpHelper interface first	✓		✓			
ip udpHelper interface even	✓		✓			
ip udpHelper interface sequential	✓		✓			
ip routing	✓		✓			
ip rip display	✓		✓	✓	✓	✓
ip rip mode	✓		✓	✓	✓	✓
ip rip compatibilityMode	✓		✓			
ip rip cost	✓		✓			
ip rip poisonReverse	✓		✓			
ip rip routeAggregationMode	✓		✓			

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
ip rip password	✓		✓			
ip rip addAdvertisement	✓		✓			
ip rip removeAdvertisement	✓		✓			
ip rip policy summary	✓		✓			
ip rip policy detail	✓		✓			
ip rip policy define	✓		✓			
ip rip policy modify	✓		✓			
ip rip policy remove	✓		✓			
ip rip statistics	✓		✓	✓	✓	✓
ip ping	✓		✓	✓	✓	✓
ip advancedPing	✓		✓	✓	✓	✓
ip traceRoute	✓		✓	✓	✓	✓
ip advancedTraceRoute	✓		✓	✓	✓	✓
ip statistics	✓		✓	✓	✓	✓
Ch 17 VRRP						
ip vrrp summary	✓		✓			
ip vrrp detail	✓		✓			
ip vrrp define	✓		✓			
ip vrrp modify	✓		✓			
ip vrrp remove	✓		✓			
ip vrrp mode	✓		✓			
ip vrrp neighbor	✓		✓			
ip vrrp statistics	✓		✓			

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
Ch 18 IP Multicast						
ip multicast dvmrp interface summary	✓		✓			
ip multicast dvmrp interface detail	✓		✓			
ip multicast dvmrp interface mode	✓		✓			
ip multicast dvmrp interface metric	✓		✓			
ip multicast dvmrp tunnels summary	✓		✓			
ip multicast dvmrp tunnels define	✓		✓			
ip multicast dvmrp tunnels remove	✓		✓			
ip multicast dvmrp tunnels address	✓		✓			
ip multicast dvmrp tunnels threshold	✓		✓			
ip multicast dvmrp tunnels metric	✓		✓			
ip multicast dvmrp routeDisplay	✓		✓			
ip multicast dvmrp cacheDisplay	✓		✓			
ip multicast dvmrp default	✓		✓			
ip multicast igmp interface summary	✓		✓			
ip multicast igmp interface detail	✓		✓			
ip multicast igmp interface TTL	✓		✓			
ip multicast igmp snooping	✓		✓			
ip multicast igmp querying	✓		✓			
ip multicast cache	✓		✓			
ip multicast traceRoute	✓		✓			

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
Ch 19 OSPF						
ip ospf areas display	✓		✓			
ip ospf areas defineArea	✓		✓			
ip ospf areas modifyArea	✓		✓			
ip ospf areas removeArea	✓		✓			
ip ospf areas addRange	✓		✓			
ip ospf areas modifyRange	✓		✓			
ip ospf areas removeRange	✓		✓			
ip ospf defaultRouteMetric display	✓		✓			
ip ospf defaultRouteMetric define	✓		✓			
ip ospf defaultRouteMetric remove	✓		✓			
ip ospf interface summary	✓		✓			
ip ospf interface detail	✓		✓			
ip ospf interface statistics	✓		✓			
ip ospf interface mode	✓		✓			
ip ospf interface priority	✓		✓			
ip ospf interface areaID	✓		✓			
ip ospf interface cost	✓		✓			
ip ospf interface delay	✓		✓			
ip ospf interface hello	✓		✓			
ip ospf interface retransmit	✓		✓			
ip ospf interface dead	✓		✓			
ip ospf interface password	✓		✓			
ip ospf linkStateData databaseSummary	✓		✓			

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
ip ospf linkStateData router	✓		✓			
ip ospf linkStateData network	✓		✓			
ip ospf linkStateData summary	✓		✓			
ip ospf linkStateData external	✓		✓			
ip ospf neighbors display	✓		✓			
ip ospf neighbors add	✓		✓			
ip ospf neighbors remove	✓		✓			
ip ospf routerID	✓		✓			
ip ospf partition display	✓		✓			
ip ospf partition modify	✓		✓			
ip ospf stubDefaultMetric display	✓		✓			
ip ospf stubDefaultMetric define	✓		✓			
ip ospf stubDefaultMetric remove	✓		✓			
ip ospf virtualLinks summary	✓		✓			
ip ospf virtualLinks detail	✓		✓			
ip ospf virtualLinks statistics	✓		✓			
ip ospf virtualLinks define	✓		✓			
ip ospf virtualLinks remove	✓		✓			
ip ospf virtualLinks areaID	✓		✓			
ip ospf virtualLinks router	✓		✓			
ip ospf virtualLinks delay	✓		✓			
ip ospf virtualLinks hello	✓		✓			
ip ospf virtualLinks retransmit	✓		✓			
ip ospf virtualLinks dead	✓		✓			

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
ip ospf virtualLinks password	✓		✓			
ip ospf policy summary	✓		✓			
ip ospf policy detail	✓		✓			
ip ospf policy define	✓		✓			
ip ospf policy modify	✓		✓			
ip ospf policy remove	✓		✓			
ip ospf statistics	✓		✓			
Ch 20 IPX						
ipx interface display	✓		✓			
ipx interface define	✓		✓			
ipx interface modify	✓		✓			
ipx interface remove	✓		✓			
ipx interface SAPadvertising	✓		✓			
ipx interface RIPadvertising	✓		✓			
ipx route display	✓		✓			
ipx route secondary	✓		✓			
ipx route static	✓		✓			
ipx route remove	✓		✓			
ipx route flush	✓		✓			
ipx server display	✓		✓			
ipx server static	✓		✓			
ipx server remove	✓		✓			
ipx server flush	✓		✓			
ipx server secondary	✓		✓			

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
ipx forwarding	✓		✓			
ipx rip mode	✓		✓			
ipx rip triggered	✓		✓			
ipx rip policy summary	✓		✓			
ipx rip policy define	✓		✓			
ipx rip policy modify	✓		✓			
ipx rip policy remove	✓		✓			
ipx sap mode	✓		✓			
ipx sap triggered	✓		✓			
ipx sap policy summary	✓		✓			
ipx sap policy detail	✓		✓			
ipx sap policy define	✓		✓			
ipx sap policy modify	✓		✓			
ipx sap policy remove	✓		✓			
ipx output-delay	✓		✓			
ipx statistics summary	✓		✓			
ipx statistics rip	✓		✓			
ipx statistics sap	✓		✓			
ipx statistics forwarding	✓		✓			
ipx statistics interface	✓		✓			
ipx oddLengthPadding	✓		✓			
ipx NetBIOS	✓		✓			
ipx secondary	✓		✓			

Table 6 Command Summary (continued)


Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
Ch 21 AppleTalk						
appletalk interface summary	✓		✓			
appletalk interface detail	✓		✓			
appletalk interface define	✓		✓			
appletalk interface modify	✓		✓			
appletalk interface remove	✓		✓			
appletalk interface statistics	✓		✓			
appletalk route display	✓		✓			
appletalk route flush	✓		✓			
appletalk aarp display	✓		✓			
appletalk aarp remove	✓		✓			
appletalk aarp flush	✓		✓			
appletalk zone display network	✓		✓			
appletalk zone display zone	✓		✓			
appletalk forwarding	✓		✓			
appletalk checksum	✓		✓			
appletalk sourceSocket	✓		✓			
appletalk ping	✓		✓			
appletalk statistics ddp	✓		✓			
appletalk statistics rtmp	✓		✓			
appletalk statistics zip	✓		✓			
appletalk statistics nbp	✓		✓			

Table 6 Command Summary (continued)

Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
Part VII Traffic Policy						
Ch 22 Quality of Service and RSVP						
qos classifier summary	✓		✓			
qos classifier detail	✓		✓			
qos classifier define	✓		✓			
qos classifier modify	✓		✓			
qos classifier remove	✓		✓			
qos control summary	✓		✓			
qos control detail	✓		✓			
qos control define	✓		✓			
qos control modify	✓		✓			
qos control remove	✓		✓			
qos ldap display	✓					
qos ldap enable	✓					
qos ldap disable	✓					
qos rsvp summary	✓		✓			
qos rsvp detail	✓		✓			
qos rsvp enable	✓		✓			
qos rsvp disable	✓		✓			
qos bandwidth display	✓		✓			
qos bandwidth modify	✓		✓			
qos excessTagging display	✓		✓			
qos excessTagging enable	✓		✓			
qos excessTagging disable	✓		✓			

Table 6 Command Summary (continued)


Commands	3500	9000: Layer 2	9000: Layer 3	9400	3900	9300
qos statistics interval	✓		✓			
qos statistics receive	✓		✓			
qos statistics transmit	✓		✓			
Part VIII Monitoring						
Ch 23 Event Log						
log display	✓					
log devices	✓					
log services	✓					
Ch 24 Roving Analysis						
analyzer display	✓	✓	✓	✓	✓	✓
analyzer add	✓	✓	✓	✓	✓	✓
analyzer remove	✓	✓	✓	✓	✓	✓
analyzer start	✓	✓	✓	✓	✓	✓
analyzer stop	✓	✓	✓	✓	✓	✓



SYSTEM-LEVEL FUNCTIONS

Chapter 3 **System Environment**

Chapter 4 **Module Environment**



3

SYSTEM ENVIRONMENT

This chapter provides guidelines and other key information about how to use `system` commands to:


- Set and modify general system parameters. Important considerations and options are also provided where applicable
- Configure management access to the system (through one of two serial connection types)
- Configure management access through the serial port. (For information about commands for configuring an out-of-band management interface, see Chapter 5.)



For more information about administering your system environment, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



In addition, this chapter describes the `script` and `logout` options from the top-level menu.

system display

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Generates a system configuration display that includes software and hardware revision numbers, module status information, and warning messages for certain system conditions.

Valid Minimum Abbreviation

sy d

Important Consideration

- A message appears in the display if any module fails a diagnostic test at start-up.

Fields in the System Display

Field	Description
Diagnostics	Whether a module has passed or failed diagnostics
Memory size (AP, FP, Flash, and Buffer)	Memory capacities of the system processors
POV	Power on verification
SysBoot	Boot software revision
ExtDiags	Extended diagnostics version number
Part number	Each module's 3Com part identification number
Product number	Each module's 3Com 3C product identification number
Rev	Unique number assigned to the hardware build by 3Com
Serial number	Each module's unique serial number
Slot number	Slot position of each hardware module
System ID	Unique number that is assigned to the system by 3Com
System name	64-character (maximum) user-defined alphanumeric name that uniquely identifies the system on your network
System up time	Time since the last system reboot
Time in service	Total operational time since the system was manufactured
Type of module	Type of physical ports
Version, build date, and time	System software version number, and date and time when the software was built



You configure the system parameters for the CoreBuilder® 9000 system through the Enterprise Management Engine (EME). See the CoreBuilder 9000 Enterprise Management Engine User Guide for a complete list and detailed explanation of the CoreBuilder 9000 system commands.

✓ 3500

9000

9400

3900

9300

system fileTransfer Sets the file transfer protocol to either Trivial File Transfer Protocol (TFTP) or File Transfer Protocol (FTP). Use this protocol to retrieve or store files across the network for system functions such as scripts, snapshots, software updates, and nvData save and restore.

Valid Minimum Abbreviation

sy f

Options

Prompt	Description	Possible Values	[Default]
File transfer protocol	File transfer protocol for the system	<div><div>■</div>TFTP</div> <div><div>■</div>FTP</div>	TFTP

system console
webHelpConfig

Sets the Uniform Resource Locator (URL) for access to the Web Management Help system.

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy co wh

Options

Prompt	Description	Possible Values	[Default]
Enter Web help installation URL	URL where the Web Management Help system files are located	–	–

✓ 3900
✓ 9300

Prompt	Description	Possible Values	[Default]
Web management	Whether remote access to the Web Management system is allowed	<ul style="list-style-type: none"> ■ enabled ■ disabled 	enabled

system console consoleAccess

Controls remote access via Telnet or modem to the system console.

Valid Minimum Abbreviation

sy co co

Options

- 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Prompt	Description	Possible Values	[Default]
Console access	Whether remote access to the system console is allowed	<ul style="list-style-type: none">■ enabled■ disabled	enabled

Enables or disables the control key combination (default: Ctrl+X) that allows you to reboot the system from the Administration Console.

Valid Minimum Abbreviation

sy co ct

Options

Prompt	Description	Possible Values	[Default]
Control keys	Whether you want to enable or disable the reboot control key combination	<ul style="list-style-type: none"> ■ enabled ■ disabled 	enabled

✓ 9300

system console password

Sets one of the password levels for the Administration Console. There are three levels of password for the Administration Console.

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy co p

Important Considerations

- The Administration Console supports three levels of access:
 - One for only browsing or viewing (read)
 - One for configuring network parameters (write)
 - One for full system administration (administer)
- When you log on for the first time, press Return or Enter at the password prompt because the initial passwords that are stored in the nonvolatile memory of the system are null for all access levels.
- To change passwords, you must enter the Console at the administer access level.
- The system does not display the password in the field as you type.
- Set a password for each access level that you want to configure.

Options

Prompt	Description	Possible Values	[Default]
Access level	Level of access for the person logging on to the system	<ul style="list-style-type: none">■ read■ write■ administer	read
Password	Text string typed by the person logging on	<ul style="list-style-type: none">■ A string of up to 32 case-sensitive characters■ Enter (for a null password)	–

Prompt	Description	Possible Values	[Default]
New screen height	New screen height in lines	<ul style="list-style-type: none"> 20 – 200 0 (to receive no prompts) 	24
Default value	Default screen height for future Administration Console sessions	<ul style="list-style-type: none"> y (yes) n (no) 	y

**system console
security display**

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Displays a summary of trusted IP client information.

Valid Minimum Abbreviation

sy co se di

Important Consideration

- If you do not have any trusted IP clients configured, this command displays only the first two fields.

Fields in the System Console Security Display

Field	Description
Trusted client access only	Whether the trusted IP client feature is enabled or disabled
Deny message	Text of the current message that is sent to a user who is not a trusted client
Index	Index number that is associated with the trusted IP client
Trusted IP address	IP address of the trusted IP client
Mask	Subnet mask that is associated with the trusted IP address

**system console
security define**

Gives a client trusted access to your system by adding the client IP address and subnet mask to an access list.

✓ 3500
9000
✓ 9400

✓ 3900
✓ 9300

**Valid Minimum Abbreviation**

sy co se de

Important Considerations

CAUTION: Be careful when you define trusted IP clients. If you specify an incorrect IP address or subnetwork address when you define a trusted IP client, you can affect your own ability to access the system. See the Implementation Guide for your system.

- Configure trusted IP clients in this order:
 - Define the trusted IP clients using `system console security define`.
 - Display the list of configured trusted IP clients using `system console security display` to verify that you have configured the trusted IP clients correctly.
 - Enable the system to verify trusted IP clients using `system console security access`.
- You can configure up to five IP addresses or five subnetwork addresses as trusted IP clients.
- An IP address or subnetwork address can be used to access the system only if it is on the trusted IP client list.
- Use the subnet mask to allow trusted access to all addresses on a particular subnetwork. Examples: The IP address 158.99.112.219 with a subnet mask of 255.255.255.0 allows trusted access to all addresses on the 158.99.112 subnetwork. The IP address 158.99.112.219 with a subnet mask of 255.255.255.255 allows access only by the single IP address 158.99.112.219.
- The trusted IP client information is retained after a system reboot; that is, it is saved in nvData.

Options

Prompt	Description	Possible Values	[Default]
IP address	IP address of the interface, chosen from the range of addresses that are assigned to your organization. This address is specific to your network and system.	Any valid IP address	–
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnet number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnet part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.	A valid subnet mask in accordance with the bits used for network number, subnet, and host number	Depends on specified IP address

system console
security remove

Removes an IP address from the trusted IP client access list.

Valid Minimum Abbreviation

sy co se r

Important Considerations

- If you remove a trusted IP client definition through the Administration Console, the definition is also removed in the Web Management Console, and vice versa.
- This command takes effect immediately. You are not prompted to confirm the deletion.

Options

Prompt	Description	Possible Values	[Default]
Trusted IP client index	One or more index numbers of the IP addresses that you want to remove	<ul style="list-style-type: none">■ 1 – 5■ all■ ? (for a list of selectable indexes)	1 (if only one client)

system console security access

Enables or disables whether the system verifies trusted IP clients on your system.

- ✓ 3500
- 9000
- ✓ 9400

Valid Minimum Abbreviation

sy co se a

Important Considerations

- ✓ 3900
- ✓ 9300



CAUTION: Be careful when you define trusted IP clients. If you specify an incorrect IP address or subnetwork address when you define a trusted IP client, you can affect your own ability to access the system. See the Implementation Guide for your system.

- Configure trusted IP clients in this order:
 - Define the trusted IP clients using system console security define.
 - Display the list of configured trusted IP clients using system console security display to verify that you have configured the trusted IP clients correctly.
 - Enable the system to verify trusted IP clients using system console security access.

Options

Prompt	Description	Possible Values	[Default]
Trusted client access only	Whether you want to allow or disallow your system to restrict access according to your list of trusted IP clients	<ul style="list-style-type: none">■ enabled■ disabled	disabled

**system console
security message**

Defines the text that is displayed to a prospective user when access to your system is denied.

✓ 3500

9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

sy co se m

Important Consideration

- Use system console security display to view the text of the current deny message.

Options

Prompt	Description	Possible Values	[Default]
Deny message	Text that is displayed to a prospective user whose IP address does not appear on the list of trusted users	Alphanumeric text of up to 85 characters and spaces	"You are not considered a trusted user. Please see your network administrator."

**system console
timeout timeOut**

Configures the system to disconnect remote sessions after a specified interval of no activity.

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy co t t

Important Considerations

- The default inactive time interval is 30 minutes.
- To change the timeout interval value before the system disconnects remote sessions, see “system console timeout interval” next for details.

Options

Prompt	Description	Possible Values	[Default]
Timeout state	Whether you want to enable or disable the timeout state	<ul style="list-style-type: none">■ enabled■ disabled	disabled

**system console
timeout interval**

Sets the remote timeout interval to a value from 1 minute through 60 minutes.

- ✓ 3500
- 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy co t i

Important Consideration

- To enable or disable the inactive timeout interval for remote sessions, see the preceding command, “system console timeout timeOut” for details.

Options

Prompt	Description	Possible Values	[Default]
Telnet timeout interval	Timeout interval	1 – 60 minutes	30

**system snapshot
summary**

Captures an image of all system summary display screens. This display reflects each application's status at the time that you use the snapshot feature.

- ✓ 3500
- 9000
- ✓ 9400

Valid Minimum Abbreviation

sy sn su

- ✓ 3900
- ✓ 9300

Captures an image of all system detail screens. The display reflects the current values of all fields and counters at the time that you use the snapshot feature.

Valid Minimum Abbreviation

✓ 3900
✓ 9300

system snapshot save Sends detail screens to a file on the host machine that you specify.

✓ 3500

9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

`sy sn sa`

Important Considerations

- The CoreBuilder 3500 uses the Trivial File Transfer Protocol (TFTP) or File Transfer Protocol (FTP) to transfer the files to the host, depending on the setting for the `system fileTransfer` option. The SuperStack® II Switch 3900, Switch 9300, and the CoreBuilder 9400 use TFTP to transfer files.

Before you transfer files:

- You must create the file to receive the snapshot images on an FTP or TFTP server *before* you send the images to the file.
- You supply the IP address of the host and specify the file according to the requirements of your TFTP or FTP implementation.
- Some TFTP implementations require that you store the file in the same directory in which the TFTP daemon (server) is running on the remote host.
- Because TFTP provides no user authentication, give the file *loose* permissions to make it publicly readable and writable. TFTP servers do not grant requests for file access.
- On the CoreBuilder 3500, if you use FTP for `system fileTransfer`, you must enter a login name and password if you are sending a file to an FTP server.

TFTP Procedure

- 1 Create an empty file with open write permissions on the host to store the system display images.
- 2 From the top level of the Administration Console, enter:
system snapshot save
- 3 Enter the IP address of the host on which you want to save the display images.
- 4 If your TFTP implementation requires a full path name, enter the full path of the file that is designated to contain the display images. (Some implementations allow you to specify only the file name and the system uses the default TFTP directory.)

While the system sends the files to the host, it displays the name of each display image that it transfers. When the transmission is complete, the system displays a message that the transfer is complete and displays the file name and the name of the host on which it stored the file.

FTP Procedure (3500 Only)

- 1** Create an empty file with open write permissions on the host to store the system display images.
- 2** From the top level of the Administration Console, enter:
system snapshot save
- 3** Enter the IP address of the host on which you want to save the display images.
- 4** Enter the full pathname of the file that you designated.
- 5** Enter your username and password.

While the system sends the files to the host, it displays the name of each display image that it transfers. When the transmission is complete, the system displays a message that the transfer is complete and displays the file name and the name of the host on which it stored the file.

**system
softwareUpdate**

Loads a new revision of system software.

Valid Minimum Abbreviation

sy so

Important Considerations

- The CoreBuilder 3500 uses the Trivial File Transfer Protocol (TFTP) or File Transfer Protocol (FTP) to transfer the files to the host, depending on the setting for the `system fileTransfer` option. The SuperStack II Switch 3900, Switch 9300, and the CoreBuilder 9400 use TFTP to transfer files.
- Before you attempt to install the system software, make sure that you have extended memory installed on your system.
- You can load the system software into flash memory while the system is operating. The system does not have to be powered down.
- Make sure that the FTP server or TFTP server software is running on the device from which you are installing the software.
- Make sure that you have defined an IP address on your system.
- Some FTP servers or TFTP servers do not accept the full pathname. If this is true on your server, enter the image filename only.
- On the CoreBuilder 3500, if you are using the FTP file transfer protocol, you must enter a login name and password.

Options

Prompt	Description	Possible Values	[Default]
Host IP address	IP address of the machine from which to load the software update	Any valid IP address	–
Install file name	Name of the image to be loaded	–	–

✓ 3500
9000
✓ 9400

✓ 3900
✓ 9300

- ✓ 3900
✓ 9300

- Use this command to determine if you need a newer baseline for viewing statistics.
- The system also indicates if you have not yet set a baseline on the system.

system baseline set Resets the baseline counters to zero.

✓ 3500

9000

✓ 9400

Valid Minimum Abbreviation

sy b s

Important Considerations

✓ 3900

✓ 9300

- Baselining is automatically enabled when you set a baseline.
- The system maintains the accumulated totals since power-up.
- The baseline is time-stamped.

**system baseline
requestedState**

Enables or disables a baseline.

Valid Minimum Abbreviation

sy b r

Important Considerations

- When you reenable a baseline, the counters return to the values that have accumulated since the most recent baseline that you set.
- Disabling a baseline returns the counters to the total accumulated values since the last power-up.

Options

Prompt	Description	Possible Values	[Default]
New value	Whether you want to enable or disable the baseline	<div><div>■ enabled</div><div>■ disabled</div></div>	enabled

system serialPort terminalSpeed

Sets the terminal speed of your system serial port. The terminal speed is set by changing the terminal connection port baud rates.

- ✓ 3500
- 9000
- 9400
- 3900
- 9300

Valid Minimum Abbreviation

sy se t

Important Considerations

- When you change the terminal port baud rate to something other than 9600, the new setting becomes the new default, even after you use the `system nvData reset` option.
- You must adjust the baud setting of your terminal or terminal emulator to match your system serial port's baud rate before you can reestablish communication using the terminal port.
- You can use this command through the terminal serial port or through a Telnet session. However, if you change the terminal speed while you are in a Telnet session, you must reboot the system for the change to take effect.

Options

Prompt	Description	Possible Values	[Default]
Terminal speed	Signal speed for the terminal connection	■ 19200	9600
		■ 9600	
		■ 4800	
		■ 2400	
		■ 1200	
Confirmation	Confirmation of terminal speed change	■ y (yes)	–
		■ n (no)	

Procedure (Local Connection)

- 1 To set the terminal speed for the serial port, from the top level of the Administration Console, enter:
- system serialPort terminalSpeed

- 2 Enter the terminal speed setting for the serial port. See the Options table for supported terminal speed rates.

The system response depends on the cable status.



The terminal speed is referred to as baud rate in the following messages.

If the cable is connected to the terminal port when you set the terminal speed for that port, the system displays the following message:

Changing the baud rate may cause a loss of communication since you are currently connected via the serial port.

Are you sure you want to change the baud rate? (y/n):

- If you respond **y** (yes), the serial port's baud rate is changed immediately, and you lose the ability to communicate with any devices connected to the port until you adjust the device baud setting to match.
- If you respond **n** (no), the rate does not change, and the display returns to the previous menu.

Procedure (IP Interface)

- 1 From the top level of the Administration Console, enter:
system serialPort terminalSpeed
- 2 Enter the terminal speed setting for the terminal port.



The terminal speed is referred to as baud rate in the following messages.

After you select the new terminal speed rate, the system displays the following message:

The baud rate will not change until the system is rebooted.

To have your change take effect without rebooting, perform this command via the serial port.

Are you sure you want to change the baud rate? (y/n):

- If you respond **y** (yes), the rate is not changed until you reboot.
- If you respond **n** (no), the rate does not change, and the display returns to the previous menu.

system serialPort
modemSpeed

Sets the port speed for the modem port to match your external modem baud setting.

✓ 3500
9000
9400

3900
9300

Valid Minimum Abbreviation

sy se m

Important Considerations

- After you use this command, you must establish a connection between your current Console session and the modem port before you dial in. (See “system serialPort configModem” later in this chapter for details.)
- Be sure that the baud setting of the modem port matches that of your external modem.
- The system immediately changes the modem port baud setting.

Options

Prompt	Description	Possible Values	[Default]
Modem speed	Signal speed for the connection	<div>■ 19200 ■ 9600 ■ 4800 ■ 2400 ■ 1200</div>	9600

**system serialPort
baudRate**

- 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Sets the baud rate of your system serial port.

Valid Minimum Abbreviation

`sy se b`

Important Considerations

- The default setting for the serial port is 9600. You can change the setting to match the port speed on your terminal or modem. The default setting for the serial port is 9600. You can change the setting to match the port speed on your terminal or modem.
- When you change the baud rate to something other than 9600, the new setting becomes the new default, even after you use the `system nvData reset` option.
- You must adjust the baud rate setting of your terminal or terminal emulator's terminal interface processor (tip) to match your system serial port's speed before you can reestablish communication using the terminal port.
- You can use this command through the terminal serial port or through a Telnet session. However, if you change the terminal speed while you are in a Telnet session, you must reboot the system for the change to take effect.

Options

Prompt	Description	Possible Values	[Default]
New value	Baud rate for the serial port connection	<ul style="list-style-type: none">■ 19200■ 9600■ 4800■ 2400■ 1200	9600
Confirmation	Confirmation of baud rate change	<ul style="list-style-type: none">■ y (yes)■ n (no)	–

Procedure (Local Connection)

- 1 To set the baud rate for the serial port, from the top level of the Administration Console, enter:

`system serialPort baudRate`

- 2 Enter the baud setting for the serial port. The system supports the following baud rates:

- 19200
- 9600
- 4800
- 2400
- 1200

The system response depends on the cable status. If the cable is connected to the terminal port when you set the baud rate for that port, the system displays the following message:

Changing the baud rate may cause a loss of communication since you are currently connected via the serial port.

Are you sure you want to change the baud rate? (y/n):

- If you respond **y** (yes), the serial port's baud rate is changed immediately, and you lose the ability to communicate to any devices connected to it until you adjust the device baud setting to match.
- If you respond **n** (no), the rate does not change, and the display returns to the previous menu.

Procedure (IP Interface)

- 1 From the top level of the Administration Console, enter:

```
system serialPort baudRate
```

- 2 Enter the baud setting for the terminal port.

After you select the new baud, the system displays the following message:

The baud rate will not change until the system is rebooted.

To have your change take effect without rebooting, perform this command via the serial port.

Are you sure you want to change the baud rate? (y/n):

- If you respond **y** (yes), the rate is not changed until you reboot the system.
- If you respond **n** (no), the rate does not change, and the display returns to the previous menu.

**system serialPort
serialPortMode**

Configures the system serial port to establish either a terminal connection or a modem connection.

- 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy se s

Options

Prompt	Description	Possible Values	[Default]
Serial port	Type of serial port configuration that you want	<div><div>■ console</div><div>■ modem</div></div>	console

Procedures

To change the serial port configuration from `console` to `modem`, perform the following steps:

- 1 Change the serial port configuration from `console` to `modem`.
- 2 Disconnect the console cable.
- 3 Connect the modem cable.

The system is ready for you to establish a modem connection. See “system serialPort configModem” next for details.

To change the serial port configuration from `modem` to `console`, perform the following steps:

- 1 Change the serial port configuration from `modem` to `console`.
- 2 Disconnect the modem cable.
- 3 Connect the console cable.

The system is ready for you to establish a console connection. (See “system serialPort baudRate” earlier in this chapter.)

**system serialPort
configModem**

Configures the external modem from the Administration Console.

Valid Minimum Abbreviation

sy se c

Important Considerations

- The system transmits characters that you have entered as output on the modem port. The system echoes characters that it receives as input on the modem port to the current Console session. Thus, the Console appears to be directly connected to the external modem.
- You may need to change the baud of the modem to match that of your modem port.

Procedure

- 1 From the top level of the Administration Console, enter:

```
system serialPort configModem
```

You can now enter commands that support the appropriate parameters for your network. All characters that you enter are transmitted to the modem port until you type the escape sequence in step 2.

- 2 When the modem is configured, enter the escape sequence ~] with no intervening characters or spaces.

Entering the escape sequence breaks the connection to the modem serial port and returns you to the previous Administration Console menu.

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

**system serialPort
enableModem**

Enables the external modem from the Administration Console.

Valid Minimum Abbreviation

sy se e

Important Consideration

- You must configure the external modem before you can enable it. See the configModem command description on the previous page.

✓ 3500
9000
✓ 9400

✓ 3900
✓ 9300

system name Assigns or changes the name of the system. The system name identifies the system to users on other systems in the network.

✓ 3500
9000
✓ 9400

Valid Minimum Abbreviation

sy na

Important Considerations

- Assign an easily recognizable and unique name for each system. For example, name the system according to its physical location, as in PARIS-ENGLAB1.
- Use quotation marks (") around any string that has embedded spaces.
- Use double quotation marks (") to enter an empty string.
- The new system name appears the next time that you display the system configuration.

Options

Prompt	Description	Possible Values	[Default]
New string	New or changed name for the system	<ul style="list-style-type: none">■ A string of up to 64 case-sensitive characters■ ? (to get information about the naming guidelines)	–

system time Displays and changes the system's current date and time, timezone, and daylight saving time.

✓ 3500

9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

sy t

Important Considerations

- The system's internal clock is set at the factory. You may want to reset the system date and time to match the system's physical location.
- 00 specifies the year 2000 for all 3Com products. See the 3Com Web site for more details.

Options

Field	Description
DateTime	Starting date and time in the following format: yyyy-mm-ddThh:mm:ss
Timezone	User-configured time zone (for example, GMT for Greenwich Mean Time)
dst	<ul style="list-style-type: none">■ If you enter y to the prompt, a sub-menu appears that lists the Daylight Savings Time around the world and a user specified option for start and end dates■ If you enter n, you are returned to the Time menu

system time datetime Sets the system's date and time.

- ✓ 3500
- 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

`sy t d`

Important Considerations

- The system's internal clock is set at the factory. You may want to reset the system date and time to match the system's physical location.
- `00` specifies the year 2000 for all 3Com products. See the 3Com Web site for more details.

Procedure

- 1 To change the system date or time, enter:

`da`

The system displays the current date and time and then prompts you to change the time.
- 2 Enter **y** (yes) or **n** (no) at the prompt.

If you respond **n**, the top-level menu appears.

If you respond **y**, the system prompts you for the correct date and time.
- 3 Enter the correct date and time in this format:

`ccyy-mm-ddThh:mm:ss`

Format	Description
yyyy	century and last two digits of the year (00-99)
<i>first</i> mm	date (1 – 31)
dd	month (1 – 12)
T	Time designator
hh	hour (1 – 12)
<i>second</i> mm	minute (00 – 59)
ss	second (00 – 59)

system time timezone Configures the local time zone and daylight savings time values.

✓ 3500
9000
9400

3900
9300

Valid Minimum Abbreviation

sy t timez

Important Considerations

- Displays the current time zone table, with time zone indexes and the time zone identifiers before it prompts you to select a time zone index.
- The local time zone value adjusts the server reply universal time to local time properly.
- The default time zone is Greenwich Mean Time (GMT).

Options

Prompt	Description	Possible Values	[Default]
Time zone index	Index number of the time zone that you want to configure	<ul style="list-style-type: none">■ 1 – 28*■ ? (lists the default selection and selectable values)	1 (GMT)

* Index number 28 prompts for an offset from the GMT in the following time format: ±hh:mm.

System Time Timezone Example (3500)

Select menu option (system/sntp): **timez**

Index	Time Zone
1	[GMT+0:00] GMT/WET/UT
2	[GMT-1:00] WAT
3	[GMT-2:00] AT
4	[GMT-3:00] Brasilia/Buenos Ar/GeorgeTown
5	[GMT-4:00] AST
6	[GMT-5:00] EST
7	[GMT-6:00] CST
8	[GMT-7:00] MST
9	[GMT-8:00] PST
10	[GMT-9:00] YST
11	[GMT-10:00] AHST/CAT/HST
12	[GMT-11:00] NT
13	[GMT-12:00] IDLW
14	[GMT+1:00] CET/FWT/MET/MEWT/SWT
15	[GMT+2:00] EET
16	[GMT+3:00] BT
17	[GMT+4:00] ZP4
18	[GMT+5:00] ZP5
19	[GMT+5:30] Bombay/Calcutta/Madras/New Dehli/Colombo
20	[GMT+6:00] ZP6
21	[GMT+7:00] WAST
22	[GMT+8:00] CCT
23	[GMT+9:00] JST
24	[GMT+9:30] Darwin/Adelaide
25	[GMT+10:00] EAST/GST
26	[GMT+11:00] Magadan/Solomon Is/N. Caledonia
27	[GMT+12:00] IDLE/NZST/NZT
28	Input an offset from GMT

Select timezone index {1-28|?} [1]:

system time dst Sets daylight savings time.

✓ 3500
9000
9400

Valid Minimum Abbreviation

`sy t ds`

Important Consideration

- Displays the daylight savings time periods for various parts of the world.

Procedure

- 1 To set daylight savings time, enter:

`ds`

The system displays the following prompt:

Do you want to set the Daylight Saving Time?(n,y)[n]:

- 2 Enter `y` (yes) or `n` (no) at the prompt.

If you respond `n`, the Time menu appears.

If you respond `y`, the system displays the following:

1 First Sunday in April to last Sunday in October (North America)

2 Last Sunday in March to last Sunday in October (Europe, parts of Asia)

3 Last Sunday in October to last Sunday in March (Parts of Australia)

4 Last Sunday in October to the Sunday on/after March 15th (New Zealand, parts of Australia)

5 Enter a start and an end dates for the current year

Select daylight saving time option {1-5|?} [1]:

The format for option 5 is: `ccyy-mm-ssThh:mm:ss`

Example: `1999-05-20T12:30:34`

- 3 Enter a daylight saving time option.

system nvData save

✓ 3500
9000
✓ 9400

Stores nonvolatile (NV) data on a server. The CoreBuilder 3500 uses the Trivial File Transfer Protocol (TFTP) or File Transfer protocol (FTP) to transfer the files to the host, depending on the setting for the `system fileTransfer` option. The SuperStack II Switch 3900, Switch 9300, and the CoreBuilder 9400 use TFTP to transfer files.

✓ 3900
✓ 9300

Valid Minimum Abbreviation

`sy nv s`

Important Considerations (TFTP)

- To store NV data, you must first create two files on the TFTP server *before* you send the data:
 - **Control file** — Use any filename that is meaningful to you.
Example: `ctrlfile`
 - **NV data file** — Use the control filename plus the `.nvd` extension.
Example: `ctrlfile.nvd`
- When the system saves NV data, it writes it to a disk file on a host computer (that is, a server) using TFTP or FTP. You can then retrieve the data from the disk file with the `restore` option.
- Some TFTP implementations require that you store the files in the same directory in which the TFTP daemon (server) is running on the remote host.
- Some TFTP implementations require a full path, while other implementations allow you to specify only the file name, and the system saves the file to the default TFTP directory. Consult your network administrator or TFTP documentation for details about your host system's TFTP implementation.
- Because TFTP provides no user authentication, give *loose* permissions to the control file and the NV data file on the remote host (that is, make the files publicly readable and writable). TFTP servers do not grant requests for file access.

Important Consideration (FTP and TFTP)

- During the save procedure, the current configuration can be altered. To detect this event, the software runs checksum on the NVRAM before and after the save.

If the checksum is different, you are notified and prompted to save the configuration again. In abnormal situations, this reiteration can continue indefinitely, so you are given the option to terminate the save. You are also prompted for a retry request after a network (TFTP) I/O failure.

Options

Prompt	Description	Possible Values	[Default]
Host IP address	IP address of the host on which you want to store the data file	Any valid IP address	Previous IP address used
NV control file	Full path of the file where you want to store the NV data	–	–
File label (optional)	Meaningful description of the file	32-character text string	–

Procedure

- 1 To save NV data, from the top level of the Administration Console, enter:

```
system nvData save
```

The system prompts you for information about saving the data. To accept the value in brackets, press Return. Any entry for IP address, filename, and user name becomes the new default.

- 2 Enter the IP address of the TFTP or FTP server.
- 3 If you are using TFTP and your implementation requires a full path, enter the full pathname of the control file. (Some implementations allow you to specify only the file name, and the server uses the default TFTP directory.)



*The system prompt says **nv Control file**, so enter the name of the control file without the **NVD** extension.*

- 4 Optionally, enter a label for the file.

Example:

```
Host IP Address [158.101.100.1]: 158.101.112.34
NV Control file (full pathname): nvdata
Enter an optional file label: Labdata
```

If the information is incorrect or if a connection cannot be made with the specified host, the system displays a message similar to this one:

```
Login incorrect.  
Error: Transfer Timed Out  
Error - I/O error while writing nonvolatile data
```

If a session is successfully opened, a system message notifies you of the success or failure of your save, as in the following examples:

```
Success  System NV data successfully stored on host 158.101.112.34.  
Failure  Saving system...transfer timed out.  
Error - I/O error while writing nonvolatile data. Do you wish  
to retry the save using the same parameters? (n,y) [y].
```

5 To save the data as proposed, enter **y**

If you enter **n**, the NV data is not saved and the previous menu appears.

The text of the failure message depends on the problem that the system encountered while it saved the NV data.

At the end of a successful save, the system display returns to the previous menu.

**system nvData
restore**

Restores the NV data that was previously saved to a file.

Valid Minimum Abbreviation

sy nv r

Important Considerations

- Before you attempt to restore the data to a system that has a different system ID, be aware that the following types of NV data may cause problems when they are restored:
 - Management IP addresses (defined in IP interface configurations) are saved as NV data and restored. To avoid duplicate IP address problems, you may need to change the IP address of defined interfaces before you connect the restored system to the network.
 - Statically configured Ethernet addresses are saved as NV data. Verify that you have no duplicate addresses when you restore the NV data.

Options

Prompt	Description	Possible Values	[Default]
Host IP address	IP address of the host on which the NV data file resides	Any valid IP address	Previous IP address that was used
NV control file	Location of the NV data file	<ul style="list-style-type: none">■ file name■ full path	Previous nv control file that was used
Do you wish to continue?	Confirmation of the operation. (You may not want to reboot because resetting nonvolatile data may leave the system in an inconsistent state, so the system reboots after each reset.)	<ul style="list-style-type: none">■ n (no)■ y (yes)	y

Procedure

- 1 To restore the NV data, from the top level of the Administration Console, enter:

system nvData restore

The system prompts you to enter information for restoring the NV data that was saved to a file.

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Press Return at any prompt to accept the current or default value in brackets.

- 2 Enter the IP address of the host on which the NV data file resides.
- 3 If you are using TFTP and your implementation requires a full path name, enter the full NV data file path and filename. Some implementations allow you to specify only the file name; the system uses the default TFTP directory.

When you restore system NV data, the software compares the system IDs, module types, and module revisions (if applicable) between the saved configuration and the system on which you are restoring the image.

- If the system finds an exact match between system IDs, module types, and module revisions, the system displays a reminder message and prompts you for verification before performing the restoration (see step 4).
- If there is *not* an exact match between system IDs, module types, and module revisions, the system displays a warning message and prompts you as follows:

```
WARNING - mismatch between saved system IDs (27DA00) and
current system (28E100)
Do you want to disregard this and continue the restore (n,y)
[y]:
```

If you want to continue the restoration, enter **y** (yes). If you do not want to continue, enter **n** (no).

- 4 At the next prompt, to have the system NV data restored as requested, enter **y** (yes). To terminate the restoration, enter **n** (no).

For example:

```
Restoring nonvolatile data may leave the system
in an inconsistent state and therefore a reboot is necessary
after each restore.
Do you wish to continue? (y/n): y
```

- If you enter **y**, the system displays the following messages:

```
Restoring nonvolatile data...done
Nonvolatile data successfully restored
```

The system automatically reboots itself after it restores the NV data.

- If you enter **n**, the restoration ends and the previous menu appears.

**system nvData
examine**

Displays the header information of the NV data file.

Valid Minimum Abbreviation

sy nv e

Important Considerations

- Some TFTP implementations allow you to specify only the file name, and the system uses the default TFTP directory.
- If a session is successfully opened, the system displays the header information that corresponds to the file name that you entered.

Options

Prompt	Description	Possible Values	[Default]
Host IP address	IP address of the host on which the NV data file resides	Any valid IP address	Previous IP address used
NV control file	Location of the NV data file	<ul style="list-style-type: none">■ file name■ full path	System NV data file

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

system nvData reset

Resets the system values to the factory defaults. You can then reconfigure the system from its original settings.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy nv rese

Important Considerations

- You are not permitted to perform an NV data reset from a Telnet session. An NV data Reset over Telnet clears all of your configurable information, including the IP interface of the box, and prevents you from managing the system without a direct console connection.
- If you enter **n** (no) when you are prompted to confirm the reset, the system displays the previous menu.



CAUTION: As a precaution, consider saving the existing NV data to a file before you reset all values to the factory defaults. Resetting NV data means that NV memory is set back to the factory defaults (except for the serial port baud rate, modem baud rate, and system boot parameters). Before you proceed, be sure that you want to reset your NV data.

Options

Prompt	Description	Possible Values	[Default]
Do you wish to continue?	Confirmation of the reset operation. (You may not want to reboot because resetting nonvolatile data may leave the system in an inconsistent state, so the system reboots after each reset.)	<ul style="list-style-type: none">■ n (no)■ y (yes)	y

system
clearDiagBlock

Prevents diagnostic information about failed modules from appearing in system display screens.

✓ 3500
9000
✓ 9400

Valid Minimum Abbreviation

sy cl

✓ 3900
✓ 9300

Important Consideration

- After you enter this command, the system immediately removes diagnostic information about failed modules from the SNMP MIB *swSysDiagnosticsGroup*.

Options

Prompt	Description	Possible Values	[Default]
Clear the diagnostic block	Confirmation of your decision to clear the diagnostic information	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y

system diagErrLog

Displays hardware diagnostic errors that have been saved in the flash memory. When the system is initializing, if the diagnostic software detects errors, and if the system completes initializing, the detected errors are written to flash memory and stored in a dynamic error log.

✓ 3500

9000

✓ 9400

Valid Minimum Abbreviation`mo dia`

✓ 3900

✓ 9300

Important Consideration

- The error messages are saved to flash memory until you power down the system or clear the error log with the `system clearDiagBlock` command.

system sntp display Displays Simple Network Time Protocol (SNTP) information.

✓ 3500

9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

`sy snt di`

Important Considerations

- SNTP handles the synchronization of system clocks in the network to the national time standards via distributed time servers.
- Your system provides the SNTP client.
- The display has two types of information:
 - **Configuration information** — User configurable parameters appear.
 - **Servers information** — Information returned by the server appears if you have defined SNTP servers. Otherwise, the `No Servers are defined` message appears.

Fields in the System SNTP Display

Field	Description
Configuration Information	
State	SNTP state. It is either <code>enabled</code> or <code>disabled</code>
PollInterval	Interval for the client to send requests to a specific server
Tolerance	Threshold for updating the local system time
Servers Information	
Server	Server IP address
Mode	Client's SNTP operating mode. This field always displays <code>Unicast</code> .
Version	Version number of the responding server (for example, 4 represents version 4, which is suitable for IPv4, IPv6, and OSI). The client version number is 3.
Stratum	8-bit integer that indicates the stratum level of the local clock (for example, a stratum value of 3 indicates a secondary reference via SNTP).
Poll	Maximum interval between successive messages
Delay	Roundtrip propagation delay from the server's reply
LastPktRcv	Date and time stamp of the last packet that was received from the specific server.

system sntp define

Specifies up to three Simple Network Time Protocol (SNTP) server IP addresses.

- ✓ 3500
- 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy snt de

Important Considerations

- You can define up to three SNTP servers for backup purposes.
- Your system provides the SNTP client.
- The system indicates that it is adding the IP address to the SNTP database. The server is assigned an index number.

Options

Prompt	Description	Possible Values	[Default]
Server's IP address	IP address of a server to add to the SNTP database	Valid IP address (except 0.0.0.0)	–

Prompt	Description	Possible Values	[Default]
Index	Index number of the server that you want to modify	Available Server index number	–
Server address	IP address of each configured server	Valid IP address	–

system sntp remove

Removes a Simple Network Time Protocol (SNTP) server IP address from the SNTP server list.

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy snt r

Options

Prompt	Description	Possible Values	[Default]
Index	Index number of the server that you want to remove	Available Server index number	–
Server address	IP address of each configured server	Valid IP address	–

system sntp state Enables or disable the Simple Network Time Protocol (SNTP) state for the system.

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy snt s

Options

Prompt	Description	Possible Values	[Default]
SNTP state	Whether you want to implement SNTP on the system	<ul style="list-style-type: none">disabledenabled	disabled

system sntp pollInterval

Sets a poll interval value. This value determines how often the Simple Network Time Protocol (SNTP) client sends a request to the SNTP server.

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy snt p

Important Consideration

- The default pollInterval value is once an hour (3600 seconds). The value 86400 (the pollInterval limit) is the number of seconds in a day.

Options

Prompt	Description	Possible Values	[Default]
Request poll interval	In seconds, the poll interval	64 – 86400 seconds	3600

✓ 3900
✓ 9300

Sets a tolerance threshold that is used to update the local system time.

Valid Minimum Abbreviation

sy sint to

Important Consideration

- If the difference between the server time and the local time exceeds the specified tolerance threshold, the client drops the server time and maintains the current local system time unchanged.

Options

Prompt	Description	Possible Values	[Default]
Time tolerance	Time threshold value, in seconds, that is used to update the local system time	0 – 3600 seconds	900

system reboot Reboots the system.

- ✓ 3500
- 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sy r

Important Considerations

- This command disconnects the present Administration Console session and starts another session whether your system is connected to the Administration Console by an external modem or through an rlogin or Telnet session.
- To view diagnostic information during reboots, connect your system through the Console serial port.

Options

Prompt	Description	Possible Values	[Default]
Reboot the system?	Confirmation that you want to reboot	<ul style="list-style-type: none">■ n (no)■ y (yes)	–

✓ 3500
9000
✓ 9400

✓ 3900
✓ 9300

script Executes a command file that you have written to expedite and automate Administration Console tasks.

Valid Minimum Abbreviation

sc

Important Considerations

- Any command that you can enter in the Administration Console can be part of a script. You can even script your entire system setup so that you can repeat the exact setup on other systems.
- You create scripts in an ASCII-based line editor, such as *EMACS* or *vi*. Scripts must be stored on the TFTP server. For the CoreBuilder 3500, you can select TFTP or FTP as the file transfer method. See “system fileTransfer” earlier in this chapter for more details.
- Some TFTP implementations require that you store the script file in the same directory in which the TFTP daemon (server) is running on the remote host.
- Because TFTP provides no user authentication, make the file permissions *loose* so that the public can read and write to the file. TFTP servers do not grant requests for file access.

Procedure

- 1 From the top level of the Administration Console, enter:

script

The system prompts you for information about where you have stored the script that you want to run: host IP address and file path

Press Return at any prompt to accept the current or default value in brackets.

- 2 Enter the path name to the script file. If you are using TFTP, see “system snapshot save” earlier in this chapter for more details about pathname requirements.

The task that you scripted runs in the Administration Console.

Example Script (3500)

This example scripts these tasks to initially configure your system:

- Changes the modem port baud
- Sets the system name
- Assigns an IP address for management
- Verifies the IP connection by pinging the system
- Enables Spanning Tree
- Sets up SNMP trap reporting

```
# This script performs some start-up configurations.
#
# Set the modem serial port baud.
#
system serialPort modemSpeed
4800          # modem serial port baud
#
# Set the system name
#
system name
Eng_CoreBuilder_4
#
# Assign an IP address to the system.
#
ip interface define
158.101.112.99 # IP address for the system
255.255.0.0    # subnet mask
1              # VLAN interface index

ip interface summary all
#
# Validate access to management workstation
#
ip ping
158.101.112.26 # management workstation address
#
# Enable the Spanning Tree Protocol
#
bridge stpState enabled
#
# Configure my node as an SNMP trap destination
#
snmp trap add
158.101.112.26 # management workstation address
all           # turn on all traps
q             # no more trap destinations
#
snmp trap display
#
```

✓ 3500
9000
✓ 9400

✓ 3900
✓ 9300

logout Terminates a Telnet session or returns control to the password prompt in a serial port session.

Valid Minimum Abbreviation

logo

Important Consideration

- Press Escape to return to the top level before you log out.

4

MODULE ENVIRONMENT

This chapter describes how to use `module` commands for modules that are installed in the CoreBuilder® 9000 7-slot, 8-slot, and 16-slot chassis to:

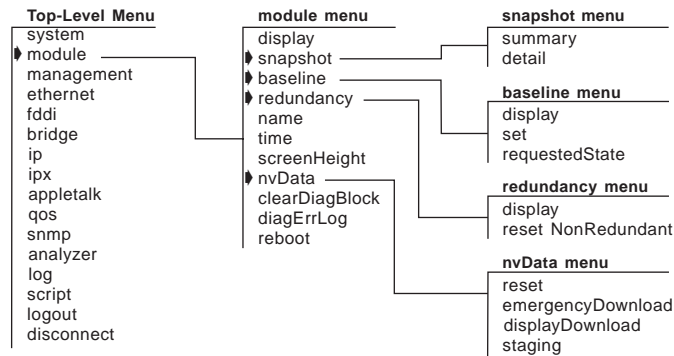
- Display the module configuration and status
- Administer a statistics baseline and module redundancy
- Set the module name and the console screen height
- View the date and time
- Manage nonvolatile data (nvData)
- Clear the module diagnostic block
- Reboot a module



For more information about administering your module parameters, see the CoreBuilder 9000 Implementation Guide.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



The redundancy option appears if you have a CoreBuilder 9000 8-slot and 16-slot chassis and if you have one or two GEN Switch Fabric Modules installed.

In addition to the `module` options, you must also use the `disconnect` option from the top-level menu to return to the Enterprise Management Engine (EME).

module display

Generates software and hardware revision numbers, module status information, and warning messages for certain module conditions.

3500
✓ 9000
9400

Valid Minimum Abbreviation

mo d

3900
9300

Important Considerations

- The module display provides the configuration information for the module to which you are currently connected. To display configuration information for another module, disconnect from the current module through the Enterprise Management Engine (EME) and then connect to the module that you want to display.
- After you have logged in to the EME, you can access the module that you want to manage.
- Diagnostic messages appear in the display only if a module fails any of the tests at start-up.

Fields in the Module Display

Field	Description
24 port Gigabit Switching Fabric	32-character alphanumeric name that uniquely identifies this module
3C number	Each module's 3Com 3C product identification number
Built	Date and time when this software version was built
Diagnostics	Whether a module has passed or failed diagnostics. Additional diagnostic messages describe a failure.
Memory size (AP, FP, Flash, Buffer)	Memory capacities of the processors
Module ID	ID number that the system assigns to that module
Rev	Unique number assigned to the hardware build by 3Com
Serial	Each module's unique serial number
Slot	Location of the module in the chassis (slot.subslot)
System name	Name of the system
System up time	Time that this module has been up and running
Time in service	Total operational time since the module was manufactured
Version	Software release number

**module snapshot
summary**

Captures an image of all the module's display screens. The values in each screen reflect the current values of all fields and counters at the time that you use the snapshot feature.

3500

✓ 9000

9400

Valid Minimum Abbreviation`mo sn su`

3900

9300

Important Consideration

- If a feature or protocol has only one display option (`display`), the module includes the same image in the snapshot of both the summary and the detail display images.

module snapshot
detail

Captures an image of all module detail display screens. The display screens contain the current values of all fields and counters at the time that you use the snapshot feature.

3500
✓ 9000
9400

Valid Minimum Abbreviation

mo sn de

3900
9300

Important Consideration

- If a feature or protocol has only one display option (`display`), the module includes that image with both the summary and detail display images.

**module baseline
display**

Displays when the current baseline was last set.

Valid Minimum Abbreviation

`mo ba dis`

Important Considerations

- Use this command to determine if you need a newer baseline for viewing statistics.
- The display indicates if you have not set the baseline on a module.

3500

✓ 9000

9400

3900

9300

module baseline set Resets the baseline counters to zero and time-stamps the baseline.

3500
✓ 9000
9400

Valid Minimum Abbreviation

mo ba set

Important Considerations

- Baselining is automatically enabled when a baseline is set.
- The module maintains the accumulated totals since power-on.
- After you disconnect from a module on which you set a baseline, the baseline is disabled. You must reconnect to that module and use the `module baseline requestedState` option to reenable the baseline.

**module baseline
requestedState**

Enables or disables a baseline.

Valid Minimum Abbreviation

`mo ba req`

Important Considerations

- When you reenable a baseline, the counters return to the values that have accumulated since the most recent baseline that you set.
- Disabling a baseline returns the counters to the total accumulated values since the last power on.
- After you disconnect from a module on which you set a baseline, the baseline is disabled. You must reconnect to that module and use the `module baseline requestedState` option to reenable the baseline.

Options

Prompt	Description	Possible Values	[Default]
Baseline	Whether you want to enable or disable the baseline	<ul style="list-style-type: none">■ enabled■ disabled	disabled

module redundancy Establishes a fault-tolerant environment for your CoreBuilder 9000 system.

3500
✓ 9000
9400

Valid Minimum Abbreviation

mo red

Important Considerations

- You must be using a CoreBuilder 9000 8-slot or 16-slot chassis.
- The Redundancy option appears on the module menu if you have one or two switch fabric modules installed. If you only have one switch fabric module installed in the chassis, the status of the second switch fabric slot is `Not Responding`.

Options

Prompt	Description	Possible Values	[Default]
Display	Module redundancy configuration and status	–	–
reset nonRedundant	Whether the module’s non-redundant indicator resets	–	–

3500
✓ 9000
9400

3900
9300

module name Assigns or changes an easily recognizable and unique module name to help you manage it.

Valid Minimum Abbreviation

mo nam

Important Considerations

- Assign an easily recognizable and unique name for each module. For example, name the module according to its physical location, such as CB9000-ENGLAB1.
- Use quotation marks (") around any string with embedded spaces.
- The new module name appears the next time that you display the configuration.

Options

Prompt	Description	Possible Values	[Default]
New name	New or changed name for the module	<ul style="list-style-type: none">■ A string of up to 32 case-sensitive characters■ ? (for information about the naming guidelines)	Current system and module name

module time Displays the module's current date and time.

3500
✓ 9000
9400

Valid Minimum Abbreviation

mo ti

Important Considerations

- You cannot change the system time from the module. You can only change the date and time from the Enterprise Management Engine (EME).
- The CoreBuilder 9000 module's internal clock is initialized when the module is shipped from the factory. You may want to reset the EME date and time to match the system's physical location.

Module Time Example

CB9000@slot10.1 [12-E/FEN-TX-L3] (module): **time**
The current module time is 05/20/98 04:37:57 PM.

module screenHeight

Changes the Administration Console's screen height to increase or decrease the space available for displaying information.

3500✓ **9000****9400****3900****9300****Valid Minimum Abbreviation****mo scr****Important Considerations**

- The setting controls the way that the module displays statistical summaries and other information that results from your use of the menus, not the way that the module displays the menus themselves.
- Each time that the screen output reaches the designated screen height, the module prompts you to press a key to display more information. Set the screen height to infinite (0) if you do not want the modules to display this prompt. At 0, however, the screen output can scroll beyond the screen, depending on your screen size.
- Most terminal screens are 24 lines.

Options

Prompt	Description	Possible Values	[Default]
New screen height	New screen height in lines	<ul style="list-style-type: none">■ 1 – 200■ 0 (for infinite height)	24
Set this value as the default?	Default screen height for future Administration Console sessions	<ul style="list-style-type: none">■ y (yes)■ n (no)	y

module nvData reset Resets the module's nonvolatile data (NV) values to the factory defaults.

3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

mo nv res

Important Considerations

- At times you may want to reset the values to the factory defaults so that you can reconfigure the module from its original settings.
- Resetting the NV data means that all NV memory is set back to the factory defaults. Before you proceed, be sure that you want to reset your NV data. Rebooting a module returns you to the Enterprise Management Engine (EME) prompt, so that you must reconnect to the module.

Prompts

Prompt	Description	Possible Values	[Default]
Do you wish to continue?	Confirmation prompt. Resetting nonvolatile data may leave the module in an inconsistent state; a reboot is necessary after each reset.	<ul style="list-style-type: none">■ n (no)■ y (yes)	y

Prompt	Description	Possible Values	[Default]
Staging setting	Whether you want to enable or disable the NV staging flag	<ul style="list-style-type: none">■ off■ on	off

**module nvData
displayDownload**

Displays emergency download information for your module.

Valid Minimum Abbreviation

mo nv dis

Important Consideration

- The download display shows the following information:
 - File Type
 - File Name
 - Server IP

3500
✓ 9000
9400

3900
9300

**module nvData
staging**

Enables either default module settings or retention of nonvolatile data settings when you hot swap a module.

3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

mo nv sta

Important Considerations

- If you hot swap a module and the staging flag is set to `on`, the new module adopts the nonvolatile data settings from the old module.
- If you hot swap a module and the staging flag is set to `off`, the new module uses the module default settings for the new module.

Options

Prompt	Description	Possible Values	[Default]
Staging setting	Whether you want to enable or disable the NV staging flag	<ul style="list-style-type: none">■ off■ on	off

module

clearDiagBlock

Prevents diagnostic information about failed modules from accumulating in module display screens.

3500

✓ 9000

9400

Valid Minimum Abbreviation

mo cle

Important Considerations

- The module immediately removes diagnostic information about failed modules from the SNMP MIB *swSysDiagnosticsGroup*.
- If you enter **y** (yes), the module immediately removes the diagnostic information about failed modules from the SNMP MIB *swSysDiagnosticsGroup*.
- If you enter **n** (no), the module displays the previous menu.

3900

9300

module diagErrLog

Displays hardware diagnostic errors that have been saved in the flash memory. When the system is initializing, if the diagnostic software detects errors, and if the system completes initializing, the detected errors are written to flash memory and stored in a dynamic error log.

3500✓ **9000****9400****3900****9300****Valid Minimum Abbreviation**`mo dia`**Important Consideration**

- The error messages are saved to flash memory until you power down the system or clear the error log with the `system clearDiagBlock` command.

module reboot Reboots the specified module.

3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`mo reboot`

Important Considerations

- Rebooting a module returns you to the Enterprise Management Engine (EME) prompt, so that you must reconnect to the module.
- If you enter `y`, the module reboots.
- If you enter `n`, the previous menu appears on the screen.

disconnect Disconnects you from the Administration Console and returns you to the Enterprise Management Engine (EME) module.

3500

✓ 9000

9400

Valid Minimum Abbreviation

`disc`

3900

9300

Important Consideration


- Disconnecting from the Administration Console does not disconnect you from a Telnet session.



ESTABLISHING MANAGEMENT ACCESS

Chapter 5 Out-of-Band Management

Chapter 6 Simple Network Management Protocol (SNMP)



5

OUT-OF-BAND MANAGEMENT

The Internet Protocol (IP) is a standard networking protocol that is used for communications among various networking devices. To gain access to the system using the Transmission Control Protocol/Internet Protocol (TCP/IP) or to manage the system using the Simple Network Management Protocol (SNMP), you must set up an IP interface to manage your system, either in-band (with your regular network traffic) or out-of-band (with a dedicated network).

- **In-Band Management** — Set up an IP routing interface and at least one virtual local area network (VLAN). See Chapter 14 for information about how to define a VLAN.
- **Out-of-Band Management** — Assign an IP address and subnet mask for the out-of-band Ethernet port on your system through the `management` menu. This chapter focuses on out-of-band management. The out-of-band Ethernet port is the 10BASE-T port on the system processor module. It is not associated with a port number. (See Chapter 16 for background information about IP addresses and subnet masks.)

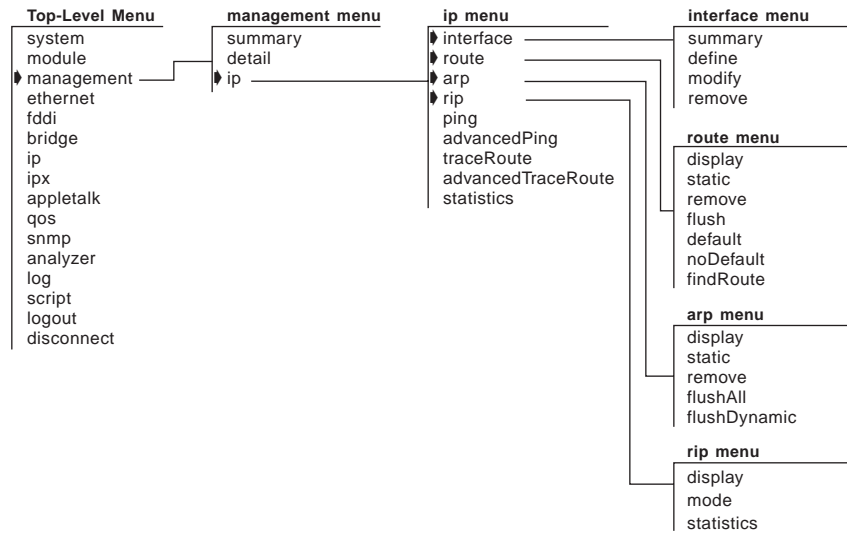
This chapter provides guidelines and other key information about how to set up an out-of-band management interface for your system.



The CoreBuilder® 9000 and SuperStack® II Switch 3900 use in-band management only. For more information about management interfaces, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



management summary

Displays Ethernet summary information about the out-of-band system management port.

- ✓ 3500
- 9000
- ✓ 9400
- 3900
- ✓ 9300

Valid Minimum Abbreviation

m sum

Important Considerations

- The management summary and management detail displays contain the same fields as the Ethernet summary and Ethernet detail displays.
- Fields that do not apply to the management port contain n/a in the management summary and management detail displays.

Fields in the Management Summary Display

Field	Description
actualFlowControl	Actual flow control setting (for Gigabit Ethernet ports). When autonegotiation is completed, the value is the autonegotiated setting. When autonegotiation is disabled, the value is the user-selected flow control value.
actualPortMode	Actual operating port mode. When autonegotiation is completed, the value shown is the autonegotiated setting. When autonegotiation is disabled, the value is the user-selected port mode.
autoNegMode	Autonegotiation mode configured for port. Possible values are enabled or disabled.
autoNegState	Current negotiation state. Possible values are disabled, configuring, completed, and failed.
linkStatus	Boolean value indicating the current state of the physical link status for this port (either enabled or disabled)
macAddress	MAC address of this port
noRxBuffers	Number of frames that were discarded because no buffer space was available
portLabel	User-defined label name. The maximum length of the string is 32 characters, including the null terminator.
portState	Current software operational state of this port. Possible values are on-line and off-line.
portType	Specific description of this port's type. Value for port type: 10/100BASE-TX (RJ-45)
reqFlowControl	If autonegotiation is disabled, a configurable parameter that sets the flow control option on the ports. If autonegotiation is enabled, flow control values are ignored.

Field	Description
reqPortMode	If autonegotiation is disabled, a configurable parameter that sets the port mode on Ethernet ports that have port mode options. If autonegotiation is enabled, port mode values are ignored.
rxBytes	Number of bytes received by this port, including framing characters
rxErrs	Sum of all receive errors that are associated with this port (summary report only)
rxFrames	Number of frames that were copied into receive buffers by this port
txBytes	Number of bytes that were transmitted by this port, including framing characters
txErrs	Sum of all transmit errors that are associated with this port (summary report only)
txFrames	Number of frames that were transmitted by this port
txQOverflows	Number of frames that were lost because transmit queue was full

management detail

Displays Ethernet detailed information about the out-of-band system management port.

✓ 3500
9000
✓ 9400

Valid Minimum Abbreviation

m det

Important Considerations

- The management summary and management detail displays contain the same fields as the Ethernet summary and Ethernet detail displays.
- Fields that do not apply to the management port contain n/a in the management summary and management detail displays.

3900
✓ 9300

Fields in the Management Detail Display

Field	Description
actualFlowControl	Actual flow control setting for the port. When autonegotiation is completed, the value is the autonegotiated setting. When autonegotiation is disabled, the value is the user-selected flow control value.
actualPortMode	Actual operating port mode. When autonegotiation is completed, the value shown is the autonegotiated setting. When autonegotiation is disabled, the value is the user-selected port mode.
alignmentErrs	Number of frames received by this port that are not an integral number of octets in length and do not pass the FCS check
autoNegMode	Autonegotiation mode configured for port. Possible values are enabled or disabled.
autoNegState	Current negotiation state. Possible values are disabled, configuring, completed, and failed.
carrierSenseErr	Number of frames discarded because the carrier sense condition was lost while attempting to transmit a frame from this port
collisions	Number of collisions detected on this port
excessCollision	Number of frames that could not be transmitted on this port because the maximum allowed number of collisions was exceeded
excessDeferrals	Number of frames that could not be transmitted on this port because the maximum allowed deferral time was exceeded
fcsErrs	Number of frames received by this port that are an integral number of octets in length but do not pass the FCS check

Field	Description
lateCollisions	Number of times that a collision was detected on this port later than 512 bit-times into the transmission of a frame
lengthErrs	Number of frames received by this port that are longer than 1518 bytes or shorter than 64 bytes
linkStatus	Boolean value indicating the current state of the physical link status for this port (either <code>enabled</code> or <code>disabled</code>)
macAddress	MAC address of this port
multiCollisions	Number of times that multiple collisions were detected on this port.
noRxBuffers	Number of frames that were discarded because no buffer space was available
portLabel	User-defined label name. The maximum length of the string is 32 characters, including the null terminator.
portState	Current software operational state of this port. Possible values are <code>on-line</code> and <code>off-line</code> .
portType	Specific description of this port's type. Values for each port type: <code>10/100BASE-TX (RJ-45)</code> , <code>100BASE-FX (SC)</code> , <code>1000BASE-SX (SC)</code> , <code>1000BASE-LX (SC)</code> .
reqFlowControl	If autonegotiation is disabled, a configurable parameter that sets the flow control option on the port. If autonegotiation is enabled, flow control values are ignored.
reqPortMode	If autonegotiation is disabled, a configurable parameter that sets the port mode on Ethernet ports that have port mode options. If autonegotiation is enabled, port mode values are ignored.
requestedState	Configurable parameter that enables or disables this port. The default is <code>enabled</code> .
rxByteRate	Average number of bytes received per second by this port during the most recent sampling period
rxBytes	Number of bytes received by this port, including framing characters
rxDiscards	Number of received frames that were discarded because there was no higher layer to receive them or because the port was disabled
rxErrs	Sum of all receive errors that are associated with this port (summary report only)
rxFrameRate	Average number of frames that were received per second by this port during the most recent sampling period. Sampling periods are 1 second long and not configurable.
rxFrames	Number of frames that were copied into receive buffers by this port

Field	Description
rxInternalErrs	Number of frames that were discarded because of an internal error during reception
rxMulticasts	Number of multicast frames that were delivered to a higher-level protocol or application by this port
rxPeakByteRate	Peak value of ethernetPortByteReceiveRate for this port since the station was last initialized
rxPeakFrameRate	Peak value of ethernetPortFrameReceiveRate for this port since the station was last initialized
rxUnicasts	Number of unicast frames that were delivered by this port to a higher-level protocol or application
txByteRate	Average number of bytes transmitted per second by the port during the most recent sampling period
txBytes	Number of bytes that were transmitted by this port, including framing characters
txDiscards	Number of transmitted frames that were discarded because the port was disabled
txErrs	Sum of all transmit errors that are associated with this port (summary report only)
txFrameRate	Average number of frames transmitted per second by this port during the most recent sampling period. Sampling periods are 1 second long (not configurable).
txFrames	Number of frames that were transmitted by this port
txInternalErrs	Number of frames that were discarded because of an internal error during transmission
txMulticasts	Number of multicast frames that are queued for transmission by a higher-level protocol or application, including those not transmitted successfully
txPeakByteRate	Peak value of ethernetPortByteTransmitRate for this port since the station was last initialized
txPeakFrameRate	Peak value of ethernetPortFrameTransmitRate for this port since the station was last initialized
txQOverflows	Number of frames that were lost because transmit queue was full
txUnicasts	Number of unicast (nonmulticast) frames that are queued for transmission by a higher-level protocol or application, including frames not transmitted successfully

**management ip
interface summary**

Displays a summary table about the out-of-band system IP management interface configuration, including parameter settings.

✓ 3500
9000
9400

Valid Minimum Abbreviation

m ip i sum

3900
9300

Fields in the Management IP Interface Summary Display

Field	Description
Index	Unique number that identifies the out-of-band interface
IP address	IP address of the out-of-band interface, chosen from the range of addresses that the central agency assigned to your organization. This address is specific to your network and system.
IP routing status	Whether the interface is available to route IP traffic (enabled) or not (disabled)
RIP status	Whether RIP is dynamically configuring its routing tables (active) or on request (passive)
State	State of the IP interface, indicating whether the interface is available for communications (up) or unavailable (down).
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnet number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnet part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.
Type	Type of device that is connected to the interface

management ip
interface define

Defines the IP address of the IP management out-of-band port.

Valid Minimum Abbreviation

m ip i d

Options

✓ 3500
9000
9400

3900
9300

Prompt	Description	Possible Values	[Default]
IP address	IP address of the out-of-band interface, chosen from the range of addresses that the central agency assigned to your organization. This address is specific to your network and system.	Any valid IP address	–
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnet number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnet part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.	Any subnet mask valid for use with the current IP address	255.255.0.0, or the subnet mask value currently stored in the system

**management ip
interface modify**

Changes the configuration of an IP management interface that you have already defined.

✓ 3500
9000
9400

Valid Minimum Abbreviation

m ip i m

Important Consideration

- Use the `management ip statistics` command to periodically monitor IP activity for your system. The statistics can help determine whether you need to change the IP management interface using the `management ip interface modify` command.

3900
9300

**management ip
interface remove**

Removes an IP management interface if you no longer need it.

Valid Minimum Abbreviation

m ip i r

Important Consideration

- Use the `management ip statistics` command to periodically monitor IP activity for your system. The statistics can help determine whether you need to remove the IP management interface using the `management ip interface remove` command.

✓ 3500
9000
9400

3900
9300

management ip route display

Displays the system’s routing table to determine which routes to other IP networks are configured and whether the routes are operational.

✓ 3500
9000
9400

3900
9300

Valid Minimum Abbreviation

m ip ro di

Important Considerations

- The system prompts you for an IP address and subnet mask. This information enables you to display only a subset of routes instead of all routes. To see all entries in the table, press Return at the prompts.
- The first line in the output (the status line) indicates whether IP routing is enabled (in-band only):
 - IP interface options (such as ICMP router discovery) appear under ip interface detail and are set on a per-interface basis.
- The route table display includes a range for the routing table entries as follows:

There are *n* of *m* possible Routing Table entries.
Where *n* is the minimum and *m* is the maximum number of entries.

Options

Prompt	Description	Possible Values	[Default]
IP address	IP address that directs the system to display only those routes that match the bits set in the specified IP address (and its corresponding subnet mask) Press Enter to take the default, which displays all entries	A valid IP address	0.0.0.0
Subnet mask	Subnet mask that directs the system to display only those routes that match the bits set in the subnet mask for the specified IP address	A valid subnet mask	Based on specified IP address

Fields in the Management IP Route Display

Field	Description
Destination	IP address of the destination network, subnetwork, or host. This field can also identify a default route, which the system uses to forward packets that do not match any other routing table entry. You may want to use the default route in place of routes to numerous destinations that all have the same gateway IP address.
Gateway	Address that directs the router how to forward packets whose destination addresses match the route's IP address and subnet mask. The system forwards such packets to the indicated gateway.
Metric	Number of networks through which a packet must pass to reach a given destination. The system includes the metric in its RIP updates to allow other routers to compare routing information that is received from different sources.
Status	Status of the route. See the following route status table.
Subnet mask	Subnet mask that is associated with the IP address of the destination network, subnet, or host.

Status for Routes

Value	Description
Direct	Route is for a directly connected network
Learned	Route was learned using indicated protocol
Learned RIP	Route was learned using RIP-1 protocol
Learned RIP-Zombie	Route was learned but is partially timed out
Learned RIP2	Route was learned using RIP-2 protocol
Local	Actual interface address
Static	Route was statically configured
Timed out	Route has timed out and is no longer valid

✓ 3500
9000
9400

3900
9300

Valid Minimum Abbreviation

m i p r o s

Important Considerations

- Before you can define static routes, you must define at least one IP interface. See “ip interface define (3500/9000 Layer 3)” in Chapter 16 for more information.
- You can define up to 128 static routes.
- Static routes remain in the table until you remove them or the corresponding interface.
- Static routes take precedence over dynamically learned routes to the same destination.
- Static routes are not included in periodic Routing Information Protocol (RIP) updates sent by the system.

Options

Prompt	Description	Possible Values	[Default]
Destination IP address	IP address of the destination network, subnet, or host for this route	A valid IP address	–
Subnet mask	Subnet mask for the specified IP address	A valid subnet mask	Based on specified IP address
Gateway IP address	IP address of the gateway used by this route	A valid router address	–

management ip route
remove

✓ 3500
9000
9400

3900
9300

Deletes an existing route.

Valid Minimum Abbreviation

m ip ro r

Important Consideration

- When you enter the command, the system deletes the route immediately from the routing table. You are not prompted to confirm the deletion.

Options

Prompt	Description	Possible Values	[Default]
Destination IP address	IP address associated with the route that you want to delete	A valid IP address	–
Subnet mask	Subnet mask for the specified IP address	A valid subnet mask	Based on specified IP address

**management ip route
flush**

Deletes all learned routes from the routing table.

Valid Minimum Abbreviation

m ip ro fl

Important Considerations

- The system deletes all learned routes from the routing table immediately. You are not prompted to confirm the deletion.
- Flushing the routing table causes Routing Information Protocol (RIP) to regenerate the routing table. The system repopulates the routing table a few seconds after you flush it.

✓ 3500
9000
9400

3900
9300

management ip route
default

✓ 3500
9000
9400

3900
9300

Adds a default route to the routing table immediately.

Valid Minimum Abbreviation

m ip ro de

Important Considerations

- If you define a default route, the system uses it to forward packets that do not match any other routing table entry. The system can learn a route using the Routing Information Protocol (RIP), or you can statically configure a default route.
- If the routing table does not contain a default route, the system cannot forward a packet that does not match any other routing table entry. When the system drops the packet, it sends an ICMP destination unreachable message to the host that sent the packet.

Options

Prompt	Description	Possible Values	[Default]
Gateway IP address	IP address that is associated with the default route that you want to add (for example, 158.101.112.253)	A valid IP address	–

management ip route
noDefault

Deletes the default route.

Valid Minimum Abbreviation

m ip ro n

Important Consideration

- The system deletes the default route from the routing table immediately after you enter the command. You are not prompted to confirm this deletion.

✓ 3500
9000
9400

3900
9300

management ip route
findRoute

✓ 3500
9000
9400

3900
9300

Searches for a route in the routing table.

Valid Minimum Abbreviation

m ip route fi

Important Considerations

- This command enables you to find a route using an IP address or a host name, as long as Domain Name System (DNS) is configured.
- When you enter this command with a valid IP address or host name, the system displays the routing table entry.

Options

Prompt	Description	Possible Values	[Default]
IP address (or host name)	IP address that is associated with the route you that want to find, or a host name, if DNS is configured	A valid IP address or host name	0.0.0.0

management ip arp display

Display the contents of the Address Resolution Protocol (ARP) cache for each interface on the system.

✓ 3500
9000
9400

Valid Minimum Abbreviation

m ip ar d

Important Considerations

- The system uses the ARP cache to find the MAC addresses that correspond to the IP addresses of hosts and other routers on the same subnets. Each device that participates in routing maintains an *ARP cache*, which is a table of known IP addresses and their corresponding MAC addresses.
- The first line in the output (the status line) indicates whether IP routing is enabled (in-band only):
 - IP interface options (such as ICMP router discovery) appear under `ip interface detail` and are set on a per-interface basis. The second status line indicates the number of entries in the ARP cache.

Fields in the Management IP ARP Display

Field	Description
Hardware address	MAC address that is mapped to the IP address
IP address	IP address of the interface
Type	Type of entry, <code>static</code> or <code>dynamic</code>

management ip arp static

Defines a static ARP cache entry on the system.

Valid Minimum Abbreviation

m ip ar s

Important Consideration

- You can define up to 128 static ARP entries.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the interface for which you want to define a static ARP entry	<ul style="list-style-type: none">■ A valid index number■ ? (for a list of selectable indexes)	–
IP address	IP address to use in the entry	A valid IP address	–
MAC address	Hardware address to use in the entry (in the format xx-xx-xx-xx-xx-xx)	A valid MAC address	–

Management IP ARP Static Example

Select interface index {1-2|?} 2
Enter IP address: 158.101.12.12
Enter MAC address: 00-00-00-00-00-01

✓ 3500
9000
9400

3900
9300

**management ip arp
remove**

Deletes an entry from the ARP cache (for example, if the MAC address has changed).

✓ 3500
9000
9400

3900
9300

Valid Minimum Abbreviation

m ip ar rem

Important Considerations

- When you enter the command, the system deletes the entry immediately from the cache. You are not prompted to confirm the deletion.
- If necessary, the system subsequently uses ARP to find the new MAC address that corresponds to that IP address.

Options

Prompt	Description	Possible Values	[Default]
IP address	IP address that is associated with the entry that you want to delete	A valid IP address	–

**management ip arp
flushAll**

Deletes all entries from the ARP cache.

Valid Minimum Abbreviation

m ip ar flushA

Important Considerations

- This command applies to the CoreBuilder 3500 only; other platforms use `ip arp flush`. To flush dynamic entries only, see the “management ip arp flushDynamic” command next.
- When you enter the command, the system deletes all entries immediately from the cache. You are not prompted to confirm the deletions.

✓ 3500
9000
9400

3900
9300

**management ip arp
flushDynamic**

Deletes all dynamic (automatically learned) entries from the ARP cache.

Valid Minimum Abbreviation

m ip ar flushD

Important Considerations

- This command applies to the CoreBuilder 3500 only; other platforms use `ip arp flush`. To flush all entries, static and dynamic, see the previous “management ip arp flushAll” command.
- When you enter the command, the system deletes all dynamic entries immediately from the cache. You are not prompted to confirm the deletions.

✓ 3500
9000
9400

3900
9300

management ip rip display

✓ 3500
9000
9400

3900
9300

Displays information about the Routing Information Protocol (RIP) interfaces on the system. RIP is one of the IP Interior Gateway Protocols (IGPs). When enabled, RIP allows the system to dynamically configure its routing tables.

Valid Minimum Abbreviation

m ip ri d

Important Considerations

- The first line in the output (the status line) indicates whether IP routing is enabled (in-band only):
 - IP interface options (such as ICMP router discovery) appear under `ip interface detail` and are set on a per-interface basis. The rest of the output contains more RIP interface information.
- The two available RIP modes are as follows:
 - **Disabled** — The system ignores all incoming RIP packets and does not generate any RIP packets of its own.
 - **Learn** — The system processes all incoming RIP packets, but it does not transmit RIP updates.

Fields in the Management IP RIP Display

Field	Description
Index	Index number of the interface
RIP-1 mode	Mode for RIP-1. If you disable RIP-1, the output lists the state as <code>off</code> . Other modes are <code>learn</code> (default), <code>advertise</code> , and <code>enabled</code> .
RIP-2 mode	Mode for RIP-2. If you disable RIP-2, the output lists the state as <code>off</code> . Other modes are <code>learn</code> (default), <code>advertise</code> , and <code>enabled</code> .

management ip rip mode

On a per-interface basis, sets one of four RIP Version 1 (RIP-1) modes and one of four RIP Version 2 (RIP-2) modes on the system.

✓ 3500
9000
9400

3900
9300

Valid Minimum Abbreviation

m ip ri m

Important Considerations

- The CoreBuilder 3500 supports RIP Version 1 as well as RIP Version 2. For each interface, you select a RIP Version 1 mode and a RIP Version 2 mode. The default RIP Version 1 mode for all platforms is `learn`. The default RIP Version 2 mode for the CoreBuilder 3500 is `disabled`.
- The four available RIP modes are as follows:
 - **Disabled** — The interface ignores all incoming RIP packets and does not generate any RIP packets of its own.
 - **Learn** — The interface processes all incoming RIP packets, but it does not transmit RIP updates. This is the default RIP mode.
 - **Advertise** — The interface broadcasts RIP updates, but it does not process incoming RIP packets.
 - **Enabled** — The interface broadcasts RIP updates and processes incoming RIP packets.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interface for which you want to set the RIP mode	<ul style="list-style-type: none"> ■ Selected interfaces ■ all ■ ? (for a list of selectable indexes) 	Previous entry, if applicable
RIP mode, Version 1	Selected RIP Version 1 mode that determines how the interface handles RIP 1 packets and updates	<ul style="list-style-type: none"> ■ disabled ■ learn ■ advertise ■ enabled 	learn, or current value
RIP mode, Version 2	How the interface handles RIP 2 packets and updates	<ul style="list-style-type: none"> ■ disabled ■ learn ■ advertise ■ enabled 	disabled, or current value

Management IP RIP Mode Example

```
Select menu option (management/ip/rip): mode
Select IP interfaces (1|all|?) [1]: 1
Interface 1 - Enter RIP Version 1 mode (disabled,learn) [learn]: disabled
Interface 1 - Enter RIP Version 2 mode (disabled,learn) [learn]: disabled
```

Displays general RIP statistics.

✓ 3500
9000
9400

m ip rip s

Fields in the Management IP RIP Statistics Display

Field	Description
queries	Number of queries
routeChanges	Number of route changes

management ip ping

Tries to reach or “ping” a specified destination using the default ping options.

✓ 3500
9000
9400

Valid Minimum Abbreviation

m ip p

3900
9300

Important Considerations

- This tool is useful for network testing, performance measurement, and management. It uses the Internet Control Message Protocol (ICMP) echo facility to send ICMP echo request packets to the IP destination that you specify.
- If you need to change the default ping options, use `management ip advancedPing`.
- You can either supply the host name or IP address as part of the command string, or you can supply the information at the prompt.
- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “ip dns define” in Chapter 16 for more information.
- When the system sends an echo request packet to an IP station using ping, the system waits for an ICMP echo reply packet. Possible responses:
 - If the host is reachable, the system displays information about the ICMP reply packets and the response time to the ping.
 - If the host does not respond, the system displays the ICMP packet information and this message: `Host is Not Responding`. You may not have configured your gateway IP address.
 - If the packets cannot reach the host, the system displays the ICMP packet information and this message: `Host is Unreachable`. A host is unreachable when there is no route to that host.
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination that you want to ping	A valid host name or IP address	0.0.0.0

Management IP Ping Example

```
Select menu option (ip): ping
Enter host name/IP address [0.0.0.0]: 158.101.111.50
Press "Enter" key to interrupt.

PING 158.101.111.50: 64 byte packets
64 bytes from 158.101.111.50: icmp_seq=0.  time=16. ms
64 bytes from 158.101.111.50: icmp_seq=1.  time=19. ms
64 bytes from 158.101.111.50: icmp_seq=2.  time=24. ms

---- 158.101.111.50 PING Statistics ----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 16/20/24
```

management ip
advancedPing

Tries to reach or “ping” a host with one or more of the advanced ping options.

✓ 3500
9000
9400

Valid Minimum Abbreviation

m ip advancedP

Important Considerations

- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See the `ip dns` commands in Chapter 16 for more information.
- The Burst option, when enabled, overrides the value that is set in the Quiet or Wait option.
- The Burst option floods the network with Internet Control Message Protocol (ICMP) echo packets and can cause network congestion. Do *not* use the Burst option during periods of heavy network traffic. Use this option only as a diagnostic tool in a network that has many routers to determine if one of the routers is not forwarding packets. For example, you can set a high count value (1000 packets), and then observe the run lights on the units: the run lights blink rapidly on routers that are forwarding packets successfully, but remain unlit, or blink slowly, on routers that are not forwarding packets successfully.
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination that you want to ping.	A valid host name or IP address	0.0.0.0
Number of ICMP Request packets	Number of ICMP echo request packets that are sent to ping a host. If the destination host does not respond after it is pinged by the number of packets that you specify, the system displays a <code>Host is Unreachable</code> or <code>Host is not Responding</code> message.	1 – 9999 packets	3

Prompt	Description	Possible Values	[Default]
Packet size	Number of bytes in each ICMP echo request packet. The packet size includes both the IP and the ICMP headers.	28 – 4096 bytes	64
Burst Transmit Ping mode	When <code>enabled</code> , sends out the ICMP echo request packets as rapidly as possible. The system displays a period (.) upon receiving an ICMP echo replay packet. Use this display to determine how many packets are being dropped during the burst. This is unique to the Burst option.	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled
Quiet mode	How much packet information the system displays after a ping. When <code>enabled</code> , the system displays information about the number of packets that the system sent and received, any loss of packets, and the average time that it took a packet to travel to and from the host. When <code>disabled</code> , the system displays more detailed status information about each ICMP echo request packet.	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled
Time between sending each packet (wait)	Number of seconds that the system waits before it sends out successive ICMP echo request packets. Set this option to a high value if network traffic is heavy and you do not want to add to the network traffic with pings in fast succession.	1 – 20 seconds	1
ICMP sourceAddress	Forces the source address of the ICMP packets to be something other than the IP address of the interface from which the packet originated. You can use this option if you have more than one IP interface defined.	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Interface index	Index number of the ICMP source IP address that you want to use.	Currently defined interfaces and their indexes	0 (the router picks the best interface)

Management IP Advanced Ping Example

```
Select menu option (ip): advancedPing
Enter host IP address [0.0.0.0]: 158.101.112.56
Enter number of ICMP request packets (1-9999) [3]:
Enter packet size (bytes) (28-4096) [64]:
Enter Burst Transmit Ping mode (disabled,enabled) [disabled]:
Enter Quiet mode (disabled,enabled) [disabled]:
Enter time (sec) waits between sending each packet (1-20) [1]: 2
Configure ICMP sourceAddress? (n,y) [y]:
      Index      Interface address
        0      Best interface (default)
        1      158.101.117.151
        2      158.101.10.1
Select interface index {0-2|?} [0]: 1
Press "Enter" key to interrupt.

PING 158.101.112.56 from 158.101.117.151: 64 byte packets
64 bytes from 158.101.112.56:  icmp_seq=0.  time=26. ms
64 bytes from 158.101.112.56:  icmp_seq=1.  time=18. ms
64 bytes from 158.101.112.56:  icmp_seq=2.  time=18. ms

---- 158.101.112.56 PING Statistics ----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 18/21/26
```

management ip traceRoute

Traces a route to a destination using the default traceRoute options.

Valid Minimum Abbreviation

m ip t

Important Considerations

- TraceRoute information includes all of the nodes in the network through which a packet passes to get from its origin to its destination. It uses the IP time-to-live (TTL) field in User Datagram Protocol (UDP) probe packets to elicit an Internet Control Message Protocol (ICMP) Time Exceeded message from each gateway to a host.
- To change the default traceRoute options, use the `management ip advancedTraceRoute` command.
- You can either supply the host name or IP address as part of the command string, or you can supply the information at the prompt.
- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See Chapter 16 for more information about `ip dns` commands.
- To track the route of an IP packet, traceRoute launches UDP probe packets with a small TTL value and then listens for an ICMP Time Exceeded reply from a gateway. Probes start with a small TTL of 1 and increase the value by 1 until one of the following events occurs:
 - The system receives a Port Unreachable message, indicating that the packet reached the host.
 - The probe exceeds the maximum number of hops (default 30).
- At each TTL setting, the system launches three UDP probe packets, and the traceRoute display shows a line with the TTL value, the address of the gateway, and the round-trip time of each probe. If a probe answers from different gateways, the traceRoute feature prints the address of each responding system. If no response occurs in the 3-second timeout interval, traceRoute displays an asterisk (*) for that probe.

✓ 3500
9000
9400

3900
9300

Other characters that can be displayed include the following:

- !N — Network is unreachable
- !H — Host is unreachable
- !P — Protocol is unreachable
- !F — Fragmentation is needed
- !<n> — Unknown packet type
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination to which you want to trace a route	A valid host name or IP address	0.0.0.0

Management IP Trace Route Example

```
Select menu option (ip): traceRoute
Enter host name/IP address [0.0.0.0]: 158.101.101.40
Press "Enter" key to interrupt.

Traceroute to 158.101.101.40: 30 hops max, 28 bytes packet

1 158.101.117.254  9 ms 22 ms 5 ms
2 158.101.112.254  8 ms 22 ms 8 ms
3 158.101.96.22   7 ms 22 ms 7 ms
4 158.101.101.40  7 ms 23 ms 6 ms
```

management ip advancedTraceRoute

Traces a route to a host with one or more of the advanced traceRoute options.

✓ 3500
9000
9400

3900
9300

Valid Minimum Abbreviation

m ip advancedT

Important Considerations

- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “ip dns define” in Chapter 16 for more information.
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination that you want to ping	<ul style="list-style-type: none"> ■ A valid host name ■ IP address 	0.0.0.0
Maximum ttl	Maximum number of hops that the system can use in outgoing probe packets	1 – 255 hops	30
Destination port	Destination (or base) UDP port number that the system uses in probe packets. Set the destination UDP port number to be very high to ensure that an application at the destination is not using that port.	30000 – 65535	33434
probeCount	Maximum number of probes that the system sends at each TTL level	1 – 10	3
Wait	Wait interval (in seconds) that determines the maximum amount of time that the system waits for a response to a probe	1 – 10 seconds	3
packetSize	Number of bytes that the system sends in each UDP probe packet	28 – 4096 bytes	28

Prompt	Description	Possible Values	[Default]
sourceAddress	Source address other than the one from which the probe packets originate. This option is available if you have more than one IP interface defined on the system.	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Interface index	Index number of the ICMP source IP address that you want to use. The system lists defined interfaces and their indexes.	A selectable interface	0 (the router picks the best interface)
Numeric mode	Whether the system shows hop addresses numerically or symbolically.	<ul style="list-style-type: none"> ■ disabled ■ enabled 	default

Management IP Advanced Trace Route Example (TTL value of 10):

```

Select menu option (ip): advancedTraceRoute
Enter host IP address [158.101.101.27]:
Enter maximum Time-to-Live (ttl) (1-255) [30]: 10
Enter Destination Port number (30000-65535) [33434]:
Enter the number of probes to be sent at each ttl level (1-10) [3]:
Enter time (sec) to wait for a response (1-10) [3]:
Enter the packet size (bytes) (28-4096) [28]:
Configure TRACEROUTE sourceAddress? (n,y) [y]:
      Index      Interface address
      0          Best interface (default)
      1          158.101.117.151
      2          158.101.10.1
Select interface index {0-2|?} [0]:
Enter Numeric mode (disabled,enabled) [disabled]:
Press "Enter" key to interrupt.

```

Traceroute to 158.101.101.27: 10 hops max, 28 bytes packet

```

 1  158.101.117.254  12 ms   7 ms   5 ms
 2  158.101.112.254  51 ms   9 ms   7 ms
 3  158.101.96.22   21 ms  15 ms   6 ms
 4  158.101.101.27  18 ms  90 ms  80 ms

```

management ip statistics

Displays different types of IP statistics: general statistics and those specific to the User Datagram Protocol (UDP) or the Internet Control Message Protocol (ICMP).

✓ 3500
9000
9400

Valid Minimum Abbreviation

m ip sta

Options

Prompt	Description	Possible Values	[Default]
Statistics	Type of IP statistics that you want to display	<ul style="list-style-type: none"> ■ ip ■ udp ■ icmp ■ all 	ip

Fields in the Management IP Statistics Display

Field	Description
forwDatagrams	Number of datagrams that the IP station tried to forward
fragCreates	Number of IP datagram fragments that were generated as a result of fragmentation on this system
fragFails	Number of IP datagrams that were discarded because they needed to be fragmented but could not be (for example, because their Don't Fragment bit was set)
fragOks	Number of IP datagrams that were successfully fragmented
inAddrErrors	Number of datagrams that the IP station discarded because of an error in the source or destination IP address
inDelivers	Number of datagrams that the IP station delivered to local IP client protocols
inDiscards	Number of packet receive discards
inHdrErrors	Number of datagrams that the IP station discarded because the IP header contained errors
inReceived	Total number of IP datagrams that were received, including those with errors
osReceives	Number of packets received that are destined to higher-level protocols such as Telnet, DNS, TFTP, and FTP
osTransmits	Number of packets that were sent through the router by higher-level protocols such as Telnet, DNS, TFTP, and FTP
outDiscards	Number of packet transmit discards

Field	Description
outNoRoutes	Number of datagrams that the IP station discarded because there was no route to the destination
outRequests	Number of datagrams that local IP client protocols passed to IP for transmission
reasmFails	Number of packet reassembly failures
reasmReqs	Number of packet reassembly requests
reasmOks	Number of successful packet reassemblies
rtDiscards	Number of packets that were discarded due to system resource errors
unkProtos	Number of packets whose protocol is unknown

Fields in the Management UDP Statistics Display

Field	Description
inDatagrams	Number of UDP packets that were received and addressed to the router or broadcast address
inErrors	Number of received UDP or ICMP packets that contain header errors
noPorts	Number of UDP packets that were received but addressed to an unsupported UDP port
outDatagrams	Number of UDP packets that were sent by the router

Fields in the Management ICMP Statistics Display

Field	Description
inAddrMaskReps	Number of ICMP address mask reply frames that were received
inAddrMasks	Number of ICMP address mask request packets that were received
inDestUnreach	Number of ICMP destination unreachable packets that were received
inEchoReps	Number of ICMP echo reply packets that were received
inEchos	Number of ICMP echo request packets that were received
inParmProbs	Number of ICMP parameter problem frames that were received
inRedirects	Number of ICMP redirect packets that were received
inSrcQuenchs	Number of ICMP source quench packets that were received
inTimeExcds	Number of ICMP time exceeded packets that were received

Field	Description
inTimeStamps	Number of ICMP time stamp request packets that were received
inTimeStampsReps	Number of ICMP time stamp reply packets that were received
messages	Number of ICMP packets that were received
outAddrMaskReps	Number of ICMP address mask reply packets that were sent
outAddrMasks	Number of ICMP address mask request packets that were sent
outDatagrams	Number of UDP packets that the router sent
outDestUnreach	Number of ICMP destination unreachable packets that were sent
outEchoReps	Number of ICMP echo reply packets that were sent
outEchos	Number of ICMP echo request packets that were sent
outErrors	Number of ICMP packets sent that were dropped due to system resource errors
outMsgs	Number of ICMP packets that were sent
outParmProbs	Number of ICMP parameter problem packets that were sent
outRedirects	Number of ICMP redirect packets that were sent
outSrcQuenchs	Number of ICMP source quench packets that were sent
outTimeExcds	Number of ICMP time exceeded packets that were sent
outTimeStampReps	Number of ICMP time stamp reply packets that were sent
outTimeStamps	Number of ICMP time stamp request packets that were sent

6

SIMPLE NETWORK MANAGEMENT PROTOCOL (SNMP)

You can manage the system using a Simple Network Management Protocol (SNMP)-based external management application (called the SNMP manager) that sends requests to the system.

The SNMP agent provides access to the collection of information about the system, called Management Information Bases (MIBs). Your views of MIB information differ depending on the SNMP management method that you choose. In addition, you can configure an SNMP agent to send traps to an SNMP manager to report significant events. Access to system information through SNMP is controlled by community strings. You can use either an in-band or an out-of-band IP interface to manage the system with SNMP.

This chapter provides guidelines and other key information about how to set up SNMP in your system.

To configure SNMP for system management with SNMP:

- 1 Assign an IP address to either the system processor out-of-band Ethernet port or an in-band Ethernet port.
- 2 Set the destination IP address to which the traps are forwarded by the system agent.



For more information about setting up SNMP, see the Implementation Guide for your system.



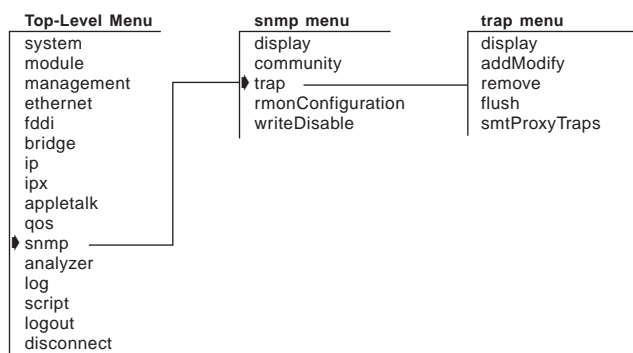
To set community strings, `snmp authentication_trap`, and `snmp extensions` on a CoreBuilder® 9000 Enterprise Switch, see the CoreBuilder 9000 Enterprise Management Engine User Guide.



You can access the Remote Monitoring (RMON) capabilities of the CoreBuilder 3500 through SNMP applications such as Transcend® Network Control Services software.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



snmp display Displays the current SNMP configurations for the community strings.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

sn d

Fields in the SNMP Display

- ✓ 3900
- ✓ 9300

Field	Description
community string	Community strings setting that controls access to the system: <ul style="list-style-type: none">■ Read-only community strings with the default public■ Read-write community strings with the default private

snmp community

Sets two SNMP community strings: read-only and read-write.

✓ 3500
9000
✓ 9400



To set the community strings for the CoreBuilder 9000, see the CoreBuilder 9000 Enterprise Management Engine User Guide.

✓ 3900
✓ 9300

Valid Minimum Abbreviation

sn c

Important Considerations

- When an SNMP agent receives an SNMP request, the agent compares the community string in the request with the community strings that are configured for the agent:
 - SNMP *get*, *get-next*, and *set* requests are valid if the community string in the request matches the agent's read-write community.
 - SNMP *get* and *get-next* requests are valid if the community string in the request matches the agent's read-only community string or read-write community string.
- You can specify any string value up to 48 characters long. Do not use embedded spaces or the # symbol.
- If you do not want to change the value of a community string, press Return or Enter at either prompt.

Options

Prompt	Description	Possible Values	[Default]
Read-only	Octet string, included in each SNMP message, that provides read-only access to system information	<ul style="list-style-type: none"> ■ public ■ A string up to 48 characters long 	public
Read-write	Octet string, included in each SNMP message, that controls read-write access to system information	<ul style="list-style-type: none"> ■ private ■ A string up to 48 characters long 	private

SNMP Community Example (3500)

```
Select menu option (snmp): community
Enter new read-only community [public]:our_app
Enter new read-write community [private]: my_mail
```


snmp trap display

Displays the SNMP traps and their currently configured destinations.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sn t d

Fields in the SNMP Trap Display

Field	Description
Trap description	Description of the system event that triggers the trap
Trap destinations configured	IP address of the system that is to receive event notification
Trap number	Identifying number of the trap that is associated with a system event
Trap numbers enabled	Traps that are active

snmp trap addModify

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Adds or modifies trap reporting destination configurations. When an event occurs, the system sends the trap that you specify here to the destination address.

Valid Minimum Abbreviation`sn t a`**Important Considerations**

- You can define up to 10 destination addresses and the set of traps that are sent to each destination address.
- No unlisted traps are transmitted.
- Specify a range of more than two trap numbers with a hyphen (-) and nonsequential trap numbers with commas.
- If the destination address that you entered is not a valid end station, if a valid IP interface is not defined on the system, or if the agent does not have a route to the destination, the agent displays this message:

`Trap address invalid or unreachable`

If you see this message, verify the IP address of the end station, that it is online, and that a route exists to the intended management station.

- See the “Device Monitoring” chapter in the *Implementation Guide* for your system for an explanation of what the individual traps mean.

Options

Prompt	Description	Possible Values	[Default]
Trap destination address	Destination IP address of the SNMP manager	A valid destination IP address	–
Trap numbers to enable	Traps that you want to direct to the SNMP Manager	<ul style="list-style-type: none"> ■ A valid trap #, range, or sequence of valid trap #s ■ all ■ ? (for a list of available trap numbers) 	–

Procedure

- 1 From the top level of the Administration Console, enter:

```
snmp trap addModify
```

The system displays the list of traps.

- 2 Enter the IP address of the SNMP manager (destination address).
- 3 Enter one or more trap numbers for that destination, **all**, or **?** to get a list of selectable values.

SNMP Trap addModify Example (3500)

Select menu option (snmp/trap): **addModify**

Trap Descriptions:

Trap #	Description
1	MIB II: Coldstart
2	MIB II: Link Down
3	MIB II: Link Up
4	MIB II: Authentication Failure
5	Bridge MIB: New Root
6	Bridge MIB: Topology Change
7	3C System MIB: System Overtemperature
8	3C System MIB: Power Supply Failure
13	3C System MIB: Address Threshold
14	3C System MIB: System Fan Failure
15	3C FDDI MIB: SMT Hold Condition
16	3C FDDI MIB: SMT Peer Wrap Condition
17	3C FDDI MIB: MAC Duplicate Address Condition
18	3C FDDI MIB: MAC Frame Error Condition
19	3C FDDI MIB: MAC Not Copied Condition
20	3C FDDI MIB: MAC Neighbor Change
21	3C FDDI MIB: MAC Path Change
22	3C FDDI MIB: Port LER Condition
23	3C FDDI MIB: Port Undesired Connection
24	3C FDDI MIB: Port EB Error Condition
25	3C FDDI MIB: Port Path Change
26	RMON MIB: Rising Alarm
27	RMON MIB: Falling Alarm
28	POLL MIB: Response Received
29	POLL MIB: Response Not Received
32	VRRP MIB: New Master
33	VRRP MIB: Authentication Failure
35	QOS MIB: QOS INTRUDER Trap

Enter the trap destination address: 158.102.31.22

Enter the trap numbers to enable (1-8,13-29,32-33,35|all|?)

[1-8,13-29,32-33,35]: 35

snmp trap remove

Removes a destination, so that no SNMP traps are reported to that destination.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

sn t r

✓ 3900

✓ 9300

Important Consideration

- When the system removes the destination address, it displays the previous menu.

snmp trap flush Removes all SNMP trap reporting destinations.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

sn t f

Important Consideration

- When you flush the SNMP trap reporting destinations, you remove all trap destination address information for the SNMP agent.

**snmp trap
smtProxyTraps**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Controls SNMP's ability to alert you, by means of an SNMP-to-SMT proxy, that a significant event is occurring in the Fiber Distributed Data Interface (FDDI) station statistics.

Valid Minimum Abbreviation

sn t s

Options

Prompt	Description	Possible Values	[Default]
SNMP-to-SMT proxy mode	Whether the SMT proxy agent is enabled or disabled	■ enabled ■ disabled	disabled

snmp

rmonConfiguration

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Configures the transmit and receive mode to monitor Ethernet and Fiber Distributed Data Interface (FDDI) statistics as follows:

- **receive** — Monitors incoming port data
- **transmitAndReceive** — Monitors incoming and outgoing port data

Valid Minimum Abbreviation

sn r

Options

Prompt	Description	Possible Values	[Default]
Transmit/receive mode	Whether RMON is configured for only incoming port data, or for both incoming and outgoing port data	<ul style="list-style-type: none">■ receive■ transmitAndReceive	Current setting


- ✓ 9300

Valid Minimum Abbreviation

sn w

Options


Prompt	Description	Possible Values	[Default]
SNMP write request mode	Whether SNMP write access is enabled or disabled	<ul style="list-style-type: none"> ■ off ■ on 	off



PHYSICAL PORT PARAMETERS

Chapter 7 Ethernet Ports

Chapter 8 Fiber Distributed Data Interface (FDDI)



7

ETHERNET PORTS

Before you configure your system, become familiar with the physical port numbering scheme on the system. Understanding the port numbering scheme enables you to:

- Manage your bridge ports (especially if you use trunking), as described in the *Implementation Guide* for your system
- Accurately define your virtual LANs (VLANs), as described in the *Implementation Guide* for your system

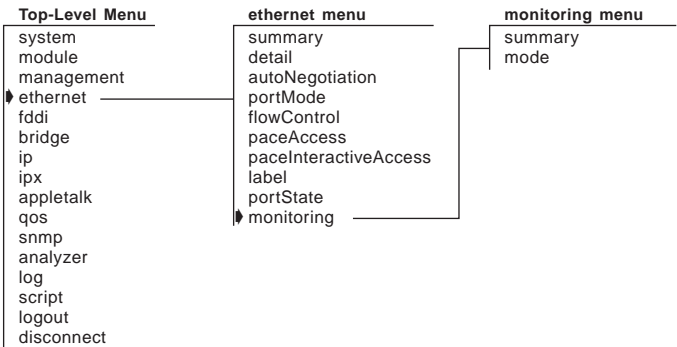
This chapter provides guidelines and other key information about how to configure Ethernet ports in your system.



For more information about port numbering and how to configure Ethernet ports, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on your system, your level of access, and the modules and other hardware options that are configured for your system. The following diagram shows the list of commands for all systems. The checklist at the beginning of each command description in this chapter shows whether your system supports the command.



ethernet summary

Displays a summary of Ethernet port information. The summary shows the port's label and status, as well as the most pertinent statistics about general port activity and port errors.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

e s

Important Considerations

- Port numbering is consecutive, regardless of module type (if you are using a system that has modules).
- Depending on the system, numbering may or may not skip an empty slot and continue with the ports that are associated with the next occupied slot. (See the *Implementation Guide* for your system for specific information about port numbering.)
- Numbering includes unused ports.
- Only one port number is assigned to a Gigabit Ethernet module in switches that use Gigabit Ethernet modules.
- The system no longer assigns port number 1 to the out-of-band management port, which does not receive a port number.
- The `rxFrames` value that the Ethernet summary command reports for a bridge port may differ from the value that the bridge port summary command reports. The Ethernet summary command counts *all* frames that are delivered to the port while the bridge port summary command reports only *valid* frames that are passed to the port. Therefore, the Ethernet summary value should exceed the bridge port summary value by the number of receive errors (`rxErrs`).
- At some prompts, you can specify the `?` option to list Ethernet ports and port numbers. The `?` option displays Selection, Port, and Port Label columns. The Selection column and Port column contain the same port numbers because they represent your physical ports.

Fields in the Ethernet Summary Display

Field	Description
actualFlowControl	Actual flow control setting. When autonegotiation is completed, the value is the autonegotiated setting. When autonegotiation is disabled, the value is the user-selected flow control value.
actualPortMode	Actual operating port mode. When autonegotiation is completed, the values shown are the autonegotiated settings. When autonegotiation is disabled, the value is the user-selected port mode.
autoNegMode	Autonegotiation mode configured for port. Possible values are <code>enable</code> or <code>disable</code> .
autoNegState	Current negotiation state. Possible values are <code>disabled</code> , <code>configuring</code> , <code>completed</code> , and <code>failed</code> .
linkStatus	Boolean value that indicates the current state of the physical link for this port (either <code>enabled</code> or <code>disabled</code>).
macAddress	MAC address of this port.
noRxBuffers	Number of frames that were discarded because no buffer space was available.
portLabel	User-defined label name. The maximum length of the string is 32 characters, including the null terminator.
portState	Current software operational state of this port. Possible values are <code>on-line</code> , <code>off-line</code> , <code>partitioned</code> , <code>tx-fault</code> , and <code>config-error</code> . The value <code>on-line</code> appears when the port is both enabled and connected to a cable. The value <code>partitioned</code> appears when the port has been disabled by the ethernet port monitoring feature.
portType	Specific description of this port's type. Values for each port type: <code>10/100BASE-T (RJ45)</code> , <code>100BASE-FX (SC)</code> , <code>1000BASE-SX (SC)</code> , <code>1000BASE-LX (GBIC)</code> , <code>1000BASE-SX (GBIC)</code> , <code>Backplane (9000)</code> .
reqFlowControl	Configurable parameter that sets the flow control option (when autonegotiation is disabled). When autonegotiation is enabled, flow control values are ignored.
reqPortMode	Configurable parameter that sets the port mode on Ethernet ports that have port mode options (when autonegotiation is disabled). When autonegotiation is enabled, port mode values are ignored.
rxBytes	Number of bytes received by this port, including framing characters.
rxErrs	Total of all receive errors that are associated with this port.

Field	Description
rxFrames	Number of frames that were copied into receive buffers by this port.
slot:channel (9000 switch fabric module)	Maps a CoreBuilder® 9000 switch fabric module port to an interface module backplane link. The “channel” designation is just a backplane trace number. For example, to troubleshoot a problem with switch fabric module port 5 (slot:channel 3:1), look at the first backplane link for slot 3.
slot:port (9000)	Module slot and port number in the CoreBuilder 9000 system.
txBytes	Number of bytes that were transmitted by this port, including framing characters.
txErrs	Sum of all transmit errors that are associated with this port (summary report only).
txFrames	Number of frames that were transmitted by this port.
txQOverflows	Number of frames that were lost because transmit queue was full.
vendorName (3500)	Vendor name for a GBIC module. Other modules display n/a.

To display Ethernet port statistics relative to a baseline, see the *Implementation Guide* for your system.

Procedure

- 1 To display summary information about Ethernet ports, enter:
ethernet summary
- 2 At the prompt (for example, (1-24|all|?)), select the ports whose information you want to display, or to display a port summary, specify ?
Indicate a range of ports with a hyphen (-). Separate nonconsecutive ports with a comma.
The system displays port information based on the ports that you specified.

ethernet detail

Displays detailed Ethernet port information including the information in the summary and additional Ethernet port statistics, such as collision counters.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

e d

Important Considerations

- Port numbering is consecutive, regardless of module type (if you are using a system that has modules).
- Depending on the system, numbering may or may not skip an empty slot and continue with the ports that are associated with the next occupied slot. (See the *Implementation Guide* for your system for specific information about port numbering.)
- Numbering includes unused ports.
- Only one port number is assigned to a Gigabit Ethernet module in switches that use Gigabit Ethernet modules.
- The system no longer assigns port number 1 to the out-of-band management port, which does not receive a port number.
- The `rxFrames` value that the Ethernet detail command reports for a bridge port may differ from the value that the bridge port detail command reports. The Ethernet detail command counts *all* frames that are delivered to the port while the bridge port detail command reports only *valid* frames that are passed to the port. Therefore, the Ethernet detail value should exceed the bridge port detail value by the number of receive errors (`rxErrs`).
- At some prompts, you can specify the `?` option to list Ethernet ports and port numbers. The `?` option displays Selection, Port, and Port Label columns. The Selection column and Port column contain the same port numbers because they represent your physical ports.

Fields in the Ethernet Detail Display

Field	Description
actualFlowControl	Actual flow control setting. When autonegotiation is completed, the value is the autonegotiated setting. When autonegotiation is disabled, the value is the user-selected flow control value.
actualPortMode	Actual operating port mode. When autonegotiation is completed, the value shown is the autonegotiated setting. When autonegotiation is disabled, the value is the port mode.
alignmentErrs (3500, 3900 and 9000)	Number of frames received by this port that are not an integral number of octets in length and do not pass the FCS check.
autoNegMode	Autonegotiation mode configured for port. Possible values are <code>enable</code> or <code>disable</code> .
autoNegState	Current negotiation state. Possible values are <code>disabled</code> , <code>configuring</code> , <code>completed</code> , and <code>failed</code> .
carrierSenseErr (3500, 3900 and 9000)	Number of frames that were discarded because the carrier sense condition was lost while transmitting a frame from this port.
excessCollision (3500, 3900 and 9000)	Number of frames that have been dropped because they experienced 15 consecutive collisions when sent from this port. This value is incremented by 1 each time that a frame experiences 15 consecutive collisions.
excessDeferrals (3500 and 9000 Layer 3)	Number of frames that were not transmitted on this port because the maximum allowed deferral time was exceeded.
fcsErrs	Number of frames received by this port that are an integral number of octets in length but do not pass the frame check sequence (FCS) test.
fragments (3900, 9000 Layer 2, 9300 and 9400)	Number of frames received by this port that were shorter than 64 bytes and had CRC or alignment errors.
jabbers (3900, 9000 Layer 2, 9300 and 9400)	Number of frames received by this port that were longer than 1518 bytes and had CRC or alignment errors.
lateCollisions (3500, 3900 and 9000)	Number of times that a collision was detected on this port later than 512 bit-times into the transmission of a frame.
lengthErrs (3500 and 9000 Layer 3)	Number of frames received by this port that are longer than 1518 bytes or shorter than 64 bytes.

Field	Description
linkStatus	Boolean value that indicates the current state of the physical link for this port (either <code>enabled</code> or <code>disabled</code>).
macAddress	MAC address of this port.
multiCollisions (3500, 3900 and 9000 Layer 3)	Number of frames that have experienced from 2 to 15 consecutive collisions <i>before successful transmission</i> from this port. If a frame also experiences a collision on the 15th attempt, it is dropped and the <code>excessCollision</code> count is increased by 1.
noRxBuffers	Number of frames that were discarded because no buffer space was available.
oversized (3900, 9000 Layer 2, 9300 and 9400)	Number of frames received by this port that were longer than 1518 bytes.
paceAccess (3900 and 9000 Layer 2)	Whether PACE® Interactive Access is <code>enabled</code> or <code>disabled</code> for this port.
portLabel	User-defined label name. The maximum length of the string is 32 characters, including the null terminator.
portState	Current software operational state of this port. Possible values are <code>on-line</code> , <code>off-line</code> , <code>partitioned</code> , <code>tx-fault</code> , and <code>config-error</code> . The value <code>on-line</code> appears when the port is both <code>enabled</code> and connected to a cable. The value <code>partitioned</code> appears when the port has been disabled by the ethernet port monitoring feature.
portType	Specific description of this port's type. Values for each port type: <code>10/100BASE-T (RJ45)</code> , <code>100BASE-FX (SC)</code> , <code>1000BASE-SX (SC)</code> , <code>1000BASE-LX (GBIC)</code> , <code>1000BASE-SX (GBIC)</code> , <code>Backplane (9000)</code> .
reqFlowControl	Configurable parameter that sets the flow control option (when autonegotiation is disabled). When autonegotiation is enabled, flow control values are ignored.
reqPortMode	Port mode on Ethernet ports that have port mode options (when autonegotiation is disabled). When autonegotiation is enabled, port mode values are ignored.
requestedState	Configurable parameter that is used to enable and disable this port. The default is <code>enabled</code> .
runts (3900, 9000 Layer 2, 9300 and 9400)	Number of frames received by this port that were shorter than 64 bytes.
rxBroadcast (3900, 9000 Layer 2, 9300 and 9400)	Number of broadcasts received by this port.

Field	Description
rxByteRate	Average number of bytes received per second by this port during the most recent sampling period.
rxBytes	Number of bytes received by this port, including framing characters.
rxDiscards (3500 and 9000 Layer 3)	Number of received frames that were discarded because there was no higher layer to receive them or because the port was disabled.
rxFrameRate	Average number of frames that were received per second by this port during the most recent sampling period. Sampling periods are 1 second long and not configurable.
rxFrames	Number of frames that were copied into receive buffers by this port.
rxInternalErrs	Number of frames that were discarded because of an internal error during reception.
rxMcastsOnly (3900, 9000 Layer 2, 9300 and 9400)	Number of multicast frames received by this port.
rxMulticasts	Number of multicast frames that were delivered to a higher-level protocol or application by this port.
rxPeakByteRate	Peak value of ethernetPortByteReceiveRate for this port since the station was last initialized.
rxPeakFrameRate	Peak value of ethernetPortFrameReceiveRate for this port since the station was last initialized.
rxUnicasts	Number of unicast (nonmulticast) frames that were delivered by this port to a higher-level protocol or application.
singleCollision (3500, 3900 and 9000)	Number of frames that have experienced only one collision before successful transmission from this port on the second attempt.
slot:port (9000)	Module slot and port number.
txBroadcasts (3900, 9000 Layer 2, 9300 and 9400)	Number of frames that were queued for transmission from this port by a higher-level protocol or application, including frames not transmitted successfully.
txByteRate	Average number of bytes that were transmitted per second by this port during the most recent sampling period.
txBytes	Number of bytes that were transmitted by this port, including framing characters.
txDiscards	Number of transmitted frames that were discarded because the port was disabled.

Field	Description
txFrameRate	Average number of frames that were transmitted per second by this port during the most recent sampling period. Sampling periods are 1 second long and not configurable.
txFrames	Number of frames that were transmitted by this port.
txInternalErrs	Number of frames that were discarded because of an internal error during transmission.
txMcastsOnly (3900, 9000 Layer 2, 9300 and 9400)	Number of multicast frames transmitted by this port.
txMulticasts	Number of multicast frames that were queued for transmission from this port by a higher-level protocol or application, including frames not transmitted successfully.
txPeakByteRate	Peak value of ethernetPortByteTransmitRate for this port since the station was last initialized.
txPeakFrameRate	Peak value of ethernetPortFrameTransmitRate for this port since the station was last initialized.
txQOverflows	Number of frames lost because transmit queue was full.
txUnicasts	Number of unicast (nonmulticast) frames that are queued for transmission by a higher-level protocol or application, including frames not transmitted successfully.
vendorName (3500)	Vendor name for a GBIC module. Other modules display n/a.

ethernet autoNegotiation

Enables or disables autonegotiation of port attributes such as duplex mode and port speed on ports that support autonegotiation.

✓ 3500
✓ 9000
✓ 9400

✓ 3900
✓ 9300

Valid Minimum Abbreviation

e a

Important Considerations

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- When you `enable` autonegotiation, the system ignores your requested `portMode` information for 10/100BASE-TX ports and your requested `flowControl` information for 1000BASE-SX ports. When you `disable` autonegotiation, the system recognizes the requested `portMode` values for ports that have `portMode` options and the requested `flowControl` values for 1000BASE-SX ports. (100BASE-FX ports and backplane ports do not support autonegotiation.)

Therefore, it is extremely important that you understand how to implement flowcontrol and `portMode` in your network. See the *Implementation Guide* for your system for more information.

Options

Prompt	Description	Possible Values	[Default]
Port	Port numbers for which you want to enable or disable autonegotiation	<ul style="list-style-type: none"> ■ A single port ■ A range of ports ■ all ■ ? (to display a port summary) 	–
Autonegotiation setting	Whether to enable or disable autonegotiation on each of the ports that you selected	<ul style="list-style-type: none"> ■ enable ■ disable 	enable

ethernet portMode

Sets the port speed (10 Mbps or 100 Mbps) and the duplex mode (full-duplex or half-duplex) on individual ports.

✓ 3500

✓ 9000

9400

✓ 3900

9300

Valid Minimum Abbreviation**e portm****Important Considerations**

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- The device that is connected to each port must be configured for the same port mode. If the port speeds differ, the link does not come up. If the duplex modes differ, link errors occur.
- Gigabit Ethernet ports do not support mode options. The value `all` refers only to ports that support port mode options.
- If you change to full-duplex mode on the port, a message indicates that collision detection will be disabled unless you configure the connected device to the same duplex mode.
- Disable autonegotiation on any port on which you are setting a specific port mode.
- 10/100BASE-TX supports the following modes and speeds:
 - 10 Mbps, full-duplex mode
 - 10 Mbps, half-duplex mode
 - 100 Mbps, full-duplex mode
 - 100 Mbps, half-duplex mode
- 100BASE-FX supports the following modes and speeds:
 - 100 Mbps, full-duplex mode
 - 100 Mbps, half-duplex mode

Options

Prompt	Description	Possible Values	[Default]
Port	Ports for which you want to change the portMode values	<ul style="list-style-type: none">■ A single port■ A range of ports■ all■ ? (to display a port summary)	–
Port mode setting	Speed and duplex mode for each of the ports that you selected	See “Important Considerations,” earlier in this section	10half (10/100BASE-TX) 100half (100BASE-FX)

Procedure

- 1 To change the port speed or duplex mode for 10/100BASE-TX ports or the duplex mode for 100BASE-FX ports, enter:

ethernet portMode

- 2 At the prompt (for example, (1-24|all|?)), enter the ports whose portMode values you want to change, or to display a port summary, specify ?

After you have selected the ports, the system prompts you to enter the port mode for the ports that you selected.

ethernet flowControl

Controls whether a Fast Ethernet or Gigabit Ethernet port can respond to or generate flow control packets.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

e f

Important Considerations

- Flow control allows a port to:
 - Decrease the frequency with which it sends packets to a receiving device, if packets are being sent too rapidly.
 - Send flow control packets to a sending device, to request that the device slow its speed of transmission to the port.
- The system does not count flow control packets in either receive or transmit statistics.

Options

Prompt	Description	Possible Values	[Default]
Port selection	Ports for which you want to set flow control characteristics	<ul style="list-style-type: none">■ A single port■ A range of ports■ all■ ? (to display a port summary)	–
Flow control setting	Flow control characteristics for each of the ports that you selected	<ul style="list-style-type: none">■ on■ off■ rxOn■ txOn	off

Flow Control Settings

Setting	Description	Available on Port Type
on	Port recognizes flow control packets and responds by pausing transmission. The port can generate flow control packets as necessary to slow incoming traffic.	Gigabit Ethernet Fast Ethernet
off	Port ignores flow control packets and does not generate them.	Gigabit Ethernet Fast Ethernet
rxOn	Port recognizes flow control packets and responds by halting transmission. The port does not generate flow control packets.	Gigabit Ethernet
txOn	Port ignores flow control packets, but it can generate them, if necessary.	Gigabit Ethernet

ethernet paceAccess***For CoreBuilder 9000: Applies to Layer 2 switching modules only.*****3500****✓ 9000****9400****✓ 3900****9300**

Configures the Ethernet ports on your system to support the PACE® Interactive Access feature, which ensures reliable timing by preventing excessive Ethernet network jitter (the variation in the timing of packet delivery that can cause garbled sound, jerky images, and delays).

Valid Minimum Abbreviation**e pa****Important Considerations**

- PACE technology is 3Com's method to provide reliable timing, optimal LAN bandwidth utilization, and data prioritization for time-sensitive multimedia and real-time applications, and data-only applications.
- PACE Interactive Access employs a "back-off" algorithm that enables your system to control traffic flow on a point-to-point link with an end station. When the network experiences congestion, the switch holds packets. PACE Interactive Access prevents an end station from "monopolizing" the link.

Options

Prompt	Description	Possible Values	[Default]
Port	Ports for which you want to set the PACE® feature	<ul style="list-style-type: none">■ A range of port numbers■ all■ ? (to display a port summary)	–
PACE setting	Whether the PACE feature is on or off for each of the ports that you selected	<ul style="list-style-type: none">■ enable■ disable	disable

ethernet
paceInteractiveAccess

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Configures the Ethernet ports on your system to support the PACE Interactive Access feature, which ensures reliable timing by preventing excessive Ethernet network jitter (the variation in the timing of packet delivery that can cause garbled sound, jerky images, and delays).

Valid Minimum Abbreviation

e pa

Important Considerations

- PACE technology is 3Com's method to provide reliable timing, optimal LAN bandwidth utilization, and data prioritization for time-sensitive multimedia and real-time applications, and data-only applications.
- PACE Interactive Access employs a "back-off" algorithm that enables your system to control traffic flow on a point-to-point link with an end station. When the network experiences congestion, the switch holds packets. PACE Interactive Access prevents an end station from "monopolizing" the link.

Options

Prompt	Description	Possible Values	[Default]
Port	Ports for which you want to set the PACE® feature	<ul style="list-style-type: none">■ A range of port numbers■ all■ ? (to display a port summary)	–
PACE setting	Whether the PACE feature is on or off for each of the ports that you selected	<ul style="list-style-type: none">■ enable■ disable	disable

ethernet label Labels the Ethernet ports to help identify the kind of device that is attached to each port (for example, LAN, workstation, or server).

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

e 1

Important Considerations

- Label Ethernet ports so that you can easily identify the devices that are attached to them (such as LANs, workstations, or servers). For example: `engineeringserver`
- A new port label appears in system displays the next time that you display information for that port.

Options

Prompt	Description	Possible Values	[Default]
Port selection	Ports for which you want to define a port label	<ul style="list-style-type: none">■ A range of port numbers■ all■ ? (to display a port summary)	–
Port label	Labels that appear the next time that you display information for the ports that you selected	String of up to 32 ASCII characters, including the null terminator	–

ethernet portState

Enables or disables Ethernet ports, controlling whether the ports send or receive frames.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

e ports

Important Consideration

- When an Ethernet port is **enabled**, frames are transmitted normally over that port. When an Ethernet port is **disabled**, the port neither sends nor receives frames.

Options

Prompt	Description	Possible Values	[Default]
Port	Ports that you want to enable or disable	<ul style="list-style-type: none"> ■ Individual ports ■ A range of port numbers ■ all ■ ? (to display a port summary) 	–
Port state	Value shown in the summary and detail displays reports: on-line for all enabled ports displayed and off-line for all disabled ports displayed	<ul style="list-style-type: none"> ■ enabled ■ disabled 	enabled

Procedure

- 1 To enable or disable an Ethernet port, from the top level of the Administration Console, enter:
ethernet portState
- 2 At the prompt (for example, (1-24|all|?)), enter the ports whose port state values you want to set, or to display a port summary, specify ?
- 3 Enter **enabled** or **disabled** for each Ethernet port.

The **portState** value shown in the summary and detail displays reports **on-line** for all enabled ports that are displayed and **off-line** for all disabled ports. The Port Status LED for each disabled port on the module indicates the disabled status.

**ethernet monitoring
summary**

3500

✓ 9000

9400

✓ 3900

9300

Displays the status of 10/100 Mbps Ethernet ports that are being monitored. The display shows the status of port statistics that are being monitored, including:

- error count
- excessive collisions
- multiple collisions
- late collisions
- runts
- fcsErrs

Valid Minimum Abbreviation**e m s****Important Consideration**

- The Ethernet monitoring feature is enabled by default, and performs these functions:
 - 1 Monitors 10/100Mbps Ethernet ports for excessive collisions, multiple collisions, late collisions, runts, and FCS errors
 - 2 Compares these error counters against user-defined thresholds
 - 3 Disables a port that reaches an error threshold
 - 4 Reports the reason that a port is disabled to the Administration Console, MIB databases, and SNMP traps
 - 5 Reenables the port after an initial backoff time interval
 - 6 Continues monitoring

ethernet monitoring mode

Enables or disables port monitoring on 10/100 Mbps Ethernet ports on the switch.

3500

✓ 9000

9400

✓ 3900

9300

Valid Minimum Abbreviation

e m m

Important Considerations

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- You can determine when a monitored port is in error and has been disabled by these port statistics:
 - The `status` value shown in the ethernet monitoring summary display reports `partitioned`.
 - The `portState` value shown in the ethernet summary and ethernet detail displays reports `partitioned`.
 - The `linkStatus` value shown in the ethernet summary and ethernet detail displays reports `disabled`.

When the monitoring feature reenables the port, port statistics resume normal values.

- The Ethernet monitoring feature is enabled by default, and performs these functions:
 - 1 Monitors 10/100Mbps Ethernet ports for excessive collisions, multiple collisions, late collisions, runts, and FCS errors
 - 2 Compares these error counters against user-defined thresholds
 - 3 Disables a port that reaches an error threshold
 - 4 Reports the reason that a port is disabled to the Administration Console, MIB databases, and SNMP traps
 - 5 Reenables the port after an initial backoff time interval
 - 6 Continues monitoring

8

FIBER DISTRIBUTED DATA INTERFACE (FDDI)


Fiber Distributed Data Interface (FDDI) is a standards-based solution that provides fast and reliable data transfer on a local area network. This chapter provides guidelines and other key information about how to configure FDDI parameters in your system.



For more information about implementing FDDI in your network, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



fddi station display

Displays FDDI station information. The system display shows the station configuration, status reporting, and the most pertinent statistics about general station activity and errors.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

fd station d

Fields in the FDDI Station Display

Field	Description
configuration	Attachment configuration for the station or concentrator. Values can be <i>Thru</i> , <i>Isolated</i> , <i>Wrap_A</i> , and <i>Wrap_B</i> .
connectPolicy	Bit string that represents the connection policies in effect on a station. How connection policies translate into bits is described in “fddi station connectPolicy” in this chapter. This value is user-defined.
ecmState	Current state of the ECM state machine.
ports	Ports numbers assigned to the FDDI module. The FDDI port numbers change depending on the configuration of your system.
remoteDisconnect	Flag indicating that the station was remotely disconnected from the network as a result of receiving an <i>fddiSMTAction</i> with the value of <i>disconnect</i> in a Parameter Management Frame (PMF). A station requires a Connect Action to rejoin the network and clear the flag.
stationID	Unique identifier for the FDDI station.
statusReporting	Whether a station generates Status Report Frames (SRFs) to report events and conditions to network management stations. This value is user-defined.
tnotify	Timer used in the Neighbor Notification protocol to indicate the interval of time between generation of Neighbor Information Frames (NIF). This value is user-defined.
traceMaxExp	Maximum propagation time for a trace on an FDDI topology. Places a lower bound on the detection time for an unrecovering ring.

**fddi station
connectPolicy**

✓ 3500
✓ 9000
9400

3900
9300

Sets the connectPolicy attribute string that represents the connection policies in effect on a station. A connection's type is defined by the types of the two ports involved in the connection.

Valid Minimum Abbreviation

`fd station c`

Important Considerations

- Port types can be A, B, M, or S.
- The system FDDI ports are type A or type B for Dual Attachment Station (DAS) ports and type M for Single Attachment Station (SAS) ports.
- By default, all connections to the system ports are valid. M-M connections are accepted so that one CoreBuilder® 3500 port can be connected to another system port.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the connection policies	<ul style="list-style-type: none">■ Any of the available ports on the installed FDDI modules■ all	–
connectPolicy	Bit string that represents the connection policies in effect on that station	See next table	–

Bit to Set for Rejecting a Station Connection

This Connection Is Rejected	If This Bit Is Set	Connection Rules
A-A	0	Undesirable peer connection that creates twisted primary and secondary rings; notify station management (SMT).
A-B	1	Normal trunk ring peer connection.
A-S	2	Undesirable peer connection that creates a wrapped ring; notify SMT.

This Connection Is Rejected	If This Bit Is Set	Connection Rules
A-M	3	Tree connection with possible redundancy. The node may not go to Thru state in Configuration Management (CFM). In a single MAC node, Port B has precedence (with defaults) for connecting to a Port M.
B-A	4	Normal trunk ring peer connection.
B-B	5	Undesirable peer connection that creates twisted primary and secondary rings; notify SMT.
B-S	6	Undesirable peer connection that creates a wrapped ring; notify SMT.
B-M	7	Tree connection with possible redundancy. The node may not go to Thru state in CFM. In a single MAC node, Port B has precedence (with defaults) for connecting to a Port M.
M-A	12	Tree connection with possible redundancy.
M-B	13	Tree connection with possible redundancy.
M-S	14	Normal tree connection.
M-M	15	Connection that allows one system port to be connected to another system port.

fddi station tNotify

Sets the timer used in the Neighbor Notification protocol to indicate the interval of time between generation of Neighbor Information Frames (NIF).

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Valid Minimum Abbreviation

fd station t

Important Considerations

- If you set the T-notify value low, your network reacts quickly to station changes, but uses more bandwidth.
- If you set the T-notify value high, less bandwidth is used, but your network does not react to station changes as quickly.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the neighbor notification timer	<ul style="list-style-type: none">■ Any of the available ports on the installed FDDI modules■ all	–
tnotify	Timer (in seconds) used in the Neighbor Notification protocol to indicate the interval of time between generation of Neighbor Information Frames (NIF)	2 – 30 seconds	30

fdi station statusReporting

Controls whether a station generates Status Report Frames (SRFs) to report events and conditions to network management stations.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

`fd station s`

Important Consideration

- If you do not have an SMT management station listening to these event reports or if you use SNMP to monitor FDDI events on all FDDI end stations, set this attribute to `disabled` so that the station does not generate SRFs.

3900
9300

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set <code>statusReporting</code>	<ul style="list-style-type: none"> ■ Any of the available ports on the installed FDDI modules ■ all 	–
<code>statusReporting</code>	Parameter that controls whether a station generates Status Report Frames (SRFs) to report events and conditions to network management stations	<ul style="list-style-type: none"> ■ enabled ■ disabled 	enabled

fddi path display Displays FDDI path information.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

fd pa d

Important Consideration

- The path display changes slightly when ports are configured as DAS ports.

3900
9300

Fields in the FDDI Path Display

Field	Description
maxTReq	Maximum time value of fddiMACT-Req that any MAC that is configured in this path uses. This value can be user-defined.
path	Current selected path.
ports	Ports numbers that are assigned to the FDDI module. The FDDI port numbers change depending on the configuration of your system.
ringLatency	Total accumulated latency of the ring that is associated with this path.
tmaxLowBound	Minimum time value of fddiMACT-Max that any MAC that is configured in this path uses. This value can be user-defined.
traceStatus	Current trace status of the path.
txxLowBound	Minimum time value of fddiMACTvxValue that any MAC that is configured in this path uses. This value can be user-defined.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the path display	<ul style="list-style-type: none">■ Any of the available ports on the installed FDDI modules■ all	–
Path	Path that you want to set. <ul style="list-style-type: none">■ A DAS port has primary and secondary paths.■ A SAS port has only a primary path.	<ul style="list-style-type: none">■ p (primary)■ s (secondary)■ all	–

fddi path tvxLowerBound

Specifies the minimum time value (in microseconds) of fddiMAC tvxValue that any MAC that is configured in this path uses.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

fd pa tv

Important Considerations

- A MAC uses its valid transmission timer (TVX) to detect and recover from certain ring errors. If a valid frame has not passed through a MAC during the time indicated by fddiMACTvxValue, the MAC reinitializes the ring.
- By adjusting the tvxLowerBound value, you specify how quickly the ring recovers from an error. The lower you set this value, the faster the network reacts to problems, but the ring might be reinitialized when there is no problem.
- The higher you set this value, the less chance of frequent reinitializations, but the network takes longer to recover from errors.
- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the tvxLowerBound	<ul style="list-style-type: none"> ■ Any of the available ports on the installed FDDI modules ■ all 	–
Path	Path that you want to set: <ul style="list-style-type: none"> ■ A DAS port has primary and secondary paths. ■ A SAS port has only a primary path. 	<ul style="list-style-type: none"> ■ p (primary) ■ s (secondary) ■ all 	–
tvxLowerBound	Minimum time value of fddiMAC tvxValue that any MAC that is configured onto this path uses	0 – 4294967295 microseconds	2500

fddi path
tmaxLowerBound

- ✓ 3500
- ✓ 9000
- 9400

- 3900
- 9300

Specifies the minimum time value (in microseconds) of fddiMAC T-Max that any MAC that is configured in this path uses. This value specifies the boundary for how high T-Req (the requested token rotation time) can be set.

Valid Minimum Abbreviation

fd pa tm

Important Consideration

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the tmaxLowerBound	<ul style="list-style-type: none">■ Any of the available ports on the installed FDDI modules■ all	–
Path	Path that you want to set. <ul style="list-style-type: none">■ A DAS port has primary and secondary paths.■ A SAS port has only a primary path.	<ul style="list-style-type: none">■ p (primary)■ s (secondary)■ all	–
tmaxLowerBound	Minimum time value of fddiMAC T-Max that any MAC that is configured onto this path uses	0 – 4294967295 microseconds	16500

fddi path maxTreq✓ **3500**✓ **9000****9400****3900****9300**

Specifies the maximum time value (in microseconds) of fddiMACT-Req that is used by any MAC that is configured in this path. T-Req is the value that a MAC bids during the claim process to determine a ring's operational token rotation time, T_Opr. The lowest T-Req bid on the ring becomes T_Opr.

Valid Minimum Abbreviation**fd pa m****Important Considerations**

- When you set T_Opr low, the token rotates more quickly, so token latency is reduced. However, more of the ring's available bandwidth is used to circulate the token.
- Higher values of T_Opr use less bandwidth to circulate the token, but they increase token latency when the ring is saturated.
- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the path maxTreq	<ul style="list-style-type: none"> ■ Any of the available ports on the installed FDDI modules ■ all 	–
Path	Path that you want to set. <ul style="list-style-type: none"> ■ A DAS port has primary and secondary paths. ■ A SAS port has only a primary path. 	<ul style="list-style-type: none"> ■ p (primary) ■ s (secondary) ■ all 	–
maxTreq	Value that a MAC bids during the claim process to determine a ring's operational token rotation time, T_Opr	0 – 4294967295 microseconds	16500

fddi mac summary

✓ 3500
✓ 9000
9400

3900
9300

Displays a summary of FDDI MAC information. A summary report displays various FDDI MAC statistics, including information about the MAC, received and transmitted frames, and received and transmitted bytes.

Valid Minimum Abbreviation

fd m s

Important Consideration

- The MAC summary display changes slightly when ports are configured as DAS ports.

Fields in the FDDI MAC Summary Display

Field	Description
currentPath	Path on which this MAC is currently located (primary, secondary, or isolated)
downstream	MAC address of this MAC's downstream neighbor
Errors	Sum of errorCount, lateCount, lostCount, and txExpiredCount
noRxBuffers	Number of frames discarded because no buffer space was available
port	Port numbers assigned to the FDDI module. The FDDI port numbers change depending on the configuration of your system.
rxBytes	Number of bytes that this MAC received
rxErrors	Number of errors that this MAC received
rxFrames	Number of frames that this MAC received
station	Unique identifier for the FDDI station.
smtAddress	Address of the MAC that was used for SMT frames
txBytes	Number of bytes that this MAC transmitted
txFrames	Number of frames that this MAC transmitted. This number does not include MAC frames.
txQOverflows	Number of frames that were discarded because the transmit queue was full
upstream	MAC address of this MAC's upstream neighbor

fddi mac detail

Displays detailed FDDI MAC information. A detail report displays various FDDI MAC statistics, including information about the MAC, received and transmitted frames, and received and transmitted bytes, as well as additional FDDI MAC statistics.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

fd m d

Important Consideration

- The MAC summary display changes slightly when ports are configured as DAS ports.

Fields in the FDDI MAC Detail Display

Field	Description
currentPath	Path on which this MAC is currently located <u>primary or secondary, isolated, concatenated, or thru</u>
downstream	MAC address of this MAC's downstream neighbor
downstreamType	PC type of this MAC's downstream neighbor
dupAddrTest	Pass or fail test for a duplicate address
duplicateAddr	Whether this address is duplicated on the FDDI ring
errorCount	Number of SMT MAC errors
frameCount	Number of frames that this MAC received
frameErrCond	Active when the frameErrorRatio is greater than or equal to frameErrorThresh.
frameErrorRatio	Ratio of the number lostCount plus the frameErrorCount divided by the frameCount plus lostCount
frameErrThresh	Threshold for determining when a MAC condition report is generated
lateCount	Number of token rotation timer expirations since this MAC last received a token
llcAvailable	Whether this MAC can send or receive LLC frames
llcService	Setting of the Logical Link Control service
lostCount	Number of frames and tokens that this MAC lost during reception
noRxBuffers	Number of frames discarded because no buffer space was available
notCopiedCond	Active when the notCopiedRatio is greater than or equal to notCopiedThresh.

Field	Description
notCopiedCount	Number of frames that were addressed to this MAC but were not copied into its receive buffers
notCopiedRatio	Ratio of notCopiedCount divided by the quantity copiedCount plus notCopiedCount
notCopiedThresh	Threshold for determining when a MAC condition report is generated
oldDownstream	Previous value of the MAC address of this MAC's downstream neighbor
oldUpstream	Previous value of the MAC address of this MAC's upstream neighbor
ringOpCount	Number of times that this MAC has entered the operational state from the nonoperational state
rmtState	State of the ring management as defined in SMT
rxByteRate	Average number of bytes per second that this MAC received during the most recent sampling period
rxBytes	Number of bytes that this MAC received
rxDiscards	Number of good frames that this MAC received and discarded before being delivered to a higher-level protocol or application. Does not include frames that were not received into receive buffers, such as missed frames.
rxFrameRate	Average number of frames per second that this MAC received during the most recent sampling period
rxFrames	Number of frames that this MAC received
rxInternalErrs	Number of frames discarded because of an internal hardware error during reception
rxMulticasts	Number of multicast frames that this MAC delivered to a higher-level protocol or application
rxPeakByteRate	Peak value of fddiMACByteReceiveRate for this MAC since the station was last initialized
rxPeakFrameRate	Peak value of fddiMACFrameReceiveRate for this MAC since the station was last initialized
rxUnicasts	Number of unicast (nonmulticast) frames that this MAC delivered to a higher-level protocol or application
smtAddress	Address of the MAC used for SMT frames
tMax	Maximum value of the target token rotation time
tMaxCapab	Maximum supported target token rotation time that this MAC can support
tNeg	Target token rotation time negotiated during the claim process
tokenCount	Number of tokens that this MAC received

Field	Description
tReq	Target token rotation time that this MAC requested
txCapab	Maximum time value of the valid transmission timer that this MAC can support
txExpiredCount	Number of times that this MAC's valid transmission timer has expired
txValue	Value of the valid transmission timer that this MAC uses
txByteRate	Average number of bytes that this MAC transmitted per second during the most recent sampling period
txBytes	Number of bytes that this MAC transmitted
txDiscards	Number of frames discarded because LLC service was not enabled or the FDDI ring was not operational
txFrameRate	Average number of frames that this MAC transmitted per second during the most recent sampling period
txFrames	Number of frames that this MAC transmitted. This number does not include MAC frames.
txInternalErrs	Number of frames discarded because of an internal hardware error during transmission
txMulticasts	Number of multicast frames queued for transmission by a higher-level protocol or application, including frames not transmitted successfully
txPeakByteRate	Peak value of fddiMACByteTransmitRate for this MAC since the station was last initialized
txPeakFrameRate	Peak value of fddiMACFrameTransmitRate for this MAC since the station was last initialized
txQOverflows	Number of frames discarded because the transmit queue was full
txUnicasts	Number of unicast frames queued for transmission by a higher-level protocol or application, including frames not transmitted successfully
upstream	MAC address of this MAC's upstream neighbor
upstreamDupAddr	Whether the address upstream of this address is duplicated on the ring

**fdi mac
frameErrorThreshold**

- ✓ 3500
- ✓ 9000
- 9400

- 3900
- 9300

Determines when the system generates a MAC condition report because too many frame errors have occurred.

Valid Minimum Abbreviation

fd m f

Important Considerations

- A frame error occurs when a frame becomes corrupted.
- A high frame error rate often indicates a faulty station on the FDDI ring or a dirty FDDI connector.
- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the mac frameErrorThreshold	<ul style="list-style-type: none">■ Any of the available ports on the installed FDDI modules■ all	–
frameErrorThreshold value	Time value set to determine when the system generates a MAC condition report because too many frame errors have occurred	0 – 4294967295 microseconds	655

fdi mac notCopiedThreshold

Sets the timing when the system generates a MAC condition report because too many frames could not be copied.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

fd m n

Important Considerations

- Not-copied frames occur when there is no buffer space available in the station (which in turn indicates congestion in the station).
- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the mac notCopiedThreshold	<ul style="list-style-type: none"> ■ Any of the available ports on the installed FDDI modules ■ all 	–
notCopiedThreshold	Time value set to determines when the system generates a MAC condition report because too many frames could not be copied	0 – 4294967295 microseconds	6550

fddi mac llcService

Sets the Logical Link Control (LLC) service so that LLC frames are sent and received on the MAC. LLC frames are all data frames that are transmitted on the network.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Valid Minimum Abbreviation

fd m l

Important Considerations

- If there is something wrong on your network, you may want to turn off data (user) traffic for a MAC by disabling LLC service. Although you have disabled data traffic from the MAC, the MAC still participates in neighbor notification and is visible to network management.
- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the mac llcService	<ul style="list-style-type: none">■ Any of the available ports on the installed FDDI modules■ all	–
llcService	Whether LLC frames are sent and received on the MAC	<ul style="list-style-type: none">■ enabled■ disabled	enabled

fddi mac path Sets the path assignment for MACs.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

fd m p

Important Considerations

- The fddiMAC path selections depend on the stationMode configuration (DAS or SAS).
- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the MAC path	<ul style="list-style-type: none"> ■ Any of the available ports on the installed FDDI modules ■ all 	–
MAC path	Path assignments for MACs	For DAS ports: <ul style="list-style-type: none"> ■ primary ■ secondary ■ isolated For SAS ports: <ul style="list-style-type: none"> ■ primary ■ isolated 	primary

fd di port display

Displays information about FDDI ports, including the type, path, and port label, as well as other FDDI port statistics, such as error counters.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

fd po d

3900
9300

Fields in the FDDI Port Display

Field	Description
connectState	Connect state of this port (disabled, connecting, standby, or active)
currentPath	Path on which this port is currently located
ebErrorCond	Whether an elasticity buffer error has been detected during the past 2 seconds
ebErrorCount	Number of elasticity buffer errors that have been detected
lctFailCount	Number of consecutive times that the link confidence test (LCT) has failed during connection management
lemCount	Number of link errors that this port detected
lemRejectCount	Number of times that the link error monitor rejected the link
lerAlarm	Link error rate estimate at which a link connection generates an alarm
lerCondition	Whether the lerEstimate is less than or equal to lerAlarm
lerCutoff	Link error rate estimate at which a link connection is broken
lerEstimate	Average link error rate. It ranges from 10 ⁻⁴ to 10 ⁻¹⁵ and is reported as the absolute value of the exponent of the link error estimate
lineState	Line state of this port
myType	Type of port connector on the port (A, B, S, M)
neighborType	Type of port connector at the other end of the physical connection (A, B, S, M)
pcmState	Current Physical Connection Management (PCM) state defined in SMT
pcWithhold	Reason for withholding the connection
pmdClass	Type of PMD entity that is associated with this port
port	Ports numbers that are assigned to the FDDI module. The FDDI port numbers change depending on the configuration of your system.
portLabel	32-character string of a user-defined name for the port

fddi port lerAlarm

Sets the link error rate (LER) value at which a link connection generates an alarm.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation**fd po lerA****Important Considerations**

- The lerAlarm value is expressed as the absolute value of the exponent (such as 1×10^{-10}). A healthy network has an LER exponent between 1×10^{-10} and 1×10^{-15} .
- If the LER value is greater than the alarm setting, then SMT sends a Status Report Frame (SRF) to the network manager software indicating a problem with a port.
- Set the lerAlarm value below these values so that you receive alarms only if your network is in poor health. The SMT Standard recommended value is 8.
- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the port lerAlarm	<ul style="list-style-type: none"> ■ Any of the available ports on the installed FDDI modules ■ all 	–
lerAlarm value	Link error rate (LER) value at which a link connection generates an alarm	4 – 15	7

fd di port lerCutoff

Sets the link error rate estimate at which a link connection is disabled. When the lerCutoff value is reached, the PHY that detected a problem is disabled.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Valid Minimum Abbreviation

fd po lerC

Important Considerations

- The lerCutoff value must be lower than the lerAlarm value so that the network management software is alerted to a problem before the PHY (port) is actually removed from the network.
- Set the lerCutoff below these values so that a port is removed only as a last resort. The SMT Standard recommended value is 7.
- The lerCutoff value is expressed as an exponent (such as 1 x 10⁻¹⁰). A healthy network has an LER exponent between 1 x 10⁻¹⁰ and 1 x 10⁻¹⁵.
- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the port lerCutoff	<ul style="list-style-type: none">■ Any of the available ports on the installed FDDI modules■ all	–
lerCutoff	Link error rate estimate at which a link connection is disabled	4 – 15	4

fddi port label

Assigns a unique name to your FDDI ports for easy identification of the devices that are attached to them (for example, workstation, server, FDDI backbone). Port labels serve as useful reference points and as an accurate means of identifying your ports for management.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

```
fd po label
```

Important Consideration

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the port label	<ul style="list-style-type: none">■ Any of the available ports on the installed FDDI modules■ all	–
Label	Name of the FDDI port used for identification	–	–

fd di port path Sets the one or more FDDI ports to be either part of the primary path or isolated from the ring.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

fd po p

Options

Prompt	Description	Possible Values	[Default]
Ports	One or more FDDI station ports for which you want to set the path	<ul style="list-style-type: none">Any of the available ports on the installed FDDI modulesall	–
Path	Sets the type of path used by the port: <ul style="list-style-type: none">isol — isolates the port from the ringpri — sets the port to be part of the primary ring	<ul style="list-style-type: none">isolpri	pri

**fddi stationMode
display**

Generates a display of FDDI stationMode information. The display shows the station mode, DAS (Dual Attachment Station) or SAS (Single Attachment Station), for each FDDI port.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

fd stationM d

3900
9300

Important Consideration

- Before the new stationMode takes effect, you must reboot your system.

Fields in the FDDI Station Mode Display

Field	Description
Ports	Ports numbers that are assigned to the FDDI module. The FDDI port numbers change, depending on the configuration of your system.
stationMode	Current FDDI stationMode, DAS or SAS, that is assigned to a specific port.

fddi stationMode
modify

Modifies the stationMode, DAS or SAS, that is assigned to a specific port number.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

fd stationM m

Important Considerations


- You cannot modify the stationMode when any of the ports in the pair are part of a trunk.

■ Before the new stationMode takes effect, you must reboot your system.

■ You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.


Options

Prompt	Description	Possible Values	[Default]
Ports	FDDI station port for which you want to set the stationMode	<div> <div>■ Any of the available ports on the installed FDDI modules</div> <div>■ all</div> </div>	–
stationMode	Mode of the FDDI port pair selected to change	<div> <div>■ DAS</div> <div>■ SAS</div> </div>	SAS
reboot	Prompt to reboot the system if you want the stationMode changes to take effect	<div> <div>■ y (yes)</div> <div>■ n (no)</div> </div>	–



BRIDGING PARAMETERS

- Chapter 9** **Bridge-Wide Parameters**
- Chapter 10** **Bridge Port Parameters**
- Chapter 11** **Trunks**
- Chapter 12** **MultiPort Link Aggregation (MPLA)**
- Chapter 13** **Resilient Links**
- Chapter 14** **Virtual LANs (VLANs)**
- Chapter 15** **Packet Filters**



9

BRIDGE-WIDE PARAMETERS

This chapter provides guidelines and other key information about how use the Administration Console to configure bridge-wide parameters.




This chapter addresses the commands in the bridge menu, except port, packetFilter, vlan, trunk, mpla, and link, which other chapters in this Command Reference Guide address.



For more information about configuring bridging and related features, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



bridge display

Displays bridge statistics and configuration information including Spanning Tree Protocol (STP) parameter values.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

b d

Fields in the Bridge Display

Field	Description
addressCount	Number of addresses in the bridge address table at the point in time in which you are viewing it. This value fluctuates but the highest value reached is recorded in the <code>PeakAddrCount</code> field.
addrTableSize	Maximum number of addresses that can be stored in the bridge address table. For CoreBuilder® switches, the value is 32K. For SuperStack® II Switches, the value is 16K.
addrThreshold	Configurable reporting threshold for the total number of addresses known on this bridge. When this threshold is reached, the system generates the SNMP trap <i>addressThresholdEvent</i> . The range of valid values for setting this object is between 1 and 1 plus the maximum table size. To configure this value for the 3500 and 9000 L3 modules, see “bridge addressThreshold” later in this chapter. This option is not available for the 3900, 9300, 9400 and 9000 L2 modules at this release.
agingTime	Configurable time period in seconds that the bridge uses to age out dynamic addresses except when a topology change has occurred. (After a topology change, the bridge uses the value shown in the <code>forwardDelay</code> field instead until it receives configuration messages without the topology change flag set). The default value for <code>agingTime</code> is 300 seconds. The acceptable range is 10 – 1,000,000. You can also enter 0 to disable aging. To configure this value, see “bridge agingTime” later in this chapter.
bridgeFwdDelay	Configurable time period in seconds that the bridge spends in <i>each of</i> two states — listening and learning — before it transitions to the forwarding state, provided that the bridge is the root bridge. (If the bridge is not the root bridge, the bridge uses the value shown in the <code>fwdDelay</code> field that is assigned to it by the root bridge.) The default value is 15 seconds. The acceptable range is 4 – 30 seconds. To configure the bridge forward delay, see “bridge spanningTree stpForwardDelay” later in this chapter.

Field	Description
bridgeHelloTime	Configurable time period in seconds that elapses between configuration messages when the bridge is the root bridge. (If the bridge is not the root bridge, the bridge uses the value shown in the <code>helloTime</code> field which is assigned to it by the root bridge.) The default value is 2 seconds. The acceptable range is 1 – 10 seconds. To configure the bridge hello time, see “bridge spanningTree stpHelloTime” later in this chapter.
bridgeIdentifier	Unique bridge identification that includes the bridge priority value and the MAC address of port 1.
bridgeMaxAge	Configurable time period in seconds that the bridge uses to discard the stored configuration message when it is operating as the root bridge. (If the bridge is not the root bridge, it uses the value shown in the <code>maxAge</code> field instead which is assigned to it by the root bridge.) The default value is 20 seconds. The acceptable range is 6 – 40 seconds. To configure the bridge maximum age, see “bridge spanningTree stpMaxAge” later in this chapter.
designatedRoot	Identity of the root bridge. It includes the root bridge’s priority value and the MAC address of port 1 on that bridge.
forwardDelay	Time period in seconds that the bridge spends in <i>each</i> of two states — listening and learning — as assigned by the root bridge. Compare with the <code>bridgeFwdDelay</code> field.
gvrpState (3500 and 9000 L3)	Status of GARP VLAN Registration Protocol (GVRP) for the entire bridge. You configure GVRP as a bridge state as well as individual port states. To configure GVRP for the bridge, see “bridge gvrpState” later in this chapter. To configure GVRP on ports, see “bridge port gvrpState” in Chapter 10.
helloTime	Time period in seconds that elapses between the configuration messages that the bridge receives from the root bridge. Compare with the <code>bridgeHelloTime</code> field.
holdTime	Minimum delay time the bridge uses between topology change BPDUs that it sends.
ipFragmentation (3500 and 9000 L3)	Shows configuration state of the IP fragmentation option. The default setting is enabled. To configure this option, see “bridge ipFragmentation” later in this chapter.
ipxTranslation (3500 and 9000 L3)	Shows configuration state of IPX SNAP translation. The default setting is disabled. To configure this option, see “bridge ipxSnapTranslation” later in this chapter.
lowLatency	(Not available at this release)

Field	Description
maxAge	Time period in seconds that the bridge uses to discard stored configuration messages. The value is determined by the root bridge. Compare with the <code>bridgeMaxAge</code> field.
mode	Reflects that the bridge operates as a transparent bridge.
peakAddrCount	Reflects the highest number of addresses that have been counted since the last address table flush. For the current size of the address table, see the <code>addressCount</code> field.
priority	Configurable STP priority value for the bridge. The default value is 0x8000. (0x signifies that 8000 is a hexadecimal number.) The acceptable range is 0x0 – 0xffff (0 – ffff). To configure a value, see “bridge spanningTree stpPriority” later in this chapter. The bridge priority is included in the bridge identifier and is considered the most significant portion because it influences root bridge selection. A lower priority increases the odds that the bridge will become the root bridge.
rootCost	Value that reflects the total cost of the best path (lowest value) from the bridge root port to the root bridge. The value sums individual port path costs. To configure path costs for ports on the bridge, see “bridge port stpCost” in Chapter 10.
rootPort	Logical port with the best path from the bridge to the root bridge.
stpGroupAddress	Address to which the bridge listens to receive configuration messages and other STP information. To modify the STP group address, see “bridge spanningTree stpGroupAddress” later in this chapter.
stpState	Whether the Spanning Tree Protocol is enabled or disabled for the bridge. The default value is disabled for all switches except CoreBuilder 9000 modules. To configure the bridge STP state, see “bridge spanningTree stpState” later in this chapter. (STP is also configured on a port-by-port basis. See “bridge port stpState” in Chapter 10.)
timeSinceLastTopology Change	Time elapsed (in hours, minutes, and seconds) since STP last reconfigured the network topology.
topologyChangeCount	Number of times that STP has reconfigured the network topology since you enabled STP or rebooted the system (whichever is less).
topologyChangeFlag	Whether the bridge topology is currently changing (<code>true</code>) or not changing (<code>false</code>)

bridge
ipFragmentation

- ✓ 3500
- ✓ 9000
- 9400

- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Determines whether the Fiber Distributed Data Interface (FDDI) and Ethernet stations that are connected to your system can communicate using IP when FDDI stations transmit packets that are too large for Ethernet. IP fragmentation divides such large FDDI packets into smaller packets that can be bridged to Ethernet LANs.

Valid Minimum Abbreviation

b ipf

Options

Prompt	Description	Possible Values	[Default]
ipFragmentation value	Whether large FDDI packets can be divided into smaller packets so that they can be bridged to Ethernet	<ul style="list-style-type: none">■ enabled■ disabled	enabled

**bridge
ipxSnapTranslation**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Translates 802.3_RAW IPX packets to FDDI_SNAP packets when they are forwarded from Ethernet to FDDI links, and vice versa when packets are forwarded from FDDI to Ethernet.

Valid Minimum Abbreviation

b ipx

Important Consideration

- When IPX SNAP Translation is disabled, the system uses standard IEEE 802.1H bridging to translate 802.3_RAW packets to FDDI_RAW packets when they are forwarded from Ethernet to FDDI, and vice versa from FDDI to Ethernet.

Options

Prompt	Description	Possible Values	[Default]
ipx SnapTranslation	Whether the system uses IPX SNAP Translation when forwarding packets between Ethernet and FDDI links	<ul style="list-style-type: none">■ enabled■ disabled	disabled

bridge
addressThreshold

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the reporting threshold for the number of Ethernet addresses that are known. When this threshold is reached, the system generates the SNMP trap called *addressThresholdEvent*.

Valid Minimum Abbreviation

b ad

Important Considerations

- The bridge address table size on CoreBuilder switches is 32K; that is, the bridge can store a maximum of 32768 addresses.
- The range of valid values for this parameter is between 1 and 1 *plus the address table size*. Setting the address threshold to the highest possible value prevents the system from generating the trap, because the value can never be reached.

Options

Prompt	Description	Possible Values	[Default]
address threshold	Threshold for the total number of addresses that are known on this bridge	1 – 32769	29491 (factory default), or current value

bridge agingTime

Sets the maximum period (in seconds) for aging out (deleting) dynamic addresses from the address table.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation**b ag****Important Considerations**

- Use this parameter to configure the system to age addresses in a timely manner, without increasing packet flooding.
- To disable the bridge aging function, set the value to 0.
- This parameter does not affect statically configured addresses.

Options

Prompt	Description	Possible Values	[Default]
aging time	Maximum period (in seconds) for aging out dynamically learned forwarding information	<ul style="list-style-type: none">■ 0 to disable■ 10 – 1,000,000 seconds	300

bridge spanningTree
stpState

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Enables or disables the Spanning Tree Protocol (STP) on your system.

Valid Minimum Abbreviation

b sp stps

Important Considerations

- The state of STP is configured in two places: the entire bridge (this command) and individual bridge ports. (See Chapter 10.) The combination of the states determines the forwarding behavior of each port, as shown in the following table:

Bridge STP State	Port STP State	Port Participates in STP?	Port Forwards Frames?
Disabled	Disabled	No	Yes, if link state is up.
	Enabled	No	Yes, if link state is up.
	Removed	No	Yes, if link state is up.
Enabled	Disabled	No	No
	Enabled	Yes	Determined by STP, provided that the link state is up.
	Removed	No	Yes, if link state is up.

- After you enable STP, the system takes several seconds to process the command before the Administration Console menu reappears.
- Although bridge-wide STP is initially disabled, default values exist for the following STP bridge parameters: priority, max age, hello time, forward delay, and group address. These values do not function until STP is enabled.
- CoreBuilder® 3500 and CoreBuilder 9000 Layer 3 modules include an ignore STP mode option. See Chapter 14 in this guide or see your system *Implementation Guide* for more information.
- Bridge-wide STP must be disabled on a CoreBuilder 9400 switch if you configure it as a MultiPoint Link Aggregation (MPLA) core switch. For more information about MPLA, see the *CoreBuilder 9400 Implementation Guide*.

Options

Prompt	Description	Possible Values	[Default]
stpState (3500, 3900, 9300, 9400)	Whether the Spanning Tree Protocol is enabled or disabled for the system	■ enabled ■ disabled	disabled (factory default), or current value
stpState (9000 L2 and L3)	Whether the Spanning Tree Protocol is enabled or disabled for the module	■ enabled ■ disabled	enabled (factory default), or current value

bridge spanningTree
stpPriority

Modifies the bridge priority, which influences the choice of the root and designated bridges.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

b sp stpp

Important Considerations

- The bridge priority is expressed as a hexadecimal value. The characters 0x signify this.
- The lower the bridge’s priority value, the more likely it is that the bridge is chosen as the root bridge or a designated bridge.
- You can change the value while STP is disabled or enabled. If you change the value while STP is disabled, the value is retained when you enable STP.

Options

Prompt	Description	Possible Values	[Default]
STP priority	Bridge-wide STP parameter	0x0 – 0xffff	0x8000 (factory default), or current value

**bridge spanningTree
stpMaxAge**

✓ 3500
✓ 9000
✓ 9400

✓ 3900
✓ 9300

Determines when the stored CPDU configuration message is discarded from the bridge's memory if the bridge is the root bridge. The current value is shown in the bridgeMaxAge field of the `bridge display`.

Valid Minimum Abbreviation

b sp stpm

Important Considerations

- If the value is too small, the STP may reconfigure the topology too often, causing temporary loss of connectivity in the network.
- If the value is too large, the network may take longer than necessary to adjust to a new STP configuration after a topology change such as the restarting of a bridge.
- A conservative value assumes a delay variance of 2 seconds per hop. The recommended and default value is 20 seconds.
- Although the possible range for stpMaxAge is 6 – 40, the available range is constrained by the following inequalities:
 - $2 \times (\text{stpForwardDelay} - 1 \text{ second}) \geq \text{stpMaxAge}$
 - $\text{stpMaxAge} \geq 2 \times (\text{stpHelloTime} + 1 \text{ second})$

Options

Prompt	Description	Possible Values	[Default]
STP max age	Value (in seconds) when the stored configuration message information is deemed too old and is discarded	6 – 40 seconds	20 (factory default), or current value

**bridge spanningTree
stpHelloTime**

Sets the time between configuration messages that the bridge generates if it is operating as the root bridge. The current value is shown in the bridgeHelloTime field of the `bridge display`.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

`b sp stph`

Important Considerations

- If the probability of losing configuration messages is high, shorten the time to make the protocol more robust.
- If the probability of losing configuration messages is low, lengthen the time to lower the overhead of the algorithm.
- The recommended Hello time is 2 seconds.
- Although the possible range for stpHelloTime is 1 – 10, the available range is constrained by the following inequality:

$$\text{stpMaxAge} \geq 2 \times (\text{stpHelloTime} + 1 \text{ second})$$

Options

Prompt	Description	Possible Values	[Default]
STP hello time	Time (in seconds) between configuration messages from the root bridge	1 – 10 seconds	2 (factory default), or current value

**bridge spanningTree
stpForwardDelay**

✓ 3500
✓ 9000
✓ 9400

✓ 3900
✓ 9300

Sets the amount of time that the bridge spends in each of the listening and learning states if it is the root bridge. The current value is shown in the bridgeFwdDelay field of the `bridge display`.

Valid Minimum Abbreviation

`b sp stpf`

Important Considerations

- This value temporarily prevents a bridge from starting to forward data packets to and from a link until news of a topology change has spread to all parts of a bridged network.
- The recommended and default value is 15 seconds.
- Setting the value too low can result in temporary loops while STP reconfigures the topology.
- Setting the value too high can lead to a longer wait while STP reconfigures the topology.
- If the configuration was not successful, the system notifies you that your changes failed, and you can try to reenter the changes.
- Although the possible range for stpForwardDelay is 4 – 30, the available range is constrained by the following inequality:
$$2 \times (\text{stpForwardDelay} - 1 \text{ second}) \geq \text{stpMaxAge}$$

Options

Prompt	Description	Possible Values	[Default]
STP forward delay	Time (in seconds) that a bridge spends in the listening state and the learning state	4 – 30 seconds	15 (factory default), or current value

bridge spanningTree
stpGroupAddress

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Sets the single address to which a bridge listens to receive Spanning Tree Protocol (STP) information. Each STP bridge on the network sends STP packets to the group address. Every STP bridge on the network receives STP packets that were sent to the group address, regardless of which bridge sent the packets. The current value is shown in the stpGroupAddress field of the `bridge display`.

Valid Minimum Abbreviation

b sp stpg

Important Considerations

- Because there is no industry standard for a group address, products from different vendors may respond to different group addresses. If STP does not seem to be working in a mixed-vendor environment, other vendors’ products may use different group addresses as their defaults. If that is true, set the STP group address to be the same across all bridges in the network.
- Before you can modify the STP group address, you must disable STP (if it is not already disabled) on the bridge. (See “bridge spanningTree stpState” earlier in this chapter.)

Options

Prompt	Description	Possible Values	[Default]
STP group address	Single address to which a bridge listens for STP information	A valid STP group address	01-80-C2-00-00-00 (factory default), or current value

bridge gvrpState *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

Enables or disables the GARP VLAN Registration Protocol (GVRP), which can help simplify management of VLAN configurations in larger networks, and determines whether the virtual LAN (VLAN) origin for a port-based VLAN is dynamic (with GVRP) or static (without GVRP).

3900

9300

Valid Minimum Abbreviation

b g

Important Considerations

- To activate GVRP in your system, you first enable it for the entire bridge (this command) and then enable it on appropriate individual bridge ports (see Chapter 10).
- To maximize the effectiveness of GVRP, it should be enabled in as many end stations and network devices as possible.
- VLANs that are created dynamically with GVRP exist only as long as a GVRP-enabled device is sending updates. If the devices no longer send updates, or if GVRP is disabled, or if the system is rebooted, all dynamic VLANs are removed.
- If you disable GVRP after it has been enabled for a period of time, the system deletes all VLAN interfaces that were learned through GVRP and leaves unchanged all VLANs that were statically configured through the Administration Console or through the Web management software.

Options

Prompt	Description	Possible Values	[Default]
GVRP state	Whether the system uses GVRP for the entire bridge	<ul style="list-style-type: none">■ enable■ disable	disabled

bridge cos enable

For CoreBuilder 9000: Applies to Layer 2 switching modules only.

- 3500
- ✓ 9000
- ✓ 9400
-
- ✓ 3900
- ✓ 9300

Enables or disables IEEE 802.1p Class of Service (CoS) on the bridge. Use this feature to help prioritize business-critical or time-sensitive traffic in your network.

Valid Minimum Abbreviation

b c e

Important Considerations

- The opportunity to be processed in the high priority queue exists only for IEEE 802.1Q tagged packets (provided that CoS is enabled) with priority values that match the high priority queue configuration. Non-tagged packets are always processed in the low priority queue, along with tagged packets with priority values that match the low priority queue configuration.
- CoS is enabled by default and initial queue assignments conform with IEEE 802.1p recommendations — that is, priorities 0 – 3 are assigned to the low priority queue and priorities 4 – 7 are assigned to the high priority queue. To modify queue assignments, see “bridge cos modify” later in this chapter.
- If you disable CoS, all tagged and non-tagged traffic is processed through the low priority queue and its buffers. (The high priority queue and buffers are shut off.)

Options

Prompt	Description	Possible Values	[Default]
CoS setting	Whether all bridge ports in the system implement Class of Service	<ul style="list-style-type: none">■ enabled■ disabled	enabled (factory default), or current value

bridge cos summary***For CoreBuilder 9000: Applies to Layer 2 switching modules only.***

3500
 ✓ 9000
 ✓ 9400

 ✓ 3900
 ✓ 9300

Displays whether Class of Service (CoS) is enabled or disabled; shows how the eight possible priority values are assigned (or, if CoS is disabled, how they were last assigned) to the two queues; and shows the rate limit that exists on the high priority queue (queue 1).

Valid Minimum Abbreviation**b c s****Important Considerations**

- By default, CoS is enabled and the eight priority values (traffic classes 0 – 7) are divided between the two queues in accordance with IEEE 802.1p recommendations — that is, queue 1 (high priority) has classes 4, 5, 6, and 7 and queue 2 (low priority) has classes 0, 1, 2, and 3.
- If CoS is disabled (indicated at the top of the display), the display reflects the most recent configuration even though it is no longer active.

Options

Prompt	Description	Possible Values	[Default]
Queue index number	Number of the queue for which you want to see information	<ul style="list-style-type: none"> ■ 1 ■ 2 ■ all 	–

Fields in the Bridge CoS Summary Display

Field	Description
Queue	Number of the queue. Queue 1 is always the high priority queue. Queue 2 is always the low priority queue.
Rate limit	Percentage of traffic allowed on the high priority queue. See the <i>Implementation Guide</i> for your system for more information about the rate limit option.
Traffic classes	Priority values assigned to each queue. The IEEE 802.1p standard specifies eight possible values (0 – 7), each of which is intended to signify a certain kind of traffic.

bridge cos modify

For CoreBuilder 9000: Applies to Layer 2 switching modules only.

- 3500
- ✓ 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Changes how the eight priority values (0 – 7) are assigned to each of the two hardware queues and changes the optional rate limit on queue 1 (the high priority queue).

Valid Minimum Abbreviation

b c m

Important Considerations

- Because you cannot configure a rate limit for queue 2, the Administration Console prompts you to enter a rate limit only if you select queue 1.
- When you assign priority values to a given queue, the system automatically assigns the remaining priority values to the other queue.
- If CoS is disabled, you can still modify the queue settings; however, they do not affect traffic until CoS is enabled.

Options

Prompt	Description	Possible Values	[Default]
Queue index	Number of the device queue whose settings you want to modify	<ul style="list-style-type: none"> ■ 1 ■ 2 ■ ? (for a list of selectable indexes) 	–
Rate limit	Throughput limit that applies only to queue 1 and is expressed as a percentage	Whole numbers from 1 – 100	100 (factory default), or current value
Class of service tags	IEEE 802.1p priority values that you want to assign to the selected queue. Use commas to separate multiple values.	<ul style="list-style-type: none"> ■ 0 – 7 ■ all ■ ? (for a list of possible values) 	For queue 1, values 4, 5, 6, 7 (factory default), or current values For queue 2, values 0, 1, 2, 3 (factory default), or current values

**bridge multicast igmp
summary****3500****✓ 9000****✓ 9400****✓ 3900****✓ 9300*****For CoreBuilder 9000: Applies to Layer 2 switching modules only.***

Displays a summary of parameters related to the Internet Group Management Protocol (IGMP) which conserves network bandwidth by directing IP multicast application traffic only to the ports that require it.

Valid Minimum Abbreviation**b mu i su****Important Consideration**

- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, IGMP commands exist under the `ip multicast igmp` menu.

Fields in the Bridge Multicast IGMP Summary Display

Field	Description
igmp snooping	Whether the IGMP snooping function is enabled or disabled for the entire system.
igmp querying	Whether the system is enabled to operate as an IGMP querier. IGMP snooping must be enabled for the querying function to operate.
igmp query source IP address	Source IP address used by the system or module in query messages if it is elected as the IGMP querier.

bridge multicast igmp
snoopMode

- 3500
- ✓ 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

For CoreBuilder 9000: Applies to Layer 2 switching modules only.

Enables or disables the snooping (listening) function of the Internet Group Management Protocol (IGMP).

Valid Minimum Abbreviation

b mu i sn

Important Considerations

- The value that you select applies to the entire system or module.
- IGMP snooping must be disabled on a CoreBuilder 9400 switch if you configure it as a MultiPoint Link Aggregation (MPLA) core switch. For more information about MPLA, see the *CoreBuilder 9400 Implementation Guide*.
- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, IGMP commands exist under the `ip multicast igmp` menu.

Options

Prompt	Description	Possible Values	[Default]
IGMP snooping	Whether your system implements IGMP snooping	<ul style="list-style-type: none">■ enabled■ disabled	enabled (factory default), or current value

**bridge multicast igmp
queryMode****3500**✓ **9000**✓ **9400**✓ **3900**✓ **9300*****For CoreBuilder 9000: Applies to Layer 2 switching modules only.***

Enables or disables the querying function of the Internet Group Management Protocol (IGMP). From all IGMP-capable devices on a given subnetwork, the one with the lowest IP address is elected as the querier.

Valid Minimum Abbreviation**b mu i querym****Important Considerations**

- The value that you select applies to the entire system.
- If you enable `igmp querymode`, but disable `igmp snoopmode`, the system or module cannot operate as an IGMP querier.
- To prevent IP multicast traffic from occupying unnecessary bandwidth, the best device to operate as the querier is the one closest to the source of IP multicast traffic. You can disable querying on select devices or manipulate IP addresses with the `bridge multicast igmp queryIpAddress` command to force this configuration.
- IGMP querying must be disabled on a CoreBuilder 9400 switch if you configure it as a MultiPoint Link Aggregation (MPLA) core switch. For more information about MPLA, see the *CoreBuilder 9400 Implementation Guide*.
- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, IGMP commands exist under the `ip multicast igmp` menu.

Options

Prompt	Description	Possible Values	[Default]
IGMP querying	Whether the system can operate as the IGMP querier if so elected	<ul style="list-style-type: none">■ enabled■ disabled	enabled (factory default), or current value

bridge multicast igmp
queryIpAddress

3500
✓ 9000
✓ 9400

✓ 3900
✓ 9300

For CoreBuilder 9000: Applies to Layer 2 switching modules only.

Configures the source address that is inserted in IGMP query packets.

Valid Minimum Abbreviation

b mu i queryi

Important Considerations

- For the CoreBuilder 9400 and SuperStack II Switch 9300 and 3900 systems, you do not need to used this command as long as you have one in-band IP interface configured; the system uses its IP address as the source IP address of query packets. Use this command only if you want the system to use a different source IP address for query packets. If there are no in-band IP interfaces configured and you want to enable querying, you must enter an IP address with this command.
- For a CoreBuilder 9000 Layer 2 switching module to offer itself as a querier, you must enter an IP address with this command.
- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, IGMP commands exist under the ip multicast igmp menu.

Options

Prompt	Description	Possible Values	[Default]
IGMP Query Source IP Address	Source address that the system uses in its IGMP queries	Any unique IP address in dotted decimal format	0.0.0.0, first in-band IP interface index, or current value (3900, 9300, 9400) 0.0.0.0, or current value (9000 L2)

**bridge multicast igmp
vlangs****3500****✓ 9000****✓ 9400****✓ 3900****✓ 9300*****For CoreBuilder 9000: Applies to Layer 2 switching modules only.***

If IGMP snooping is enabled, lists the VLAN IDs of VLANs that are carrying IP multicast traffic.

Valid Minimum Abbreviation**b mu i v****Important Consideration**

- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, IGMP commands exist under the `ip multicast igmp` menu.

bridge multicast igmp groups

- 3500
- ✓ 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

For CoreBuilder 9000: Applies to Layer 2 switching modules only.

Displays IP multicast group and associated port information for a selected VLAN.

Valid Minimum Abbreviation

b mu i g

Important Considerations

- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, IGMP commands exist under the `ip multicast igmp` menu.
- If no IP multicast group traffic is present on the selected VLAN, you see this message: `No groups exist for this VLAN`

Options

Prompt	Description	Possible Values	[Default]
VLAN ID	ID number of the VLAN for which you want to display group and port information	<ul style="list-style-type: none"> ■ A valid VLAN ID (VID) number ■ ? (for a list of selectable VIDs) 	1 (Default VLAN)

Fields in the Bridge Multicast IGMP Groups Display

Field	Description
VLAN ID	ID number of the selected VLAN.
Group	Hexidecimal equivalent of the IP multicast group address shown in the <code>IpAddress</code> column.
IpAddress	IP multicast group address of the traffic that the system or module has observed on the selected VLAN.
Ports	Ports that lead to group members.

**bridge multicast igmp
desQuerier****3500****✓ 9000****✓ 9400****✓ 3900****✓ 9300*****For CoreBuilder 9000: Applies to Layer 2 switching modules only.***

Determines whether the system or module is the designated querier for the selected VLAN.

Valid Minimum Abbreviation**b m u i d****Important Considerations**

- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, IGMP commands exist under the `ip multicast igmp` menu.
- If the system or module is not functioning as the querier, you see this message: `Device IS NOT designated querier on VLAN`
- If the system or module is functioning as the querier, you see this message: `Device IS designated querier on VLAN`

Options

Prompt	Description	Possible Values	[Default]
VLAN ID	ID number of the VLAN for which you are requesting information	<ul style="list-style-type: none">■ A valid VLAN ID (VID)■ ? (for a list of selectable VIDs)	1 (Default VLAN)

bridge multicast igmp
rPorts

- 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

For CoreBuilder 9000: Applies to Layer 2 switching modules only.

Lists the ports in the selected VLAN that lead to IP multicast routers.

Valid Minimum Abbreviation

b mu i r

Important Considerations

- The system determines which ports in a VLAN lead to multicast routers by snooping on advertisements from the following routing protocols: Distance-Vector Multicast Routing Protocol (DVMRP), Multicast Open Shortest Path First (MOSPF), and Protocol Independent Multicast (PIM).
- Router port entries age out after 100 seconds. Routing protocol advertisements are usually sent every few seconds.
- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, IGMP commands exist under the ip multicast igmp menu.

Options

Prompt	Description	Possible Values	[Default]
VLAN ID	ID number of the VLAN for which you want to list ports that lead to IP multicast routers	<ul style="list-style-type: none">■ A valid VLAN ID (VID)■ ? (for a list of selectable VIDs)	1 (Default VLAN)

bridge multicast igmp qPort

3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

For CoreBuilder 9000: Applies to Layer 2 switching modules only.

Displays the number of the port that receives incoming IGMP queries for the selected VLAN.

Valid Minimum Abbreviation

b mu i qp

Important Considerations

- If no query packets have been received within the last five minutes (approximately) when you enter this command, the system responds:
No queries are heard by the switch
- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, IGMP commands exist under the `ip multicast igmp` menu.

Options

Prompt	Description	Possible Values	[Default]
VLAN ID	ID number of the VLAN for which you want to display the port that last received query packets	<ul style="list-style-type: none"> ■ A valid VLAN ID (VID) ■ ? (for a list of selectable VIDs) 	1 (Default VLAN)

BRIDGE PORT PARAMETERS

This chapter provides guidelines and other key information about how to manage bridge ports in your system.



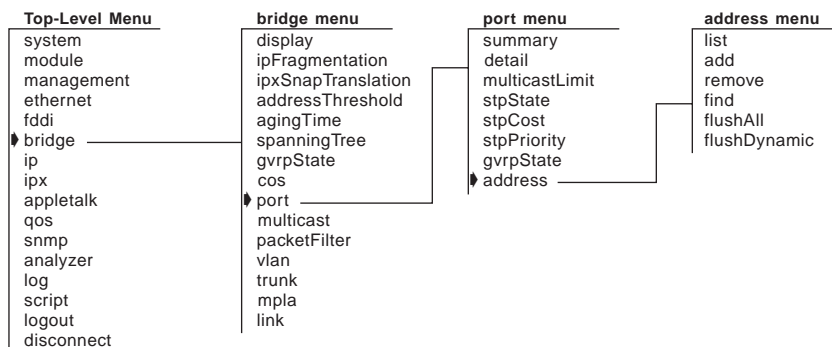
This chapter covers bridge port options only. For information about other bridge menu options, use the Table of Contents to find the appropriate chapter in this Command Reference Guide.



For more information about configuring bridge ports in your network, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



bridge port summary

Displays a summary of bridge port information, including the Spanning Tree Protocol (STP) configurations for selected bridge ports.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation**b po su****Important Considerations**

- The port numbering that is displayed is always sequential, although it depends on the placement of the modules that you have configured into your system. See the *Implementation Guide* for your system for more information about port numbering.
- For resilient links, the main and standby ports are shown in ascending order.
- When you are prompted to select ports, you can enter ? to see a matrix of information about the bridge ports. This matrix is useful, for example, if you have trunks configured but forget which port is the anchor port. You must use the anchor port number to display a port summary for a trunk.

Fields in the Bridge Port Summary Display

Field	Description
fwdTransitions	Number of times that the port has entered the forwarding state since you enabled STP or rebooted the system. This value is useful for determining the relative stability of the topology.
gvrpState (3500, 9000 L3)	Whether GARP VLAN Registration Protocol (GARP) is enabled or disabled on the port. For GVRP to function on the port, you must enable it on the port as well as the entire bridge. To configure GVRP on a port, see "bridge port gvrpState" later in this chapter. To configure GVRP on the bridge, see "bridge gvrpState" in Chapter 9.
linkState	State of the link (up or down), that is, whether it is available for communication.
portId	Port identification, which includes the port priority value (first 2 digits after "0x") and the logical port number (last 2 digits). Both are shown as hexadecimal values.

Field	Description
portNumber	Logical index number that the system assigns to the bridge port, which may not correspond with the physical port number depending on your system configuration. (For example, when you define a trunk, only the anchor port receives a portNumber.) As you add and remove logical ports, portNumbers are reassigned so that they remain consecutive.
rxDiscards	Total number of frames received on the bridge port that have been discarded. This value reflects a summary of all statistics that end with <code>Discs</code> , <code>Discards</code> , or <code>Filters</code> .
rxFrames	Total number of frames that this bridge port received from its segment. However, unlike the <code>rxFrames</code> field in the Ethernet display which counts all frames, this field does not count frames in error. Thus, this value may be lower than the value shown in the <code>rxFrames</code> field in the Ethernet display.
state	<p>Current operating state of the port:</p> <ul style="list-style-type: none">■ Blocking — The bridge continues to run STP on the port, but the bridge does not receive packets from the port, learn locations of station addresses from it, or forward packets onto it.■ Listening — The bridge continues to run STP and to transmit configuration messages on the port, but it discards packets that are received on the port and does not transmit packets that are forwarded to the port.■ Learning — STP operating state in which the bridge receives packets on the port to learn the location of some of the stations that are located on the port.■ Forwarding — The bridge receives packets on the port and forwards or does not forward them, depending on address comparisons with the bridge's source address list. Provided that the link state is up, this <code>state</code> field indicates <code>forwarding</code> even if STP is disabled for the bridge.■ Disabled — Management has disabled the port or the link state is down.

Field	Description
stp	<p>Configurable status of STP on a port. If bridge-wide STP is enabled, the port STP configuration options are:</p> <ul style="list-style-type: none">■ <code>enabled</code> — STP sets the operating state of the port (blocking, listening, etc.) according to network topology characteristics. This is the default configuration for all ports.■ <code>disabled</code> — STP is disabled and the port is disabled. The port does not participate in STP decisions, frame reception, or frame transmission.■ <code>removed</code> — STP is disabled on the port but the port can still receive or transmit frames if its link state is up. <p>If bridge-wide STP is disabled, the port STP setting has no effect; as long as its link state is up, the port forwards all valid frames. To see a matrix of port and bridge STP settings, see “bridge port stpState” in this chapter.</p>
txFrames	<p>Number of frames that this port has transmitted. This object counts a frame transmitted on the interface that corresponds to this port only if the frame is for a protocol that the local bridging function is processing (includes bridge management frames).</p>

bridge port detail

Displays detailed information about bridge ports, including the Spanning Tree Protocol (STP) configurations for the bridge port.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

b po d

Important Considerations

- The port numbering that is displayed for your ports is always sequential, although it depends on the placement of the modules that you have configured into your system. See the *Implementation Guide* for your system for more information about port numbering.
- For resilient links, the main and standby ports are shown in ascending order.
- When you are prompted to select ports, specify the ? option to see a list of information about your bridge ports. This matrix is useful, for example, if you have trunks configured but forget which port is the anchor port. You must use the anchor port number to display a port summary for a trunk.

Fields in the Bridge Port Detail Display

Field	Description
designatedBridge	Identity of the designated bridge of the LAN to which the port is attached. It is an STP port parameter.
designatedCost	Cost through this port to get to the root bridge. The designated cost of the root port is the same as the cost that is received in incoming BPDUs from the designated bridge for that LAN. It is an STP port parameter.
designatedPort	Identity of the designated port on the designated bridge.
designatedRoot	Identity of the root bridge in the LAN, which includes the root bridge's priority value and the MAC address of port 1 on that bridge.
fwdTransitions	Number of times that the port has entered the forwarding state since you enabled STP or rebooted the system. This value is useful for determining the relative stability of the topology.
gvrpState (3500, 9000 L3)	Whether the GARP VLAN Registration Protocol (GVRP) is enabled or disabled on the port. For GVRP to function on the port, you must enable it on the port as well as the entire bridge. To configure GVRP on a port, see "bridge port gvrpState" in this chapter. To configure GVRP on the bridge, see "bridge gvrpState" in Chapter 9.

Field	Description
linkState	State of the link (up or down), that is, whether it is available for communication.
pathCost	Cost to add to the total path cost when this port is the root port. To configure a port's STP cost, see "bridge port stpCost" in this chapter.
portNumber	Logical index number that the system assigns to the bridge port, which may not correspond with the physical port number depending on your system configuration. (For example, when you define a trunk, only the anchor port receives a portNumber.) As you add and remove logical ports, portNumbers are reassigned so that they remain consecutive.
portId	Port identification, which includes the port priority value (first 2 digits after "0x") and the logical port number (last 2 digits). Both are shown as hexadecimal values.
priority	Configurable STP port priority value. The default value is 0x80. (0x signifies that the value to follow is a hexadecimal number.) The acceptable range is 0x0 – 0xff. To configure port priority values, see "bridge port stpPriority" in this chapter. The port priority is included in the port ID and is considered the most significant portion because it is the first factor that determines if a port is to be the designated port when more than one bridge port is attached to the same LAN. The lowest priority is chosen. If all ports in a bridge have the same priority, then the port number is used as the determining factor.
rxAllFilters (3500 and 9000 L3)	Number of frames that the bridge port discarded due to a user-defined packet filter on its "receive all" path.
rxBlockedDiscs (3500 and 9000 L3)	Number of frames that the bridge port discarded because the receiving bridge port was not in the forwarding state.
rxErrorDiscs	Number of frames that the bridge port discarded because of internal bridge system errors (such as hardware and software address table discrepancies).
rxForwards (3500 and 9000 L3)	Total number of frames (all types) that the bridge port received and forwarded to another bridge port.
rxForwardMcasts (3900, 9300, 9400, and 9000 L2)	Number of multicast frames that the bridge port received and forwarded to another bridge port.
rxForwardUcasts (3900, 9300, 9400, and 9000 L2)	Number of unicast frames the bridge port received and forwarded to another bridge port.
rxFloodUcasts	Number of unicast frames that the port received and flooded to one or more ports.

Field	Description
rxFrames	Total number of frames that this bridge port received from its segment. However, unlike the rxFrames field in the Ethernet display which counts all frames, this field does not count frames in error. Thus, this value may be lower than the value shown in the rxFrames field in the Ethernet display.
rxInternalFilters (3500 and 9000 L3)	Number of frames discarded due to customer filters on the rxInternal path.
rxMcastExcDiscs	Number of multicast frames that were discarded when rxMcastLimit was exceeded.
rxMcastExceeds	Amount of time that rxMcastLimit has been exceeded.
rxMcastFilters (3500 and 9000 L3)	Number of frames that were discarded due to a user-defined packet filter on the "receive multicast" path of this port.
rxMcastLimit	Configurable parameter that limits the rate of multicast frames that are forwarded from a bridge port. The default value is 0, which means there is no limit. To configure this option, see "bridge port multicastLimit" in this chapter.
rxMcastLimitType (3900, 9300, 9400, and 9000 L2)	Configurable parameter that selects the type of frames on which the multicast limit operates (both multicast and broadcast frames, or broadcast frames only). The default value is McastBcast. To configure this option, see "bridge port multicastLimit" in this chapter.
rxNoDestDiscs (3900, 9300, 9400, and 9000 L2)	Number of frames that this port discarded because of an unknown VLAN ID or because the port was in a non-forwarding Spanning Tree state.
rxNoRescrDiscs (3500 and 9000 L3)	Number of frames that this port discarded due to insufficient resource availability (buffer space).
rxOtherDiscs (3900, 9300, 9400, 9000 L2, 9000 L3)	Number of frames that this port discarded because they contained either invalid (group) source addresses or source addresses that belong to this bridge (indicates network loops).
rxOtherDiscards (3500)	Number of frames that this port discarded because they contained either invalid MAC source addresses or source addresses that belong to this bridge.
rxSameSegDiscs	Number of frames that this port discarded because the destination address is known on the same network segment as the source address and, thus, the frame does not need to be bridged.
rxSecurityDiscs	Number of frames that this port discarded because they contained source addresses that were statically configured on another bridge port.
SRRingNumber (3500 and 9000 L3)	(Not available at this release)

Field	Description
SRHopLimit (3500 and 9000 L3)	(Not available at this release)
state	<p>Current operating state of the port:</p> <ul style="list-style-type: none"> ■ Blocking — The bridge continues to run STP on the port, but the bridge does not receive packets from the port, learn locations of station addresses from it, or forward packets onto it. ■ Listening — The bridge continues to run STP and to transmit configuration messages on the port, but it discards packets that are received on the port and does not transmit packets that are forwarded to the port. ■ Learning — Similar to listening, but the bridge receives packets on the port to learn the location of some of the stations that are located on the port. ■ Forwarding — The bridge receives packets on the port and forwards or does not forward them, depending on address comparisons with the bridge's source address list. Provided that the link state is up, this <code>state</code> field indicates <code>forwarding</code> even if STP is disabled for the bridge. ■ Disabled — Management has disabled the port or the link state is down.
stp	<p>Configurable status of STP on a port. Provided that bridge-wide STP is enabled, the port STP configuration states function as follows:</p> <ul style="list-style-type: none"> ■ enabled — STP sets the operating state of the port (blocking, listening, etc.) according to network topology characteristics. This is the default configuration for all ports. ■ disabled — STP is disabled and the port is disabled. The port does not participate in STP decisions, frame reception, or frame transmission. ■ removed — STP is disabled on the port but the port can still receive or transmit frames if its link state is up. <p>If bridge-wide STP is disabled, this port STP setting is meaningless; as long as its link state is up, the port forwards all frames. To configure STP on a port, see "bridge port stpState" in this chapter.</p>
txAllFilters (3500 and 9000 L3)	Number of frames that the bridge port discarded because of a user-defined packet filter on its "transmit all" path.
txBlockedDiscs (3500 and 9000 L3)	Number of frames that this bridge port discarded because the transmitting bridge port was not in the forwarding state.

Field	Description
txFrames	Number of frames that this port transmitted to its segment. This object counts a frame transmitted on the interface that corresponds to this port only if the frame is for a protocol that the local bridging function is processing (includes bridge management frames).
txMcastFilters (3500 and 9000 L3)	Number of frames that this port discarded because of a user-defined packet filter on its "transmit multicast" path.
txMtuExcDiscs (3500 and 9000 L3)	Number of frames that this port discarded because of excessive size.

bridge port multicastLimit

✓ 3500
✓ 9000
✓ 9400

✓ 3900
✓ 9300

Sets a threshold value on a bridge port that affects the per-second forwarding rate of multicast or broadcast traffic that originates on the segment connected to that port.

Valid Minimum Abbreviation

b po m

Important Considerations

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- A value of zero indicates that no limit is configured.
- If you want to configure a limit for a trunk, be sure to apply it to the trunk's anchor port (lowest-numbered port) only. However, be aware that the limit that you specify applies separately to *each link* in the trunk, even though you only enter it once — that is, it is not an aggregate.
- For a larger array of similar options in the CoreBuilder 3500 and 9000 Layer 3 modules, see the Quality of Service (QoS) chapter in this guide (Chapter 22) and in the appropriate *Implementation Guide*.

Options

Prompt	Description	Possible Values	[Default]
Bridge ports	Bridge ports for which you want to set the multicastLimit	One or more valid bridge port numbers	—
Frame type (3900, 9300, 9400, 9000 L2)	Frame type to which the limit shall apply	<ul style="list-style-type: none"> ■ BcastOnly (broadcasts only) ■ McastBcast (multicasts and broadcasts) 	McastBcast
Multicast threshold value	Configurable parameter that limits the per-second receive rate of specified traffic.	<ul style="list-style-type: none"> ■ 0 – 200 (K frames/sec) (3500 and 9000 L3) ■ 0 – 200000 (frames/sec) (3900, 9300, 9400, and 9000 L2) 	0

bridge port stpState

Sets the Spanning Tree Protocol (STP) state for one or more bridge ports. The selection is effective only if STP is enabled for the system or module.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

b po stps

Important Considerations

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- The following table explains the forwarding behavior of a port based on its bridge and port STP states:

Bridge STP State	Port STP State	Port Participates in STP?	Port Forwards Frames?
Disabled	Disabled	No	Yes, if link state is up.
	Enabled	No	Yes, if link state is up.
	Removed	No	Yes, if link state is up.
Enabled	Disabled	No	No
	Enabled	Yes	Determined by STP, provided that the port link state is up.
	Removed	No	Yes, if link state is up.

Options

Prompt	Description	Possible Values	[Default]
Ports	Ports for which you want to control the STP setting	<ul style="list-style-type: none">■ One or more valid port numbers■ ? (to display a port summary)	–
STP state	Spanning Tree Protocol state that you assign to specified ports	<ul style="list-style-type: none">■ enabled■ disabled■ removed	enabled (factory default), or current value

bridge port stpCost

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Sets the path cost that the Spanning Tree Protocol (STP) adds to the root cost field in a configuration message that the port receives. The system uses this value to determine the path cost to the root through the port. The current value is shown in the pathCost field of the `bridge port detail` display.

Valid Minimum Abbreviation`b po stpc`**Important Considerations**

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- A larger path cost value makes the LAN that is reached through the port more likely to be low in the Spanning Tree topology. The lower the LAN is in the topology, the less through traffic it carries. For this reason, you may want to assign a large path cost to a LAN that has a lower bandwidth or to one on which you want to minimize traffic.
- If your configuration is successful, the previous menu appears. If the configuration is not successful, the system notifies you that your changes failed, and you can try to reenter your changes.
- See the *IEEE 802.1D MAC Bridges* standard for recommended path cost settings.

Options

Prompt	Description	Possible Values	[Default]
Bridge ports	Bridge ports for which you want to set the path cost	One or more valid bridge port numbers	–
STP cost	Configurable bridge port STP parameter that specifies the cost to add to the total path cost when this port is the root port	1 – 65535	<ul style="list-style-type: none">■ 100 (Ethernet)■ 10 (Fast Ethernet)■ 1 (Gigabit Ethernet)

bridge port
stpPriority

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Sets the Spanning Tree Protocol (STP) bridge port priority. This value influences the choice of port when the bridge has two or more ports that have the same path cost and that are connected to the same LAN, which creates a loop. STP selects the bridge port with the lowest priority and places the remaining ports in the blocking state. The current value is shown in the priority field of the `bridge port detail` display.

Valid Minimum Abbreviation

b po stpp

Important Considerations

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- Port priority is a 1-octet value written in hexadecimal format.
- If all ports in a bridge have the same priority value, then the port number is used as the determining factor.
- If your configuration is successful, the previous menu appears. If the configuration is not successful, the system notifies you that your changes failed, and you can try to reenter your changes.

Options

Prompt	Description	Possible Values	[Default]
Bridge ports	Bridge ports for which you want to set the STP port priority	A valid bridge port number	–
STP priority	One-octet value that determines which port is the designated port when there is more than one port attached to the same LAN	0x0 – 0xff, where 0x precedes a hexadecimal value	0x80

bridge port gvrpState *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Allows the port to participate in sending and receiving GARP VLAN Registration Protocol (GVRP) updates, which can help you simplify the management of IEEE 802.1Q VLAN configurations.

Valid Minimum Abbreviation

b po g

Important Considerations

- You can use this command to configure the same setting on multiple ports simultaneously. When you specify multiple port numbers, the system prompts you to choose the setting and then applies it to all of the ports.
- To activate GVRP in your system, you must enable it for the entire bridge (see “bridge gvrpState” in Chapter 9) as well as on individual bridge ports (this command). If you enable GVRP on a port but you have not enabled it for the bridge, GVRP does not function.
- To maximize the effectiveness of GVRP, it should be enabled in as many end stations and network devices as possible.
- GVRP updates are not sent to any blocked Spanning Tree Protocol (STP) ports. GVRP operates only on ports that are in the forwarding state.
- If GVRP is enabled on the bridge and on a given port which changes to the STP forwarding state, the port automatically begins to participate in GVRP.
- VLANs that are created dynamically with GVRP exist only as long as a GVRP-enabled device is sending updates — if the devices no longer send updates, GVRP is disabled, or the system is rebooted, all dynamic VLANs are removed, but the system leaves unchanged all VLANs that were statically configured through a management interface.

Options

Prompt	Description	Possible Values	[Default]
Bridge ports	Bridge ports for which you wish to enable or disable the GVRP state	<ul style="list-style-type: none">■ enable■ disable	disable

bridge port address list

Displays the MAC addresses (canonical addresses) that are currently associated with selected bridge ports, as well as the address type (static or dynamic).

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

b po a l

Important Consideration

- If you have multiple ports that are associated with a trunk, the display groups ports that are associated with a trunk on one line (for example, 3,4,6) and lists the addresses that are associated with the trunk.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface indexes (only if in allClosed mode)	Index numbers of the VLANs to which the desired bridge ports belong	<ul style="list-style-type: none">■ One or more valid VLAN indexes■ all■ ? (for a list of selectable VLANs)	–
Bridge ports	Bridge ports for which you want to display MAC addresses	<ul style="list-style-type: none">■ One or more valid bridge port numbers■ all■ ? (to display a port summary)	–

**bridge port address
add**

Adds new MAC addresses to the selected bridge ports as statically configured addresses.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

b po a a

Important Considerations

- If you have multiple ports that are associated with a trunk, the display groups ports that are associated with a trunk on one line (for example, 3, 4, 6) and lists the addresses that are associated with the trunk.
- A statically configured address is never aged out of the address table and cannot be learned on a different port. You must first remove it from its former port.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface indexes (only if in allClosed mode)	Index numbers of the VLANs to which the desired bridge ports belong	<ul style="list-style-type: none"> ■ One or more valid VLAN indexes ■ all ■ ? (for a list of selectable VLANs) 	—
Bridge ports	Bridge ports to which you want to add certain MAC addresses	<ul style="list-style-type: none"> ■ One or more valid bridge port numbers ■ ? (to display a port summary) 	—
MAC address	MAC address that you want to add to the selected port	A valid MAC address	—

bridge port address
remove

Removes individual MAC addresses from the address table.

Valid Minimum Abbreviation

b po a r

Important Consideration

- This command is typically used to remove only static MAC addresses, because the bridge could relearn a dynamic MAC address shortly after you remove it.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface indexes (only if in allClosed mode)	Index numbers of the VLANs to which the desired bridge ports belong	<ul style="list-style-type: none">■ One or more valid VLAN indexes■ all■ ? (for a list of selectable VLANs)	–
MAC address	MAC address that you want to remove	A valid MAC address	–

**bridge port address
find**

Displays the bridge port (as well as the vlan index number if the system is in allClosed mode) that is associated with a specified MAC address.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation**b po a fi****Options**

Prompt	Description	Possible Values	[Default]
MAC address	MAC address (canonical address) that you want to find on the system	A valid MAC address	—

bridge port address
flushAll

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Removes all static and dynamic MAC addresses from the bridge ports that you select. Static MAC addresses are those that you specified using the `bridge port address add` option. Dynamic MAC addresses are those that the bridge learned automatically.

Valid Minimum Abbreviation

b po a flush

Important Consideration

- If the bridge is power cycled, reset, or rebooted, the address table is automatically flushed.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface indexes (only if in allClosed mode)	Index numbers of the VLANs to which the desired bridge ports belong	<ul style="list-style-type: none">■ One or more valid VLAN indexes■ all■ ? (for a list of selectable VLANs)	–
Bridge ports	Bridge ports for which you want to remove all addresses	<ul style="list-style-type: none">■ One or more valid bridge port numbers■ ? (to display a port summary)	–

**bridge port address
flushDynamic**

Removes all dynamic MAC addresses from the bridge ports that you select. Dynamic MAC addresses are those that the bridge learned by receiving and processing packets.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

b po a flushd

Important Consideration

- If the bridge is power cycled, reset, or rebooted, the address table is automatically flushed.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface indexes (only if in allClosed mode)	Index numbers of the VLANs to which the desired bridge ports belong	<ul style="list-style-type: none">■ One or more valid VLAN indexes■ all■ ? (for a list of selectable VLANs)	–
Bridge ports	Bridge ports for which you want to remove all addresses	<ul style="list-style-type: none">■ One or more valid bridge port numbers■ ? (to display a port summary)	–

TRUNKS

You can configure a system to aggregate multiple network links into a single *trunk*. With trunking you can create high-speed point-to-point or multipoint connections without changing or replacing existing cabling. In addition, trunking provides automatic point-to-point redundancy between two devices. Redundant links normally have one link disabled by Spanning Tree (to prevent looping); trunking utilizes both links.

This chapter provides guidelines and other key information about how to configure trunking in your system.

The system treats trunked bridge ports in the same way that it treats normal individual bridge ports. Also, all higher-level network functions — including Spanning Tree algorithms, virtual LANs (VLANs), and Simple Network Management Protocol (SNMP) management — do not distinguish a trunk from any other network port. Unlike for any other network port, the system automatically distributes traffic across the ports that are associated with a trunk. If any of the trunk's ports go down or up, the system automatically redistributes traffic across the new arrangement of operational ports.




For more trunking information, see the Implementation Guide for your system.



Trunks can work with MultiPoint Link Aggregation (MPLA). MPLA is a feature for the CoreBuilder® 9400 that increases the capacity and availability of campus LAN cores without using complex, meshed router networks. Functioning at Layer 2, MPLA provides both dual-homed link resiliency and automatic load sharing over point-to-multipoint backbone connections. MPLA increases network availability using scalable Gigabit Ethernet connections among multiple campus switches. For more information about MPLA and trunking, see the CoreBuilder 9400 Implementation Guide.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



bridge trunk
autoMap summary

Displays a list of slot numbers that have been selected to support automatic backplane trunking.

3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

b t a

Important Considerations

- Automatic backplane trunking is supported only through the switch fabric modules and managed interface modules.
- The Gigabit Ethernet (GEN) Switch Fabric Module (Model Number 3CB9FG24T) has 24 non-blocking Gigabit Ethernet ports that connect to the chassis backplane to provide high-speed, low-latency connectivity between CoreBuilder 9000 interface modules.
- The GEN Switch Fabric Module supports port trunking for 12 groups, with up to six ports in a group.

Fields in the Bridge Trunk autoMap Summary Display

Field	Description
Slot number	Port numbers selected to be in the trunk.
AutoMap status	Whether the autoMap on the slot is enabled or disabled

bridge trunk
autoMap
enable/disable

3500
✓ 9000
9400

3900
9300

Dynamic backplane trunking provides automatic backplane trunking on the switch fabric modules and managed interface modules.

Valid Minimum Abbreviation

b t a e

Important Considerations

- You can enable or disable the autoMap function on slots.
- All trunking is performed through the switch fabric module.
- Do not perform backplane trunking through the interface modules.
- When you enable autoMap on a module in a specific slot, the switch fabric module verifies that the switch fabric module and the interface module’s backplane configuration support dynamic backplane mapping.
- When you disable autoMap on a module in a specific slot, the switch fabric module verifies that the interface module’s backplane configuration is compatible with that of the switch fabric module’s backplane.

Options

Prompt	Description	Possible Values	[Default]
Slot number	Slot number to choose to enable or disable autoMap function	One slot number	–

**bridge trunk
autoMap test**

Indicates what happens when you do a reset on the switch fabric module when autoMap is enabled.

3500✓ **9000****9400****3900****9300****Valid Minimum Abbreviation****b t a t****Important Consideration**

- After you enable or disable a module for automatic backplane trunking, the switch fabric module verifies that the interface module's backplane configuration is compatible or not compatible to the switch fabric's configuration.

- If autoMap is enabled on the desired interface module, then the switch fabric module verifies that the switch fabric and interface module's backplane configuration satisfies the requirements of automatic backplane trunking. If the switch fabric module and interface module satisfy requirements, then no reset is required. If not, then the switch fabric module determines if the interface module must reconfigure to support backplane configuration requirements. If the switch fabric module or interface module must reconfigure, then a message is sent to the user that a reset is required:

Fabric reset required for trunk configuration to be effective for the module in slot x.

- If autoMap is disabled, the switch fabric module verifies that the interface module's backplane configuration is compatible to the switch fabric's configuration. If the interface module is compatible to the switch fabric module's backplane configuration, then no reset is required. If not, then the switch fabric module determines if the interface module can support the defined switch fabric module's backplane configuration.

If the interface module can, then a message is sent to the interface module with the desired backplane port configuration. Otherwise, the switch fabric module will modify its configuration and a message is sent to the user that a reset is required:

Fabric reset required for trunk configuration to be effective for the module in slot x.

**bridge trunk
summary**

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Displays summary information about configured trunks on your system. In a summary report, the system displays the trunk name and index number, the ports defined in that trunk, whether the Trunk Control Message Protocol (TCMP) is enabled or disabled, and whether the port link is up or down.

Valid Minimum Abbreviation**b t s****Fields in the Bridge Trunk Summary Display**

Field	Description
Index	Identifying number that the system assigned to the trunk. You can select all or one trunk.
Name	Trunk name that you defined.
Ports	Port numbers in the trunk.
State	Whether the trunk is up or down
TCMP	Whether the Trunk Control Message Protocol (TCMP) is enabled or disabled.

bridge trunk detail

Displays detailed trunk information in addition to the summary information.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

b t det

- ✓ 3900
- ✓ 9300

Fields in the Bridge Trunk Detail Display

Field	Description
FlowC	For Gigabit Ethernet trunks, the flow control setting (on, off, rxOn, txOn). For other media types, the field contains n/a to indicate that flow control does not apply.
Index	Identifying number for the system or module that the system assigned to the trunk.
Missing	Number of ports that are configured for the trunk, but are missing because an interface module is inaccessible.
Mode	Operating mode: 100half or 100full for Fast Ethernet and 1000full for Gigabit Ethernet.
Name	Trunk name that you defined.
Node trunk id	TCMP identifier that the system assigned to the trunk.
Node trunk id list	Node trunk identifications that each port has detected on the trunk.
Ports	Ports in the trunk. The second half of the display lists each individual port in each trunk.
Present	Number of ports that participate in the trunk.
rxBadType	Number of TCMP messages received that contain a bad TCMP Type field.
rxBadVersion	Number of TCMP messages received that contain a bad TCMP version number.
rxFrames	Number of TCMP messages that were received on each port.
rxHellos	Number of TCMP helloMessages that were received on each port.
rxOverflow	Number of times that TCMP has detected a TCMP trunk configuration that exceeds the eight-node maximum.
rxSameTrunkId	Number of times that TCMP has received a helloMessage that contains the TCMP agent's own Node trunk id (an illegal configuration).
Selected node trunk id list	Node trunk identifications that are selected for use on the trunk.
State	Whether the trunk is up or down.

Field	Description
TCMP	Whether TCMP is enabled or disabled for the trunk.
Tcmpstate	TCMP state for each port in the trunk: <ul style="list-style-type: none">■ <code>notInUse</code> — Not selected for use in the trunk■ <code>selected</code> — Selected for use in the trunk, but not yet active in the trunk■ <code>inUse</code> — Active in the trunk
Trunk state	State (<code>up</code> or <code>down</code>) of each port link in the trunk.
txFrames	Number of TCMP messages that were transmitted on each port.
txHellos	Number of TCMP helloMessages that were transmitted on each port.
Type	The network type that you assigned to the ports in the trunk (for example, <code>FDDI</code> , <code>Fast_Ethernet</code> , or <code>Gigabit_Ethernet</code>).

bridge trunk define

Defines one or more trunks on the system. When you define a trunk, you specify ports and characteristics for the trunk.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation**b t def**

✓ 3900

✓ 9300

Important Considerations

- If you have more than one media type on your system (for example, Fiber Distributed Data Interface (FDDI), Fast Ethernet, and Gigabit Ethernet), you are prompted for a media type before you are prompted for the trunk information.
- All links to be trunked must be homogeneous. For example, you cannot mix Fast Ethernet and Gigabit Ethernet links in a trunk.
- If you have already defined other trunks on your system, you cannot select ports that are part of an existing trunk.
- In general, create trunks before you define your virtual LANs (VLANs). If you create a trunk whose ports are part of existing VLANs, the VLAN bridge port configuration changes. For example, if you have the default VLAN as well as IP VLANs and you then define a trunk with ports in one of the IP VLANs, the system removes those ports from that VLAN and places them in the default VLAN. You must modify the VLAN and add the new bridge ports to the appropriate VLAN. This situation does not apply if you have only the default VLAN (all ports are part of the default VLAN).
- When you define a VLAN to include trunk ports, specify the anchor port (lowest-numbered port) that is associated with the trunk.
- Do not use Gigabit Ethernet (GEN) Interface modules (such as the 2-port 1000BASE-SX Gigabit Ethernet (GEN) Interface Module) when defining trunks.
- Enter ? to see the port summary (for example, to indicate whether there are ports associated with FDDI Dual Attach Station (DAS) pairs), and then enter the appropriate port numbers. To specify an FDDI DAS pair, specify the lowest-numbered port in the DAS pair.
- The number of trunk groups and the number of ports within a trunk group depend on your system. See the Options table.

The 3CB9FG24T switch fabric module supports up to 12 trunk groups on the CoreBuilder 9000.

- If you are working with Gigabit Ethernet modules in a SuperStack II Switch, keep in mind that each Gigabit Ethernet module uses an internal trunk resource towards the limit of four. You can trunk Gigabit Ethernet modules together (each with one port) to consolidate the Gigabit trunk resources. If you have four trunks defined and you add a Gigabit Ethernet module to the system, after a boot, the system reports that the configuration is incompatible. You must delete one of the existing trunks.
- You must reboot the module at the end of the trunk definition process. (You can define multiple trunks in one `bridge trunk define` operation.) On the CoreBuilder® 9000, rebooting a module returns you to the EME prompt, which requires you to reconnect to the module.
- The following considerations apply to the trunk clustering function, MultiPoint Link Aggregation (MPLA), in the CoreBuilder 9400 system:
 - When you configure a new switch, define multipoint aggregated links and reboot the system before you define other trunks and VLANs on the switch.
 - On a reboot, existing trunks and VLANs are deleted, and the default VLAN is restored. Trunked ports that were part of a VLAN before reboot are moved to the default VLAN. You must then redefine trunks or VLANs that you want to continue to use. See Chapter 11.

Options

Prompt	Description	Possible Values	[Default]
Mac type (if you have more than one)	Media type for the trunk.	Depends on your configuration: <ul style="list-style-type: none"> ■ FDDI ■ Fast_Ethernet ■ Gigabit_Ethernet ■ 10/100BASE-TX ■ 100BASE-FX ■ 1000BASE-SX ■ ? (for a list of selectable media types) 	–

Prompt	Description	Possible Values	[Default]
Ports	Total number of the bridge ports that you want to be part of the trunk.	<div>9000:</div> <ul style="list-style-type: none"> Layer 2 modules support up to 4 trunk groups with up to 6 ports per trunk Layer 3 modules support up to 3 trunk groups with up to 6 ports per trunk <div><i>The 6-port SAS (3-port DAS) FDDI Layer 3 supports 3 trunk groups. In SAS mode the trunks can contain up to 6 ports. In DAS mode, the trunks can contain up to 3 ports.</i></div> <ul style="list-style-type: none"> The FGA24 switch fabric module supports up to 4 trunk groups with up to 6 ports per trunk and the FGA24T and GA9 switch fabric modules support up to 12 trunk groups with up to 6 ports per trunk <div><i>(The GA9 cannot support 12 trunk groups because there are not enough ports on this module.)</i></div> <div>3500:</div> <ul style="list-style-type: none"> Supports up to 4 trunk groups with up to 8 ports per trunk <div>3900:</div> <ul style="list-style-type: none"> Supports up to 4 trunk groups with up to 6 ports per trunk <div>9300 and 9400:</div> <ul style="list-style-type: none"> Supports up to 12 trunk groups with up to 6 ports per trunk all ? (for a list of selectable ports) 	—

Prompt	Description	Possible Values	[Default]
Mode	Operating mode for the trunk.	<ul style="list-style-type: none"> ■ 100half, 100full (for Fast Ethernet) ■ 10half, 10full, 100half, 100full (for platforms that support 10 Mbps Ethernet) 	–
Flow control	Flow control setting (Ethernet only)	<ul style="list-style-type: none"> ■ on ■ off ■ rxOn (Gigabit Ethernet) ■ txOn (Gigabit Ethernet) 	off
Trunk name	Name of the trunk. Use quotation marks (") around any string with embedded spaces.	<ul style="list-style-type: none"> ■ Maximum 32 alphanumeric characters ■ ? (for a list of selectable names) 	–
TCMP	Trunk Control Message Protocol (TCMP). Performs the following functions: <ul style="list-style-type: none"> ■ Detects and corrects trunks that violate trunk configuration rules ■ Ensures orderly activation and deactivation of trunk ports 	<ul style="list-style-type: none"> ■ enabled ■ disabled 	enabled

Procedure

You can define all trunks in one `bridge trunk define` operation and then reboot. At the end of each trunk definition, the system prompts you to define another trunk.

- 1 If you have more than one media type on your system, enter the media type (for example, `Fast_Ethernet` OR `Gigabit_Ethernet`).
- 2 Enter the ports that you want to be part of the trunk.
To get information about the selectable ports, enter `?`
- 3 Enter the correct operating mode for 10/100 Ethernet ports.
If you are configuring a Fast Ethernet trunk, select the Fast Ethernet port mode (`100half`, `100full`).
- 4 For an Ethernet trunk, enter the flow control setting (`on`, `off`, `rxOn`, OR `txOn`).

- 5 Enter the trunk name, or to get information about specifying the trunk name, enter **?**
- 6 Specify whether TCMP is enabled or disabled.
The system indicates that the trunk definition is complete and allows you to define additional trunks until you reach the system trunk limit.
- 7 At the system prompt, to define another trunk enter **y** (yes) or to end the trunk sequence, enter **n** (no).

You must then reboot to enable the trunks to take effect.

Bridge Trunk Define Example (9000)

The example shows a define operation that creates two trunks.

```
Select menu option: bridge trunk define
Select mac type {Fast_Ethernet,Gigabit_Ethernet|?}: Fast_Ethernet
Select ports (14-19|all|?): 14-18
Enter mode {100half,100full|?}: 100full
Enter trunk name {?} []: "Trunk_1"
Enter TCMP state (disabled,enabled) [enabled]: enabled
Trunk definition complete
```

```
Define another trunk? (y,n) [n]: y
Select mac type {Fast_Ethernet,Gigabit_Ethernet|?}: Fast_Ethernet
Select up to 8 ports (1-12|?): 1-6
Enter trunk name {?} []: Trunk2
Enter TCMP state (disabled,enabled) [enabled]: enabled
Trunk definition complete
```

The configuration of the ports will be modified.

The system must be rebooted to complete trunk configuration.
This may take a few minutes.

```
Are you sure you want to reboot the system? (n,y) [y]: y
```

bridge trunk modify

Changes a trunk in either of two ways:

- ✓ 3500
- ✓ 9000
- ✓ 9400

- Modifies a trunk's characteristics (for example, a Fast Ethernet operating mode or the Trunk Control Message Protocol (TCMP) state).
- Adds or removes a port from the trunk, as long as you maintain at least one of the original ports in the trunk.

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

b t m

Important Considerations

- Keep at least one port that you defined in the original trunk. To completely redefine a trunk configuration, remove the trunk and define a new one.
- You cannot modify, add, or remove ports that are part of different trunks from the trunk that you are modifying.
- To avoid configuration errors, do not modify Fiber Distributed Data Interface (FDDI) station mode port pairs when any of the ports in the pair are members of a trunk.
- If you have more than one media type on your system (for example, Fast Ethernet and Gigabit Ethernet), you are prompted for a media type before you are prompted for the trunk information.
- Any changes that you make to the trunk's characteristics take effect immediately and do not interrupt trunk operations. If you add or remove a port, however, you must reboot the system to implement the change.

- In an FDDI trunk:
 - You cannot modify FDDI station mode port pairs when any of the ports in the pair are in a trunk.
 - When you modify the station mode, any FDDI ports that are associated with virtual LANs (VLANs) or a trunk are removed from the VLAN or trunk.
- Within a trunk, you cannot change certain port characteristics, such as FDDI station mode. For example, in an FDDI trunk, you cannot change a trunked DAS (dual attach station) port to an SAS (single attach station) port or an SAS port to a DAS port.
- If you change an FDDI port pair from SAS to DAS, you select the pair using just the lower of the two port numbers, just as with a trunk anchor port.

Options

Prompt	Description	Possible Values	[Default]
Trunk index	Index number of the trunk that you want to modify	<ul style="list-style-type: none">■ One or more trunks■ ? (for a list of selectable trunk indexes)	–

Prompt	Description	Possible Values	[Default]
Ports	Total number of the bridge ports that you want to be part of the trunk	<p>9000:</p> <ul style="list-style-type: none"> ■ Layer 2 modules support up to 4 trunk groups with up to 6 ports per trunk ■ Layer 3 modules support up to 3 trunk groups with up to 6 ports per trunk <p><i>The 6-port SAS (3-port DAS) FDDI Layer 3 supports 3 trunk groups. In SAS mode the trunks can contain up to 6 ports. In DAS mode, the trunks can contain up to 3 ports.</i></p> <ul style="list-style-type: none"> ■ The FGA24 switch fabric module supports up to 4 trunk groups with up to 6 ports per trunk and the FGA24T switch fabric module supports up to 12 trunk groups with up to 6 ports per trunk <p><i>(The GA9 cannot support 12 trunk groups because there are not enough ports on this module.)</i></p> <p>3500:</p> <ul style="list-style-type: none"> ■ Supports up to 4 trunk groups with up to 8 ports per trunk <p>3900:</p> <ul style="list-style-type: none"> ■ Supports up to 4 trunk groups with up to 6 ports per trunk <p>9300 and 9400:</p> <ul style="list-style-type: none"> ■ Supports up to 12 trunk groups with up to 6 ports per trunk ■ all 	Currently configured ports

Prompt	Description	Possible Values	[Default]
Mode	Operating mode for a 10/100 Ethernet trunk	<ul style="list-style-type: none">100half, 100full (for Fast Ethernet)10half, 10full, 100half, 100full (for platforms that support 10 Mbps Ethernet)	Current mode
Flow control (Gigabit Ethernet only)	Flow control setting for a Gigabit Ethernet trunk	<ul style="list-style-type: none">onoffrxOntxOn	Current value
Trunk name	Name of the trunk. Use quotation marks (") around any string with embedded spaces.	<ul style="list-style-type: none">Maximum 32 alphanumeric characters? (for a list of selectable trunk names)	Current trunk name
TCMP	Trunk Control Message Protocol (TCMP). Performs the following functions: <ul style="list-style-type: none">Detects and corrects trunks that violate trunk configuration rulesEnsures orderly activation and deactivation of trunk ports	<ul style="list-style-type: none">enableddisabled	enabled (factory default), or current TCMP state

Procedure

To modify trunk information for a bridge, follow these steps:

- 1** Enter the trunk index number, or to display the selectable trunks, enter `?`
The system shows the media type for the trunk (for example, `Fast Ethernet, Gigabit Ethernet, Or FDDI`).
- 2** At the prompt, enter the ports that you want to be part of the trunk, or to display a port summary, enter `?`
The maximum number of ports per trunk is 8 (for the CoreBuilder 3500 and the CoreBuilder 9000 Layer 3 modules).
- 3** To change the 10/100 operating mode, enter the new operating mode, or to display information about the selectable values, enter `?`
For Fast Ethernet, you can select 100 Mbps, running in half-duplex or full-duplex mode. All ports in the trunk are set to the specified operating mode.
- 4** To change the flow control setting for a Gigabit Ethernet trunk only, enter a new flow control setting
- 5** To change the name of the trunk, enter the new name, or to view information on how to specify a trunk name, enter `?`
The name can have up to 32 characters. Use quotation marks around any character string that has embedded spaces.
- 6** Enter the TCMP state. The system default is `enabled`.
If you modified the port information, the system displays a message to inform you that the port configuration will change and then displays a reboot prompt.
- 7** At the system prompt, to reboot the system, enter `y` (yes) and implement the new trunk information, or to return to the previous menu, enter `n` (no).
Entering `n` (no) cancels the trunk changes. The system reports that it is unable to continue with the trunk configuration.

Bridge Trunk Modify Example (9000)

Select menu option: **bridge trunk modify**

Select trunk index {1-3|?}: ?

Selectable trunks

selection	ports	name
1	7,8,12	trunk1
2	1,2,4,19	trunk2
3	3,14,17,18	trunk3

Select trunk index {1-3|?}: **2**

Fast Ethernet

Select ports (1,2,5,6,15,16,19|all|?) [1,2,19]: **1,2**

Enter trunk name {?} [trunk2]:

Enter TCMP state (disabled,enabled) [enabled]:

The configuration of the ports will be modified.

Are you sure you want to reboot the system? (n,y) [y]:

bridge trunk remove

Removes a previously defined trunk. You can remove one or more trunks with this command.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation**b t r****Important Considerations**

- The number of trunk groups and the number of ports within a trunk group depend on your system. See the Options table. However, because each Gigabit Ethernet module uses an internal trunk resource towards the limit of four (Gigabit Ethernet only), keep in mind how many trunk resources you have when you remove trunks. For example, if you have a trunk with two Gigabit Ethernet ports (which consolidates two Gigabit trunk resources into one) as well as three other trunks, and you then try to remove the Gigabit Ethernet trunk, you will exceed the trunk resource limit. (The Gigabit Ethernet ports use two trunk resources.) The system reports that it is unable to remove the trunk because the trunk resource limit would be exceeded.
- Removing a trunk requires a module reboot. For CoreBuilder 9000 modules, rebooting a module returns you to the EME prompt, which requires you to reconnect to the module.

Options

Prompt	Description	Possible Values	[Default]
Trunk index	Index number of the trunk that you want to remove	<ul style="list-style-type: none">■ One or more valid trunk index number■ all■ ? (for a list of selectable trunk indexes)	–

Bridge Trunk Remove Example (9000)

```
Select menu option: bridge trunk remove  
CB9000@slot10.1 [12-E/FEN-TX-L3] (bridge/trunk) remove  
Select trunk index(s) {1-2|all|?}: 2
```

The configuration of the ports will be modified.

The module must be rebooted to complete trunk configuration.
This may take a few minutes.

Are you sure you want to reboot the system? (n,y) [y]:**y**

12

MULTIPOINT LINK AGGREGATION (MPLA)

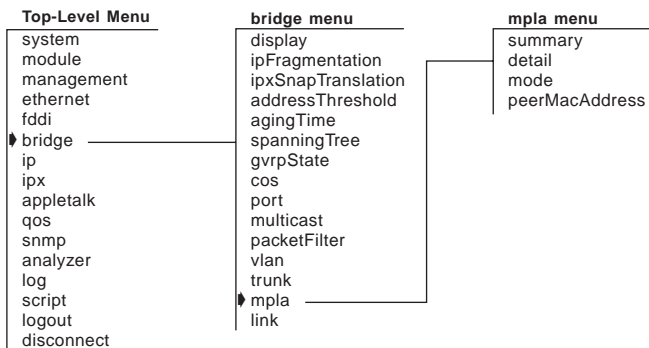
MultiPoint Link Aggregation (MPLA) increases the capacity and availability of campus LAN cores without using complex, meshed router networks. Functioning at Layer 2, MPLA provides both dual-homed link resiliency and automatic load sharing over point-to-multipoint backbone connections. MPLA increases network availability using scalable Gigabit Ethernet connections among multiple campus switches.



For more information about MPLA and trunking, see the CoreBuilder 9400 Implementation Guide.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



3900
9300

Fields in the Bridge MPLA Summary Display

Field	Description
Mode	Whether MPLA is enabled on the switch. Possible values are <code>enabled</code> and <code>disabled</code> . The default is <code>disabled</code> .
Peer Switch Interface State	The state (<code>up</code> or <code>down</code>) of the out-of-band management port on the other (peer) switch in the MPLA core.

bridge mpla detail Displays the trunk state and node trunk IDs for the switch ports.

3500 **Valid Minimum Abbreviation**

9000 b mp d

✓ **9400**

Fields in the Bridge MPLA Detail Display

Field	Description
Mode	Whether MPLA is enabled on the switch. Possible values are enabled and disabled. The default is disabled.

bridge mpla mode

Enables or disables the MultiPoint Link Aggregation feature on the switch.

3500

9000

✓ 9400

3900

9300

Valid Minimum Abbreviation

b mp m

Important Considerations

- Use only CoreBuilder 9400 systems as MPLA core switches.
The core of a multipoint aggregated link must contain two 9400 switches, whose out-of-band management ports also must be directly connected.
- Use only Switch 3900 devices as edge switches.
Each MPLA edge switch must have at least one physical link to each core switch. Multiple trunked links may connect an edge and core switch, for added bandwidth
- Use only Gigabit Ethernet links between MPLA core switches and edge switches.
- All links from an edge switch to the MPLA core switches must be aggregated (trunked) at the edge switch.
- While the Trunk Control Message Protocol (TCMP) is optional in point-to-point trunks, you must configure it to run on all of the point-to-multipoint links between MPLA edge switches and core switches.
- You can enable these features in MPLA *edge* switches, but not in MPLA *core* switches:
 - Spanning Tree
 - IGMP snooping
 - Resilient links
 - Roving analysis port
- When you configure a new MPLA core switch, define MPLA configurations and reboot the switch before you define other trunks and VLANs on the switch.

Procedure

- 1** To enable MultiPoint Link Aggregation on the switch, use the `bridge mpla mode enable` command.

To disable MultiPoint Link Aggregation on the switch, use the `bridge mpla mode disable` command.
- 2** Select the ports that you want to be part of the multipoint aggregated link using the `bridge trunk define` command, as described in Chapter 11.
- 3** Reboot the switch to implement the multipoint aggregated link selection.

On reboot, existing trunks and VLANs are deleted, and the default VLAN is restored.

bridge mpla peerMacAddress

Specifies the MAC address of the out-of-band management port of the *attached* CoreBuilder 9400 switch in the MPLA core (the *peer* core switch).

3500

9000

✓ 9400

3900

9300

Valid Minimum Abbreviation

b mp p

Important Considerations

- You execute this command on each of the two CoreBuilder 9400 switches in the MPLA core.
- In each core switch, you use this command to specify the MAC address for the out-of-band management port of the *attached* peer switch.



You must use this management port to connect a crossover Ethernet cable to the out-of-band management port on the switch you are presently configuring.

- The input format for this MAC address is 00-00-00-00-00-00

RESILIENT LINKS

Resilient links protect your network against the failure of an individual link or device by providing a secondary backup link that is inactive until it is needed.

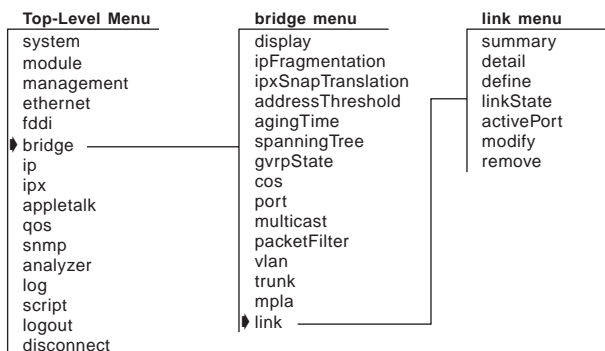
This chapter provides guidelines and other key information about how to configure resilient links in your system.



For more information about resilient links, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



bridge link summary***For CoreBuilder 9000: Applies to Layer 2 switching modules only.*****3500****✓ 9000****✓ 9400****✓ 3900****✓ 9300**

Displays summary information about configured resilient links on your system. In a summary report, the system displays the index number, link name, and whether the link is up or down.

Valid Minimum Abbreviation**b l s****Fields in the Bridge Link Summary Display**

Field	Description
Index	Number that the system assigned to the resilient link pair. You can select all resilient link pairs or one resilient link pair.
Name	Name of the defined resilient link pair.
State	Whether the resilient link pair is up or down.

bridge link detail

For CoreBuilder 9000: Applies to Layer 2 switching modules only.

- 3500
- ✓ 9000
- ✓ 9400
-
- ✓ 3900
- ✓ 9300

Displays detailed link information in addition to the summary information.

Valid Minimum Abbreviation

b 1 det

Fields in the Bridge Link Detail Display

Field	Description
Active Port	Port that carries network traffic
Enable State	Whether the resilient pair is enabled or disabled
Index	Number that the system assigned to the resilient link pair. You can select all resilient link pairs or one resilient link pair.
Main Link	Link state (up or down) of the main link
Main Port	Main resilient link port
Name	Name of the defined link
Standby Link	Link state (up or down) of the standby link
Standby Port	Standby resilient link port to which traffic shifts if the main resilient link port fails
State	Whether the resilient link pair is up or down

bridge link define***For CoreBuilder 9000: Applies to Layer 2 switching modules only.***

Defines one or more links on the system. When you define a link, you specify ports and characteristics for the link.

Valid Minimum Abbreviation**b 1 def****Important Considerations**

- Connect the network cable to the resilient link ports after you reboot the system; failure to do so may create a bridge loop in your network.
- In general, create resilient links before you define your virtual LANs (VLANs). If you plan to create resilient links for part of a VLAN, create the resilient links before you create the VLAN.
- When you create a resilient link that includes ports that are part of a VLAN, those ports are removed from the VLAN. You must modify the VLAN and add the new bridge port to the appropriate VLAN. This situation does not apply to the default VLAN (all ports are part of the default VLAN).
- If you have already defined other resilient links or trunks on your system, you cannot select ports that are part of an existing resilient link pair or a trunk.
- You must reboot the system at the end of the link definition process. (You can define multiple links in one `define` operation.)
- The resilient link port pair uses a single MAC address for frames sourced by this pair.
- The resilient link name can be up to 32 alphanumeric characters.

3500**✓ 9000****✓ 9400****✓ 3900****✓ 9300**

Options

Prompt	Description	Possible Values	[Default]
Resilient link name	Name of the link. Use quotation marks around any character string that contains spaces	Maximum 32 alphanumeric characters	–
Main Port	Main port that you want to be part of the link.	Any of the available ports on the system	–
Standby Port	Standby port that you want to be part of the link.	Any of the available ports on the system	–
Define another link?	Whether you want to define another link.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n
Reboot the system?	Resilient links that you define do not take effect until you reboot the system.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y

Procedure

You can define all links in one `bridge link define` operation and then reboot. At the end of each link definition, the system prompts you to define another link.

- 1 Enter the link name, or to get information about specifying the link name, enter `?`
- 2 Select the port that you want to be the main port.
To get information about the selectable ports, enter `?`
- 3 Select the port that you want to be the standby port.
To get information about the selectable ports, enter `?`
- 4 At the system prompt, to define another link enter `y` (yes), or to end the link define sequence, enter `n` (no).

You must reboot for the links to take effect.

bridge link linkState***For CoreBuilder 9000: Applies to Layer 2 switching modules only.***

Sets the linkState value (enabled or disabled) for a specific resilient link.

3500**✓ 9000****✓ 9400****✓ 3900****✓ 9300****Valid Minimum Abbreviation****b 1 1****Important Considerations**

- When the `bridge link linkState` option is enabled, the resilient link transmits or receives frames.
- When the `bridge link linkState` option is disabled, the resilient link no longer transmits or receives frames.

Options

Prompt	Description	Possible Values	[Default]
Resilient link index	Index number of the resilient link that you want to modify	<ul style="list-style-type: none"> ■ Depends on configured links ■ ? (for a list of selectable link indexes) ■ A valid link number ■ all 	–
linkState value	Whether you want the resilient link for the selected link index to transmit and receive frames	<ul style="list-style-type: none"> ■ enable ■ disable 	–

bridge link activePort

3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

For CoreBuilder 9000: Applies to Layer 2 switching modules only.

Sets either the main port or the standby port as the active port. The active port carries the network traffic.

Valid Minimum Abbreviation

b l a

Options

Prompt	Description	Possible Values	[Default]
Resilient link index	Index number of the link whose active port you want to set	<div><div>■</div> Depends on configured links</div> <div><div>■</div> ? (for a list of selectable link indexes)</div> <div><div>■</div> A valid link number</div> <div><div>■</div> all</div>	–
Active port state	Port that you want to carry network traffic	<div><div>■</div> main</div> <div><div>■</div> standby</div>	–

bridge link modify***For CoreBuilder 9000: Applies to Layer 2 switching modules only.***

Modifies the link name, as well as the main port and standby port, of a defined resilient link.

3500**✓ 9000****✓ 9400****✓ 3900****✓ 9300****Valid Minimum Abbreviation****b l m****Important Considerations**

- Connect the network cable to the resilient link port after you reboot the system.
- In general, create links before you define your Virtual LANs (VLANs). If you plan to create resilient links for part of a VLAN, create the resilient links before you create the VLAN.
- When you create a resilient link that includes ports that are part of a VLAN, those ports are removed from the VLAN. You must modify the VLAN and add the new bridge port to the appropriate VLAN. This situation does not apply to the default VLAN (all ports are part of the default VLAN).
- If you have already defined other links or trunks on your system, you cannot select ports that are part of an existing link or a trunk.
- You must reboot the system at the end of the link definition process. (You can define multiple links in one `define` operation.)
- The resilient link port pair uses a single MAC address for frames sourced by this pair.
- The resilient link name can be up to 32 alphanumeric characters.

Options

Prompt	Description	Possible Values	[Default]
Resilient link name	New resilient link name. Use quotation marks around any character string that has embedded spaces.	Maximum 32 alphanumeric characters	–
Main port	New port to be the main port of the defined resilient link.	Any of the available ports on the system	–
Standby port	New port to be the standby port of the defined resilient link.	Any of the available ports on the system	–
Define another link?	Whether you want to define another link.	■ y (yes) ■ n (no)	n
Reboot the system?	Resilient links that you define do not take effect until you reboot the system.	■ y (yes) ■ n (no)	y

bridge link remove***For CoreBuilder 9000: Applies to Layer 2 switching modules only.***

Removes a previously defined resilient link pair. You can remove one or more resilient link pairs with this command.

3500**✓ 9000****✓ 9400****✓ 3900****✓ 9300****Valid Minimum Abbreviation****b l r****Important Consideration**

- Removing a link requires that you reboot the system.

Options

Prompt	Description	Possible Values	[Default]
Resilient link index	Index number of the resilient link pair that you want to remove	<ul style="list-style-type: none">■ Depends on configured links■ ? (for a list of selectable link indexes)	–

Bridge Link Remove Example

Select menu option: **bridge link remove**

Select link index(s) (1-2|all|?): **2**

The configuration of the ports will be modified.

The system must be rebooted to complete resilient link configuration.

This may take a few minutes.

Are you sure you want to reboot the system? (n,y) [y]: **y**

14

VIRTUAL LANs (VLANs)

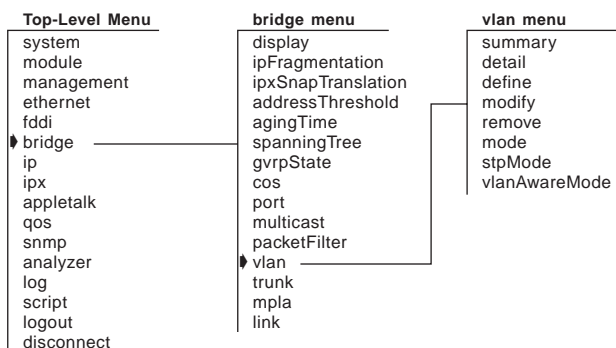
A virtual LAN (VLAN) is a logical definition of a network work group. It is roughly equivalent to a broadcast domain. A *VLAN interface* is your system's point of attachment to a given VLAN. A VLAN and a VLAN interface are analogous to an IP subnetwork and an IP interface.



For more information about VLANs, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



The `bridge vlan stpMode` command is available only when you enable `allClosed` mode on the CoreBuilder® 3500 system or the CoreBuilder 9000 Layer 3 switching modules. The command does not appear when you are using the default VLAN mode (`allOpen`) or when you use `allClosed` mode on a Layer 2 system or module.

bridge vlan summary

Displays a summary of VLAN information. In a summary report, the system displays the ports and protocols that are assigned to each VLAN.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviations

b v s (in allOpen mode on Layer 2 or Layer 3 switches and modules)

b v su (in allClosed mode on Layer 3 switches and modules)

Important Considerations

- The summary display lists the physical ports that are associated with each VLAN interface. It does not indicate bridge port characteristics (for example, trunked ports). See “bridge vlan detail” next for this information.
- The VLAN mode (shown in the Type field) affects VLANs as follows:
 - For the CoreBuilder 3500, CoreBuilder 9400, SuperStack® II Switch 3900, and SuperStack II Switch 9300, the VLAN mode affects all configured VLANs on the system. For the entire system, the default VLAN mode is `allOpen`.
 - For the CoreBuilder 9000, the VLAN mode affects all VLANs that are associated with a particular module (the switch fabric module, or all VLANs on a switching module). For each module, the default VLAN mode is `allOpen`.
- As of Release 3.0.0, the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules support one of three origins for a VLAN:
 - If you explicitly create the VLAN with a `bridge vlan define` operation, the origin of the VLAN is `static`.
 - If you create a router port IP interface (for which the system automatically creates a router port VLAN), the origin of the router port VLAN is `router`.
 - If you enable dynamic port-based VLAN configuration via the GARP VLAN Registration Protocol (GVRP), the origin is `GVRP`.

- GVRP is based on IEEE 802.1Q and allows for dynamic configuration of port-based VLANs. GVRP can help you simplify the management of VLAN configurations in larger networks. Use the command `bridge port gvrpState` to explicitly enable GVRP on the participating bridge ports *and* use the command `bridge gvrpState` to enable the bridge GVRP state for the entire system. The bridge GVRP state enables you to control GVRP on the system without losing the per-port GVRP state. By default, the GVRP state for the entire system is `disabled` and the GVRP state for each bridge port is `disabled`.
- The system prompts you for a VLAN interface index number before it displays the summary information.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index numbers of the VLAN interfaces for which you want summary information	<ul style="list-style-type: none"> ■ One or more selectable VLAN interface index numbers ■ all ■ ? (for a list of selectable indexes) 	1 (if you have only one VLAN)

Fields in the Bridge VLAN Summary Display

Field	Description
Index	System-assigned index number that identifies a VLAN of the identified origin. Statistics appear in the display for the VLAN that you specify.
Name	Character string of from 0 through 32 bytes that identifies the VLAN. The default VLAN always uses the name <code>Default</code> .

Field	Description
Origin	<p>For all Layer 2 systems or switching modules, the VLAN origin is always <code>static</code>, which indicates that the user created the VLAN. For the CoreBuilder® 3500 or CoreBuilder 9000, the origin indicates one of the following:</p> <ul style="list-style-type: none"> ■ <code>static</code> — The VLAN was created statically (user-configured by using the <code>bridge vlan define</code> command). ■ <code>router</code> — The VLAN was created automatically by a router port IP interface (of router origin). You create a router port IP interface using the <code>ip interface define</code> command with the interface type <code>port</code>. You cannot modify or remove a router port VLAN. ■ <code>GVRP</code> — The VLAN was created dynamically from a GVRP update (<code>GVRP</code>). You must enable the GVRP state for the entire system as a bridge-wide parameter <i>and</i> for the participating bridge ports as a bridge-port parameter.
Ports	<p>Index numbers of the bridge ports that belong to the VLAN, or the bridge port that belongs to the router port IP interface.</p> <p>On the CoreBuilder 9000, the list of ports includes the front-panel ports and the appropriate backplane ports. Example: On a 12-port Layer 3 module, the list of ports includes ports 1 – 12 and port 13, which is the module's backplane port.</p>
Type (VLAN mode)	<p>Either <code>allOpen</code> or <code>allClosed</code>. VLANs in <code>allOpen</code> mode share a single address table for all configured VLANs; in <code>allClosed</code> mode, each VLAN has its own unique address table. Standard bridging rules apply based on the table addresses that are assigned to the specific VLAN. A router port IP interface requires that you put the system in <code>allClosed</code> mode.</p>
VID	<p>Unique, user-defined integer (VLAN ID) that identifies this VLAN. It is used by management operations. You can assign or modify a VID that is associated with a static VLAN; you cannot modify the VID selected automatically after you define a router port IP interface, nor can you change the VID of the default VLAN. The default VLAN requires a VID of 1.</p>
VLAN Aware Mode (3500 and 9000 Layer 3)	<p>Whether the VLAN aware mode (tagging mode) is <code>allPorts</code> or <code>taggedVlanPorts</code>. The default for CoreBuilder 3500 Release 2.0 or later is <code>allPorts</code>; <code>allPorts</code> is also the default as of CoreBuilder 9000 software Release 3.0. The value <code>taggedVlanPorts</code> is a compatibility mode for VLANs configured prior to CoreBuilder 3500 Release 2.0 and for VLANs configured on CoreBuilder 9000 Layer 3 modules prior to CoreBuilder 9000 Release 3.0.</p>

bridge vlan detail

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Displays per-port information such as tagging in addition to the VLAN summary information. For the CoreBuilder 3500 and the CoreBuilder 9000 Layer 3 switching modules, this command also displays VLAN statistics.

Valid Minimum Abbreviation`b v det`**Important Considerations**

- The default VLAN always uses VLAN ID (VID) 1 and the name `Default`. For Layer 3 systems and modules, it also uses the protocol type `unspecified`.
- The VLAN ID (VID) is used as the 802.1Q tag if tagging is enabled for a port.
- The VLAN statistics for the CoreBuilder 3500 and CoreBuilder 9000 Layer 3 switching modules are valid *only* under one of the following conditions:
 - If the VLANs are defined for the same protocol type (or for the type called `unspecified`) and do not share any ports. Example: IP VLAN1 has ports 1 through 6 and IP VLAN2 has ports 7 through 12.
 - If the VLANs are explicitly defined for different protocol types. In this case, the VLANs may share ports. Example: An IP VLAN and an IPX VLAN both use ports 2 through 4.
- As of Release 3.0.0, the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules support one of three origins for a VLAN:
 - If you explicitly create the VLAN with a `bridge vlan define` operation, the origin of the VLAN is `static`.
 - If you create a router port IP interface (for which the system automatically creates a router port VLAN), the origin of the router port VLAN is `router`.
 - If you enable dynamic port-based VLAN configuration via the GARP VLAN Registration Protocol (GVRP), the origin is `GVRP`.

- GVRP is based on IEEE 802.1Q and allows for dynamic configuration of port-based VLANs. GVRP can help you simplify the management of VLAN configurations in larger networks. Use the command `bridge port gvrpState` to explicitly enable GVRP on the participating bridge ports *and* use the command `bridge gvrpState` to enable the bridge GVRP state for the entire system. The bridge GVRP state enables you to control GVRP on the system without losing the per-port GVRP state. By default, the GVRP state for the entire system is `disabled` and the GVRP state for each bridge port is `disabled`.
- The system prompts you for a VLAN interface index number before it displays the detail information.
- Either you can use network-based IP VLANs (by supplying Layer 3 address information when you configure a VLAN for IP), or you create the IP VLAN and then define multiple IP interfaces per VLAN. See Chapter 16.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index numbers of the VLAN interfaces for which you want detailed information	<ul style="list-style-type: none"> ■ One or more selectable VLAN interface index numbers ■ all ■ ? (for a list of selectable indexes) 	1 (if you have only one VLAN)

Fields in the Bridge VLAN Detail Display

Field	Description
Ignore STP mode (3500 and 9000 Layer 3)	Whether a VLAN can ignore STP blocked ports and let routing traffic pass through. Possible values: <code>enabled</code> and <code>disabled</code> .
Index	System-assigned index number that identifies a VLAN. Statistics appear for the VLAN that you specify.
Layer 3 addresses (3500 and 9000 Layer 3)	Information that is used to set up flood domains for overlapping IP VLAN subnetworks (network-based VLANs).
Name	Character string 0 through 32 bytes that identifies the VLAN. The default VLAN always uses the name <code>Default</code> .

Field	Description
Origin	<p>For all Layer 2 systems or switching modules, the VLAN origin is always <code>static</code>, which indicates that the VLAN was created by the user. For the CoreBuilder® 3500 or CoreBuilder 9000, the origin indicates one of the following:</p> <ul style="list-style-type: none"> ■ <code>static</code> — The VLAN was created statically (user-configured by using the <code>bridge vlan define</code> command). ■ <code>router</code> — The VLAN was created automatically by the router port IP interface (of router origin). You create a router port IP interface using the <code>ip interface define</code> command with the interface type <code>port</code>. You cannot modify or remove a router port VLAN. ■ <code>GVRP</code> — The VLAN was created dynamically from a GVRP update (<code>GVRP</code>). You must enable the GVRP state for the entire system as a bridge-wide parameter <i>and</i> for the participating bridge ports as a bridge-port parameter.
Ports/Port	<p>Index numbers of the bridge ports that belong to each VLAN. In the second part of the detail display, the Port column lists the ports for the VLAN individually and indicates ports that are trunked or have tagging.</p> <p>On the CoreBuilder 9000, the list of ports includes the front-panel ports and the appropriate backplane ports. Example: On a 12-port Layer 3 module, the list of ports includes ports 1 – 12 and port 13, which is the module's backplane port.</p>
Protocol (3500 and 9000 Layer 3)	Protocol suites for the VLAN. VLANs that are associated with router port IP interfaces always have IP as the protocol type. The default VLAN always uses the protocol type <code>unspecified</code> .
rxBcastBytes (3500 and 9000 Layer 3)	Number of received broadcast bytes
rxBcastFrames (3500 and 9000 Layer 3)	Number of received broadcast frames
rxMcastBytes (3500 and 9000 Layer 3)	Number of received multicast bytes
rxMcastFrames (3500 and 9000 Layer 3)	Number of received multicast frames
rxUcastBytes (3500 and 9000 Layer 3)	Number of received unicast bytes
rxUcastFrames (3500 and 9000 Layer 3)	Number of received unicast frames

Field	Description
Tag type (3500 and 9000 Layer 3)	Whether tagging is set to none or 802.1Q (IEEE 802.1Q tagging)
txBcastBytes (3500 and 9000 Layer 3)	Number of transmitted broadcast bytes
txBcastFrames (3500 and 9000 Layer 3)	Number of transmitted broadcast frames
txMcastBytes (3500 and 9000 Layer 3)	Number of transmitted multicast bytes
txMcastFrames (3500 and 9000 Layer 3)	Number of transmitted multicast frames
Type (VLAN Mode)	Either allOpen or allClosed. VLANs in allOpen mode share a single address table for all configured VLANs. In allClosed mode, each VLAN has its own unique address table. Standard bridging rules apply based on the table addresses that are assigned to the specific VLAN. Router port IP interfaces require allClosed mode.
VID	Unique, user-defined integer (VLAN ID) that identifies this VLAN. It is used by management operations. You can assign or modify a VID that is associated with a static VLAN; you cannot modify the VID selected automatically after you define a router port IP interface, nor can you change the VID of the default VLAN. The default VLAN requires a VID of 1.
VLAN Aware Mode (Layer 3 only)	Whether the VLAN aware mode (tagging mode) is allPorts or taggedVlanPorts. The default for CoreBuilder 3500 Release 2.0 or later is allPorts; allPorts is also the default as of CoreBuilder 9000 software Release 3.0. The value taggedVlanPorts is a compatibility mode for VLANs configured prior to CoreBuilder 3500 Release 2.0 and for VLANs configured on CoreBuilder 9000 Layer 3 modules prior to CoreBuilder 9000 Release 3.0.

**bridge vlan define
(3500/9000 Layer 3)**

✓ 3500
✓ 9000
9400

3900
9300

**For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Creates a VLAN on the CoreBuilder 3500 system or a CoreBuilder 9000 Layer 3 module. When you explicitly configure a VLAN on the system, you assign information such as a VLAN ID (VID), a set of bridge ports, and, optionally, a protocol type and IEEE 802.1Q tagging.

For details about this command on the SuperStack II Switch 3900, the Switch 9300, the CoreBuilder 9400, and CoreBuilder 9000 Layer 2 modules, see “bridge vlan define (3900/9300/9400/ 9000 Layer 2)” next.

Valid Minimum Abbreviation

b v def

Important Considerations

- If you have previously defined your IP VLANs with Layer 3 address information (that is, network-based VLANs), you can redefine your IP VLANs without Layer 3 address information. To define multiple IP interfaces per IP VLAN, use the `ip interface define` command. (See Chapter 16.)
- The VLAN that you create with this command is a static VLAN. To establish routing between static IP VLANs, define your IP VLANs and then use the `ip interface define` command to define an IP routing interface. As of Release 3.0, you can also specify the interface type `vlan` to create one or more IP routing interfaces for a static IP VLAN.
- If you have a router port IP interface on the system, you *cannot* specify the port that belongs to the router port IP interface when you explicitly define an IP VLAN (or a VLAN that includes the IP protocol). A router port IP interface is an alternative to static VLANs and allows routing versus bridging. You create a router port IP interface by entering the `ip interface define` command with the interface type `port` and a single bridge port. A router port IP interface requires `allClosed` mode. See Chapter 16 for more information.

- You must specify a VID in the range from 2 through 4094. You can no longer define a VLAN other than the default VLAN with a VID of 1. VID 1 is reserved for the default VLAN only as of Release 3.0.0. (As of Release 3.0.0, the default VLAN always uses the name `Default` and the protocol type `unspecified`.) If you delete the default VLAN, you can redefine it with VID 1 only.
- You cannot delete a VLAN that has a routing interface associated with it.
- If you plan to use the trunking feature or the MPLA feature, define the appropriate trunks *before* you define your VLANs. See Chapter 11 for more information.
- If you plan for your VLAN to include trunk ports, specify the anchor port (lowest-numbered port) that is associated with the trunk. For example, if ports 1 through 3 are associated with a trunk, specify `1` to define the VLAN to include all of the physical ports in the trunk (ports 1 through 3). If you have not defined trunks, specify one or more port numbers, or `a11` to assign all ports to the VLAN interface.
- If a port is shared by another VLAN, verify that if tagging is the only distinguishing characteristic between the VLANs, the specified tag type is not in conflict with the port's tag type in another VLAN (that is, there is only one port that is tagged `none`).
- Do not use this command if you want GVRP to dynamically create IEEE 802.1Q port-based VLANs. Instead, explicitly enable the GVRP state for the participating ports and enable the GVRP state for the entire system. To set the per-port GVRP state, use the `bridge port gvrpState`. (See Chapter 10.) To set the bridge-wide GVRP state, use the `bridge gvrpState` command. (See Chapter 9.)
- Whether you are bridging or routing, you can select more than one protocol suite per VLAN and specify one protocol at each of the prompts. Use the protocol type of `unspecified` to create a port-based VLAN.
- The IPX protocol type `IPX-802.2-SNAP` is available for both the CoreBuilder 3500 system and the CoreBuilder 9000 Layer 3 switching modules.
- For the CoreBuilder 9000, keep the following considerations in mind:
 - When you define a VLAN on a switching module (and other switching modules in the system also define this VLAN), you must define the VLAN on both the switching module *and* on the switch fabric module.

- When you define the VLAN on the Layer 3 switching module, you must specify any front-panel ports in the VLAN as well as the module's backplane port. The specified backplane port must also be tagged if you have more than one VLAN and plan to communicate with VLANs on other modules on the CoreBuilder 9000 through the switch fabric module.
- When you define the VLAN on the switch fabric module, you must specify which switch fabric module backplane port is connected to the module backplane port. The switch fabric module backplane port must also be tagged if you have more than one VLAN.
- When you use a Layer 3 switching module to establish routing between VLANs on other switching modules, you can configure the backplane port of the Layer 3 switching module as part of the VLANs and then define a routing interface for each VLAN. One VLAN equals one network or subnetwork.
- For configurations that include FDDI ports, if you plan for your VLAN to include FDDI DAS ports, you must specify the lowest-numbered port in the DAS pair when defining the ports in the VLAN. See Chapter 8.
- Specify ? to see the port summary (for example, to see whether ports are associated with a trunk), and then enter the appropriate port numbers.
- The VID is used as the IEEE 802.1Q tag if tagging is enabled for a port.

Options

Prompt	Description	Possible Values	[Default]
VID	Unique, user-defined integer used by management operations	<ul style="list-style-type: none"> ■ If the default VLAN exists, 2 – 4094 ■ If the default VLAN does not exist, 1 to redefine the default VLAN, or 2 – 4094 for other VLANs 	Next available VID

Prompt	Description	Possible Values	[Default]
Bridge ports	<p>Index numbers of the bridge ports that belong to the VLAN. If you include trunked ports, specify the anchor port of the trunk. On the CoreBuilder 9000, the list of ports includes the front-panel ports and the module's backplane port.</p> <p>When you define a VLAN that includes the IP protocol type, you cannot specify a port that is owned by a router port IP interface.</p>	<ul style="list-style-type: none"> ■ One or more of the ports that are available to be assigned to the VLAN ■ all ■ ? (for a list of selectable ports) 	–
Protocol suite (for VLANs other than the default)	<p>One or more protocol suites that you want to specify for the VLAN</p> <p>The default VLAN always uses the protocol type unspecified.</p>	<ul style="list-style-type: none"> ■ IP ■ IPX ■ Apple (for AppleTalk) ■ XNS ■ DECnet ■ SNA ■ Vines ■ X.25 ■ NETBEUI ■ unspecified (Default VLAN or a port-based VLAN) ■ IPX-II ■ IPX-802.2 ■ IPX-802.3 ■ IPX-802.2-SNAP 	unspecified (factory default)

Prompt	Description	Possible Values	[Default]
Layer 3 address configuration (IP VLAN only)	Whether you want to define Layer 3 information for the IP VLAN Since this is the last release to support Layer 3 address information in IP VLANs, avoid this mechanism and instead define multiple IP interfaces for the VLAN with <code>ip interface define</code> commands.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y
Layer 3 address and mask (IP VLAN only)	Fields (IP network address and subnet mask) you can use to set up flood domains for overlapping IP VLAN subnetworks. This is the last release to support Layer 3 address information in IP VLANs.	Any valid IP network address and subnet mask	—
Per-port tagging	Whether you want to configure IEEE 802.1Q VLAN tagging. You are prompted to answer for each port that you selected.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y
Tag type	Whether you want to configure no tagging or IEEE 802.1Q tagging (the VID) for each port.	<ul style="list-style-type: none"> ■ none ■ 802.1Q 	none
VLAN name (for VLANs other than the default)	Unique, user-defined name that identifies members of the VLAN. If you use spaces, put quotation marks around the VLAN name.	Up to 32 ASCII characters or spaces	—

Procedure

- 1 Enter the VLAN identification (VID) number in the range 2 – 4094.
- 2 Select the bridge ports.
- 3 Select one or more protocol suites.

If you select an IP protocol suite, proceed with step 4. If you did not choose an IP protocol suite for this interface, proceed to step 5.

- 4 Specify whether you want to specify Layer 3 address information (**n** or **y**). The default is **y**. Specify **n** if possible and instead define multiple IP interfaces for this VLAN using `ip interface define` commands. (See Chapter 16.) If you still want to specify Layer 3 address information for an IP VLAN:
 - a Enter **y** for Layer 3 addressing.
 - b Enter the Layer 3 network address.
 - c Enter the Layer 3 subnet mask. To accept the default or current value in brackets [], press Return or Enter.
- 5 Specify whether you want per-port tagging (**n** or **y**). The default is **y**.
- 6 If you specified per-port tagging, enter the tag type for the indicated port (**none** or **802.1Q**).
- 7 If you have defined more than one port, you are prompted again for a tag type for each port.
- 8 Enter the VLAN name.

Bridge VLAN Define Example (9000 Layer 3)

This example shows the steps necessary to define a protocol-based VLAN for IPX 802.3 on a Layer 3 switching module. In this example, only the backplane port (port 13) of the module has IEEE 802.1Q tagging; the front-panel ports in this VLAN are not tagged. Because you have tagged the module's backplane port, you must also tag the corresponding switch fabric module port of the switch fabric module for that VLAN. (Use the EME to connect to the switch fabric module and configure the VLAN.)

```
CB9000@slot2.1 [12-E/FEN-TX-L3] (bridge/vlan): define
Enter VID (2-4094) [5]: 5
Select bridge ports (1-13|all|?): 1-3,13
Enter protocol suite
(IP,IPX,Apple,XNS,DECnet,SNA,Vines,X25,NetBEUI,unspecified,
IPX-II,IPX-802.2,IPX-802.3): IPX-802.3
Enter protocol suite ('q' to quit)
(IP,Apple,XNS,DECnet,SNA,Vines,X25,NetBEUI,IPX-II,IPX-802.2,
IPX-802.3): q
Configure per-port tagging? (n,y) [y]: y
Enter port 1 tag type (none,802.1Q): none
Enter port 2 tag type (none,802.1Q): none
Enter port 3 tag type (none,802.1Q): none
Enter port 13 tag type (none,802.1Q): 802.1Q
Enter VLAN Name {?} [ ]: IPX1
```

Bridge VLAN Define Example (3500)

This example shows the steps necessary to define an IP VLAN with IEEE 802.1Q tagging on some ports. (Instead of supplying Layer 3 address information when you define the VLAN, you can define multiple IP interfaces for this VLAN.) This VLAN has *trunk ports*.

```
Select menu option: bridge vlan define
Enter VID (1-4094) [2]: 2
Select bridge ports (1-4,6,9-13|all|?) [3,6]: ?

Default selection: [3,6]

Selectable bridge ports
```

selection	ports	label
1	1	
2	2	
3	3,5	CampusLk1
4	4	
6	6-8	CampusLk2
9	9	
10	10	
11	11	
12	12	
13	13	

```
Select bridge ports (1-4,6,9-13|all|?) [3,6]: 3,6,9
Enter protocol suite
(IP,IPX,Apple,XNS,DECnet,SNA,Vines,X25,NetBEUI,unspecified,IPX-II,IPX-802.2
IPX-802.3,IPX-802.2-SNAP): IP
Enter protocol suite ('q' to quit)
(IPX,Apple,XNS,DECnet,SNA,Vines,X25,NetBEUI,IPX-II,IPX-802.2,IPX-802.3,
IPX-802.2-SNAP): q
Configure layer 3 address? (n,y) [y]: n
Configure per-port tagging? (n,y) [y]: y
Enter port 3,5 tag type (none,802.1Q) [none]: none
Enter port 6-8 tag type (none,802.1Q) [none]: 802.1Q
Enter port 9 tag type (none,802.1Q): none
Enter VLAN Name {?} [ ]: IP1
```

bridge vlan define (3900/9300/9400/ 9000 Layer 2)

3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300



Creates a port-based VLAN on standalone systems or the CoreBuilder 9000 Layer 2 modules. When you configure a port-based VLAN, you assign a VLAN ID (VID), a set of bridge ports, and, optionally, IEEE 802.1Q tagging.

For details about this command on the CoreBuilder 3500 and CoreBuilder 9000 Layer 3 modules, see “bridge vlan define (3500/9000 Layer 3)” earlier in this chapter.

Valid Minimum Abbreviation

b v def

Important Considerations

- On the SuperStack II Switch 3900 or 9300, you can define a maximum of 127 port-based VLANs on a single system.
- By default, all ports are defined to be part of the default VLAN, which always uses a VID of 1 and the name Default as of Release 3.0.0. If you delete the default VLAN, you can redefine it with VID 1 only.
- You cannot delete a VLAN that has an IP interface associated with it.
- The VID is used as the IEEE 802.1Q tag for a port if tagging is enabled.
- On the CoreBuilder 9000, the list of ports includes the front-panel ports and both backplane ports (even though only the lower-numbered backplane port is enabled by default). On the SuperStack II Switch 3900, the list of ports includes the 24 or 36 10/100 ports and any Gigabit Ethernet ports in use.
- For the CoreBuilder 9000, keep the following considerations in mind:
 - When you define a VLAN on a switching module and other switching modules in the system also define this VLAN, you must define the VLAN on both the switching module *and* on the switch fabric module.
 - When you define the VLAN on the Layer 2 switching module, you must specify any front-panel ports in the VLAN as well as the module's lower-numbered backplane port. The specified backplane port must also be tagged if you have more than one VLAN.
 - When you define the VLAN on the switch fabric module, you must specify the switch fabric module backplane port that is connected to the switching module's backplane port. The switch fabric module port must also be tagged if you have more than one VLAN.

Options

Prompt	Description	Possible Values	[Default]
VID	Unique, user-defined integer used by global management operations	<ul style="list-style-type: none"> ■ If the default VLAN exists, 2–4094 ■ If the default VLAN does not exist, 1 to redefine the default VLAN, or 2–4094 for other VLANs 	Next available VID
Bridge ports	Index numbers of the bridge ports that belong to the VLAN. If you include trunked ports, specify the anchor port of the trunk. See “Important Considerations” for information about the list of ports.	<ul style="list-style-type: none"> ■ One or more of the ports that are available to be assigned to the VLAN ■ all ■ ? (for a list of selectable ports) 	–
Per-port tagging	Whether you want to configure 802.1Q VLAN tagging. You are prompted to answer for each port that you selected.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y
Tag type	Whether you want no tagging or IEEE 802.1Q tagging (the VID). For a port shared by another VLAN, verify that the specified tag type is not in conflict with the port’s tag type in another VLAN.	<ul style="list-style-type: none"> ■ none ■ 802.1Q 	none
VLAN name (for VLANs other than the default)	Unique, user-defined name that identifies members of the VLAN. If you use spaces, put quotation marks around the VLAN name.	Up to 32 ASCII characters or spaces	–

Procedure

Press Return or Enter to accept the default or existing values that appear in brackets [].

- 1 Enter the VLAN identification (VID) number.
- 2 Enter one or more port numbers. To assign all ports to the VLAN, enter **all**
- 3 Configure the per-port tagging.

- 4 Enter the tag type for each port in the VLAN.
- 5 Enter the VLAN name.

Bridge VLAN Define Example (9000 Layer 2)

This example shows a port-based VLAN that includes tagged front-panel ports and a tagged backplane port (port 21). These ports are tagged because they overlap with ports that belong to other VLANs:

- Because the front-panel ports are tagged, any attached devices must be IEEE 802.1Q enabled.
- Because the backplane port is tagged, the corresponding switch fabric module port must also be tagged in the VLAN definition on the switch fabric module. (You connect to the switch fabric module and define the VLAN to include the appropriate tagged switch fabric module port, based on the slot that contains the switching module.)

```
CB9000@slot 10.1 [20-E/FEN-TX-L2] (bridge/vlan): define
Enter VID (2-4094) [3]: 3
Select bridge ports (1-22|all|?): 1-5,21
Configure per-port tagging? (n,y) [y]: y
Enter port 1 tag type (none,802.1Q): 802.1Q
Enter port 2 tag type (none,802.1Q): 802.1Q
Enter port 3 tag type (none,802.1Q): 802.1Q
Enter port 4 tag type (none,802.1Q): 802.1Q
Enter port 5 tag type (none,802.1Q): 802.1Q
Enter port 21 tag type (none,802.1Q): 802.1Q
Enter VLAN Name {?} [ ]: vlantag3
```

Bridge VLAN Define Example (3900)

This example shows a port-based VLAN that includes tagged ports.

```
Select menu option (bridge/vlan): define
Enter VID (2-4094) [2]: 2
Select bridge ports (1-39|all|?): 3-5
Configure per-port tagging? (n,y) [y]: y
Enter port 3 tag type (none, 802.1Q) [none]: 802.1Q
Enter port 4 tag type (none, 802.1Q) [none]: 802.1Q
Enter port 5 tag type (none, 802.1Q) [none]: 802.1Q
Enter VLAN name {?} [ ]: Sales
```

**bridge vlan modify
(3500/9000 Layer 3)**

✓ 3500
✓ 9000
9400

3900
9300

**For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Changes an existing port-based, protocol-based, or network-based VLAN definition on the CoreBuilder 3500 system or CoreBuilder 9000 Layer 3 module.

To use this command on the SuperStack II Switch 3900 or Switch 9300, the CoreBuilder 9400, and CoreBuilder 9000 Layer 2 modules, see “bridge vlan modify (3900/9300/9400/ 9000 Layer 2)” next.

Valid Minimum Abbreviation

b v modi

Important Considerations

- Before you modify the port assignments for a VLAN, always enter **?** to review the system port summary. If the VLAN includes trunk ports, you must specify the anchor (lowest numbered) port in each trunk. If there are no trunk ports, enter one or more port numbers, or enter **a11** to assign all ports to the VLAN.
- For the CoreBuilder 3500, if you want to modify your VLAN to include FDDI DAS ports, you must specify the lowest-numbered port in the DAS pair.
- If you modify the default VLAN, you can only change the member ports or the tag status. You cannot change the name or the VID or the protocol type of unspecified.
- If you modify the tagging type of a backplane port on a switching module, make sure that you modify the tagging type of the corresponding port on the switch fabric module.
- To modify a VLAN to support more than one protocol suite for the VLAN, specify one protocol at each of the prompts.
- Select the bridge ports that you want to be part of the modified VLAN, or specify **?** to display a port summary with the selectable bridge ports.
- If tagging is enabled for a port, the software uses the VID as the 802.1Q tag.

- If you modify the tagging for a port shared by another VLAN, and tagging is the only distinguishing characteristic between the VLANs, verify that the new tag type does not conflict with the port's tag type in another VLAN. (A shared port can use a tag type of `none` for only one of its VLANs; for all other VLANs to which it belongs, the shared port must use IEEE 802.1Q tagging.)

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	System-assigned index number that identifies a VLAN	<ul style="list-style-type: none"> ■ Selectable VLAN index ■ all ■ ? (for a list of selectable indexes) 	1 (if you have only the default VLAN)
VID (for VLANs other than the default)	Unique, user-defined integer used by global management operations	2 – 4094	Current VID
Bridge ports	Index numbers of the bridge ports that belong to the VLAN. To add trunked ports, specify the anchor port of the trunk. You cannot add a port owned by a router port IP interface.	<ul style="list-style-type: none"> ■ One or more index numbers of the ports that are available to be assigned to the VLAN ■ all ■ ? (for a list of selectable ports) 	Current ports in the VLAN

Prompt	Description	Possible Values	[Default]
Protocol suite (for VLANs other than the default)	One or more protocol suites that you want to specify for the VLAN	<ul style="list-style-type: none"> ■ IP ■ IPX ■ Apple (for AppleTalk) ■ XNS ■ DECnet ■ SNA ■ Vines ■ X.25 ■ NETBEUI ■ unspecified ■ IPX-II ■ IPX-802.2 ■ IPX-802.3 ■ IPX-802.2-SNAP (3500 only) 	Current protocol type
Modify Layer 3 address (IP VLAN)	<p>Whether you want to modify the Layer 3 information for the VLAN</p> <p>Avoid this mechanism and instead define multiple IP interfaces per VLAN with <code>ip interface define</code> commands.</p>	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y
Layer 3 address and mask (IP VLAN)	Optional fields (IP network and mask) used to set up flood domains for overlapping IP VLAN subnetworks	Any valid IP network address and mask	Current address and mask
Per-port tagging	Whether you want to modify the per-port 802.1Q VLAN tagging. You are prompted to answer for each port that you specified.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y
Tag type	Either no tagging or IEEE 802.1Q tagging (the VID)	<ul style="list-style-type: none"> ■ none ■ 802.1Q 	Current tag type for each port
VLAN name (for VLANs other than the default)	Unique, user-defined name that identifies members of the VLAN. If you use spaces, put quotation marks around the VLAN name.	Up to 32 ASCII characters or spaces	Current name

Procedure

To modify information for a VLAN, follow these steps:

- 1 Select the VLAN interface index.
- 2 For a VLAN other than the default VLAN, enter the VLAN identification (VID) number.
- 3 Specify the index numbers of the bridge ports.
- 4 For a VLAN other than the default VLAN, specify one or more protocol suites.

If you have selected the IP protocol suite, proceed with step 5. If you did not define an IP protocol suite for this VLAN, proceed to step 7.

- 5 Specify whether you want to modify Layer 3 address information (**n** or **y**). Since this is the last release to support Layer 3 address information, specify **n** if possible and instead define multiple IP interfaces for this VLAN using `ip interface define` commands. (See Chapter 16.) If you still want to modify Layer 3 address information for an IP VLAN:
 - a Enter **y** for Layer 3 addressing.
 - b Enter the Layer 3 network address.
 - c Enter the Layer 3 subnet mask. To accept the default or current value in brackets [], press Return or Enter.
- 6 Specify whether you want to modify per-port tagging.
- 7 If you want to modify per-port tagging, enter the new tag type for the port (**none** or **802.1Q**).
- 8 If you have specified that you want to modify more than one port, enter a tag type for each port.
- 9 For a VLAN other than the default VLAN, enter a new VLAN name or keep the current name.

The VLAN name can include up to 32 ASCII characters, including spaces. If you include spaces, put quotation marks around the VLAN name.

Bridge VLAN Modify Example (9000 Layer 3)

This example shows the steps to modify the per-port tagging for a protocol-based VLAN on a Layer 3 module. In this example, front-panel port 5 is changed to have IEEE 802.1Q tagging, and its associated device is IEEE 802.1Q enabled.

```
CB9000@slot2.1 [12-E/FEN-TX-L3] (bridge/vlan): modify  
Select VLAN interface index {1-5|?}: 5
```

```

Enter VID (2-4094) [5]: 5
Select bridge ports (1-13|all|?) [1-5,13]: 1-5,13
Enter protocol suite
(IP,IPX,Apple,XNS,DECnet,SNA,Vines,X25,NetBEUI,unspecified,
IPX-II,IPX-802.2,IPX-802.3) [IPX-802.3]: IPX-802.3
Enter protocol suite ('q' to quit) (IP,IPX, Apple, XNS,
DECnet,SNA,Vines,X25,NetBEUI,IPX-II,IPX-802.2): q
Modify per-port tagging? (n,y) [y]: y
Enter port 1 tag type (none,802.1Q) [none]: none
Enter port 2 tag type (none,802.1Q) [none]: none
Enter port 3 tag type (none,802.1Q) [none]: none
Enter port 4 tag type (none,802.1Q) [none]: none
Enter port 5 tag type (none,802.1Q) [none]: 802.1Q
Enter port 13 tag type (none,802.1Q) [802.1Q]: 802.1Q
Enter VLAN Name {?} [IPX]: IPX1

```

Bridge VLAN Modify Example (3500)

This example shows the steps to modify the member ports and per-port tagging for an IP VLAN.

```

Select menu option: bridge vlan modify
Select VLAN interface index {1-2|?}: 2
Enter VID (2-4094) [2]: 2
Select bridge ports (1-4, 6, 9-13|all|?) [3,6,9]: 9,11
Enter protocol suite
(IP,IPX,Apple,XNS,DECnet,SNA,Vines,X25,NetBEUI,unspecified,
IPX-II,IPX-802.2,IPX-802.3, IPX-802.2-SNAP) [IP]: IP
Enter protocol suite ('q' to quit)
(IPX,Apple,XNS,DECnet,SNA,Vines,X25,NetBEUI,IPX-II,
IPX-802.2, IPX-802.3, IPX-802.2-SNAP): q
Modify layer 3 address? (n,y) [y]: n
Modify per-port tagging? (n,y) [y]: y
Enter port 9 tag type (none,802.1Q) [none]: 802.1Q
Enter port 11 tag type (none,802.1Q) [none]: 802.1Q
Enter VLAN Name {?} [IP1]: IP1

```

bridge vlan modify (3900/9300/9400/ 9000 Layer 2)

Changes a port-based VLAN definition on the indicated system Layer 2 module. See “Important Considerations” for information on when changes take effect.

3500
✓ 9000
✓ 9400



To use this command on the CoreBuilder 3500 or CoreBuilder 9000 Layer 3 modules, see the “bridge vlan modify (3500/9000 Layer 3)” earlier in this chapter.

✓ 3900
✓ 9300

Valid Minimum Abbreviation

b v modi

Important Considerations

- You need not reboot the system for the changes to take effect. However, depending on the number of VLANs that are affected, the system may take several minutes to return control to you.
- If you modify the tagging type of a backplane port on a switching module, make sure that you modify the tagging type of the corresponding port on the switch fabric module.
- If tagging is enabled for a port, the software uses the VID as the 802.1Q tag.
- If you modify the default VLAN, you can only change the member ports or the tag status. You cannot change the name or the VID.
- If you modify the tagging for a port shared by another VLAN, verify that the new tag type does not conflict with the port's tag type in another VLAN. (A shared port can use a tag type of `none` for only one of its VLANs; for all other VLANs to which it belongs, the shared port must use IEEE 802.1Q tagging.)

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	System-assigned index number that identifies a VLAN	<ul style="list-style-type: none"> ■ Selectable VLAN index ■ all ■ ? (for a list of selectable indexes) 	1 (if you have only the default VLAN)

Prompt	Description	Possible Values	[Default]
VID (for VLANs other than the default)	Unique, user-defined integer used by management operations	2 – 4094	Current VID
Bridge ports	Index numbers of the bridge ports that belong to the VLAN. To add trunked ports, specify the anchor port of the trunk.	<ul style="list-style-type: none"> ■ One or more index numbers of the ports that are available to be assigned to the VLAN ■ all ■ ? (for a list of selectable ports) 	Current ports in VLAN
Per-port tagging	Whether you want to configure 802.1Q VLAN tagging. You are prompted to answer for each port that you selected.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	y
Tag type	Either no tagging or IEEE 802.1Q tagging (the VID)	<ul style="list-style-type: none"> ■ none ■ 802.1Q 	Current tag type for each port
VLAN name (for VLANs other than the default)	Unique, user-defined name that identifies members of the VLAN. If you use spaces, put quotation marks around the VLAN name.	Up to 32 ASCII characters or spaces	Current name

Procedure

- 1 Enter the VLAN interface index.
- 2 For a VLAN other than the default VLAN, enter a VLAN identification (VID) number or keep the default in brackets.
- 3 Specify the index numbers of the bridge ports.
- 4 Specify whether you want to modify per-port tagging.
- 5 If you modify per-port tagging, enter the new tag type for the port (**none** or **802.1Q**).
- 6 If you have defined more than one port, enter a tag type for each port.
- 7 For a VLAN other than the default VLAN, enter a new VLAN name or keep the current name.

The VLAN name can include up to 32 ASCII characters, including spaces. If you include spaces, put quotation marks around the VLAN name.

Bridge VLAN Modify Example (9000 Layer 2)

This example shows the removal of two ports from a port-based VLAN that includes tagged front-panel ports and a tagged backplane port (port 21).

```
CB9000@slot 10.1 [20-E/FEN-TX-L2] (bridge/vlan): modify
Select VLAN interface index {1-3|?}: 3
Enter VID (2-4094) [3]: 3
Select bridge ports (1-22|all|?) [1-5,21]: 1-3,21
Configure per-port tagging? (n,y) [y]: y
Enter port 1 tag type (none,802.1Q) [802.1Q]: 802.1Q
Enter port 2 tag type (none,802.1Q) [802.1Q]: 802.1Q
Enter port 3 tag type (none,802.1Q) [802.1Q]: 802.1Q
Enter port 21 tag type (none,802.1Q) [802.1Q]: 802.1Q
Enter VLAN Name {?} [vlan3]: vlantag3
```

Bridge VLAN Modify Example (3900)

This example shows default VLAN changes in the ports and per-port tagging type.

```
Select menu option (bridge/vlan): modify
Select VLAN interface index {1-2|?}: 1
Select bridge ports (1-27|all|?) [1-27]: 2-6
Modify per-port tagging? (n,y) [y]:
Enter port 2 tag type (none,802.1Q) [none]: 802.1Q
Enter port 3 tag type (none,802.1Q) [none]: 802.1Q
Enter port 4 tag type (none,802.1Q) [none]: 802.1Q
Enter port 5 tag type (none,802.1Q) [none]: 802.1Q
Enter port 6 tag type (none,802.1Q) [none]: 802.1Q
```

bridge vlan remove Deletes a VLAN definition.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

b v r

Important Considerations

- When you remove a VLAN on a CoreBuilder 9000 Layer 2 or Layer 3 module, the system prompts you to verify that you want to wait the several minutes that it may take for the removal to be complete.
- You cannot remove a VLAN that is associated with any type of routing interface (for example, a router port VLAN created by a router port IP interface or a protocol-based VLAN associated with a particular router interface).

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	System-assigned index number that is associated with the VLAN	<ul style="list-style-type: none"> ■ A selectable VLAN index ■ all ■ ? (for a list of selectable indexes) 	–
Continue verification (9000 Layer 2 and Layer 3)	Whether you want to continue with the VLAN removal, even though the removal may take a few minutes to complete	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y

Bridge VLAN Remove Example (3500)

```
Select menu option: bridge vlan remove
Select VLAN interface indexes (1-2|all|?): ?

Selectable vlans

selection  VID    ports                name
1          1      1-13                Default
2          2      3,5-9,11            IP1

Select VLAN interface indexes (1-2|all|?):2
```

bridge vlan mode

Determines whether data with a unicast MAC address can be forwarded between VLANs.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation**b v mode****Important Considerations**

- Select a VLAN mode as follows:
 - **allOpen** — Use this less restrictive mode if you do not have security issues concerning the forwarding of data between VLANs. It is the default VLAN mode for all VLANs that you create. It permits data with a unicast MAC address to be forwarded between VLANs. The allOpen mode implies that the system uses a single bridge address table for all of the VLANs on the system.
 - **allClosed** — Use this restrictive mode if you are concerned about security between VLANs. Data cannot be forwarded between VLANs but can still be routed between VLANs. This mode implies that each VLAN that you create has its own address table.
- For the CoreBuilder 3500 system and CoreBuilder 9000 Layer 3 modules, if you are using allClosed mode and STP (with multiple routes to a destination), you can also use the command “bridge vlan stpMode” to disable STP blocking for a specified VLAN.
- For the CoreBuilder 9000, set a VLAN mode for each switching module and the switch fabric module.
- Changing this mode removes all VLANs and redefines the default VLAN.
- Before you issue this command to change the mode, you must remove all routing interfaces, including router port IP interfaces. If routing interfaces are defined, the system displays this message:


```
could not change configured VLAN mode - interface in use by client.
```

Options

Prompt	Description	Possible Values	[Default]
VLAN mode	Selected VLAN mode for the entire system	<ul style="list-style-type: none"> ■ allOpen ■ allClosed 	allOpen (factory default), or current value

bridge vlan stpMode

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

If allClosed mode is enabled, allows the system to ignore the Spanning Tree Protocol (STP) state for a specified VLAN interface or all interfaces, for either routing or bridging.

Valid Minimum Abbreviation

b v st

Important Considerations

- This mode is valid only if the VLAN mode is set to allClosed.
- To disable the STP state on a *per-port* basis with either allOpen or allClosed mode, use the bridge port stpState command. See Chapter 10.
- If you have configured router port IP interface (and therefore have a router port VLAN), the ignore STP mode is enabled and cannot be changed. You cannot select a router port VLAN with this command.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	System-assigned index number that is associated with the VLAN	<ul style="list-style-type: none">■ Any selectable VLAN index number■ all■ ? (for a list of selectable indexes)	–
STP state	Whether you want to ignore the STP state for the VLAN index	<ul style="list-style-type: none">■ disabled■ enabled	disabled

Bridge VLAN STP Mode Example (3500)

```
selection  VID    ports          name
1          1      1-13          Default
2          2      3,5-9,11      IP1

Select VLAN interface index(es) (1-2|all|?):2
Ignore STP state for VLAN index: 2 (disabled,enabled) [disabled]:enabled
```

bridge vlan vlanAwareMode

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

For compatibility purposes, allows the system to observe previous VLAN resource usage and tagged-frame ingress rules for CoreBuilder 3500 serial-port upgrades from Release 1.2.0 to 2.1.0 or 3.0.0 (or CoreBuilder 9000 upgrades from Release 2.0.0 to 3.0.0).

Valid Minimum Abbreviation

b v v

Important Considerations

- Use this command *only* if you upgrade your system and the system reports an error after reaching the VLAN resource limit during a power up with a serial-port console connection. During the upgrade, the difference in resource usage and modes of tagging could cause the later release to use more VLAN resources than did the earlier release, thereby causing a decrease in the total number of allowable VLANs.
- If the system reaches the VLAN resource limit during the upgrade, it displays an error message to identify the index of the VLAN that it was unable to create. The system removes all bridge ports from the VLAN that it could not restore from NV data but does maintain the previously stored NV data.
- The difference in VLAN resource usage is based on the following:
 - In CoreBuilder 3500 Release 1.2.0 (and CoreBuilder 9000 Release 2.0.0), all bridge ports were *not* VLAN aware (tagging aware) unless they were assigned to a VLAN that has one or more tagged ports. This behavior is associated with the VLAN aware mode of `taggedVlanPorts`. If you see the VLAN resource error message, you can restore your VLANs by issuing this command and setting the VLAN aware mode to `taggedVLANPorts`. If VLANs are already defined, the system prompts you to reboot the system to put the new mode into effect.
 - As of CoreBuilder 3500 Release 2.0.0, (and CoreBuilder 9000 Release 3.0.0), all bridge ports become VLAN aware after a software update or after an NV data reset and do not have to be explicitly tagged in order to forward tagged frames. This behavior is associated with the default VLAN aware mode of `allPorts`. If you do not see the VLAN internal resource error message, maintain the VLAN aware mode of `allPorts`.

- The VLAN aware mode reflects the difference in tagged-frame ingress rules between releases. Therefore, even if the system can accommodate the number of VLANs from the earlier release, be aware that it begins using different ingress rules for tagged frames.
 - The CoreBuilder 3500 tagged-frame ingress rules vary for 1.2.0, 2.0.0, and 3.0.0. For more information, see the *CoreBuilder 3500 Implementation Guide*.
 - The CoreBuilder 9000 tagged-frame ingress rules vary for 2.0.0 to 3.0.0. For more information, see the *CoreBuilder 9000 Implementation Guide*.

Options

Prompt	Description	Possible Values	[Default]
VLAN aware mode	Whether all ports are tagging aware or only tagged ports are tagging aware	<ul style="list-style-type: none">■ allPorts■ taggedVlanPorts	allPorts
Reboot system?	Since changing the mode requires you to reboot, whether you want to reboot the system or cancel the request	<ul style="list-style-type: none">■ n (no)■ y (yes)	y (yes)

Bridge VLAN Aware Mode Example (3500)

```
Select menu option (bridge/vlan): vlanAwareMode
VLAN-aware mode (taggedVlanPorts,allPorts) [allPorts]:
taggedVLANPorts
Changing the VLAN-aware mode will reboot the system -
continue? (n,y) [y]: y
```


PACKET FILTERS

This chapter provides guidelines and other key information about how to administer bridge packet filters in your system, including the following tasks:

- Listing and displaying packet filters
- Creating, deleting, editing, and loading packet filters
- Assigning and unassigning packet filters
- Managing port groups

Independently configurable packet filtering is provided for the packet processing paths on each bridge port of the system. After you create a packet filter, you can assign the filter to the transmit or the receive paths of any bridge port or group of bridge ports.

The filter executes a series of test operations on the packet's contents and, if the result is zero, it stops (filters) the packet. If the end result is non-zero, the filter lets the packet pass.



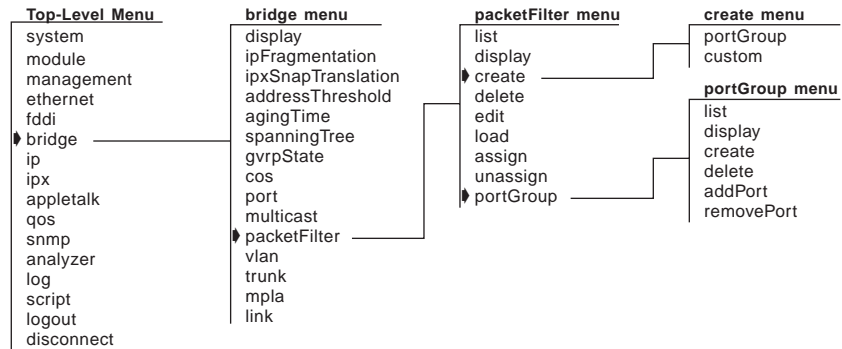
For more information about implementing packet filters on your network, see the Implementation Guide for your system.



For the CoreBuilder® 9000 platform, the commands in this chapter apply to Layer 3 switching modules only.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



**bridge packetFilter
list**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Lists the currently defined packet filters.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

b pa li

Bridge Packet Filter List Example (3500)

Select menu option (bridge/packetFilter): **list**

```
Packet Filter 1 - rejdifportgrp
Port 11, txA, rxA
```

In the example, the system has one packet filter, with a filter id of 1 and a defined name of `rejdifportgrp`. This filter is loaded onto port 11. The filter is assigned to both the transmit all (`txA`) path and the receive all (`rxA`) path.

bridge packetFilter display

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the contents of the specified packet filter.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

b pa di

Important Considerations

- Possible values for filters (*n*) depend on the number of created or loaded filters on the system.
- The packet filter id and name are displayed, followed by a list of the packet filter instructions.

Options

Prompt	Description	Possible Values	[Default]
Select filter	Identifier (id number) of the filter that you want to display	<ul style="list-style-type: none"> ■ 1 – <i>n</i> ■ ? (for a list of all filters) 	Current filter id

Sample Bridge Packet Filter Display (3500)

Select menu option (bridge/packetFilter): **display**

Select filter {1|?} [1]:

```

Packet Filter 1 - rejdiffportgrp
  name                "rejdiffportgrp"
  pushDPGM
  pushSPGM
  and
  pushLiteral.1      0x00000000
  ne
  
```

**bridge packetFilter
create portGroup*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Creates the portGroup (rejdiffportgroup) standard hardware filter.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation**b pa c p****Important Considerations**

- The portGroup (rejdiffportgroup) packet filter rejects a frame if the destination and source ports are not in the same group.
- “Creating” a hardware filter means that the code for the filter is copied from firmware into non-volatile memory.
- To verify that the filter has been created, use the `bridge packetFilter list` command. To see the contents of the portGroup filter, use the `bridge PacketFilter display` command.
- The system only creates the packet filter definition. You must still assign ports and masks to port groups, as described for “bridge packetFilter portGroup create” later in this chapter, and assign the standard filter to ports and filtering paths, as described for “bridge packetFilter assign” later in this chapter.
- At present, portGroup is the only filter supported in hardware.

Bridge Packet Filter Create Port Group Example

This example shows the user creating the portGroup filter.

```
Select menu option (bridge/packetFilter): create portgroup  
Packet filter 1 stored.
```

**bridge packetFilter
create custom**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Creates a custom packet filter using the built-in editor.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

b pa c c

Important Considerations

- You can create custom filters to add filtering logic based on the content of the packet.
- The built-in editor is a simple one-line-at-a-time editor that supports a short list of EMACS-style editing commands.
- Save your work periodically with Ctrl+w. When you press Esc to exit the built-in editor, the system examines the filter's syntax. If the syntax is correct, the filter is loaded into the switch's non-volatile memory. Incorrect syntax filters are not loaded into non-volatile memory and are not saved across editor sessions.
- After you create the filter, edit it using "bridge packetFilter edit" as described later in this chapter.
- The alternative to creating a custom packet filter using the built-in editor is to create the packet filter on an external system and transfer it across the network into the switch. See "bridge packetFilter load" later in this chapter.
- You can also use the Filter Builder component of the Web Management application to create custom filters.
 - On CoreBuilder 3500 systems, you can load the filter on to the switch directly from Filter Builder.
 - On CoreBuilder 9000 system, you must save the filter to an ASCII file and then download the file to the switch manually using TFTP and the "bridge packetFilter load" command described later in this chapter.
- The system only creates the packet filter definitions. You must still assign the standard filter to ports and filtering paths, as described for "bridge packetFilter assign" later in this chapter.

3900
9300

Create Custom Bridge Packet Filter Example (3500)

After you enter the custom filter editor, the system displays the editor commands, as shown here.

Select menu option (bridge/packetFilter): **create custom**

Editor Commands

```
Buffer: list = Ctrl-l
Line:   next = Ctrl-n, previous = Ctrl-p
Cursor: start = Ctrl-a, end = Ctrl-e, left = Ctrl-b, right = Ctrl-f
Insert: line = Enter
Delete: previous = Ctrl-h (BSP), current = Ctrl-d (DEL), line = Ctrl-k
Mode:   insert/overstrike toggle = Ctrl-o
Save:   Ctrl-w
Exit:   Esc
```

You now enter packet filter language statements that define the packet filter algorithm. See the *Implementation Guide* for your system for information about developing the packet filters.

bridge packetFilter delete

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Deletes the selected packet filter.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

b pa de

Important Considerations

- You cannot delete a filter if it is assigned. Before you can delete the filter, you must unassign the filter from the assigned ports.
- Possible values for filters (*n*) depend on the number of created or loaded filters on the system.
- To find the id of the filter, list the filters using the `bridge packetFilter list` command.

Options

Prompt	Description	Possible Values	[Default]
Select filter	Identifier (id number) of the filter that you want to delete	<ul style="list-style-type: none"> ■ 1 – <i>n</i> ■ ? (for a list of all identifiers) 	Current filter number
Delete packet filter?	Whether you want to delete the selected packet filter	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y

Bridge Packet Filter Delete Examples (3500)

Select menu option (bridge/packetFilter): **delete**

Select filter {1|?} [1]: **1**

Delete packet filter (n,y) [y]: **y**

Packet filter 1 has been deleted.

If the filter is assigned, it cannot be deleted. The system responds as follows to the delete command:

Select menu option (bridge/packetFilter): **delete**

Select filter {1|?} [1]: **1**

The selected filter is assigned

This problem prevents the deletion of this filter.

bridge packetFilter
edit

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies an existing packet filter using the built-in editor.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

b pa e

Important Considerations

- The built-in editor is a simple one-line-at-a-time editor that supports a short list of EMACS-style editing commands.
- The system displays the editor commands that you use to edit the packet filters. You can edit packet filter language statements that define the packet filter algorithm. See the *Implementation Guide* for your system for information about developing the packet filters.
- Save your work periodically with Ctrl+w. To complete the editing process, press Esc. The system replaces the filter or creates a new filter, depending on your response to the prompts.
- When you exit, the system examines the filter's syntax. If the syntax is correct, the filter is loaded into the switch's non-volatile memory.
- Possible values for filters (*n*) depend on the number of created or loaded filters on the system.

Options

Prompt	Description	Possible Values	[Default]
Select filter	Identifier (id) number of the filter that you want to edit	<ul style="list-style-type: none">■ 1 – <i>n</i>■ ? (for a list of all identifiers)	Most recent filter edited
Replace existing filter?	Whether to replace the selected filter	<ul style="list-style-type: none">■ n (no)■ y (yes)	y
Store as new filter?	Whether to create a new filter	<ul style="list-style-type: none">■ n (no)■ y (yes)	y

Replace Existing Filter Example (3500)

```
Select menu option (bridge/packetFilter): edit
Select filter {1|?} [1]:
Editing packet filter 1.
```

Editor Commands

```
Buffer: list = Ctrl-l
Line:   next = Ctrl-n, previous = Ctrl-p
Cursor: start = Ctrl-a, end = Ctrl-e, left = Ctrl-b, right = Ctrl-f
Insert: line = Enter
Delete: previous = Ctrl-h (BSP), current = Ctrl-d (DEL), line = Ctrl-k
Mode:   insert/overstrike toggle = Ctrl-o
Save:   Ctrl-w
Exit:   Esc
```

Edit buffer has been saved

```
name                "rejdifportgrp"
Replace existing filter (n,y) [y]: y
Packet filter 1 has been replaced.
```

Store as New Filter Example (3500)

```
Select menu option (bridge/packetFilter): edit
Select filter {1-2|?} [1]: 1
Editing packet filter 1.
```

Editor Commands

```
Buffer: list = Ctrl-l
Line:   next = Ctrl-n, previous = Ctrl-p
Cursor: start = Ctrl-a, end = Ctrl-e, left = Ctrl-b, right = Ctrl-f
Insert: line = Enter
Delete: previous = Ctrl-h (BSP), current = Ctrl-d (DEL), line = Ctrl-k
Mode:   insert/overstrike toggle = Ctrl-o
Save:   Ctrl-w
Exit:   Esc
```

Edit buffer has been saved

```
name                "BlockGeoB"
Replace existing filter (n,y) [y]: n
Store as new filter (n,y) [y]: y
Packet filter 3 stored.
```

**bridge packetFilter
load**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Transfers a packet filter file from another host machine to the switch to which you are currently connected.

Valid Minimum Abbreviation

b pa lo

Important Considerations

- On the CoreBuilder 3500, before you use the `packetFilter load` command, select the required file transfer protocol (TFTP or FTP) using the `system fileTransfer` command.
- On the CoreBuilder 3500, this single command transfers a file from another host and loads it in a one-step process.
- On the CoreBuilder 9000, loading a filter from another host is a two-step process. You must first `download` the packet filter source file to the Enterprise Management Engine (EME) using TFTP, then `connect` to the Layer 3 module, then enter this `bridge packetFilter load` command.

The syntax of the EME download command is:

```
download module <slot.subslot> filter <IP address> <filename>
```

You must use TFTP to download on the CoreBuilder 9000. FTP does not work.

When you enter `bridge packetFilter load`, the CoreBuilder 9000 does not prompt you for any options. Instead, the module simply looks for the downloaded filter on the EME. If it finds it, it loads it. If it does not find it, it prints the message `Filter not found`.

- TFTP or FTP hosts may place restrictions on which files and pathnames are valid. See your host administrator or host documentation for TFTP and FTP information.
- `bridge packetFilter load` verifies the syntax of the filter. If the syntax is correct, it stores the filter into non-volatile memory. If the syntax is incorrect, you are prompted to enter the built-in editor so that you can fix the filter.

Options (3500)

Prompt	Description	Possible Values	[Default]
Host IP address	IP address of the machine from which you want to transfer the filter	Any valid IP address	current IP address
File pathname	Path and file name of the filter to transfer	<ul style="list-style-type: none"> ■ ? (for a list of criteria for entering the pathname) ■ Up to 128 characters 	path and file name last loaded

Bridge Packet Filter Load Example (3500)

The system transfers the specified filter and displays a confirmation message:

```
Select menu option (bridge/packetFilter): load
Host IP address: 158.101.112.191
File pathname {?}: /tftpboot/srackley/joe.fil
Packet filter 2 stored.
```

Bridge Packet Filter Load Example (9000)

The user has copied the source text for the “reject multicast traffic” filter, rejmulticast.fil, from the Filter Builder application to the TFTP application’s root directory on host 159.101.8.112. (You must use TFTP; FTP does not work.)

The user then logs on to the CoreBuilder 9000 EME and issues the download command to transfer the filter file to the EME. Note that the user specifies the type of download (filter) and for which module (6.01) the filter is destined.

```
CB9000> download module 6.01 filter 159.101.8.112 rejmulticast.fil

File transfer request pending.
Downloading file from external file server to eme - 000000289
Downloading file from eme to module 6.1 - 000000289
File transfer completed successfully.
```

The user next connects to the module and loads the filter.

```
CB9000> connect 6.01
```

```
Menu options (Corebuilder 9000-94DC8): -----
list                - List all packet filters
display             - Display a packet filter
create              - Create a packet filter
delete              - Delete a packet filter
edit                - Edit a packet filter
load                - Load a packet filter
assign              - Assign a packet filter
unassign            - Unassign a packet filter
portGroup           - Administer port groups
```

```
Type "q" to return to the previous menu or ? for help.
```

```
-----
CB9000@slot6.1 [12-E/FEN-TX-L3] (): bridge packetFilter load
Packet filter 1 stored.
```

Lastly, the user lists the loaded filters with `bridge packetfilter list` and confirms the contents of the filter with `bridge packetfilter display`.

```
CB9000@slot6.1 [12-E/FEN-TX-L3] (bridge/packetFilter): list
```

```
Packet Filter 1 - rejMulticast
No port assignments
```

```
Menu options (Corebuilder 9000-94DC8): -----
-----
list                - List all packet filters
display             - Display a packet filter
create              - Create a packet filter
delete              - Delete a packet filter
edit                - Edit a packet filter
load                - Load a packet filter
assign              - Assign a packet filter
unassign            - Unassign a packet filter
portGroup           - Administer port groups
```

```
Type "q" to return to the previous menu or ? for help.
```

```
-----
CB9000@slot6.1 [12-E/FEN-TX-L3] (bridge/packetFilter): display 1
```

```
Packet Filter 1 - rejMulticast
name                "rejMulticast"
pushField.b         0
pushLiteral.b       0x01
and
not
```

bridge packetFilter assign

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Assigns a selected packet filter to a port or set of ports (port group).

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

b pa a

Important Considerations

- When you assign a packet filter to one or more ports, you must assign a processing path. The path (transmit all, transmit multicast, receive all, receive multicast, or receive internal) of a port can have only one packet filter assigned to it; however, you can assign a single packet filter to multiple paths and ports.
- If you try to assign a filter to a port that already has a filter assigned, the system displays a warning message and the assignment fails.
- After you assign the filter, ports and paths are removed from the list of possible values (which are listed in the Options table).
- Possible values for filters (*n*) depend on the number of created or loaded filters on the system.
- Possible values for bridge ports (*n*) depend on the number of existing bridge ports on the system.

Options

Prompt	Description	Possible Values	[Default]
Select filter	Identifier (id number) of the filter that you want to assign	<ul style="list-style-type: none"> ■ 1 – <i>n</i> ■ ? (for a list of valid filter identifiers) 	Current valid selected filter
Select bridge ports	Number of the bridge port to which you want to assign the selected filter	<ul style="list-style-type: none"> ■ 1 – <i>n</i> ■ all ■ ? (for a list of valid ports) 	Current valid selected bridge port

Prompt	Description	Possible Values	[Default]
Select path(s)	Identifier of the path to which you want to assign the selected filter	<ul style="list-style-type: none">■ txA■ txM■ rxA■ rxM■ rxI■ all■ ? (for a list of valid paths)	Current valid selected path

Bridge Packet Filter Assign Examples (3500)

```
Select menu option (bridge/packetFilter): assign
Select filter {1|?} [1]:
Select bridge port(s) (1-12|all|?) [4-6]: all
Select path(s) (txA,txM,rxA,rxM,rxI|all|?): txA
```

To specify multiple ports, use the hyphen (-) to indicate ranges, and commas to indicate individual, non-contiguous ports. To specify multiple paths, separate the paths with commas.

```
Select menu option (bridge/packetFilter): assign
Select filter {1|?} [1]:
Select bridge port(s) (1-6|all|?): 1-3,6
Select path(s) (txA,txM,rxA,rxM,rxI|all|?): txA,rxA,rxI
```

bridge packetFilter unassign

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Unassigns selected packet filter from one or more ports.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

b pa u

Important Considerations

- The packet filter that you want to unassign must have been assigned to at least one port.
- Possible values for filters (n) depend on the number of created or loaded filters on the system.
- Possible values for bridge ports (n) depend on the number of existing bridge ports on the system.
- After you unassign the filter, ports and paths are added to the list of possible values (which are listed in the Options table).

Options

Prompt	Description	Possible Values	[Default]
Select filter	Identifier (id number) of the filter that you want to unassign	<ul style="list-style-type: none"> ■ 1 – n ■ ? (for a list of valid filter identifiers) 	Current valid selected filter
Select bridge ports	Numbers of one or more bridge ports from which you want to unassign the selected filter	<ul style="list-style-type: none"> ■ 1 – n ■ all ■ ? (for a list of valid ports) 	Current valid selected bridge port
Select path(s)	Identifiers of one or more paths from which you want to unassign the selected filter	<ul style="list-style-type: none"> ■ txA ■ txM ■ rxA ■ rxM ■ rxI ■ all ■ ? (for a list of valid paths) 	Current valid selected path

Bridge Packet Filter Unassign Examples (3500)

The unassignment is from the transmit all (`txA`) paths on port 1.

Select menu option (bridge/packetFilter): **unassign**

Select filter {1|?} [1]: **1**

Select bridge port [1]: **1**

Select path(s) (`txA`,`rxA`|all|?) [`txA`,`rxA`]: **txA**

**bridge packetFilter
portGroup list**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays a list of currently defined port groups.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

b p a p l

Bridge Packet Filter Port Group List Example

Select menu option (bridge/packetFilter/portGroup): **list**

```
Port Group 1 - Marketing
  Port group mask - bit 15
Port Group 2 - Sales
  Port group mask - bit 32
```

In the example, the system has two port groups defined: Marketing and Sales. The display shows the group id, group name (if any), and group mask.

**bridge packetFilter
portGroup display**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays a port group.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

b pa p di

Important Consideration

- Possible values for port groups (n) depend on the number of user-defined port groups on the system.

Options

Prompt	Description	Possible Values	[Default]
Select port group	Number of the port group to display	<ul style="list-style-type: none">■ 1 – n■ ? (for a list of valid port groups)	Current port group

Sample Bridge Packet Filter Port Group Display (3500)

```
Select menu option (bridge/packetFilter/portGroup): display  
Select port group {1-2|?} [2]: 2
```

```
Port Group 2 - Sales  
Port 5
```

```
Port 6
```

bridge packetFilter portGroup create

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Creates a port group.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

b p a p c

Important Considerations

- You can create up to 32 port groups, one for each bit in the 32-bit port group mask.
- The `portGroup create` command only creates port group associations. You must create and assign a filter to a port group to affect filtering. See “bridge packetFilter create portGroup” and “bridge packetFilter assign” earlier in this chapter.
- Possible values for bridge ports (n) depend on the number of bridge ports on the system.

Options

Prompt	Description	Possible Values	[Default]
Select port group mask	Mask that you want to assign to the port group	<ul style="list-style-type: none"> ■ 1 – 32 ■ ? (for a list of masks) 	–
Select port group name	Name of the port group that you want to create Use quotation marks around any string with embedded spaces. Use "" to enter an empty string	<ul style="list-style-type: none"> ■ Up to 32 alphanumeric characters ■ ? (for name criteria) 	–
Select bridge port	Number of the bridge port that you want to add to the new group	<ul style="list-style-type: none"> ■ 1 – n ■ all ■ ? (for a list of valid ports) 	–

Bridge Packet Filter Port Group Create Example (3500)

```
Select menu option (bridge/packetFilter/portGroup): create
Select port group mask {1-32|?}: 15
Select port group name {?}[]: Marketing
Port Group 1 - Marketing - has been created
Select bridge port(s) (1-6|all|?): 1,3,4

Select menu option (bridge/packetFilter/portGroup): create
Select port group mask {1-14,16-32|?}: 32
Select port group name {?} []: Sales
Port Group 2 - Sales - has been created
```

bridge packetFilter portGroup delete

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Deletes a selected port group.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

b pa p de

Important Considerations

- When you delete port groups from the system, those groups are no longer available for use in packet filters.
- When you delete a port group, the remaining port group IDs are automatically renumbered to maintain consecutive numbering.
- Possible values for port groups (n) depend on the number of user-defined port groups on the system.

Options

Prompt	Description	Possible Values	[Default]
Select port group	Number of the port group to delete	<ul style="list-style-type: none"> ■ 1 – n ■ ? (for a list of groups) 	Current port group
Delete port group?	Whether to delete the selected port group	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y

Bridge Packet Filter Port Group Delete Example

```
Select menu option (bridge/packetFilter/portGroup): delete
Select port group {1-2|?} [2]: 1
Delete port group (n,y) [y]: y
Port Group 1 - Marketing - has been deleted.
```

bridge packetFilter
portGroup addPort

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Adds ports to an existing port group.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

b pa p a

Important Considerations

- You add ports to an existing group by entering port identifiers at the prompts. At least one port group must exist before you can add ports.
- The maximum number of ports that a port group can contain is 32, which is the maximum number of ports on a switching system.
- Possible values for port groups (*m*) depend on the number of user-defined port groups on the system.
- Possible values for bridge ports (*n*) depend on the number of existing bridge ports on the system.

Options

Prompt	Description	Possible Values	[Default]
Select port group	Number of the port group to which you want to add a bridge port	<ul style="list-style-type: none">■ 1 – <i>m</i>■ ? (for a list of groups)	Current port group
Select bridge port	Number of the bridge port that you want to add to the selected port group	<ul style="list-style-type: none">■ 1 – <i>n</i>■ all■ ? (for a list of groups)	–

Bridge Packet Filter Port Group Add Port Examples

```
Select menu option (bridge/packetFilter/portGroup): add
Select port group {1-2|?} [2]: 2
Select bridge port(s) (1-6|all|?): 2
```

When you display port group 2, the display shows that port 2 is added:

```
Select menu option (bridge/packetFilter/portGroup): display
Select port group (1-2|all|?) [2]:

Port Group 2 - Sales
Port 2
Port 6

Port 5
```

**bridge packetFilter
portGroup
removePort**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes ports from a port group.

Valid Minimum Abbreviation

b pa p r

Important Considerations

- At least one group must exist before you can remove a port from a port group.
- Possible values for port groups (m) depend on the number of user-defined port groups on the system.
- Possible values for bridge ports (n) depend on the number of existing bridge ports on the system.

Options

Prompt	Description	Possible Values	[Default]
Select port group	Number of the port group from which you want to remove a bridge port	<ul style="list-style-type: none"> ■ 1 – m ■ ? (for a list of groups) 	Current port group
Select bridge port	Number of the bridge port that you want to remove from the selected port group	<ul style="list-style-type: none"> ■ 1 – n ■ all ■ ? (for a list of ports) 	–

Bridge Packet Filter Port Group Remove Port Examples

Select menu option (bridge/packetFilter/portGroup): **remove**

Select port group {1-2|?} [2]: **2**

Select bridge port(s) (1-6|all|?): **6**


Displaying port group 2 shows that port 6 is removed:

Select menu option (bridge/packetFilter/portGroup): **display**

Select port group (1-2|all|?) [2]:

Port Group 2 - Sales
Port 2

Port 5



ROUTING PROTOCOLS

Chapter 16 **Internet Protocol (IP)**


Chapter 17 **Virtual Router Redundancy Protocol (VRRP)**

Chapter 18 **IP Multicast**

Chapter 19 **Open Shortest Path First (OSPF)**

Chapter 20 **IPX**

Chapter 21 **AppleTalk**



INTERNET PROTOCOL (IP)

To route packets using the Internet Protocol (IP), you:

- Establish an IP routing interface
- Decide which IP options and routing protocols you want to use
- Enable IP routing

An IP routing interface defines the relationship between an IP virtual LAN (VLAN) and the subnetworks in the IP network. Each routing IP VLAN interface is associated with one VLAN that supports IP. The system has one interface defined for each subnet that is directly connected to it.

You can also choose between two different routing models when you establish an IP routing interface:

- VLAN-based routing
Because bridging is faster in normal circumstances, the system first tries to determine if it can bridge the frame before routing it.
- Router port-based routing
The system first tries to route packets that belong to recognized protocols, and then bridges all other packets. If the network or a portion of the network is devoted to routing IP frames, this model makes network traffic more efficient.

This chapter provides guidelines and other key information about how to configure IP in your system. This chapter addresses the commands in the `ip` menu except for `multicast` and `ospf`, which other chapters in this *Command Reference Guide* explain.

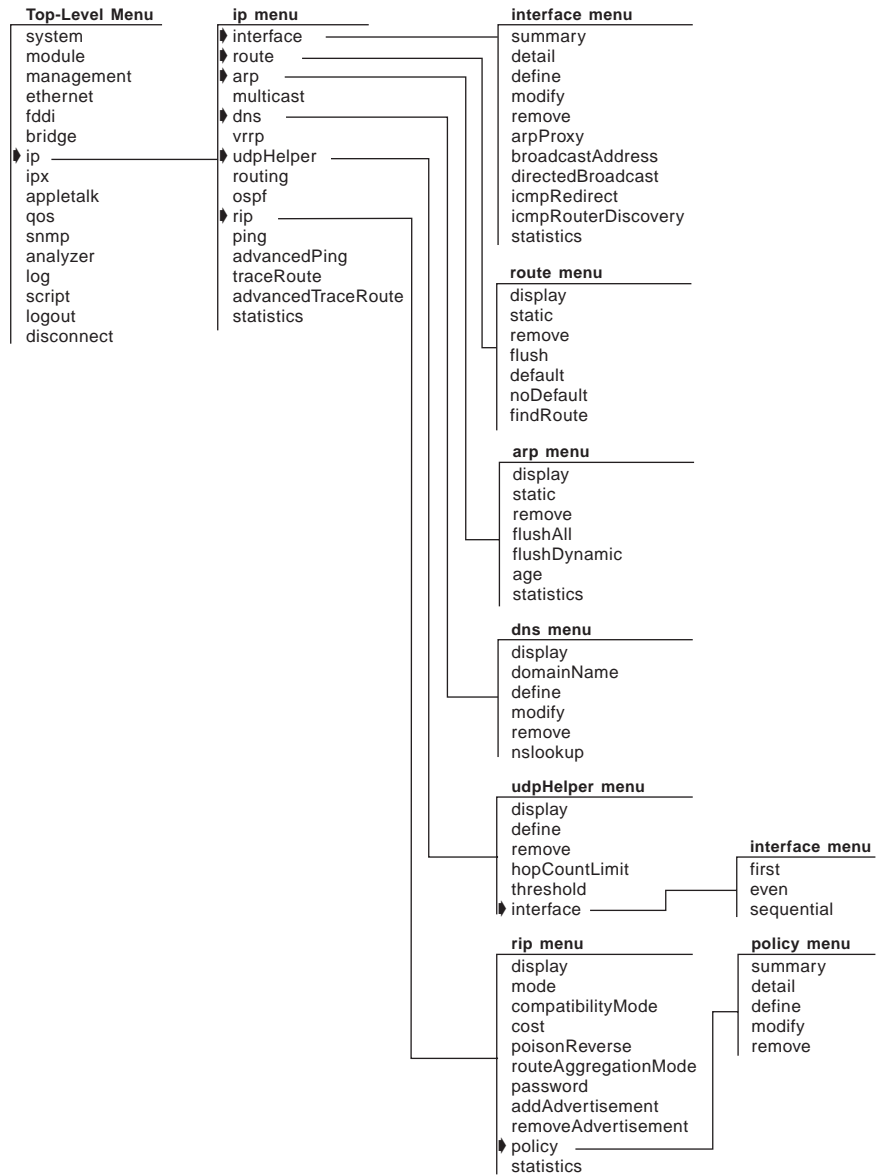
For information about IP multicast, see Chapter 18. For information about Open Shortest Path First (OSPF) routing using IP, see Chapter 19.



For more information about IP routing, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



ip interface summary***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays summary information about the IP interfaces that are configured on the system.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation`ip i su`**Important Considerations**

- When you enter the command, you are prompted for an interface index number even if you have only one interface defined.
- The first line in the output (the status line) indicates whether IP routing is enabled:
 - For CoreBuilder 9000 Layer 3 modules, it also indicates whether ICMP router discovery is enabled on the system.
 - For the CoreBuilder 3500, IP interface options (such as ICMP router discovery) appear under “ip interface detail” and are set on a per-interface basis.
- The Type field differs according to platform:
 - In the CoreBuilder 3500, which provides port-based routing and VLAN-based routing, the Type field displays whether the IP interface is VLAN-based or router port-based.
 - In all other platforms, which provide VLAN-based routing, the Type field displays whether the IP interface is used for VLAN traffic or for system management.
- The last (rightmost) field in the display differs according to platform:
 - In the CoreBuilder 3500, the ID field displays either the logical port number that is associated with a router port-based IP interface, or the VLAN interface index number that is associated with the IP interface.
 - In all other platforms, the VLAN index field displays the VLAN interface index number that is associated with the IP interface.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the IP interface whose summary information you want to display	<ul style="list-style-type: none">One or more selectable interface indexesall? (for a list of selectable interface indexes)	–

Fields in the IP Interface Summary Display

Field	Description
Index	Index number of the IP interface whose summary information you want to display
IP address	IP address of the interface, chosen from the range of addresses that the central agency assigned to your organization. This address is specific to your network and system.
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnetwork number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnetwork part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.
State	State of the IP interface. It indicates whether the interface is available for communications (up) or unavailable (down).
Type	<ul style="list-style-type: none">Type of interface: VLAN-based or router port-based (3500)Type of interface: VLAN or system (all other platforms)
ID (3500)	<ul style="list-style-type: none">Logical port number of the router port-based IP interface or the VLAN index that is associated with the IP interface
VLAN index (3900, 9000, 9300, 9400)	<ul style="list-style-type: none">VLAN index number that is associated with the IP interface

ip interface detail

Displays detailed information about the specified interfaces or all interfaces.

✓ 3500
9000
9400

Valid Minimum Abbreviation

`ip i det`

3900
9300

Important Consideration

- When you enter the command, you are prompted for an interface index number even if you have only one interface defined.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the IP interface whose summary information you want to display	<ul style="list-style-type: none"> ■ One or more configured indexes ■ all ■ ? (for a list of selectable indexes) 	–

Fields in the IP Interface Detail Display

Field	Description
ARP proxy	Whether ARP proxy is enabled or disabled for the specified interface.
Broadcast address	Broadcast address for the specified interface.
Directed broadcast	Whether the forwarding of a directed broadcast (all 1s in the host portion of the address) is enabled or disabled for the specified interface. (A directed broadcast is a packet that is sent to a specific network or series of networks.)
ICMP redirect	Whether ICMP redirect is enabled or disabled for the specified interface.
ICMP router discovery	Whether the ICMP Router Discovery is enabled or disabled for the specified interface
Index	Index number that is associated with the interface.
IP address	IP address of the interface, chosen from the range of addresses that the central agency assigned to your organization. This address is specific to your network and system.

Field	Description
Preference	Whether there is a preference being used for the specified interface. If ICMP router discovery is enabled, the system uses the routing interface with the highest preference level.
State	State of the IP interface. It indicates whether the interface is available for communications (<code>up</code>) or unavailable (<code>down</code>).
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnetwork number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnetwork part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.
Type	Type of interface: VLAN-based (<code>VLAN</code>) or router port-based (<code>port</code>).
Index	Index number of the IP VLAN that is associated with the IP interface.
MaxAdvInterval	Maximum advertisement interval between ICMP router discovery advertisements (in seconds).
MinAdvInterval	Minimum advertisement interval between ICMP router discovery advertisements (in seconds).
Holdtime	Length of time that ICMP router discovery advertisements are held valid.
State	State of the IP interface. It indicates whether the interface is available for communications (<code>up</code>) or unavailable (<code>down</code>).
ID	<ul style="list-style-type: none">■ Logical port number of the IP interface (if the Type field displays <code>port</code>)■ VLAN index number that is associated with the IP interface (if the Type field displays <code>VLAN</code>)

IP Interface Detail Example (3500)

Select menu option (ip/interface): **detail**

Select IP interfaces (1|all|?) [1]: **1**

IP routing is disabled

Index	IP address	Subnet mask	State	Type	ID
1	158.101.31.21	255.255.255.0	Down	Port	1

Index	ARP proxy	Broadcast address	Directed broadcast	ICMP redirect
1	enabled	255.255.255.255	enabled	enabled

Index	ICMP router discovery	Preference	MaxAdvInterval	MinAdvInterval	Holdtime
1	disabled	n/a	n/a	n/a	n/a

**ip interface define
(3500/9000 Layer 3)****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Defines an IP interface.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**`ip i def`**Important Considerations**

- When you define an IP interface, you must decide whether you want the interface to use router port-based routing or VLAN-based routing.
 - Router port-based routing directs the system to attempt to route the frame before it attempts to bridge the frame.

When you set up a router port-based IP interface, the system automatically creates a virtual LAN (VLAN) for the interface. The system assigns the next available VLAN index number to this VLAN.

- VLAN-based routing directs the system to attempt to bridge the frame before it attempts to route the frame.

When you set up a VLAN-based IP interface, you must first define a VLAN and select IP as a protocol supported by the VLAN, as described in Chapter 14.



If you define a router port, you do not have to define the VLAN first; the corresponding single-port VLAN is automatically defined.

- Port-based routing uses allClosed mode; VLAN-based routing uses either allClosed or allOpen mode. If you attempt to set up a router port-based IP interface in allOpen mode, the system notifies you with a message that it will change the VLAN mode to allClosed and recreate the default VLAN, clearing your existing VLANs in the process. Then the system prompts you to continue. (See the port-based router example at the end of this command description.)
- You cannot define a port-based IP interface on a port that is already a member of a VLAN-based IP interface. To change from one type of interface to another, you must redefine all IP interfaces and VLANs that are associated with that port.



CAUTION: Using different routing models (port-based or VLAN-based) in the same network without careful planning can adversely affect your network operations. Be sure that you understand the potential effects of router port-based and VLAN-based routing on your network. See the Implementation Guide for the CoreBuilder 3500 and for the CoreBuilder 9000 for detailed information about IP interfaces and VLANs.

Options

Prompt	Description	Possible Values	[Default]
IP address	IP address of the interface, chosen from the range of addresses that the central agency assigned to your organization. This address is specific to your network and system.	A valid IP address in the range of addresses that are assigned to your organization	–
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnetwork number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnetwork part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.	A valid subnet mask in accordance with the bits that are used for network number, subnetwork, and host number	Depends on specified IP address
Interface type	Whether to use router port-based routing or VLAN-based routing.	<ul style="list-style-type: none"> ■ port ■ vlan 	vlan
VLAN mode (for router port-based routing)	Whether the system removes all VLANs and recreates the default VLAN to enable port-based routing.	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Bridge port (for router port-based routing)	Port to use for port-based routing (may designate only one port).	<ul style="list-style-type: none"> ■ 1 – n ■ ? (for a list of selectable ports) 	–
VLAN interface index (for VLAN-based routing)	Index number of the IP VLAN that is associated with the IP interface; for a VLAN-based IP interface, you must assign this number. (Not applicable if you have more than one VLAN)	<ul style="list-style-type: none"> ■ A selectable VLAN index ■ ? (for a list of selectable VLAN indexes) 	Next available index number

IP Interface Define Example (Port-based Routing)

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter interface type (vlan/port) [vlan]: port
VLAN mode must be changed to allClosed to support this
interface.
This removes all VLANs, then re-creates the Default VLAN.
continue? (n,y) [y]: y
Select bridge port (1-6|?): 1
```

IP Interface Define Example (VLAN-based Routing)

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter interface type (vlan/port) [vlan]: vlan
Enter VLAN interface index {3|?} [3]: 3
```

ip interface define (3900/9300/9400/ 9000 Layer 2)

Defines an IP interface.

Valid Minimum Abbreviation

`ip i def`

Important Consideration

- Before you define the IP (routing) interface, first define a virtual LAN (VLAN) and select IP as a protocol that the VLAN supports, as described in Chapter 16.

Options

Prompt	Description	Possible Values	[Default]
IP address	IP address of the interface, chosen from the range of addresses that the central agency assigned to your organization. This address is specific to your network and system.	A valid IP address in the range of addresses that are assigned to your organization	–
Subnet mask	32-bit number that uses the same format and representation as an IP address. The subnet mask determines which bits in the IP address are interpreted as the network number, the subnetwork number, and the host number. Each IP address bit that corresponds to a 1 in the subnet mask is in the network/subnetwork part of the address. Each IP address bit that corresponds to a 0 is in the host part of the IP address.	A valid subnet mask in accordance with the bits that are used for network number, subnetwork, and host number	Depends on specified IP address
VLAN interface index	Index number of the IP VLAN that is associated with the IP interface. (Not applicable if you have more than one VLAN)	<ul style="list-style-type: none"> ■ A selectable VLAN index ■ ? (for a list of selectable VLAN indexes) 	Current value

IP Interface Define Example

```
Enter IP address: 158.101.1.1
Enter subnet mask [255.255.0.0]: 255.255.255.0
Enter VLAN interface index {2|?}[2]:2
```

ip interface modify

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Changes the configuration of an interface that you have already defined.

- ✓ 3500
- ✓ 9000
- ✓ 9400

Valid Minimum Abbreviation

ip i m

- ✓ 3900
- ✓ 9300

Important Consideration

- On the CoreBuilder 3500, you cannot modify the port number (router port-based routing) after it has been defined because of the associated virtual LAN (VLAN); you must remove the interface and then redefine it.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number that is associated with the interface that you want to modify. (Not applicable if you have more than one interface)	<ul style="list-style-type: none">■ A selectable IP interface index■ ? (for a list of selectable indexes)	Current value
IP address	IP address of the interface that you want to modify.	A valid IP address in the range of addresses that are assigned to your organization	Current IP address
Subnet mask	Subnet mask for the interface that you want to modify.	A valid subnet mask in accordance with the bits that are used for network number, subnetwork, and host number	Current subnet mask
VLAN interface index (for VLAN-based routing)	Index number of the IP VLAN that is associated with the IP interface; for a VLAN-based IP interface, you must assign this number. (Not applicable if you have more than one VLAN)	<ul style="list-style-type: none">■ A selectable VLAN index■ ? (for a list of selectable VLAN indexes)	Current value

ip interface remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Removes an IP interface from the system's routing table.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip i re`

✓ 3900

✓ 9300

Important Considerations

- Before you remove the interface, remove any static entries in the routing table or the Address Resolution Protocol (ARP) cache.
- On the CoreBuilder 3500, if you remove a router port-based IP interface, the system removes the virtual LAN (VLAN) that is associated with it as well.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number that is associated with the interfaces that you want to remove (Not applicable if you have more than one interface)	<ul style="list-style-type: none">■ One or more selectable interface indexes■ all■ ? (for a list of selectable interface indexes)	Current value

ip interface arpProxy ***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

✓ 3500
✓ 9000
9400

On a per-interface basis, enables or disables ARP proxy, which helps end stations on a subnetwork reach remote subnetworks that do not have routing capabilities or a default gateway configured.

Valid Minimum Abbreviation

ip i a

Important Considerations

- When ARP proxy is enabled and an end station sends an Address Resolution Protocol (ARP) request for a remote network, the system determines if it has the best route and then answers the ARP request by sending its own MAC address to the end station. The end station then sends the frames for the remote destination to the system, which uses its own routing table to reach the destination on the other network.
- When an interface is defined, the default ARP proxy state is `enabled`.
- The end stations must view the entire network configuration as one network (that is, by using a smaller subnet mask).
- Evaluate prolonged use of ARP proxy because it has some drawbacks, including increased ARP traffic and a need for larger ARP tables to handle the mapping of IP addresses to MAC addresses.

Options

Prompt	Description	Possible Values	[Default]
Interface	Index number for the interface for which you want to enable or disable ARP proxy. (Not applicable if you have more than one interface)	<ul style="list-style-type: none">■ One or more selectable interface indexes■ all■ ? (for a list of selectable interface indexes)	Current value
ARP proxy state	Whether you want to implement ARP proxy on an interface. The system prompts you for a state for each interface.	<ul style="list-style-type: none">■ enabled■ disabled	Current value

IP Interface ARP Proxy Example (3500)

```
Select menu option (ip/interface): arpproxy  
Select IP interfaces (1,2|?|all):2  
  Interface 2 - Enter proxy state (disabled, enabled)  
[enabled]: enabled
```

ip interface
broadcastAddress

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

On a per-interface basis, assigns the broadcast address that the system uses to forward the received directed broadcasts and advertise Routing Information Protocol (RIP) packets.

Valid Minimum Abbreviation

ip i b

Important Considerations

- You assign the broadcast address on a per-interface basis.
- When an IP interface is configured, its default broadcast address is 255.255.255.255.
- The broadcast address that you specify affects the RIP advertisement address that is used for the RIP interface. You see the specified broadcast address as the advertisement address under the RIP menus. See “ip rip display” later in this chapter for information about the RIP interface display.
- You cannot change the broadcast address for an interface if you have added any RIP advertisement addresses to that interface. See “ip rip addAdvertisement” later in this chapter for more information.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces to which you want to assign a broadcast address (Not applicable if you have more than one interface)	<ul style="list-style-type: none">■ One or more interface indexes■ all■ ? (for a list of selectable interface indexes)	Current value
Broadcast address per interface	Broadcast address that you want to assign to an interface	A valid address	Current address

ip interface directedBroadcast

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Specifies whether the forwarding of a directed broadcast (all 1s in the host portion of the address) is enabled or disabled for a specified interface. A *directed broadcast* is a packet that is sent to a specific network or series of networks.

Valid Minimum Abbreviation

ip i di

Important Considerations

- You define the directed broadcast state on a per-interface basis.
- When the state is `enabled` and the system determines that the destination is different from the interface that is receiving the directed broadcast, the system uses the broadcast address that is defined for this interface to forward the directed broadcast.
- You can disable the forwarding of a directed broadcast if security is an issue.
- By default, the directed broadcast state is `enabled`.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index numbers of the interfaces to which you want to enable or disable the forwarding of a directed broadcast. (Not applicable if you have more than one interface)	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Current value
Directed broadcast state	Whether you want to implement direct broadcast on an interface. The system prompts you for a state for each interface.	<ul style="list-style-type: none"> ■ enabled ■ disabled 	Current value

IP Interface Directed Broadcast Example (3500)

```
Select menu option (ip/interface): directedBroadcast
Select IP interfaces (1,2|all|?):2
Interface 2 - Enter directed broadcast state
(disabled, enabled) [enabled]:
```

**ip interface
icmpRedirect**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables or disables the transmission of an Internet Control Message Protocol (ICMP) redirect to the sender of a frame to indicate that there is a better gateway available to handle the frame than this routing interface.

Valid Minimum Abbreviation

`ip i icmpre`

Important Considerations

- The software determines whether there is a better path for the frame by determining whether the source interface is the same as the destination interface and whether the frame's sender is on a directly connected network. If the software determines that a received frame has a better path available through another gateway:
 - It sends an ICMP redirect message back to the originator of the frame indicating the better gateway to use in the future
 - It routes the frame to the gateway
- ICMP redirect can be set on a per-interface basis.
- For better performance or if you have applications that ignore ICMP redirects, disable the ability of the interface to send ICMP redirects.
- If you have two interfaces that belong to virtual LANs (VLANs) that share a given port and you want to completely disable ICMP redirects for that port, disable the redirects for each interface that shares that port. If you disable it for only one interface and enable it for the other, you may not get the performance improvement that you want.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces to which you want to enable or disable the transmission of an ICMP redirect to the sender of a frame. (Not applicable if you have more than one interface)	<ul style="list-style-type: none">■ One or more selectable interface indexes■ all■ ? (for a list of selectable interface indexes)	Current value
ICMP redirect state	Whether you want to implement ICMP redirect state on an interface. The system prompts you for a state for each interface.	<ul style="list-style-type: none">■ enabled■ disabled	disabled, or current value

**ip interface
icmpRouterDiscovery**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables or disables Internet Control Message Protocol (ICMP) router discovery, which enables hosts that are attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers and determine which router to use for a default gateway.

Valid Minimum Abbreviation

```
ip i icmpro
```

Important Considerations

- ICMP router discovery can be set on a per-interface basis.
- When you enable the state for an interface, the system prompts you for a preference. (See RFC 1256.) By default, this preference level is 0. Use the preference to control the use of certain routers as the default router. The host uses the router with the highest preference level.
- An appropriately configured end station can locate one or more routers on the LAN to which it is attached. The end station then automatically installs a default route to each of the routers that are running Internet Control Message Protocol (ICMP) router discovery. You do not need to manually configure a default route. ICMP redirect messages subsequently channel the IP traffic to the correct router.
- You can configure only certain end stations to work with the ICMP router discovery protocol. See the documentation for your workstation to determine whether you can configure it to work with this protocol.
- You can configure and display three timers for ICMP router discovery on the CoreBuilder 3500:
 - **Maximum advertisement interval** — The maximum time interval between advertisements.
 - **Minimum advertisement interval** — The minimum time interval between advertisements.
 - **Advertisement holdtime** — The length of time that advertisements are held valid.



The ranges for minimum advertisement interval depend on the set values for maximum advertisement interval and the holdtime range depends on the input values for both the maximum and minimum advertisement intervals.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to enable or disable ICMP router discovery. (Not applicable if you have more than one interface)	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Current value
Router discovery state	Whether you want to implement ICMP router discovery on an interface. The system prompts you for a state for each interface.	<ul style="list-style-type: none"> ■ enabled ■ disabled 	disabled, or current value
Preference	If you select <code>enabled</code> , the host interprets an unsigned integer as a 32-bit signed two's-complement integer that represents the preference level to associate with the interface. Higher values produce higher preference levels. The minimum value is reserved so that the address is not used as a default router address, only for specific IP destinations.	minimum value (hex 80000000) -2 ³¹ to 2 ³¹	0
Maximum advertisement interval	Maximum interval between advertisements.	4 – 1800 seconds	600
Minimum advertisement interval	Minimum interval between advertisements.	3 – 600 seconds	450
Advertisement holdtime	Length of time that advertisements are held valid.	600 – 9000 seconds	1800

IP Interface ICMP Router Discovery Example (3500)

Select menu option (ip/interface): **icmprouterdiscovery**

Select IP interfaces (1|all|?) [1]: **1**

Interface 1 - Enter router discovery state (disabled,enabled) [disabled]: **enabled**

Interface 1 - Enter router discovery preference [0]:

Interface 1 - Enter maximum advertisement interval (4-1800) [600]:

Interface 1 - Enter minimum advertisement interval (3-600) [450]:

Interface 1 - Enter advertisement holdtime (600-9000) [1800]:

ip interface statistics***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays IP interface statistics on a per-interface basis.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ip i st`**Important Consideration**

- The system prompts you for an interface index number even if you have only one interface defined.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the interface whose statistics you want to display	<ul style="list-style-type: none"> ■ One or more configured interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	–

Fields in the IP Interface Statistics Display

Field	Description
inAddrErrors	Number of datagrams that the IP station discarded because of an error in the source or destination IP address
inCsumErrors	Number of datagrams that were dropped because of a checksum error
inDelivers	Number of datagrams that the IP station delivered to local IP client protocols
inDiscards	Number of packet receive discards
inForwards	Total number of packets that were forwarded (that is, routed through hardware or software or both)
inHdrErrors	Number of datagrams that the IP station discarded because the IP header contained errors
inReceives	Total number of IP datagrams that were received, including those with errors
inSameSegment	Number of packets that were received on an interface and that need to be forwarded out on the same interface

Field	Description
inTtlExceeds	Number of packets that were received on an interface and that need to be forwarded, but that have an IP header TTL value of less than 2
outDiscards	Number of packet transmit discards
outForwards	Total number of packets that a router has forwarded to an outbound interface (that is, routed through hardware or software or both)

ip route display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Displays the system's routing table to determine which routes to other IP networks are configured and whether the routes are operational.

Valid Minimum Abbreviation

`ip route di`

Important Considerations

- For the CoreBuilder 3500 only, the system prompts you for an IP address and subnet mask. As a result, you can display only a subset of routes instead of all routes. To see all entries in the table, simply press Enter at these prompts.
- The first line in the output (the status line) indicates whether IP routing is enabled:
 - For the CoreBuilder 9000 Layer 3 module, it also indicates whether Internet Control Message Protocol (ICMP) router discovery is enabled on the system.
 - For the CoreBuilder 3500, IP interface options (such as ICMP router discovery) appear under "ip interface detail" earlier in this chapter and are set on a per-interface basis.

Options (3500 only)

Prompt	Description	Possible Values	[Default]
IP address	IP address (and its corresponding subnet mask) for which to display only those routes that match the bits set in it	<ul style="list-style-type: none"> ■ A valid IP address ■ 0.0.0.0 (displays all entries) 	0.0.0.0
Subnet mask	Subnet mask for the specified IP address for which to display only those routes that match the bits set in it	A valid subnet mask of a specified IP address	Current value

Fields in the IP Route Display

Field	Description
Destination	IP address of the destination network, subnetwork, or host. This field can also identify a default route, which the system uses to forward packets that do not match any other routing table entry. You may want to use the default route in place of routes to numerous destinations that all have the same gateway IP address.
Subnet mask	Subnet mask that is associated with the IP address of the destination network, subnetwork, or host.
Metric	Associated cost of sending a packet to the destination. The system includes the metric in its RIP and OSPF updates to allow other routers to compare routing information received from different sources.
Gateway	Address that directs the router how to forward packets whose destination addresses match the route's IP address and subnet mask. The system forwards such packets to the indicated gateway.
Status	Status of the route. See the following status table.
TTL	Time To Live — Time remaining before the route expires or is reset.

Status for Routes

Field	Description
Direct	Route is for a directly connected network
Learned	Route was learned using indicated protocol
Learned RIP-Zombie	Route was learned but is partially timed out. This condition is applied to all learned routes reached by an interface gateway which is in the down state.
Learned RIP2	Route was learned using RIP-2 protocol
Local	Actual interface address
Static	Route was statically configured
Timed out	Route has timed out and is no longer valid

ip route static *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a static route.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip route s`

✓ 3900

✓ 9300

Important Considerations

- Before you can define static routes, you must define at least one IP interface. See “ip interface define (3500/9000 Layer 3)” earlier in this chapter for more information.
- For the CoreBuilder 3500, you can define up to 256 static routes.
- For the other platforms, you can define up to 64 static routes.
- Static routes remain in the table; you must remove them before you can remove the corresponding interface.
- Static routes take precedence over dynamically learned routes to the same destination
- Static routes are included in periodic Routing Information Protocol (RIP) updates that the system sends.

Options

Prompt	Description	Possible Values	[Default]
Destination IP address	IP address of the destination network, subnetwork, or host for this route	A valid IP address	–
Subnet mask	Subnet mask for the specified IP address	A valid subnet mask	Based on specified IP address
Gateway IP address	IP address of the gateway that this route uses	A valid router address	–

ip route remove

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Deletes an existing route.

Valid Minimum Abbreviation

ip route r

Important Consideration

- When you enter the command, the system deletes the route immediately from the routing table. You are not prompted to confirm the deletion.

Options

Prompt	Description	Possible Values	[Default]
Destination IP address	IP address of the route that you want to delete	A valid IP address	–
Subnet mask	Subnet mask for the specified IP address	A valid subnet mask	Based on specified IP address

ip route flush *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes all learned routes from the routing table.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip route fl`

✓ 3900

✓ 9300

Important Considerations

- The system flushes all learned routes from the routing table immediately. You are not prompted to confirm the deletion.
- Flushing the routing table does not cause the Routing Information Protocol (RIP) to update the routing table. You must change the metric to update the routing table.

ip route default

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Adds a default route to the routing table immediately.

Valid Minimum Abbreviation

ip route de

Important Considerations

- If you define a default route, the system uses it to forward packets that do not match any other routing table entry. The system can learn a route through the routing protocol, or you can statically configure a default route.
- The system can learn a default route.
- If the routing table does not contain a default route, the system cannot forward a packet that does not match any other routing table entry. When the system drops the packet, it sends an Internet Control Message Protocol (ICMP) `destination unreachable` message to the host that sent the packet.
- On the CoreBuilder 3500 or the CoreBuilder 9000 Layer 3 module, you establish a static sink default route, so that the system can advertise itself as a default router. The static sink default route is not used in any of the system's forwarding decisions because it does not have a valid next-hop gateway, but it can be advertised to all of the system's neighbors (unless you establish IP policies to prevent the advertisement).

Defining a static sink default route causes the route to be advertised through any IP protocols that you have configured on the system (for example, Open Shortest Path First (OSPF) and RIP). For more information about static sink default routes, see the *Implementation Guide* for the CoreBuilder 3500 or for the CoreBuilder 9000.

Options

Prompt	Description	Possible Values	[Default]
Gateway IP address	IP address of the route that you want to add as the default	<ul style="list-style-type: none">■ A valid IP address■ 0.0.0.0 (static sink default route)	–

ip route noDefault *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes the default route.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip route n`

✓ 3900

✓ 9300

Important Consideration

- The system deletes the default route from the routing table immediately after you enter the command. You are not prompted to confirm the deletion.

ip route findRoute

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

3900

9300

Searches for a route in the routing table.

Valid Minimum Abbreviation

ip route fi

Important Considerations

- This command enables you to find a route using an IP address or a host name, as long as the Domain Name System (DNS) is configured.
- When you enter this command with a valid IP address or host name, the system displays the routing table entry.

Options

Prompt	Description	Possible Values	[Default]
IP address (or host name)	IP address of the route that you want to find, or a host name, if DNS is configured	<ul style="list-style-type: none">■ A valid IP address■ A valid host name	0.0.0.0, or current value

ip arp display ***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays the contents of the Address Resolution Protocol (ARP) cache for each interface on the system.

Valid Minimum Abbreviation

`ip ar d`

Important Considerations

- The system uses the ARP cache to find the MAC addresses that correspond to the IP addresses of hosts and other routers on the same subnetworks. Each device that participates in routing maintains an *ARP cache*, which is a table of known IP addresses and their corresponding MAC addresses.
- The first line in the output (the status line) indicates whether IP routing is enabled:
 - For the CoreBuilder 9000 Layer 3 module, it also indicates whether Internet Control Message Protocol (ICMP) router discovery is enabled on the system.
 - For the CoreBuilder 3500, IP interface options (such as ICMP router discovery) appear under “ip interface detail” earlier in this chapter and are set on a per-interface basis. The second status line indicates the number of entries in the ARP cache.

Fields in the IP ARP Display

Field	Description
Circuit	Circuit identifier
Hardware address	MAC address that is mapped to the IP address
I/F	Index number of the associated interface
IP address	IP address of the interface
Type	Type of entry — <i>static</i> or <i>dynamic</i>

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

ip arp static

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines a static Address Resolution Protocol (ARP) cache entry on the system.

Valid Minimum Abbreviation

ip ar s

Important Considerations

■ For the CoreBuilder 3500, you can define up to 128 static ARP entries.

■ For the other platforms, you can define up to 64 entries.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the interface for which you want to define a static ARP entry	<div>■ A selectable interface index</div> <div>■ ? (for a list of selectable interface indexes)</div>	–
IP address	IP address to use in the entry	A valid IP address	–
MAC address	Hardware address to use in the entry	A valid MAC address in the format xx-xx-xx-xx-xx-xx	–

IP ARP Static Example

Select interface index {1-2|?} 2

Enter IP address: 158.101.12.12

Enter MAC address: 00-00-00-00-00-01

ip arp remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes an entry from the Address Resolution Protocol (ARP) cache (for example, if the MAC address has changed).

Valid Minimum Abbreviation

`ip ar rem`

Important Considerations

- When you enter the command, the system deletes the entry from the cache immediately. You are not prompted to confirm the deletion.
- If necessary, the system subsequently uses ARP to find the new MAC address that corresponds to that IP address.

Options

Prompt	Description	Possible Values	[Default]
IP address	IP address for the entry that you want to delete	A valid IP address	–

ip arp flushAll *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes all entries from the Address Resolution Protocol (ARP) cache.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip ar flushA`

✓ 3900

✓ 9300

Important Considerations

- To flush dynamic entries only, see “ip arp flushDynamic” next in this chapter.
- When you enter the command, the system deletes all entries from the cache immediately. You are not prompted to confirm the deletion.

ip arp flushDynamic ***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Deletes all dynamic (learned) entries from the Address Resolution Protocol (ARP) cache.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip ar flushD`

✓ 3900

✓ 9300

Important Considerations

- To flush *all* entries, static and dynamic, see the previous “ip arp flushAll” option.
- When you enter the command, the system deletes all dynamic entries from the cache immediately. You are not prompted to confirm the deletion.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

ip arp age *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Sets the age time for dynamic Address Resolution Protocol (ARP) cache entries.

Valid Minimum Abbreviation

ip ar a

Important Considerations

- The *age time* determines how long, in minutes, that the dynamic entries remain in the ARP cache before they are removed.
- By default, the system flushes the entry from the cache when it reaches the age time.
- A value of 0 indicates no age time, and the entry remains in the table until you remove it with the `ip arp remove` option or flush the ARP cache with the appropriate `flush` option.

Options

Prompt	Description	Possible Values	[Default]
Age time	Time that dynamic entries remain in the ARP cache	<ul style="list-style-type: none">■ 1 – 1440 minutes■ 0 (to disable aging)	15 (factory default), or current value

ip arp statistics***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

✓ 3500

✓ 9000

9400

3900

9300

Displays detailed information about the specified interfaces or all interfaces.

Valid Minimum Abbreviation`ip ar status`**Important Considerations**

- Your system tracks the number of times that a particular Address Resolution Protocol (ARP) event occurs.
- If a port that has multiple IP interfaces associated with it receives an ARP frame that is discarded because of an address mismatch, the inReceives and inDiscards statistics are incremented for the *first* interface of all the interfaces that are associated with the port.
- The system supports baselining for ARP statistics.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the IP interface from which to select ARP statistics	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	1

Fields in the IP ARP Statistics Display

Field	Description
inDiscards	<p>Received ARP frames that have been discarded due to one of the following reasons:</p> <ul style="list-style-type: none"> ■ Frame had a source address that did not match any directly connected IP interface that was associated with the port on which it was received ■ Frame contained an invalid header ■ Frame was not an ARP request or an ARP reply
inReceived	ARP frames (requests, replies, and discards) that were received on an IP interface

Field	Description
inReplies	ARP reply frames that were received on an IP interface
inRequests	ARP request frames that were received on an IP interface
outIfdown	Failure of the system to send one of the following three frames because the state of the IP interface was down: <ul style="list-style-type: none">■ ARP request■ ARP reply■ IP frame to be forwarded (pending ARP resolution)
outMemErrors	Failure of the system to allocate memory to transmit either an ARP request or an ARP reply
outReplies	ARP replies that were transmitted from an IP interface
outRequests	ARP requests that were transmitted from an IP interface

ip dns display *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Displays the current domain name and the name servers that are associated with it.

Valid Minimum Abbreviation

`ip d di`

Important Considerations

- The *Domain Name System (DNS)* client provides DNS lookup functionality to the CoreBuilder IP ping and traceRoute features. You can specify a host name rather than an IP address when you perform various operations (for example, when you use `ping` or `traceRoute` to contact an IP station).
- With the DNS commands, you specify one or more name servers that are associated with a domain name. Each name server maintains a list of IP addresses and their associated host names. When you use `ping` or `traceRoute` with a host name, the DNS client attempts to locate the name on the name servers that you specify. When the DNS client locates the name, it resolves it to the associated IP address.
- See UNIX Network File System (NFS) documentation for information about how to create and maintain lists of domain names and IP addresses on the name servers.

Fields in the IP DNS Display

Field	Description
Domain name	Name of the domain name (up to 79 alphanumeric characters)
Name server	Name server that is associated with the domain

ip dns domainName

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Changes the name of a currently defined domain.

- ✓ 3500
- ✓ 9000
- ✓ 9400
- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

ip d do

Important Considerations

- You can specify a domain name with up to 79 alphanumeric characters.
- Use single quotation marks (' ') around any string that has embedded spaces. Use double quotation marks (" ") to enter an empty string.

Options

Prompt	Description	Possible Values	[Default]
Domain name	Name of the domain. The name can be up to 79 characters long.	<ul style="list-style-type: none">■ A valid domain name■ ? (to get information about specifying a domain name)	– (or current name)

ip dns define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a new name server IP address to associate with the current domain name.

Valid Minimum Abbreviation

`ip d de`

Important Considerations

- When the system accepts the new IP address, it displays a message like the following:

Server's IP address xxxxx is added to the DNS database

- The system assigns an index number to the new IP address. Use this index number to modify or remove this IP address.

Options

Prompt	Description	Possible Values	[Default]
Name server IP address	IP address of the name server that you want to define	A valid IP address	–

ip dns modify

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Modifies a currently defined name server IP address.

Valid Minimum Abbreviation

ip d m

Important Considerations

- When you enter the command, the system displays the list of name server addresses and the index number that is associated with each.
- The system assigns an index number to the new IP address. Use this index number to modify this IP address.

Options

Prompt	Description	Possible Values	[Default]
Index	Index number of the name server IP address that you want to modify	<ul style="list-style-type: none">■ A selectable server index number■ ? (for a list of selectable server indexes)	–
Name server IP address	New IP address of the name server that you want to use	A valid IP address	–

ip dns remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Deletes a previously defined name server IP address.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

`ip d r`

Important Consideration

- When you enter the command, the system displays the list of name server addresses and the index number that is associated with each.

Options

Prompt	Description	Possible Values	[Default]
Index	Index number of the name server IP address that you want to remove	<ul style="list-style-type: none">■ A selectable server index number■ ? (for a list of selectable server indexes)	–

ip dns nslookup

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Maps an IP address to a host name or a host name to an IP address on a name server.

Valid Minimum Abbreviation

ip d n

Important Considerations

- Specify a host name or IP address at the prompt.
- Enter a string of up to 255 characters.
- Use single quotation marks (' ') around any string with embedded spaces. Use double quotation marks (" ") to enter an empty string.

Options

Prompt	Description	Possible Values	[Default]
IP address or host name	IP address or host name that you want to map	<ul style="list-style-type: none">■ A host name of up to 255 characters■ A valid IP address	–

ip udpHelper display***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

✓ 3500
✓ 9000
9400

Displays the BOOTP (bootstrap protocol) hop count and the threshold configuration. Also lists the ports with their IP forwarding addresses that are defined in your system.

Valid Minimum Abbreviation

ip u di

3900
9300

Important Considerations

- With UDP Helper, you can send User Datagram Protocol (UDP) packets between routed networks. UDP Helper provides support for UDP services such as BOOTP and DHCP (Dynamic Host Configuration Protocol), which rely on the BOOTP relay agent.
- When you configure the logical BOOTP port, you can boot hosts through the router. UDP Helper also provides a relay agent for DHCP broadcasts. UDP packets that rely on the BOOTP relay agent are modified and then forwarded through the router.
- BOOTP (including DHCP) uses UDP port 67.
- With UDP Helper, you can configure the amount of time that a UDP packet is forwarded between subnetworks. The system discards UDP packets based on the hop count and the seconds value only for BOOTP and DHCP.

Fields in the IP udpHelper Display

Field	Description
UDP port	UDP port number — usually the value 67
Forwarding address	Forwarding address that is used for UDP packets

ip udpHelper define

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines port numbers or IP forwarding addresses for the UDP Helper.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Valid Minimum Abbreviation

ip u de

Important Considerations

- You can have up to 63 combinations of port numbers and IP forwarding addresses per router.
- You can have multiple IP address entries for the same ports.

Options

Prompt	Description	Possible Values	[Default]
UDP port number	Port number for UDP	1 – 65535	67 (factory default), or current value
IP forwarding address	Forwarding addresses that are used for UDP packets	A valid IP address	–

ip udpHelper remove***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***✓ **3500**✓ **9000****9400****3900****9300**

Removes a port number or IP forwarding address that has been defined for UDP Helper.

Valid Minimum Abbreviation`ip u r`**Important Consideration**

- The system immediately removes the port numbers and IP forwarding addresses that you specified. You are not prompted to confirm the deletion.

Options

Prompt	Description	Possible Values	[Default]
UDP port number	UDP port number that you want to remove	1 – 65535	67 (factory default), or current value
IP forwarding address	Forwarding addresses that you want to remove	A valid IP address	–

ip udpHelper
hopCountLimit

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the maximum hop count to specify how many steps the system uses to forward a packet through the router.

Valid Minimum Abbreviation

ip u h

Options

Prompt	Description	Possible Values	[Default]
BOOTP hop count limit	Maximum number of hops to allow for UDP packet forwarding	0 – 16	4 (factory default), or current value

**ip udpHelper
threshold**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the maximum number of times that the system forwards a packet to the network.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

`ip u t`

3900
9300

Important Consideration

- By default, there is no threshold (0).

Options

Prompt	Description	Possible Values	[Default]
BOOTP relay threshold	Maximum number of times that the system forwards a packet to the network	0 – 65535	0 (factory default), or current value

**ip udpHelper
interface first*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Configures UDP Helper to support overlapped IP interfaces by using the first interface.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation`ip u i f`

3900

9300

Important Considerations

- *Overlapped* IP interfaces are multiple logical interfaces that are defined for a single physical port. You can specify how UDP Helper forwards packets from overlapped IP interfaces with one of three interface options (`first`, `even`, or `sequential`).
- The value `first` directs the system to use the first overlapped IP interface as the source network for forwarded packets.
- The system implements your selection immediately. You can view the UDP Helper configuration when you configure the forwarding address.

**ip udpHelper
interface even**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Configures UDP Helper to support overlapped IP interfaces by evenly distributing interfaces.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

`ip u i e`

3900

9300

Important Considerations

- The value `even` directs the system to hash the client's MAC address to determine the source network for forwarded packets. This arrangement evenly distributes the interface among those on the network.
- The system implements your selection immediately. You can view the UDP Helper configuration when you configure the forwarding address.

**ip udpHelper
interface sequential**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Configures UDP Helper to support overlapped IP interfaces by distributing the interfaces sequentially.

Valid Minimum Abbreviation

`ip u i s`

Important Considerations

- The value `sequential` directs the system to assign each overlapped IP interface, in turn, as the source network for forwarded packets.
- The system implements your selection immediately. You can view the UDP Helper configuration when you configure the forwarding address.

ip routing *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Controls whether the system forwards or discards IP packets that are addressed to other hosts.

Valid Minimum Abbreviation

`ip routi`

Important Considerations

- When you enable IP routing, the system acts as a standard IP router: it forwards IP packets from one subnetwork to another when required.
- When you disable IP routing, the system discards any IP packets that are not addressed directly to one of its defined IP interfaces.
- By default, IP routing is disabled on the system.

Options

Prompt	Description	Possible Values	[Default]
IP routing state	Whether IP routing is implemented on the system	<ul style="list-style-type: none">■ enabled■ disabled	disabled

ip rip display **For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

✓ 3500
✓ 9000
✓ 9400

Displays information about the Routing Information Protocol (RIP) interfaces on the system. RIP is one of the IP Interior Gateway Protocols (IGPs). When RIP is enabled, the system dynamically configures its routing tables.

✓ 3900
✓ 9300

Valid Minimum Abbreviation

`ip ri d`

Important Considerations

- The output for this display differs according to platform.
- The first line in the output (the status line) indicates whether IP routing is enabled:
 - For the CoreBuilder 9000 Layer 3 module, it also indicates whether Internet Control Message Protocol (ICMP) router discovery is enabled on the system.
 - For the CoreBuilder 3500, IP interface options (such as ICMP router discovery) appear under “ip interface detail” earlier in this chapter and are set on a per-interface basis. The rest of the output contains more RIP interface information.
- The four available RIP modes are as follows:
 - **Disabled** — The system ignores all incoming RIP packets and does not generate any RIP packets of its own.
 - **Learn** — The system processes all incoming RIP packets, but it does not transmit RIP updates.
 - **Advertise** (3500 and 9000 only) — The system broadcasts RIP updates, but it does not process incoming RIP packets.
 - **Enabled** (3500 and 9000 only) — The systems broadcasts RIP updates and processes incoming RIP packets.
- An advertising router sends a RIP message every 30 seconds with both the IP address and a *metric* (the distance to the destination from that router) for each destination. Each router through which a RIP packet must travel to reach a destination equals one *hop*.

Fields in the IP RIP Display

Field	Description
Advertisement Addresses (3500 and 9000 only)	List of available advertisement addresses. The list is used for RIP-2 updates only if the RIP-1 compatibility mode is enabled. RIP-1 always uses advertisement addresses.
Compatibility Mode (3500 only)	Whether RIP 1 compatibility mode is enabled or disabled (by default, disabled).
Cost (3500 and 9000 only)	RIP cost for the interface (by default, 1).
Index	Index number of the interface.
Poison Reverse (3500 and 9000 only)	Whether poison reverse mode is enabled or disabled (by default, enabled).
RIP-1 Mode	Mode for RIP-1. If you disable RIP-1, the output lists the state as <code>off</code> . Other modes are <code>learn</code> (default), <code>advertise</code> , and <code>enabled</code> .
RIP-2 Mode	Mode for RIP-2. If you disable RIP-2, the output lists the state as <code>off</code> . Other modes are <code>learn</code> (default), <code>advertise</code> , and <code>enabled</code> .
Route Aggregate (3500 only)	Whether Route Aggregation mode is enabled or disabled

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

ip rip mode *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

On a per-interface basis, sets one of four RIP Version 1 (RIP-1) modes on the system. For all platforms except the CoreBuilder 9000, also allows you to set RIP Version 2 (RIP-2) modes.

Valid Minimum Abbreviation

ip ri m

Important Considerations

- Platforms except the CoreBuilder 9000 support RIP Version 1 as well as RIP Version 2. For each interface, you select a RIP Version 1 mode and a RIP Version 2 mode. The default RIP Version 1 mode for all platforms is `learn`. The default RIP Version 2 mode is `learn`.
- The four available RIP modes are as follows:
 - **Disabled** — The interface ignores all incoming RIP packets and does not generate any RIP packets of its own.
 - **Learn** — The interface processes all incoming RIP packets, but it does not transmit RIP updates. This is the default RIP mode.
 - **Advertise** (3500 and 9000 only) — The interface broadcasts RIP updates, but it does not process incoming RIP packets.
 - **Enabled** (3500 and 9000 only) — The interface broadcasts RIP updates and processes incoming RIP packets.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the RIP mode	<ul style="list-style-type: none">■ One or more selectable interface indexes■ all■ ? (for a list of selectable interface indexes)	Previous entry, if applicable

Prompt	Description	Possible Values	[Default]
RIP mode, Version 1	Selected RIP Version 1 mode that determines how the interface handles RIP 1 packets and updates	<ul style="list-style-type: none"> ■ disabled ■ learn ■ advertise (3500/9000) ■ enabled (3500/9000) 	learn (factory default), or current value
RIP mode, Version 2 (not 9000)	Selected RIP mode that determines how the interface handles RIP 2 packets and updates	<ul style="list-style-type: none"> ■ disabled ■ learn ■ advertise (3500 only) ■ enabled (3500 only) 	learn (factory default), or current value

IP RIP Mode Example

Select IP interfaces (1,2|all|?): **1**

Interface 1 - Enter RIP Version 1 mode
(disabled,learn,advertise,enabled) [learn]: **disabled**

Interface 1 - Enter RIP Version 2 mode
(disabled,learn,advertise,enabled) [learn]: **enabled**

ip rip compatibilityMode

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

On a per-interface basis, sets the RIP Version 1 compatibility mode.

Valid Minimum Abbreviation

ip ri com

Important Considerations

- The RIP-1 compatibility mode determines how the software sends periodic RIP-2 updates. (For RIP-1, the software never uses the multicast address; it uses the advertisement list.)
 - When the system is configured to advertise RIP-2 packets and compatibility mode is disabled, the software uses the multicast address of 224.0.0.9 when sending periodic updates. This latest industry recommendation reduces the load on hosts that are not configured to listen to RIP-2 messages.
 - When the system is configured to advertise RIP-2 packets and compatibility mode is enabled, the software uses the advertisement list for RIP-2 updates.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the RIP compatibility mode	<ul style="list-style-type: none">■ One or more selectable interface indexes■ all■ ? (for a list of selectable interface indexes)	Previous entry, if applicable
RIP-1 compatibility mode	Selected RIP Version 1 compatibility mode that determines how the system handles RIP-2 updates	<ul style="list-style-type: none">■ disabled■ enabled	disabled

ip rip cost *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

On a per-interface basis, sets the RIP cost.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`ip ri cos`

Important Considerations

- The default cost value is 1, which is appropriate for most networks.
- The system uses the cost number, between 1 and 15, to calculate route metrics. Unless your network has special requirements, assign a cost of 1 to all interfaces.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the RIP cost	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable
RIP cost	Selected RIP cost for the interface	1 – 15	1 (factory default), or current value

ip rip poisonReverse **For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Enables or disables RIP Poison Reverse mode on the system.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip ri poi

3900
9300

Important Considerations

- Your system always implements *Split Horizon*, a scheme that aims to avoid the problems that are associated with reverse-route updates (that is, the updates that are sent to a neighboring router that include the routes that are learned from that router). The scheme omits the routes that are learned from one neighbor in the updates that are sent to that neighbor (the reverse routes). Poison reverse works with Split Horizon as follows:
 - When you enable *Poison Reverse* for use with the Split Horizon scheme (the default), the system advertises reverse routes in updates, but sets the metrics to 16 (infinity). Setting the metric to infinity breaks the loop immediately when two routers have routes that point to each other.
 - When you disable Poison Reverse for the Split Horizon scheme, reverse routes are simply not advertised.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the poison reverse mode	<ul style="list-style-type: none">■ One or more selectable interface indexes■ all■ ? (for a list of selectable interface indexes)	Previous entry, if applicable
Poison Reverse mode	Whether you want to implement poison reverse for the selected interface	<ul style="list-style-type: none">■ disabled■ enabled	Current value

ip rip routeAggregation Mode

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the route aggregation mode.

Valid Minimum Abbreviation

ip ri ro

Important Considerations

- *Route aggregation mode* determines which route table entries are sent during a RIP Version 2 update.
 - If route aggregation mode is `enabled`, RIP-2 can function like RIP-1 and “collapse” route table entries for all subnets of a directly connected network. For example, if route aggregation is `enabled`, and the system is advertising subnets 150.100.31.0 and 150.100.32.0, only the entry for network 150.100.0.0 is sent in the update. With RIP Version 2, you *must* enable route aggregation mode if you want the interface to collapse the route table entries and function like RIP-1.
 - If route aggregation mode is `disabled` (the default), a RIP-2 update sends all routing table entries.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to set the route aggregation mode	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable
Route aggregation mode	Whether you want to implement route aggregation on the selected interface	<ul style="list-style-type: none"> ■ disabled ■ enabled 	Current value

ip rip password

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Sets the IP RIP-2 password so that you can choose the IP interfaces that can put RIP-2 updates into their routing tables.

Valid Minimum Abbreviation

ip ri pa

Important Considerations

- If the sending interface has an IP RIP-2 password, the receiving interface must have the same IP RIP-2 password. If the receiving interface has a different password or a null password, its routing table is not updated.
- If you are using RIP-1, do not use the password option.
- You cannot use the ASCII string `none` as the password. This string is reserved to indicate the default password, which is a null value.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the IP interfaces that you want to allow to receive route updates	<ul style="list-style-type: none">■ One or more selectable interface indexes■ all■ ? (for a list of selectable interface indexes)	–
Password	Combination of characters that you set as the RIP-2 password	<ul style="list-style-type: none">■ up to 16 alphanumeric characters■ null password	null password

IP RIP Password Example

Select menu option (ip/rip): **password**
Select IP interfaces (1,2|all|?): **1**
Interface 1 - Enter password {?} [none]: **wings**

ip rip addAdvertisement

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Adds an advertisement address to an IP RIP interface.

Valid Minimum Abbreviation

`ip ri a`

Important Considerations

- The system uses the specified advertisement address to advertise routes to other stations on the same network. It uses this address for sending updates. (RIP-2 updates depend on the setting of RIP compatibility mode.)
- Advertisement addresses are handled differently based on RIP-1 and RIP-2.
 - For the CoreBuilder 3500, each interface that you define initially uses the default broadcast address (255.255.255.255) as the advertisement address. With RIP-1 updates, the address that you specify becomes the new RIP-1 advertisement address if you change the broadcast address. If you subsequently use RIP-2 (configure the interface to send RIP-2 advertisements) and have the RIP-1 compatibility mode disabled, the multicast address is used for updates.
 - For the CoreBuilder 9000, each interface that you define initially uses the directed broadcast address as the RIP advertisement address (all 1s in the host field).
- You can specify up to 64 advertisement addresses in separate iterations.
- On the CoreBuilder 3500:
 - After you add an advertisement address, you cannot subsequently change the broadcast address.
 - If you are using RIP-2 for the interface, you must enable RIP compatibility mode if you want the system to use the advertisement list instead of the multicast address for RIP updates. See “ip rip compatibilityMode” earlier in this chapter for more information.
- To add an advertisement address on other platforms, you must remove the directed broadcast address if you only want the address that you added to be used for RIP advertisements.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to add the advertisement address	<ul style="list-style-type: none">■ One or more selectable interface indexes■ ? (for a list of selectable interface indexes)	Previous entry, if applicable
Advertisement address	Selected IP address to add to the list of advertisement addresses	A valid IP address	–

**ip rip remove
Advertisement*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Removes an advertisement address from the list of RIP advertisement addresses for an interface.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation`ip ri re`

3900

9300

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interfaces for which you want to remove the advertisement address	<ul style="list-style-type: none">■ One or more selectable interface indexes■ ? (for a list of selectable interface indexes)	Previous entry, if applicable
Advertisement address	Advertisement address that you want to remove	An address from the advertisement list	–

ip rip policy summary

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

3900

9300

Displays summary information about RIP routing policies.

Valid Minimum Abbreviation

ip ri pol s

Important Considerations

- Your system has one unified IP routing table. Route policies enable you to control the flow of routing information between the network, the protocols, and the unified routing table on your system.
- Route policies are classified as follows:
 - *Import policies* import routing information from what RIP learns from a router/neighbor to the unified routing table. (You can also import routing information from Open Shortest Path First (OSPF).)
 - *Export policies* send information from the routing table to RIP and RIP routers, which controls what is going out on the wire to the RIP domain. (You can also export from the routing table to OSPF.)
- The system tracks policies that you define in both RIP and OSPF, so the indexes that are assigned to your policies may have gaps (for example, if you have RIP policies 1 and 2 and OSPF policies 3-6, the next policy that is available for RIP or OSPF is 7).

Fields in the IP RIP Policy Summary Display

Field	Description
Action	Action for the route — accept or reject
Index	Index number of the policy
Protocol	Protocol (for example, RIP)
Route	Route affects the policy
Source	Source router (all is from all routers)
Type	Whether the policy is an import or export policy
Weight	Administrative weight — 1 through 16

ip rip policy detail *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays detailed information about RIP routing policies.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

`ip ri pol det`

3900

9300

Important Considerations

- This display contains the summary information and two additional fields: Interface and Metric.
- Route policies are classified as follows:
 - *Import policies* import routing information from what RIP learns from a router/neighbor to the unified routing table. (You can also import routing information from Open Shortest Path First (OSPF).)
 - *Export policies* send information from the routing table to RIP and RIP routers, which controls what is going out on the wire to the RIP domain. (You can also export from the routing table to OSPF.)

Fields in the IP RIP Policy Detail Display

Field	Description
Action	Action for the route — <code>accept</code> or <code>reject</code>
Index	Index number of the policy
Interface	Interface that is associated with the policy (all applies to all interfaces)
Metric	Assigned metric, a value 0 through 16 for RIP-1 or RIP-2 (metrics can use options +, -, /, *, and %)
Protocol	Protocol (for example, <code>RIP</code>)
Route	Route that the policy affects
Source	Source router (all is from all routers)
Type	Whether the policy is an import or export policy
Weight	Administrative weight — 1 through 16

ip rip policy define***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Defines an import or export route policy for RIP.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ip ri pol def`**Important Considerations**

- Route policies are classified as follows:
 - *Import policies* import routing information from what RIP learns from a router/neighbor to the unified routing table. (You can also import routing information from Open Shortest Path First (OSPF).)
 - *Export policies* send information from the routing table to RIP and RIP routers, which controls what is going out on the wire to the RIP domain. (You can also export from the routing table to OSPF.)
- The system assigns an index number to each policy and takes into account all route policies set on the system, RIP and OSPF (You can define up to 128 routing policies total, shared between OSPF and RIP policies).
- Certain conditions are associated with import and export policies. See the import and export policy tables that follow the Options table for lists of the conditions.
- You can set up an IP RIP or OSPF import or export policy to accept or advertise the default route, as long as the default route exists in the routing table. When you define a policy, you are always prompted for the route subnet mask after the route address, regardless of whether you specify the wildcard route address of 0.0.0.0. For more information about the default route and routing policies, see the *CoreBuilder 3500 Implementation Guide* or the *CoreBuilder 9000 Implementation Guide*.

Options

Prompt	Description	Possible Values	[Default]
Policy type	Type of policy	<ul style="list-style-type: none"> ■ import ■ export 	import
Origin protocols	Which protocol advertises the route (for export policies only)	<ul style="list-style-type: none"> ■ directory ■ static ■ rip ■ ospf ■ all 	static
Source address	Router's IP address	<ul style="list-style-type: none"> ■ A valid IP address ■ 0.0.0.0 ■ all 	0.0.0.0
Route address	Associated route IP address	<ul style="list-style-type: none"> ■ A valid IP address ■ 0.0.0.0 ■ all 	0.0.0.0
Route subnet mask	Subnet mask for the route (for example, 255.255.0.0)	A valid mask	Based on route
IP interfaces	Index number of the interface indexes for which you want to define a routing policy	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	all, or previous entry, if applicable
Policy action	Whether to accept or reject the route	<ul style="list-style-type: none"> ■ accept ■ reject 	accept
Metric adjustment	For accept conditions only, increase or decrease in the converted route metric by the specified value. Options: + (add) - (subtract) * (multiple metric by value) / (use new metric as divisor) % (modulus, remainder of division operation as integer)	0 – 16, with or without options	0, which does not change the metric

Prompt	Description	Possible Values	[Default]
Administrative weight	Metric value for this policy (higher values have higher priority)	1 – 16	1

RIP Import Policy Conditions for Specified Interfaces

Source Router	Route (address/mask)	Action	Description
Specified router	Specified route/mask	accept	Accept specified route from specified source router on specified interfaces with or without metric adjustments (+, -, *, /, %).
Specified router	all (0.0.0.0)	accept	Accept all routes from specified router on specified interfaces with or without metric adjustments (+, -, *, /, %).
all (all routers)	Specified route/mask	accept	Accept specified route on specified interfaces with or without metric adjustments (+, -, *, /, %).
all	all	accept	Accept all routes on specified interfaces with or without metric adjustments (+, -, *, /, %).
Specified router	Specified route/mask	reject	Reject specified route from specified router on specified interfaces. (Metrics are not applicable.)
Specified router	all	reject	Reject all routes from specified router on specified interfaces.
all	Specified route/mask	reject	Reject specified route from all routers on specified interfaces.
all	all	reject	Reject all routes on specified interfaces.

RIP Export Policy Conditions for Specified Interfaces

Protocol	Source Router	Route	Action	Description
RIP, OSPF, static	Specified router or all routers	Specified route/mask	accept	Advertise RIP/OSPF/static specified route from specified source router on specified interfaces with or without metric adjustments (+, -, *, /, %).
RIP, OSPF, static	Specified router or all routers	all (0.0.0.0)	accept	Advertise all RIP/OSPF/static routes from specified router on specified interfaces with or without metric adjustments (+, -, *, /, %).
RIP, OSPF, static	Specified router or all routers	Specified route/mask	reject	Do not advertise the RIP/OSPF/static specified route on specified interfaces.
RIP, OSPF, static	Specified routers or all routers	all	reject	Do not advertise all RIP/OSPF/static routes on specified interfaces.

Example of Import Policy

```
Select menu option (ip/rip/policy): define
Enter policy type (import,export) [import]: import
Enter source address [0.0.0.0]:
Enter route address [0.0.0.0]: 158.101.135.40
Enter route subnet mask [255.255.0.0]:
Select IP interfaces (1,2|all|?) [1]: 1
Enter policy action (accept,reject) [accept]:
Enter metric adjustment ([+,-,*,/,%]0-16) [0]:
Enter administrative weight (1-16) [1]:
```

Example of Export Policy

```
Select menu option (ip/rip/policy): define
Enter policy type (import,export) [import]: export
Enter origin protocols (dir,sta,rip,ospf|all|?) : rip
Enter source address [0.0.0.0]:
Enter route address [0.0.0.0]:
Select IP interfaces (1,2|all|?) [1]: all
Enter policy action (accept,reject) [accept]:
Enter metric adjustment ([+,-,*,/,%] 0-16) [0]:
Enter administrative weight (1-16) [1]:
```

ip rip policy modify

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies an existing route policy for RIP.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip ri pol m

3900
9300

Important Considerations

- Route policies are classified as follows:
 - *Import policies* import routing information from what RIP learns from a router/neighbor to the unified routing table. (You can also import routing information from Open Shortest Path First (OSPF).)
 - *Export policies* send information from the routing table to RIP and RIP routers, which controls what is going out on the wire to the RIP domain. (You can also export from the routing table to OSPF.)
- The system assigns an index number to each policy that you define. This index takes into account all route policies set on the system, RIP and OSPF, so the assigned index can be higher than you may expect.

Options

Prompt	Description	Possible Values	[Default]
Policy type	Type of policy	<ul style="list-style-type: none">■ import■ export	import
Origin protocols (export)	Whether or not the route is a static route (for export policies only)	<ul style="list-style-type: none">■ RIP■ OSPF■ all	–
Source address	IP address of the source router	<ul style="list-style-type: none">■ A valid IP address■ 0.0.0.0■ all	0.0.0.0
Route address	Route that is associated with the source network	<ul style="list-style-type: none">■ A valid IP address■ 0.0.0.0■ all	0.0.0.0
Route subnet mask	Subnet mask that is associated with the route	A valid mask	Based on source network (for example, 255.255.0.0)

Prompt	Description	Possible Values	[Default]
IP interfaces	Index number of the interface for which you want to define a routing policy.	<ul style="list-style-type: none"> ■ One or more selectable interface indexes ■ all ■ ? (for a list of selectable interface indexes) 	Previous entry, if applicable
Policy action	Whether the route is accepted or rejected	<ul style="list-style-type: none"> ■ accept ■ reject 	accept
Metric adjustment	Used with accept, increase or decrease in the converted route metric by the specified value Options include: + (add) - (subtract) * (multiple metric by value) / (use new metric as divisor) % (modulus, take remainder of division operation expressed as an integer)	0 – 16	0, which does not change the metric
Administrative weight	Metric value for this policy (higher values have higher priority over lower-numbered values associated with the route)	1 – 16	1

ip rip policy remove

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Deletes a previously defined route policy.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Valid Minimum Abbreviation

ip rip pol r

Important Considerations

- The system assigns an index number to each policy that you define. This index takes into account all route policies that are set on the system, RIP and OSPF, so the assigned index can be higher than you may expect.
- When you remove a policy, the associated index is available for future use.

Options

Prompt	Description	Possible Values	[Default]
Policy index	Index number that is associated with the policy that you want to delete	<ul style="list-style-type: none">■ One or more selectable policy indexes■ all■ ? (for a list of selectable policy indexes)	–

ip rip statistics***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays general RIP statistics.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation`ip rip s`

✓ 3900

✓ 9300

Fields in the IP RIP Statistics Display

Field	Description
queries	Number of queries
routeChanges	Number of route changes

ip ping *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Tries to reach or “ping” a specified destination using the default ping options.

Valid Minimum Abbreviation

`ip p`

Important Considerations

- This tool is useful for network testing, performance measurement, and management. It uses the ICMP echo facility to send Internet Control Message Protocol (ICMP) echo request packets to the IP destination that you specify.
- If you need to change the default ping options, use the `ip advancedPing` option. (The command description for `ip advancedPing` lists the default ping options.)
- You can either supply the host name or IP address as part of the command string, or you can supply the information at the prompt.
- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “ip dns domainName” earlier in this chapter for more information.
- When the system sends an echo request packet to an IP station using ping, the system waits for an ICMP echo reply packet. Possible responses:
 - If the host is reachable, the system displays information about the ICMP reply packets and the response time to the ping.
 - If the host does not respond, the system displays the ICMP packet information and this message: `Host is Not Responding`. You may not have configured your gateway IP address.
 - If the packets cannot reach the host, the system displays the ICMP packet information and this message: `Host is Unreachable`. A host is unreachable when there is no route to that host.
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination that you want to ping	<ul style="list-style-type: none">■ A valid host name■ IP address	0.0.0.0, or current value

IP Ping Example

```
Select menu option (ip): ping
Enter host name/IP address [0.0.0.0]: 158.101.111.50
Press "Enter" key to interrupt.

PING 158.101.111.50: 64 byte packets
64 bytes from 158.101.111.50: icmp_seq=0. time=16. ms
64 bytes from 158.101.111.50: icmp_seq=1. time=19. ms
64 bytes from 158.101.111.50: icmp_seq=2. time=24. ms

---- 158.101.111.50 PING Statistics ----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip (ms) min/avg/max = 16/20/24
```

ip advancedPing *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Tries to contact a host with one or more of the advanced ping options.

Valid Minimum Abbreviation

`ip advancedP`

Important Considerations

- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “ip dns domainName” earlier in this chapter for more information.
- The `burst` option, when enabled, overrides the value set in the `quiet` or `wait` option.
- The `burst` option floods the network with Internet Control Message Protocol (ICMP) echo packets and can cause network congestion. Do *not* use the `burst` option during periods of heavy network traffic. Use this option only as a diagnostic tool in a network that has many routers to determine if one of the routers is not forwarding packets. For example, you can set a high count value (1000 packets), and then observe the run lights on the units: the run lights blink rapidly on routers that are forwarding packets successfully, but remain unlighted, or blink slowly, on routers that are not forwarding packets successfully.
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination that you want to ping.	<ul style="list-style-type: none">■ A valid host name■ IP address	0.0.0.0
Number of ICMP Request packets	Number of ICMP echo request packets that are sent to ping a host. If the destination host does not respond after it is pinged by the number of packets that you specify, the system displays a <code>Host is Unreachable</code> or <code>Host is not Responding</code> message.	1 – 9999 packets	3

Prompt	Description	Possible Values	[Default]
Packet size	Number of bytes in each ICMP echo request packet. The packet size includes both the IP and the ICMP headers.	28 – 4096 bytes	64
Burst Transmit Ping mode	How rapidly to send out ICMP echo request packets. When <i>enabled</i> , sends out the ICMP echo request packets as rapidly as possible. The system displays a period (.) upon receiving an ICMP echo replay packet. Use this display to determine how many packets are being dropped during the burst. This is unique to the burst option.	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled
Quiet mode	How much packet information to display after a ping. When <i>enabled</i> , the system displays information about the number of packets that the system sent and received, any loss of packets, and the average time that it took a packet to travel to and from the host. When <i>disabled</i> , the system displays more detailed status information about each ICMP echo request packet.	<ul style="list-style-type: none"> ■ disabled ■ enabled 	disabled
Time between sending each packet (wait)	Number of seconds that the system waits before it sends out successive ICMP echo request packets. Set this option to a high value if network traffic is heavy and you choose not to add to the network traffic with pings in fast succession.	1 – 20 seconds	1
ICMP sourceAddress	Whether to force the source address of the ICMP packets to be something other than the IP address of the interface from which the packet originated. You can use this option if you have more than one IP interface defined.	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Interface index	Index number of the ICMP source IP address that you want to use. The system lists currently defined interfaces and their indexes.	A selectable interface index	0 (the router picks the best interface)

IP Advanced Ping Example

```
Select menu option (ip): advancedPing
Enter host IP address [0.0.0.0]: 158.101.112.56
Enter number of ICMP request packets (1-9999) [3]:
Enter packet size (bytes) (28-4096) [64]:
Enter Burst Transmit Ping mode (disabled,enabled) [disabled]:
Enter Quiet mode (disabled,enabled) [disabled]:
Enter time (sec) waits between sending each packet (1-20) [1]: 2
Configure ICMP sourceAddress? (n,y) [y]:
      Index      Interface address
      0          Best interface (default)
      1          158.101.117.151
      2          158.101.10.1
Select interface index {0-2|?} [0]: 1
Press "Enter" key to interrupt.

PING 158.101.112.56 from 158.101.117.151: 64 byte packets
64 bytes from 158.101.112.56:  icmp_seq=0.  time=26. ms
64 bytes from 158.101.112.56:  icmp_seq=1.  time=18. ms
64 bytes from 158.101.112.56:  icmp_seq=2.  time=18. ms

---- 158.101.112.56 PING Statistics ----
3 packets transmitted, 3 packets received, 0% packet loss
round-trip (ms)  min/avg/max = 18/21/26
```

ip traceRoute *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Traces a route to a destination using the default traceRoute options.

✓ 3500

✓ 9000

✓ 9400

Valid Minimum Abbreviation

`ip t`

✓ 3900

✓ 9300

Important Considerations

- TraceRoute information includes all of the nodes in the network through which a packet passes to get from its origin to its destination. It uses the IP time-to-live (TTL) field in UDP probe packets to elicit an Internet Control Message Protocol (ICMP) Time Exceeded message from each gateway to a host.
- To change the default traceRoute options, use `ip advancedTraceRoute`. (The command description for “`ip advancedTraceRoute`” lists the default traceRoute options.)
- You can either supply the host name or IP address as part of the command string, or you can supply the information at the prompt.
- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “`ip dns domainName`” earlier in this chapter for more information.
- To track the route of an IP packet, traceRoute launches User Datagram Protocol (UDP) probe packets with a small TTL value and then listens for an ICMP Time Exceeded reply from a gateway. Probes start with a small TTL of 1 and increase the value by 1 until one of the following events occurs:
 - The system receives a `Port Unreachable` message, which indicates that the packet reached the host.
 - The probe exceeds the maximum number of hops (default 30).

- At each TTL setting, the system launches three UDP probe packets, and the traceRoute display shows a line with the TTL value, the address of the gateway, and the round-trip time of each probe. If a probe answers from different gateways, the traceRoute feature prints the address of each responding system. If no response occurs in the 3-second timeout interval, traceRoute displays an asterisk (*) for that probe.

Other characters that can be displayed include the following:

- !N — Network is unreachable
- !H — Host is unreachable
- !P — Protocol is unreachable
- !F — Fragmentation is needed
- !<n> — Unknown packet type
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination to which you want to trace a route	<ul style="list-style-type: none"> A valid host name IP address 	0.0.0.0

IP Trace Route Example

```

Select menu option (ip): traceRoute
Enter host name/IP address [0.0.0.0]: 158.101.101.40
Press "Enter" key to interrupt.

Traceroute to 158.101.101.40: 30 hops max, 28 bytes packet

1 158.101.117.254  9 ms 22 ms 5 ms
2 158.101.112.254  8 ms 22 ms 8 ms
3 158.101.96.22   7 ms 22 ms 7 ms
4 158.101.101.40  7 ms 23 ms 6 ms
  
```

ip advancedTraceRoute

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Traces a route to a host with one or more of the advanced traceRoute options.

Valid Minimum Abbreviation

`ip advancedT`

Important Considerations

- When you specify a host name, the host name and its associated IP address *must* be configured on a network name server. Also, you must add the IP address on the name server to the list of name server addresses that are associated with the network domain name. See “ip dns domainName” earlier in this chapter for more information.
- To interrupt the command, press Enter.

Options

Prompt	Description	Possible Values	[Default]
Host name or IP address	Host name or IP address of the destination that you want to ping.	<ul style="list-style-type: none"> ■ A valid host name ■ IP address 	0.0.0.0
Maximum ttl	Maximum number of hops that the system can use in outgoing probe packets.	1 – 255 hops	30
Destination port	Destination (or base) UDP port number that the system uses in probe packets. Set the destination UDP port number to be very high to ensure that an application at the destination is not using that port.	30000 – 65535	33434
Probe count	Maximum number of probes that the system sends at each TTL level.	1 – 10	3
Wait	Maximum amount of time that the system waits for a response to a probe.	1 – 10 seconds	3
Packet size	Number of bytes that the system sends in each UDP probe packet.	28 – 4096 bytes	28
Source address	Source address other than the one from which the probe packets originate. This option is available if you have more than one IP interface defined on the system.	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the ICMP source IP address that you want to use The system lists defined interfaces and their indexes	A selectable interface index	0 (the router picks the best interface)
Numeric mode	Whether the system shows hop addresses numerically or symbolically	<div>■ disabled</div> <div>■ enabled</div>	disabled

IP Advanced Trace Route Example (TTL value of 10):

```

Select menu option (ip): advancedTraceRoute
Enter host IP address [158.101.101.27]:
Enter maximum Time-to-Live (ttl) (1-255) [30]: 10
Enter Destination Port number (30000-65535) [33434]:
Enter the number of probes to be sent at each ttl level (1-10) [3]:
Enter time (sec) to wait for a response (1-10) [3]:
Enter the packet size (bytes) (28-4096) [28]:
Configure TRACEROUTE sourceAddress? (n,y) [y]:
      Index      Interface address
      0          Best interface (default)
      1          158.101.117.151
      2          158.101.10.1
Select interface index {0-2|?} [0]:
Enter Numeric mode (disabled,enabled) [disabled]:
Press "Enter" key to interrupt.

Traceroute to 158.101.101.27: 10 hops max, 28 bytes packet

 1  158.101.117.254  12 ms    7 ms    5 ms
 2  158.101.112.254  51 ms    9 ms    7 ms
 3  158.101.96.22   21 ms   15 ms    6 ms
 4  158.101.101.27  18 ms   90 ms   80 ms

```

ip statistics *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Displays different types of IP statistics: general statistics and those specific to the User Datagram Protocol (UDP) or the Internet Control Message Protocol (ICMP).

Valid Minimum Abbreviation

ip sta

Options

Prompt	Description	Possible Values	[Default]
Statistics	Type of IP statistics that you want to display	<ul style="list-style-type: none"> ■ ip ■ udp ■ icmp ■ all 	ip

Fields in the IP Statistics Display

Field	Description
forwDatagrams	Number of datagrams that the IP station tried to forward
fragCreates	Number of IP datagram fragments that were generated as a result of fragmentation on this system
fragFails	Number of ip datagrams that were discarded because they needed to be fragmented but could not be (for example, because their Don't Fragment bit was set)
fragOks	Number of IP datagrams that were successfully fragmented
inAddrErrors	Number of datagrams that the IP station discarded because of an error in the source or destination IP address
inDelivers	Number of datagrams that the IP station delivered to local IP client protocols
inDiscards	Number of packet receive discards
inHdrErrors	Number of datagrams that the IP station discarded because the IP header contained errors
inReceived	Total number of IP datagrams that were received, including those with errors
osReceives	Number of packets that were received that are destined to higher-level protocols such as Telnet, DNS, TFTP, and FTP
osTransmits	Number of packets that were sent through the router by higher-level protocols such as Telnet, DNS, TFTP, and FTP
outDiscards	Number of packet transmit discards

Field	Description
outNoRoutes	Number of datagrams that the IP station discarded because there was no route to the destination
outRequests	Number of datagrams that local IP client protocols passed to IP for transmission
reasmFails	Number of packet reassembly failures
reasmReqs	Number of packet reassembly requests
reasmOks	Number of successful packet reassemblies
rtDiscards	Number of packets that were discarded due to system resource errors
unkProtos	Number of packets whose protocol is unknown

Fields in the UDP Statistics Display

Field	Description
inDatagrams	Number of UDP packets that were received and addressed to the router or broadcast address
inErrors	Number of received UDP packets that contain header errors
noPorts	Number of UDP packets that were received but addressed to an unsupported UDP port
outDatagrams	Number of UDP packets that the router sent

Fields in the ICMP Statistics Display

Field	Description
inAddrMaskReps	Number of ICMP address mask reply frames that were received
inAddrMasks	Number of ICMP address mask request packets that were received
inDestUnreach	Number of ICMP destination unreachable packets that were received
inErrors	Number of received ICMP packets that contain header errors
inEchoReps	Number of ICMP echo reply packets that were received
inEchos	Number of ICMP echo request packets that were received
inParmProbs	Number of ICMP parameter problem frames that were received
inRedirects	Number of ICMP redirect packets that were received
inSrcQuenchs	Number of ICMP source quench packets that were received

Field	Description
inTimeExcds	Number of ICMP time exceeded packets that were received
inTimeStamps	Number of ICMP time stamp request packets that were received
inTimeStampsReps	Number of ICMP time stamp reply packets
messages	Number of ICMP packets that were received
outAddrMaskReps	Number of ICMP address mask reply packets that were sent
outAddrMasks	Number of ICMP address mask request packets that were sent
outDestUnreach	Number of ICMP destination unreachable packets that were sent
outEchoReps	Number of ICMP echo reply packets that were sent
outEchos	Number of ICMP echo request packets that were sent
outErrors	Number of ICMP packets that were sent that were dropped due to system resource errors
outMsgs	Number of ICMP packets that were sent
outParmProbs	Number of ICMP parameter problem packets that were sent
outRedirects	Number of ICMP redirect packets that were sent
outSrcQuenchs	Number of ICMP source quench packets that were sent
outTimeExcds	Number of ICMP time exceeded packets that were sent
outTimeStampReps	Number of ICMP time stamp reply packets that were sent
outTimeStamps	Number of ICMP time stamp request packets that were sent

17

VIRTUAL ROUTER REDUNDANCY (VRRP)

Virtual Router Redundancy Protocol (VRRP) provides fault-tolerant routing on a LAN by eliminating the single point of failure that exists when hosts are configured with a static default gateway. This chapter provides guidelines and other key information about configuring VRRP on your system.



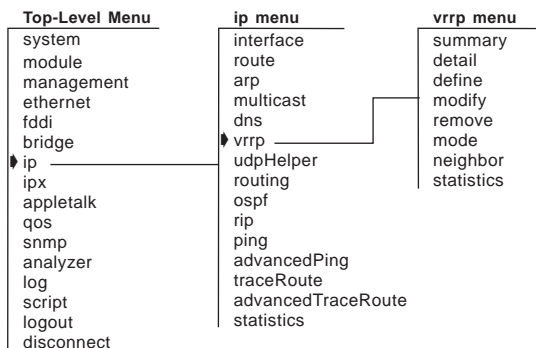
For more information about VRRP, see the Implementation Guide for your system.



For the CoreBuilder® 9000 platform, the commands in this chapter apply to Layer 3 switching modules only.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



ip vrrp summary *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Displays summary information about configured virtual routers on your system.

Valid Minimum Abbreviation

`ip v s`

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) for which you want to display virtual router information	<ul style="list-style-type: none"> One or more valid IP VLAN index numbers all ? (for a list of selectable indexes) 	–
Virtual router ID	ID of the virtual router for which you want to display summary information	<ul style="list-style-type: none"> Valid virtual router ID (1 – 255) ? (for list of selectable IDs) 	ID of virtual router that is defined on the VLAN

Fields in the IP VRRP Summary Display

Field	Description
Address	IP address of the virtual router
Auth	Whether the VRRP router uses simple password authentication. If password authentication is configured, the VRRP router discards any VRRP packet that does <i>not</i> have a matching authentication string.
Error	Last type of invalid advertisement received, or none.
Interval	Time, in seconds, between virtual router advertisements. The Master router advertises all IP addresses that are associated with the virtual router. Backup routers on the VRID consider the Master down if two advertisement intervals pass with no advertisement from the Master.
Ports	Ports that are defined on the virtual LAN (VLAN) and that are associated with the virtual router
Preempt	Whether a backup virtual router preempts a Master with a lower priority. <code>yes</code> allows preemption; <code>no</code> prohibits it.

Field	Description
Pri	Priority of the the virtual router. Represented by a value from 0 through 255. Used in Master router election. Value of 255 indicates that the router owns the IP addresses that are associated with the virtual router. 0 indicates that the current Master has stopped participating in VRRP.
State	<p>Current state of the VRRP router. One of the following:</p> <ul style="list-style-type: none"> ■ Master — In this state, the router is the active forwarding router for all IP addresses that are associated with the virtual router. ■ Backup — In this state, the router monitors the availability of the Master router. If the Master router fails, the Backup router assumes forwarding responsibility for all IP addresses that are associated with the virtual router. ■ Initialize — Transitional state between Backup and Master states. Typically indicates that the virtual router has been configured but not enabled, or that the virtual router mode has been set to disabled. <p>In this state, the router waits for a Startup event. When the router receives the Startup event, it broadcasts an ARP request that contains the virtual router MAC address for all IP addresses that are associated with the virtual router and transitions to the Master state. If the Startup event is not received, it transitions to the Backup state.</p>
Type	Type of virtual router: primary or backup
VLAN Index	Index number of the virtual LAN (VLAN) on which the virtual router is defined
VRID	Virtual Router ID (0 – 255) . Must be unique on the LAN

Sample IP VRRP Summary Display

```
Select menu option (ip/vrrp): summary
Enter VLAN interface index (2|?) [2]:
Enter virtual router ID (1|?) [1]:
```

```
VLAN Index: 2   Ports: 7-12,14
```

VRID	Address	Type	State	Interval	Pri	Preempt	Auth	Error
1	158.101.175.228	Primary	Master	1 sec.	255	Yes	pass	none

ip vrrp detail *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays summary information and detailed statistics for the specified virtual router.

Valid Minimum Abbreviation

`ip v det`

Important Consideration

- Displays both summary information and the VRRP router statistics table for locally configured virtual routers, whether they are in the Master, Backup, or Initialize state.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) for which you want to display virtual router information	<ul style="list-style-type: none"> ■ One or more valid IP VLAN index numbers ■ all ■ ? (for a list of selectable indexes) 	–
Virtual router ID	ID of the virtual router for which you want to display summary information	<ul style="list-style-type: none"> ■ Valid virtual router ID (0 – 255) ■ ? (for list of selectable IDs) 	ID of virtual router that is defined on the VLAN

Fields in the IP VRRP Detail Display

Field	Description
Address	IP address of the virtual router
addrListErrors	Total number of VRRP advertisements that were received that do not match the address list defined for the virtual router
advertReceived	Total number of VRRP advertisements that this virtual router has received
advIntErrors	Total number of VRRP advertisement packets that were received for which the advertisement interval is different than the one that is configured for the virtual router

Field	Description
Auth	Whether the VRRP router uses simple password authentication. If password authentication is configured, the VRRP router discards any VRRP packet that does <i>not</i> have a matching authentication string.
authFailures	Total number of VRRP advertisements that this virtual router has received that did not have the correct simple text authentication password
becomeMaster	Total number of times that this virtual router has changed to the Master state
Error	Last type of invalid advertisement received, or <i>none</i> .
Interval	Time, in seconds, between virtual router advertisements. The Master router advertises all IP addresses that are associated with the virtual router. Backup routers on the VLAN consider the Master down if two advertisement intervals pass with no advertisement from the Master.
InvalidAuthType	Total number of VRRP advertisements that the virtual router has received with the Authentication Type not equal to the locally configured authentication method
invalidPktTypeRx	Number of VRRP advertisements with an invalid value in the Type field that this virtual router has received
ipTtlErrors	Total number of VRRP advertisements with IP TTL (Time-to-Live) not equal to 255 that this virtual router has received
MasterIpAdd	IP address of the Master for this virtual router.
Ports	Ports that are defined on the virtual LAN (VLAN) and that are associated with the virtual router
Preempt	Whether the router preempts a Master with a lower priority. <i>Yes</i> allows preemption; <i>no</i> prohibits it.
Pri	Priority of the virtual router. Represented by a value from 0 through 255. Used in Master router election. Value of 255 indicates that the router owns the IP addresses that are associated with the virtual router. 0 indicates that the current Master has stopped participating in VRRP.
PrimaryIpAddr	IP address which VRRP advertisements use as the source of the IP packet.
priorityZeroRx	Total number of VRRP advertisements with a priority of 0 that this virtual router has received. The priority of zero (0) indicates that the current Master has stopped participating in VRRP. Used to trigger Backup routers to quickly transition to Master without having to wait for the current Master to time out.

Field	Description
priorityZeroTx	Total number of VRRP advertisements with a priority of 0 that this virtual router has sent. The priority of zero (0) indicates that this virtual router was acting as Master but stopped participating in VRRP. Used to trigger backup routers to quickly transition to Master without having to wait for the current Master to time out.
State	<p>Current state of the VRRP router. One of the following:</p> <ul style="list-style-type: none">■ Master — In this state, the router is the active forwarding router for all IP addresses that are associated with the virtual router.■ Backup — In this state, the router monitors the availability of the Master router. If the Master router fails, the Backup router assumes forwarding responsibility for all IP addresses that are associated with the virtual router.■ Initialize — Transitional state between Backup and Master states. Typically indicates that the virtual router has been configured but not enabled, or that the virtual router mode has been set to disabled. <p>In this state, the router waits for a Startup event. When the router receives the Startup event, it broadcasts an ARP request that contains the virtual router MAC address for all IP addresses that are associated with the virtual router and transitions to the Master state. If the Startup event is not received, it transitions to the Backup state.</p>
Type	Type of virtual router: <code>primary</code> or <code>backup</code>
versionErrors	Total number of VRRP advertisements with an unknown or unsupported version number that this virtual router has received
VLAN Index	Index number of the virtual LAN (VLAN) on which the virtual router is defined
VRID	Virtual Router ID. Number that identifies the virtual router on the LAN

Sample IP VRRP Detail Display

Select menu option (ip/vrrp): **detail**
Enter VLAN interface index (2|?) [2]:
Enter virtual router ID (1|?) [1]:

VLAN Index: 2 Ports: 7-12,14

VRID	Address	Type	State	Interval	Pri	Preempt	Auth	Error
1	158.101.175.228	Primary	Master	1 sec.	255	Yes	pass	none
VIDX	VRID	becomeMaster	advertReceived	ckSumErrors	versionErrors			
2	1	1	0	0	0			0
VIDX	VRID	advIntErrors	securViolations	ipTtlErrors	priorityZeroRx			
2	1	0	0	0	0			0
VIDX	VRID	priorityZeroTx	invalidPktTypeRx	addrListErrors	unknownAuthType			
2	1	0	0	0	0			0
VIDX	VRID	authTypeErrors						
2	1	0						

ip vrrp define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a virtual router on the system.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

`ip v def`

3900

9300

Important Considerations

Primary Routers

- Authentication passwords can be up to eight alphanumeric characters.
- You can define *one* Primary router per VLAN.
- Primary routers own the IP addresses that you associate with a virtual router.
- When you define a Primary virtual router, the possible VLANs that you can select are the IP VLANs on the router that have no virtual routers configured.
- The virtual router ID (VRID) must be unique across all locally attached LAN segments and unique for the local router.
- When you define a Primary virtual router, you cannot use the VRID of a virtual router that is already defined on the system or the VRID of a neighboring VRRP router.

Backup Routers

- Backup routers back up the primary router of a specified virtual router and assume Master state responsibilities for the virtual router should the primary router fail.
- When you define a Backup virtual router, you cannot use the VRID of a primary router that is defined on the system. You cannot define a Primary and Backup VRRP router for the same virtual router on the same routing device.

Address Mode

- In `auto-learn` mode, systems learn the IP addresses to associate with the specified VRID.
- In `IP address` mode, the system prompts you to select the interface index from a list.
- After a reboot, the address learning process restarts for each virtual router in `auto-learn` address mode.
- When you define a Primary virtual router, selecting `auto-learn` as the address mode automatically adds all IP addresses that are associated with the selected VLAN to the primary virtual router.

- When you define a Primary router on a VLAN that contains a single interface, the single interface is automatically chosen as the primary address when you select `IP-address` as the Address mode.
 - When you define a Backup virtual router, selecting `auto-learn` as the address mode configures the Backup router to learn the IP addresses that are associated with the virtual router by means of VRRP advertisements from the Primary router. The Primary router must be up for backup routers to auto-learn the addresses that are associated with the specified VRID.
 - When you define Backup virtual routers, the auto-learn address mode option enables auto address learning for the specified VRID. If a new interface is added to the VLAN on a primary virtual router, the new IP address is sent out in VRRP advertisements so that the Backup routers in auto-learn mode can learn the new address without having to manually add the new address to each backup router.
- Advertisement Intervals*
- The smaller the advertisement interval, the smaller the failover time if the master fails.
 - The advertisement interval must be the same across the set of VRRP routers that are associated with a single VRID. Backup routers must have the same advertisement interval as the Master router.

Options

Prompt	Description	Possible Values	[Default]
Virtual router type	Type of virtual router that you want to define	<ul style="list-style-type: none"> ■ Primary ■ Backup 	Primary
VLAN interface index	Index number of the virtual LAN (VLAN) on which you want to define the virtual router	<ul style="list-style-type: none"> ■ Index number of an IP virtual LAN (VLAN) that is defined on the system. ■ ? (for a list of selectable indexes) 	Index number of first available VLAN
VRID	Virtual router identifier. Identifies the virtual router that you want to define on the LAN.	1 – 255	1
Address mode	Method by which the virtual router you want to define learns its IP addresses	<ul style="list-style-type: none"> ■ auto-learn ■ IP address 	auto-learn

Prompt	Description	Possible Values	[Default]
Advertise interval	Time between virtual router advertisements.	1 – 255 seconds	1
Preempt mode	Whether a higher priority backup router may preempt a lower priority master	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Authentication type	Whether a password is needed to access the virtual router	<ul style="list-style-type: none"> ■ none ■ pass 	none
Password	Character string to authenticate access to virtual router	up to eight alphanumeric characters	–

IP VRRP Define Example

```

Select menu option (ip/vrrp): define
Enter virtual router's type (Primary,Backup) [Primary]:
Enter VLAN interface index {2-5|?}: 2
Enter VRID (1-255) [1]: 2
Enter address mode (auto-learn,IP-address) [auto-learn]:
Enter the advertise interval in sec (1-255) [1]:
Enter virtual router preempt mode (no,yes) [yes]:
Enter Authentication Type (none,pass) [pass]: pass
Enter 8 characters password {?}: echoe

```

ip vrrp modify**For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Modifies an existing virtual router.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation`ip v modi`

3900

9300

Important Considerations

- Authentication passwords can be up to eight alphanumeric characters.
- You can define *one* Primary router per VLAN.
- Primary routers own the IP addresses that you associate with a virtual router.
- When you define a Primary virtual router, the possible VLANs that you can select are the IP VLANs on the router that have no virtual routers configured.
- The virtual router ID (VRID) must be unique across all locally attached LAN segments and unique for the local router.
- When you define a Primary virtual router, you cannot use the VRID of a virtual router that is already defined on the system or the VRID of a neighboring VRRP router.
- Backup routers back up the primary router of a specified virtual router and assume Master state responsibilities for the virtual router should the primary router fail.
- When you define a Backup virtual router, you cannot use the VRID of a primary router that is defined on the system. You cannot define a Primary and Backup VRRP router for the same virtual router on the same routing device.
- In `auto-learn` mode, systems learn the IP addresses to associate with the specified VRID.
- In `IP address` mode, the system prompts you to select the interface index from a list.
- After a reboot, the address learning process restarts for each virtual router in `auto-learn` address mode.
- When you define a Primary virtual router, selecting `auto-learn` as the address mode automatically adds all IP addresses that are associated with the selected VLAN to the primary virtual router.

*Primary Routers**Backup Routers**Address Mode*

- When you define a Primary router on a VLAN that contains a single interface, the single interface is automatically chosen as the primary address when you select `IP-address` as the Address mode.
 - When you define a Backup virtual router, selecting `auto-learn` as the address mode configures the Backup router to learn the IP addresses that are associated with the virtual router by means of VRRP advertisements from the Primary router. The Primary router must be up for backup routers to auto-learn the addresses that are associated with the specified VRID.
 - When you define Backup virtual routers, the auto-learn address mode option enables auto address learning for the specified VRID. If a new interface is added to the VLAN on a primary virtual router, the new IP address is sent out in VRRP advertisements so that the Backup routers in auto-learn mode can learn the new address without having to manually add the new address to each backup router.
- Advertisement Intervals*
- The smaller the advertisement interval, the smaller the failover time if the master fails.
 - The advertisement interval must be the same across the set of VRRP routers that are associated with a single VRID. Backup routers must have the same advertisement interval as the Master router.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) on which you want to define the virtual router	<ul style="list-style-type: none"> ■ Index number of an IP virtual LAN (VLAN) that is defined on the system. ■ ? (for a list of selectable indexes) 	Index number of first available VLAN
VRID	Virtual router identifier. Identifies the virtual router that you want to define on the LAN.	1 – 255	1
Virtual router type	Type of virtual router that you want to define	<ul style="list-style-type: none"> ■ Primary ■ Backup 	Primary
Address mode	Method by which the virtual router you want to define learns its IP addresses	<ul style="list-style-type: none"> ■ auto-learn ■ IP address 	auto-learn

Prompt	Description	Possible Values	[Default]
Advertise interval	Time between virtual router advertisements.	1 – 255 seconds	1
Preempt mode	Whether a higher priority backup router may preempt a lower priority master	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	y
Authentication type	Whether a password is needed to access the virtual router	<ul style="list-style-type: none"> ■ none ■ pass 	none
Password	Character string to authenticate access to virtual router	up to eight alphanumeric characters	–

IP VRRP Modify Example

```

Select menu option (ip/vrrp): modify
Enter VLAN interface index {2-3|?}: 2
Enter virtual router ID {1|?} [1]:
Enter virtual router's type (Primary,Backup) [Primary]:
Enter address mode (auto-learn,IP-address) [auto-learn]: IP-address
Old Ip Association address list:
  VRID    VIDX    Address
   1       2    158.101.175.228
  Interface 158.101.175.228 will be selected as your primary address.
Enter the advertise interval in sec (1-255) [1]:
Enter virtual router preempt mode (no,yes) [yes]: no
Enter Authentication Type (none,pass): none
Enter virtual router state (enabled,disabled) [enabled]:

```

ip vrrp remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Removes one or more existing virtual routers from the system.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`ip v r`

Important Consideration

- If you attempt to remove a virtual router that is in the Master state, you are prompted to confirm the operation:
 - If you enter **no**, the system does not remove the virtual router.
 - If you enter **yes**, the system removes the virtual router, which sends an advertisement to the other virtual routers that one of them must assume Master responsibilities immediately.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) on which you want to define the virtual router	<ul style="list-style-type: none"> ■ Index number of a IP virtual LAN (VLAN) defined on the system ■ ? (for a list of selectable indexes) 	Index number of first available VLAN
VRID	Virtual router identifier. Identifies the virtual router that you want to define on the LAN	1 – 255	1

IP VRRP Remove Example

Select menu option (ip/vrrp): **remove**

Enter VLAN interface index (2-3|all|?): **2**

Enter virtual router ID (1|?) [1]:

ip vrrp mode *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Enables or disables a configured virtual router.

Valid Minimum Abbreviation

`ip v mode`

3900
9300

Important Considerations

- You must configure the virtual router before you can enable it.
- You cannot modify or remove a virtual router that is enabled; you must disable the virtual router before you can change or delete the virtual router.

Options

Prompt	Description	Possible Values	[Default]
VLAN interface index	Index number of the virtual LAN (VLAN) on which you want to define the virtual router	<ul style="list-style-type: none">■ Index number of a IP virtual LAN (VLAN) defined on the system■ all■ ? (for a list of selectable indexes)	Index number of first available VLAN
VRID	Virtual router identifier. Identifies the virtual router that you want to define on the LAN	1 – 255	1
Virtual router mode	Explicitly turns on or turns off a configured virtual router	<ul style="list-style-type: none">■ enabled■ disabled	disabled

IP VRRP Mode Example

```
Select menu option: ip vrrp mode
Enter VLAN interface index (2-3|all|?): all
Enter virtual router ID (1-2|all|?): all
Vrid 1 - Enter virtual router mode (enabled,disabled)
[disabled]: enabled
Vrid 2 - Enter virtual router mode (enabled,disabled)
[disabled]: enabled
```

ip vrrp neighbor *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Displays a list of neighboring virtual routers.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ip v n

3900

9300

Important Considerations

- Any locally defined virtual router is not displayed.
- If the Address and MasterRouterAddr fields contain the same IP address, the listed virtual router is in the Master state.

Fields in the IP VRRP Neighbor Display

Field	Description
VLAN Index	Index number of the VLAN on which the virtual router is defined
VRID	Virtual Router ID. Number that identifies the virtual router on the LAN
Address	IP address of the neighbor virtual router, which may be a Master or Backup router
MasterRouterAddr	IP address of the Master virtual router
Interval	Time, in seconds, between virtual router advertisements
Priority	Priority among the backup routers to become the Master virtual router
Auth	Authentication type: whether a password is needed to access the virtual router
Config	Whether the virtual router has been locally configured

ip vrrp statistics

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays general VRRP statistics for the virtual router.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ip v st

3900

9300

Fields in the IP VRRP Statistics Display

Field	Description
ckSumErrors	Total number of VRRP advertisements with an invalid VRRP checksum value that this virtual router has received
versionErrors	Total number of VRRP advertisements with an unknown or unsupported version number that this virtual router has received.
vriderrors	Total number of VRRP advertisements with an invalid VRID number that this virtual router has received

IP MULTICAST

This chapter provides guidelines and other key information about how to configure and manage IP multicast routing commands from the Administration Console of the CoreBuilder® 3500 and CoreBuilder 9000 Layer 3 switching modules.



For the CoreBuilder 9000 platform, the commands in this chapter apply to Layer 3 switching modules only.



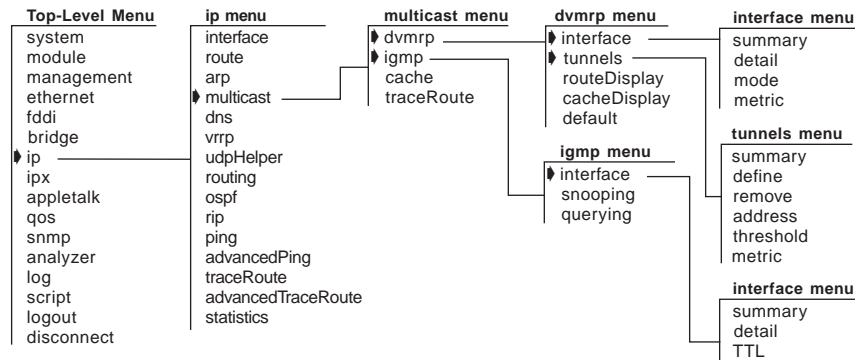
For more information about IP multicast technology, concepts, and implementation procedures, see the Implementation Guide for your system.



For IGMP commands in Layer 2 switching systems (CoreBuilder 9400, CoreBuilder 9000 Layer 2 switching modules, SuperStack® II Switch 3900, and SuperStack II Switch 9300), see Chapter 9.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured on your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



ip multicast dvmrp
interface summary

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays summary information about IP interfaces that may or may not be operating as IP multicast routing interfaces using the Distance-Vector Multicast Routing Protocol (DVMRP).

Valid Minimum Abbreviation

ip m d i s

Fields in the IP Multicast DVMRP Interface Summary Display

Field	Description
Index	Number associated with the interface for identification purposes
Address	IP address of the interface
Metric	Numeric DVMRP metric or "cost" that you assign to the interface
State	Role that the interface plays in IP multicast delivery. One or more of the following descriptors may appear: <ul style="list-style-type: none">■ <code>querier</code> — The interface is functioning as the IGMP Querier for its subnetwork.■ <code>non-querier</code> — The interface is <i>not</i> functioning as the IGMP Querier for its subnetwork.■ <code>leaf</code> — There are no routers downstream of this interface; IP multicast group members may reside on this subnetwork.■ <code>non-leaf</code> — The interface is a branch in the IP multicast delivery tree. There are one or more IP multicast routing interfaces downstream of this interface.■ <code>one-way</code> — Traffic is moving downstream only.■ <code>disabled</code> — DVMRP is disabled on the interface.■ <code>up</code> — The IP interface is available to support network communication.■ <code>down</code> — The IP interface is not available to support network communication.

ip multicast dvmrp interface detail

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays information about IP interfaces that run the Distance-Vector Multicast Routing Protocol.

Valid Minimum Abbreviation

ip m d i s

Fields in the IP Multicast DVMRP Interface Detail Display

Field	Description
Index	Number associated with the interface for identification purposes
Address	IP address of the interface
Metric	Numeric DVMRP metric or “cost” that you assign to the interface
State	<p>Role that the interface plays in IP multicast delivery. One or more of the following descriptors may appear:</p> <ul style="list-style-type: none"> ■ <code>querier</code> — The interface is functioning as the IGMP Querier for its subnetwork. ■ <code>non-querier</code> — The interface is <i>not</i> functioning as the IGMP Querier for its subnetwork. ■ <code>leaf</code> — There are no routers downstream of this interface; IP multicast group members may reside on this subnetwork. ■ <code>non-leaf</code> — The interface is a branch in the IP multicast delivery tree. There are one or more IP multicast routing interfaces downstream of this interface. ■ <code>one-way</code> — Traffic is moving downstream only. ■ <code>disabled</code> — DVMRP is disabled on the interface. ■ <code>up</code> — The IP interface is available to support network communication. ■ <code>down</code> — The IP interface is not available to support network communication.
Group	IP multicast group addresses of the traffic that is being received and forwarded on that interface.
Peer, Port	IP address of the upstream router. The additional information to the right relates to the version of DVMRP that is running and the port in the local interface that connects to the peer router.

ip multicast dvmrp
interface mode

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables or disables the Distance-Vector Multicast Routing Protocol (DVMRP) per routing interface. This protocol facilitates router-to-router communication for building source-rooted spanning trees that deliver IP multicast traffic to IP multicast group members.

Valid Minimum Abbreviation

ip m d i m

Important Considerations

- When DVMRP is enabled on an interface, the interface is configured with the default value of 1 for the metric, which you can modify at any time. See “ip multicast dvmrp interface metric” later in this chapter.
- If DVMRP is enabled on any interface, IGMP snooping should also be enabled in the system. See “ip multicast igmp snooping” later in this chapter.
- If DVMRP is disabled, the interface cannot participate in building spanning trees for IP multicast. However, as long as IGMP snooping is enabled, the interface forwards appropriate IP multicast traffic to downstream group members. If IGMP snooping is disabled, then the interface only forwards IP multicast traffic with addresses in the reserved range.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the interface for which you want to enable or disable DVMRP	<ul style="list-style-type: none">■ A valid IP interface index number■ all■ ? (for a list of selectable indexes)	–
DVMRP mode	Whether DVMRP mode is enabled or disabled	<ul style="list-style-type: none">■ enabled■ disabled	disabled (factory default), or current value

**ip multicast dvmrp
interface metric**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the DVMRP metric on an interface for which DVMRP is enabled.

Valid Minimum Abbreviation

`ip m d i m`

Important Considerations

- Use this command if you want to modify the metric value of 1 that the system assigns to an interface when you define it, even if DVMRP is not yet enabled.
- The metric affects the shape of the IP multicast spanning tree when there are multiple paths to the same downstream destination. The lower cost path is the preferred path.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the routing interface for which you want to modify the default metric	<ul style="list-style-type: none">■ A valid IP interface index number■ ? (for a list of selectable index numbers)	–
metric	DVMRP cost for the interface	<ul style="list-style-type: none">■ 1 – 32	1 (factory default), or current value

**ip multicast dvmrp
tunnels summary**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Summaries key information about DVMRP tunnels that you have configured in your system. Tunnels enable IP multicast spanning trees to be constructed through and beyond areas of the network (routers) that do not support IP multicast routing. The two tunnel end points must lie in different systems and subnetworks.

Valid Minimum Abbreviation

ip m d t s

Important Considerations

- The index number shown in the DVMRP tunnel summary display is the tunnel index number. When you define a DVMRP tunnel, the system assigns a tunnel index number to it, which is different from the routing interface index number. Tunnel index numbers provide a way to identify individual tunnels, which is necessary because multiple tunnel end points can be configured on the same routing interface. Tunnel index numbers are also needed so that you can remove tunnels without removing the interface with which it is associated.
- When you remove a tunnel, the system does not dynamically re-order remaining tunnels in the DVMRP tunnel summary display. For example, if you had three tunnels with tunnel index numbers 1, 2, and 3 and you then removed tunnel 2, the display lists the remaining tunnels with their original tunnel index numbers (1 and 3, in this example). The system assigns tunnel index 2 to the next *new* tunnel that you define. After 2 is used, the system can assign tunnel index 4 for the next new tunnel, and so on.
- You can define multiple IP multicast tunnel end points on the same local routing interface, but each must lead to a different remote interface. You cannot define multiple IP multicast tunnels between the same two end points (interfaces).

Fields in the IP Multicast DVMRP Tunnels Summary Display

Field	Description
Index	Tunnel index number, which is different from the routing interface index number that is shown under <code>Index</code> in other displays.
Local address	IP address of the local interface that serves as one of two multicast tunnel end points.
Remote address	IP address of the remote interface (a different system, a different subnetwork) that serves as the other multicast tunnel end point.
Metric	DVMRP cost of the tunnel. The system assigns a value of 1 when you define the tunnel, but you can modify that value at any time (see “ip multicast dvmrp tunnels metric”). This value can be different from the metric that you assigned to the interface itself (see “ip multicast dvmrp interface metric”).
TTL	Time-to-live (TTL) threshold of the tunnel. The system assigns a value of 1 when you define the tunnel, but you can modify that value at any time (see “ip multicast dvmrp tunnels threshold”). This value can be different from the TTL threshold that you assigned to the interface itself (see “ip multicast igmp interface TTL”).
State	Role that the interface in the multicast delivery tree. For possible entries and definitions, see “ip multicast dvmrp interface summary” earlier in this chapter.

**ip multicast dvmrp
tunnels define**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines one end point of a DVMRP tunnel. The other tunnel end point lies on an IP multicast routing interface on a different system and subnetwork. One or more unicast routers lie between these tunnel end points.

Valid Minimum Abbreviation

`ip m d t d`

Important Considerations

- IP multicast tunnels are not required in all networks. Configure a tunnel only if you need to have IP multicast traffic forwarded through one or more routers that do not understand IP multicast protocols and would therefore filter IP multicast packets. Because IP multicast packets are encapsulated in unicast format at the tunnel entrance point, the interim routers in the tunnel forward the packets onward toward the other tunnel exit point.
- Think of an IP multicast tunnel end point as being layered on top of a regular DVMRP routing interface. Therefore, before you can define a multicast tunnel end point in your system, you must first define at least one IP virtual LAN (VLAN), define at least one IP interface, and enable DVMRP on the interface.
- The remote tunnel end point must lie on a different system and subnetwork.
- You must define the tunnel on both end points — that is, on both the local system and the remote system — even though you specify the address of the remote interface in the local system.
- When you define a tunnel with local and remote addresses, the system automatically assigns the value 1 as both the tunnel metric and the tunnel TTL threshold, as shown in the IP multicast DVMRP tunnel summary display. You can change these values through menu options.
- IP multicast interfaces and tunnels have similar characteristics, such as TTL threshold and metric. The characteristics of a tunnel do not have to match the characteristics of the interface on which it is configured.
- You can define multiple tunnel end points on the same local routing interface in your system, but these tunnels must lead to different remote routing interfaces.

Options

Prompt	Description	Possible Values	[Default]
interface	Index number of the interface on which you want to create a DVMRP tunnel end point	<ul style="list-style-type: none">■ A valid IP interface index number■ ? (for a list of selectable indexes)	—
Remote address	IP address of the remote multicast tunnel end point. Use standard dotted decimal notation.	A valid IP interface on a different system and subnetwork	—

**ip multicast dvmrp
tunnels remove**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Deletes a DVMRP tunnel end point from the system.

Valid Minimum Abbreviation

ip m d t r

Important Considerations

- To remove a tunnel, specify its tunnel index number. This number is different from the routing interface index number. Reference the DVMRP tunnel summary display prior to deleting a tunnel.
- If you try to remove an IP interface in your system, and you have a DVMRP tunnel defined on that interface, the system warns you with an error message. Before you can remove the IP interface, you must remove the DVMRP tunnel.
- When you remove a tunnel, the system does not dynamically re-order remaining tunnels in the DVMRP tunnel summary display. For example, if you had three tunnels with tunnel index numbers 1, 2, and 3 and you then removed tunnel 2, the display lists the remaining tunnels with their original tunnel index numbers (1 and 3, in this example). The system assigns tunnel index 2 to the next *new* tunnel that you define. After 2 is used, the system can assign tunnel index 4 for the next new tunnel, and so on.

Options

Prompt	Description	Possible Values	[Default]
Multicast tunnel index	Index number of the multicast tunnel that you want to remove from the system	<ul style="list-style-type: none">■ A valid DVMRP tunnel index number■ ? (for a list of selectable tunnel index numbers)	–

ip multicast dvmrp tunnels address

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the remote IP address that is defined in an existing DVMRP tunnel.

Valid Minimum Abbreviation

ip m d t a

Important Consideration

- The remote address that you specify must represent a routing interface on a different system and subnetwork.

Options

Prompt	Description	Possible Values	[Default]
tunnel	Index number of the tunnel for which you modify the remote tunnel end point	<ul style="list-style-type: none"> ■ A valid DVMRP tunnel index number in the system ■ ? (for a list of selectable tunnel index numbers) 	–
remote address	A valid IP address on a different system and subnetwork. Use the 0.0.0.0 format.	A valid IP address	current value

ip multicast dvmrp
tunnels threshold

- ✓ 3500
- ✓ 9000
- 9400

- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Modifies the time-to-live (TTL) threshold on an existing DVMRP tunnel.

Valid Minimum Configuration

ip m d t t

Important Consideration

- When you first define a tunnel, the system automatically assigns the value 1 as the TTL threshold for the tunnel (which is different from the interface TTL threshold). Use this command to modify the TTL threshold value on any existing tunnel.

Options

Prompt	Definition	Possible Values	[Default]
tunnel	Index number of the existing DVMRP tunnel on which you want to modify the TTL threshold	<ul style="list-style-type: none">■ A valid DVMRP tunnel index number■ ? (for a list of selectable tunnel index numbers)	–
threshold	Value that determines whether IP multicast packets are forwarded. The interface compares the packet TTL to the TTL threshold	1 – 32	1 (factory default), or current value

ip multicast dvmrp tunnels metric

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the metric or “cost” of an existing DVMRP tunnel.

Valid Minimum Configuration

`ip m d t m`

Important Consideration

- When you first define a tunnel, the system automatically assigns the value 1 as the metric or “cost” of the tunnel (which is different from the interface metric). Use this command to modify the metric value on any existing tunnel.

Options

Prompt	Definition	Possible Values	[Default]
tunnel	Index number of the existing DVMRP tunnel on which you want to modify the metric	<ul style="list-style-type: none"> ■ A valid DVMRP tunnel index number ■ ? (for a list of selectable tunnel index numbers) 	–
metric	DVMRP cost for the tunnel. This value affects the shape of the IP multicast spanning tree when there are multiple paths to the same downstream destination. The lower cost path is chosen first.	1 – 32	1 (factory default), or current value

ip multicast dvmrp
routeDisplay

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays IP multicast route information that your system has learned from using the Distance-Vector Multicast Routing Protocol (DVMRP). The system uses this information to forward IP multicast traffic that it receives.

Valid Minimum Abbreviation

ip m d r

Fields in the IP Multicast DVMRP Route Display

Field	Description
Origin	IP address of the subnetwork that contains an IP multicast source, followed by a forward slash and subnetwork mask.
Gateway	IP address of the routing interface that lies upstream of the local system on the path back towards an IP multicast source. If the source subnetwork is connected directly to your system, this field contains a dash (--).
Metric	<p>Number of hops from your system back to the origin subnetwork. This value is <i>not</i> the DVMRP interface or tunnel metric, which are shown under <code>Metric</code> in other displays.</p> <p>Occasionally, instead of a numeric value, you may see <code>NR</code>, meaning “network unreachable.” Your system may have trouble computing the hop count because of factors such as an upstream router being temporarily congested. This condition is usually resolved in a short period of time.</p>
Tmr	Amount of time (in seconds) since each entry was last reset.
Parent	The interface that connects to the upstream router (Gateway). Because DVMRP forms a loopless spanning tree to reach all hosts for a given IP multicast group, your system always chooses a single parent interface. Either an <code>I</code> or a <code>T</code> precedes the index number. An <code>I</code> indicates that the index is an interface index number. A <code>T</code> indicates that the index is a tunnel index number.
Children	Interfaces that communicate with downstream routers or local subnetworks. The system forwards incoming IP multicast traffic through these interfaces. Either an <code>I</code> or a <code>T</code> precedes each index number. An <code>I</code> precedes an interface index number. A <code>T</code> precedes a tunnel index number.

ip multicast dvmrp cacheDisplay

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the DVMRP cache, which is a collection of information about the IP multicast packets that have traveled through the system.

Valid Minimum Abbreviation

ip m d c

Options

Prompt	Description	Possible Values	[Default]
Multicast source address	Source for which you want to view cache information	<ul style="list-style-type: none"> ■ Depends on your network ■ 255.255.255.255 for all sources 	255.255.255.255 (factory default), or current value
Multicast group address	Multicast group for which you want to view cache information	<ul style="list-style-type: none"> ■ Depends on your network ■ 255.255.255.255 for all groups 	255.255.255.255 (factory default), or current value

Fields in the IP Multicast DVMRP Cache Display

Field	Description
Source	<p>Information about IP multicast sources:</p> <ul style="list-style-type: none"> ■ Entries preceded by angle brackets (>) are subnetworks that contain sources. ■ Entries without angle brackets are the IP addresses of source devices.
Group	IP multicast group address of packets coming from the source and subnetwork to the left.
CTmr	Time since the cache entry was originally recorded. Time is noted in hours (h), minutes (m), and seconds (s).
Age	Value that indicates the remaining life for the cache entry. Time is recorded in minutes (m) and seconds (s). The system assigns a life of approximately 7 minutes to each entry. When the age of the entry decreases to zero, the entry either disappears or is refreshed.
PTmr	Time remaining before the system sends a prune message to an upstream router. Time is shown in minutes (m) and seconds (s). When traffic is actively flowing, a dash (-) indicates that no prune message has been sent upstream.

Field	Description
inVif	<p>Interface that receives incoming IP multicast traffic from the spanning tree for the source, subnetwork, and group listed on the left.</p> <p>The interface is presented as an index number and either an <code>I</code> or a <code>T</code> precedes the index number. An <code>I</code> precedes a routing interface index number. A <code>T</code> precedes a tunnel index number.</p> <p>A <code>P</code> after the index number indicates that a prune message has been sent to an upstream router.</p> <p>The entry <code><none></code> may appear if the system is not able to build the cache entry correctly. This temporary condition corrects itself quickly.</p>
outVif	<p>Interfaces to which traffic from the inVif is being forwarded.</p> <p>Each interface is presented as an index number and either an <code>I</code> or a <code>T</code> precedes each index number. An <code>I</code> precedes a routing interface index number. A <code>T</code> precedes a tunnel index number.</p> <p>A <code>p</code> after an index number indicates that the upstream router has pruned this branch of the delivery tree and no multicast packets are being forwarded through this local interface. Eventually this entry disappears from the cache display.</p> <p>Either no entry or <code><none></code> appears in this column if the system is not able to build the cache entry correctly. This temporary condition corrects itself quickly.</p>
Ports	<p>Physical ports that correspond to the interfaces that are listed in the outVifs field. The Ports field shows a dash (--) when there are no outgoing interfaces and when the outgoing interfaces are tunnels.</p>

ip multicast dvmrp default

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Configures a default route for IP multicast traffic on a DVMRP interface. This interface advertises itself as a default route to neighboring DVMRP routers.

Valid Minimum Abbreviation

ip m d d

Important Considerations

- A default route metric of 0 means that the default route function is not activated on the interface (interface does not advertise 0.0.0.0 to DVMRP routers). Values other than 0 means that the default route function is activated and these values represent the “cost” of the default route.
- Definitions of default route modes:
 - **all** — The interface advertises the default route plus all other known routes to neighboring DVMRP routers.
 - **only** — The interface advertises only the default route to neighboring DVMRP routers.

If the system learns a default route, it propagates it no matter which mode is set on a given interface.

- The system allows you to configure an interface as a DVMRP default route, even when DVMRP is disabled on the interface. If DVMRP is disabled, the interface does not advertise itself as a default route.

Options

Prompt	Definition	Possible Values	[Default]
interface	Index number of the routing interface on which you want to configure a default route	<ul style="list-style-type: none"> ■ A valid interface index number ■ ? (for a list of selectable indexes) 	1 (factory default), or current value
default route metric	Value that you assign to the default route as the “cost” of that route	0 – 32	0 (factory default), or current value
default route advertise mode	Routes that the interface advertises to neighboring DVMRP routers	<ul style="list-style-type: none"> ■ all ■ only 	all (factory default), or current value

ip multicast igmp
interface summary

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Summarizes key information about IGMP interfaces.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip m i i s

3900
9300

Fields in the IP Multicast IGMP Interface Summary Display

Field	Description
Index	Number assigned to the routing interface to its right.
Address	IP address of a routing interface in the system
TtlThreshold	Time-to-live (TTL) threshold that is assigned to the interface. This threshold affects IP multicast packets only.
Protocol	Multicast routing protocol that registers with IGMP. In release 3.0 software, there is one supported routing protocol (DVMRP).
Querier	IP address of the IGMP querier in the subnetwork to which the interface belongs. If the interface is functioning as the IGMP querier, this field shows self.

ip multicast igmp interface detail

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Supplements the IP multicast IGMP interface summary display with group and port information.

Valid Minimum Abbreviation

ip m i i d

Fields in the IP Multicast IGMP Interface Detail Display

Field	Description
Index	Number assigned to the routing interface to its right for identification purposes.
Address	IP address of a routing interface in the system that is identified by the index number to its left.
TtlThreshold	Time-to-live (TTL) threshold that is assigned to the interface. This threshold affects IP multicast packets only.
Protocol	Multicast routing protocol that registers with IGMP. In release 3.0 software, there is one supported routing protocol (DVMRP).
Querier	IP address of the IGMP querier in the subnetwork to which the interface belongs. If the interface is functioning as the IGMP querier, this field shows <code>Self</code> .
group	IP multicast group address for which packets have been received or forwarded
port(s)	Physical port numbers that are associated with the interface listed in the <code>Address</code> field that see incoming or outgoing traffic.

ip multicast igmp
interface TTL

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the time-to-live (TTL) threshold of a given routing interface. The interface compares the TTL value in each IP multicast packet against its TTL threshold. If the packet TTL is greater than the threshold TTL, the interface decrements the packet TTL by 1 and forwards the packet, provided that no other restrictions exist.

Valid Minimum Abbreviation

ip m i i t

Important Considerations

- Because IGMP is enabled by factory default, the system assigns a TTL threshold value of 1 as soon as you create an IP interface.
- This TTL threshold affects IP multicast packets only.

Options

Prompt	Description	Possible Values	[Default]
IP interfaces	Index numbers of the interfaces for which you want to modify the TTL threshold	<ul style="list-style-type: none">■ One or more valid interface index numbers■ ? (for a list of selectable indexes)	–
TTL threshold	Value you want to assign to the specified interfaces	0 – 255	1 (factory default), or current value

**ip multicast igmp
snooping**

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables or disables the system's ability to understand the Internet Group Management Protocol (IGMP) and snoop on IGMP packets to determine if IP multicast group members exist downstream from routing interfaces and therefore if the system should forward group traffic on those interfaces.

Valid Minimum Abbreviation`ip m i s`**Important Considerations**

- Your selection applies to all interfaces in the system.
- 3Com recommends that you keep IGMP snooping enabled at all times. It adds little processing overhead to the system and enhances the efficiency of your network if IP multicast traffic is present.

Options

Prompt	Description	Possible Values	[Default]
snooping mode	Whether the system can observe, record, and react to IGMP packets and set filters on appropriate ports in an interface	<ul style="list-style-type: none">■ enabled■ disabled	enabled (factory default), or current value

ip multicast igmp
querying

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables or disables the system’s ability to operate as the Internet Group Management Protocol (IGMP) querier if so elected by other IGMP-capable devices in the subnetwork. The IGMP querier is always the device with the lowest IP address.

Valid Minimum Abbreviation

ip m i q

Important Considerations

- Your selection applies to all interfaces in the system.
- The most efficient bandwidth usage is achieved by having the device that is closest to the source of IP multicast traffic operate as the querier for a given subnetwork.

Options

Prompt	Description	Possible Values	[Default]
query mode	Whether the system can offer itself as a candidate for election as the IGMP querier	<ul style="list-style-type: none">■ enabled■ disabled	enabled (factory default), or current value

✓ 3500

✓ 9000

9400

3900

9300

ip multicast cache

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays information about IP multicast traffic that has been observed on the system. For more detailed information, review the DVMRP cache. (See “ip multicast dvmrp cacheDisplay” earlier in this chapter.)

Valid Minimum Abbreviation

ip m c

Important Consideration

■ Although the Administration Console menu description is protocol independent multicast cache, this cache is not related to the multicast routing protocol called *Protocol Independent Multicast (PIM)*.

Options

Prompt	Description	Possible Values	[Default]
Multicast source address	Source for which you want to view cache information	<div><div>■ Depends on your network</div><div>■ 255.255.255.255 for all sources</div></div>	255.255.255.255 (factory default), or current value
Multicast group address	Multicast group for which you want to view cache information	<div><div>■ Depends on your network</div><div>■ 255.255.255.255 for all groups</div></div>	255.255.255.255 (factory default), or current value

Fields in the IP Multicast Cache Display

Field	Description
source	Subnetwork that contains a source device that is sending traffic addressed to the IP multicast group listed in the <code>group</code> field.
group	IP multicast group address of packets coming from the subnetwork listed to its left.
inVif	Index number of the interface that receives incoming IP multicast group traffic. Either an <code>I</code> or a <code>T</code> precedes the index number. An <code>I</code> indicates a regular IP multicast interface. A <code>T</code> indicates that the interface also operates as a DVMRP tunnel.
outVif	Index numbers of the interfaces to which traffic from the inVif is being forwarded.
inPorts	Physical port that corresponds to the interface that is listed in the inVifs field.
outPorts	Physical ports that correspond to the interfaces that are listed in the outVifs field.

✓ 3500

✓ 9000

9400

3900

9300

**ip multicast
traceRoute**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Provides a method for tracing the path that an IP multicast packet takes from a source to a particular receiver. Unlike unicast IP traceroute, multicast traceroute works in the reverse and requires a special packet type and implementation in routing devices.

Valid Minimum Abbreviation

`ip m t`

Important Considerations

- This command traces the path backwards from a specific receiving device to a specific source device. When you use this command, the receiver is assumed to be the system to which you are connected.
- This command produces a display that shows IP addresses of the interfaces that span from your system back to the source that you specify. The display also shows the number of hops back to those interfaces, the multicast routing protocols used, and the amount of time it takes to reach each hop from the receiver.
- All interim devices must support IP multicast traceroute for you to see a complete path on the display.

Options

Prompt	Description	Possible Values	[Default]
source IP address	IP address of the source device that sends traffic to a specific IP multicast group address	Any valid IP address for IP multicast source devices in your network	–
multicast group address	The IP multicast group address that the source is using for a particular application. This is useful when all applications come from the same source.	Any valid IP multicast group address used by source devices in your network	–

19

OPEN SHORTEST PATH FIRST (OSPF)

This chapter describes commands that you can use to configure Open Shortest Path First (OSPF) routing on your system.



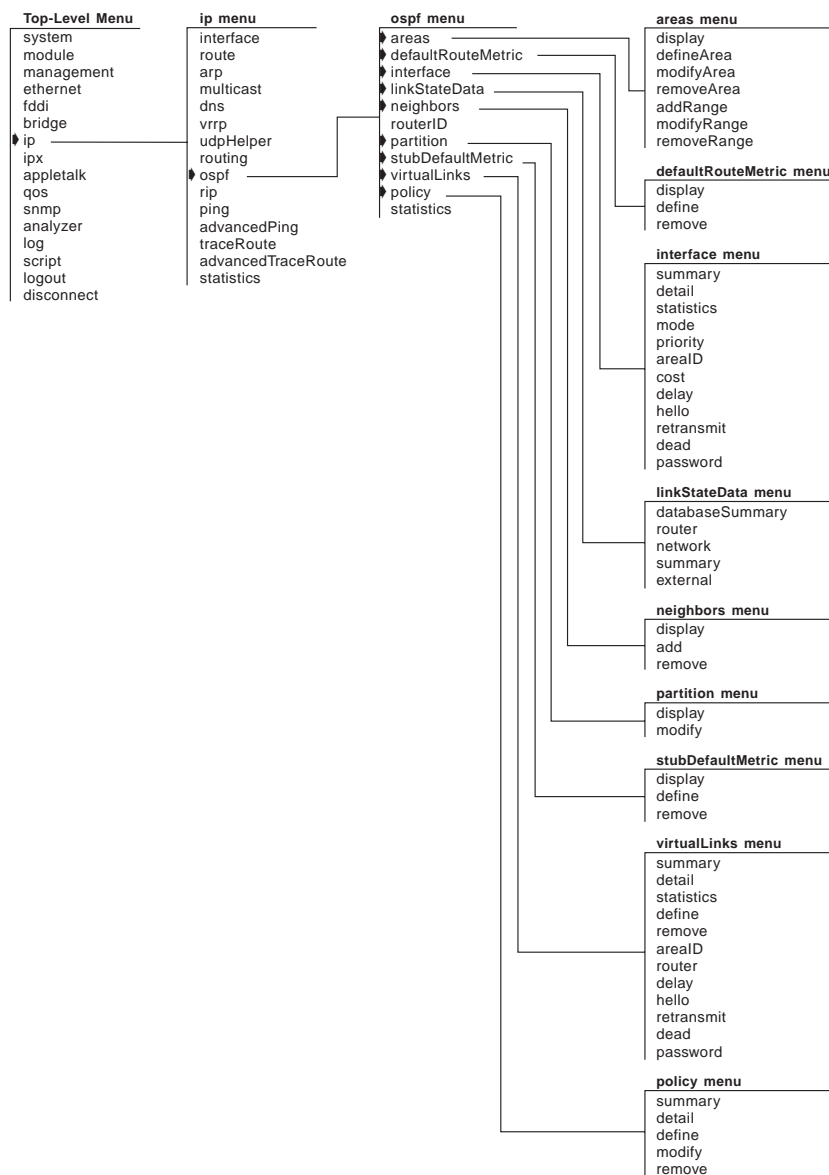
For more information about administering OSPF routing on your network, see the Implementation Guide for your system.



For the CoreBuilder[®] 9000, the commands in this chapter apply to Layer 3 switching modules only.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



ip ospf areas display

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays a list of existing OSPF areas.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a di

3900
9300

Fields in the IP OSPF Areas Display

Field	Description
Advertise	Whether the network range is advertised (y) or not (n)
AreaID	Area identifier
Indx	Entry index number for the area
IP Address	Network portion of IP address range
Mask	IP address range subnet mask
Stub	Whether the area is a stub area (y) or not (n)

**ip ospf areas
defineArea**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Defines an OSPF area.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a de

Important Considerations

- The backbone area 0.0.0.0 is configured by default.
- The area ID must be unique for the autonomous system.
- On the CoreBuilder 3500, you can define a maximum of eight areas.

Options

Prompt	Description	Possible Values	[Default]
Area ID	In the form n.n.n.n (where 0 <= n <= 255); functions as an area identification number to the OSPF autonomous system	Up to 255.255.255.255	–
Stub area	Whether this area is a stub area	<div><div>■ y (yes)</div><div>■ n (no)</div></div>	n (factory default), or current value

ip ospf areas
modifyArea

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Modifies an existing OSPF area.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a modifya

Options

3900
9300

Prompt	Description	Possible Values	[Default]
Area	Index number of the area that you want to modify	<div><div>■ Valid area index number</div><div>■ ? (for a list of selectable indexes)</div></div>	–
Area ID	In the form n.n.n.n (where 0 <= n <= 255); functions as an area identification number to the OSPF autonomous system	Up to 255.255.255.255	–
Stub area	Whether this area is a stub area	<div><div>■ y (yes)</div><div>■ n (no)</div></div>	n (factory default), or current value

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500
✓ 9000
9400

```
ip o a removea
```

Options

Prompt	Description	Possible Values	[Default]
Area	Index number of the area that you want to remove	<ul style="list-style-type: none"> Valid area index number all ? (for a list of selectable indexes) 	First available index number

**ip ospf areas
addRange**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Adds a range to an existing OSPF area.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a a

Options

3900
9300

Prompt	Description	Possible Values	[Default]
Area	Index number of the area to which you want to add the range	<ul style="list-style-type: none">Valid area index number? (for a list of selectable indexes)	–
IP address	IP address of the range that you want to add to the area	Up to 255.255.255.255	–
Subnet mask	Subnet mask of the range that you want to add to the area	Variable, based on address range class	Variable, based on address range class
Advertise range	Whether to advertise area range	<ul style="list-style-type: none">y (yes)n (no)	y

**ip ospf areas
modifyRange*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Modifies an OSPF area range.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ip o a modifyr`**Options**

Prompt	Description	Possible Values	[Default]
Area	Index number of the area that contains the range that you want to modify	<ul style="list-style-type: none"> ■ Valid area index number ■ ? (for a list of selectable indexes) 	–
IP address of range	Existing range that you want to modify (in the form of an IP address)	Up to 255.255.255.255	–
IP address	Range (in the form of an IP address)	Up to 255.255.255.255	Current value
Subnet mask	Subnet mask of the range that you want to modify	Variable, based on address range class	Current value
Advertise range	Whether to advertise the area range	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	Current value

IP OSPF Areas Modify Range Example

```

Select area {1-2|?}: 1
Enter IP address of range to modify: 3.3.3.1
Enter IP address [3.3.3.1]: 2.2.2.2
Enter subnet mask [255.0.0.0]: 255.255.0.0
Advertise this area range (yes,no) [yes]: y

```


ip ospf areas
removeRange

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Removes an OSPF area range.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o a remove

Options

Prompt	Description	Possible Values	[Default]
Area	Index number of the area that contains the range that you want to delete	<div><div>■</div>Valid area index number</div> <div><div>■</div>? (for a list of selectable indexes)</div>	–
IP address	IP address of the range that you want to delete	Up to 255.255.255.255	–

**ip ospf
defaultRouteMetric
display**

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the cost of a default route.

Valid Minimum Abbreviation

`ip o d di`

Important Considerations

- If a default metric is not defined, the router does not advertise itself as the default router.
- By default, the default route metric is not defined.

Field in the IP OSPF Default Route Metric Display

Field	Description
Default route metric	Cost (metric) that is associated with the default route. A higher cost indicates a slower route, for example, because it entails more hops or less bandwidth.

ip ospf
defaultRouteMetric
define

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines the default route metric for the router.

Valid Minimum Abbreviation

ip o d de

Important Considerations

- If a default metric is not defined, the router does not advertise itself as the default router.
- By default, the default route metric is not defined.
- Defining is default route metric is useful when the configuration supports multiple paths to the same destination. It provides a way to signify which of the paths is to be preferred.

Options

Prompt	Description	Possible Values	[Default]
Default route metric	Cost (metric) that is associated with the default route	1 – 65535	–

**ip ospf
defaultRouteMetric
remove**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes the default route metric.

Valid Minimum Abbreviation

ip o d r

Important Considerations

- If a default metric is not defined, the router does not advertise itself as the default router.
- By default, the default route metric is not defined.
- The default route metric is removed immediately after you enter the command. You are not prompted to confirm the deletion.

✓ 3500

✓ 9000

9400

3900

9300

ip ospf interface
summary

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays summary information for the system's OSPF interface configuration.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o i su

3900
9300

Fields in the IP OSPF Interface Summary Display

Field	Description
ArealD	OSPF area to which the interface belongs
Dead Intvl	Time interval (in seconds) before OSPF declares that a neighbor is dead
Hello Intvl	OSPF Hello packet transmit interval (in seconds) for the interface
Indx	Interface entry index; same number as the IP interface index
Password	Password that is associated with the OSPF interface
Pri	OSPF router priority for the interface
Rxmit Intvl	LSA retransmit interval (in seconds)
Xmit Cost	Interface transmit cost
Xmit Delay	Interface transmit delay (in seconds)

**ip ospf interface
detail*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays summary and detailed information for the system's OSPF interface configuration.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation`ip o i det`

3900

9300

Important Consideration

- The display also indicates whether IP routing and Internet Control Message Protocol (ICMP) router discovery are enabled and gives the OSPF router ID.

Fields in the IP OSPF Interface Detail Display

Field	Description
ArealD	OSPF area to which the interface belongs
BDR	IP interface of the backup designated router (BDR)
Dead Intvl	Time interval (in seconds) before OSPF declares that a neighbor is dead
DR	IP interface of the designated router (DR)
Hello Intvl	OSPF Hello packet transmit interval (in seconds) for the interface
Indx	Index number that corresponds to the IP interface for which OSPF information is displayed
IP address	IP address of the OSPF interface
Notes	When RouterID appears, the interface address is being used as the OSPF router ID
Password	Password that is associated with the OSPF interface
Pri	OSPF router priority for the interface
Rxmit Intvl	LSA retransmit interval (in seconds)

Field	Description
State	Interface state: <ul style="list-style-type: none">■ Disabled — OSPF is not enabled on the interface.■ Down — Interface is down, but OSPF is enabled on it.■ Loopback — Interface is a loopback interface.■ Waiting — Router is trying to determine the identity of the DR and BDR on the network.■ PTP — Interface is operational and connects to either a point-to-point network or a virtual link. The router attempts to form adjacency with the neighboring router.■ DRother — Interface is on a multiaccess network where this router is not the designated router or backup designated router.■ BDR — Router is the backup designated router on the attached network.■ DR — Router is the designated router on the attached network.
Xmit Cost	Interface transmit cost
Xmit Delay	Interface transmit delay (in seconds)

ip ospf interface statistics

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays statistics that are associated with specified OSPF interfaces.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

ip o i st

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the interface for which you want to display statistics	<ul style="list-style-type: none"> Valid interface index number all ? (for a list of selectable indexes) 	–

Fields in the IP OSPF Interface Statistics Display

Field	Description
adjacencyDown	Number of times that OSPF adjacencies have gone down
adjacencyUp	Number of times that OSPF adjacencies have been formed
authError	Number of packets discarded due to OSPF authentication errors Interpretation: <ul style="list-style-type: none"> A non-zero value is bad and means that packets from some OSPF routers are being discarded due to authentication errors. This statistic is incremented under the following circumstances: <ul style="list-style-type: none"> If the OSPF packet authentication type is something other than simple password (i.e., cryptographic authentication is not supported in the current implementation). If the OSPF packet contains a password but the interface does not have a password configured. If the OSPF packet has a simple password that does not match the password defined for the OSPF interface.
computeDR	Number of times that the designated router has been computed
lsaXsumError	Number of LSA checksum errors that were detected
mismatchAreaID	Number of interface area ID mismatches that were detected
mismatchAreaType	Number of interface area type mismatches that were detected

Field	Description
mismatchDead	<p>Number of router dead interval mismatches that were detected</p> <p>Interpretation:</p> <ul style="list-style-type: none">■ A non-zero value is bad and means that some OSPF routers on the interface are configured with a different dead interval than this router. This prevents the router from becoming a neighbor with these other routers. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none">■ When an OSPF Hello packet is received and the dead interval it defines is different from the dead interval configured on the OSPF interface.
mismatchHello	Number of Hello packet interval mismatches that were detected
mismatchMask	Number of subnet mask mismatches that were detected
packetXsumError	Number of packet checksum errors since interface has come up
receiveDD	<p>Number of database description packets that were received from valid OSPF neighbors.</p> <p>Interpretation:</p> <ul style="list-style-type: none">■ A non-zero value is OK. <p>Database description packets are sent when forming adjacencies with valid neighbors. A large number of receiveDD packets in a network whose configuration has not changed could indicate that adjacencies are being torn down and re-established.</p> <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none">■ When an OSPF database descriptor packet from a valid OSPF neighbor is received.
receivedUnknown	Number of unknown LSAs that were received

Field	Description
receiveError	<p>Number of general receive errors.</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value indicates that OSPF packets are being dropped and that this could be causing routing problems. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF Hello packet is received and the packet length is too short. ■ When an OSPF Hello packet is received that has the same router ID as the router receiving the packet. ■ When an OSPF database descriptor packet is received and the packet length is too short. ■ When an OSPF link state request (LSR) packet is received and the packet length is too short. ■ When processing an LSR packet, if the area is not configured on the interface. ■ When an OSPF link state update (LSU) packet is received and the packet length is too short. ■ When processing an LSU packet, if there are more than 500 advertisements the packet is not processed. ■ When an OSPF link state acknowledgement (LSAck) packet is received and the packet length is too short. ■ When processing an LSAck packet, if the area described by the packet is not known by the router receiving the packet. ■ When processing any OSPF packet, if the packet length is less than the OSPF header length then it must have been truncated and the packet is dropped. ■ When an OSPF packet is received on an interface that is not running OSPF. ■ When an OSPF packet is received over a virtual link, but the virtual link is down or not configured. ■ When an OSPF packet is received (over a non-virtual link) from a source whose IP network does not match the IP network of the interface on which it was received. ■ When an OSPF packet is received on a Non-Broadcast Multiple Access network from an unknown neighbor. ■ When an OSPF packet is received whose version is not OSPF version 2.
receiveHello	Number of Hello packets that were received
receiveLSAck	Number of LSA acknowledgments that were received
receiveLSR	Number of LSA request packets that were received

Field	Description
receiveLSU	Number of link state update packets that were received
transmitDD	<p>Number of database description packets that were transmitted</p> <p>Interpretation:</p> <ul style="list-style-type: none">■ A non-zero value is OK. <p>Database description packets are sent when forming adjacencies with valid neighbors. A large number in a network whose configuration has not changed could indicate that adjacencies are being torn down and re-established.</p> <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none">■ When an OSPF database descriptor packet is transmitted.
transmitError	<p>Number of general transmit errors</p> <p>Interpretation:</p> <ul style="list-style-type: none">■ A non-zero value indicates that an OSPF packet could not be sent either out a particular interface, or to a particular destination. This could prevent OSPF from running properly within the autonomous system and lead to routing problems. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none">■ When an OSPF Hello, LSU, or LSAck is being sent as a multicast packet on a non-broadcast multiple access network.
transmitHello	Number of Hello packets that were transmitted
transmitLsAck	Number of LSA acknowledgments that were transmitted
transmitLSR	Number of LSA request packets that were transmitted
transmitLSU	Number of link state update packets that were transmitted

**ip ospf interface
mode*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Enables or disables OSPF on specified IP interfaces.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o i m

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more IP interfaces on which you want to enable or disable OSPF	<ul style="list-style-type: none"> One or more valid IP interface index numbers all ? (for a list of selectable indexes) 	all (factory default), or current value
OSPF mode	Whether to disable or enable OSPF on the specified IP interface	<ul style="list-style-type: none"> disabled enabled 	disabled (factory default), or current value

ip ospf interface
priority

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Assigns interface priority to the OSPF router.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o i pr

Important Consideration

- The interface priority of an OSPF router determines its status as a designated router.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more IP interfaces to which you want to assign a priority	<ul style="list-style-type: none">■ One or more valid IP interface index numbers■ all■ ? (for a list of selectable indexes)	all (factory default), or current value
Priority	Interface priority: <ul style="list-style-type: none">■ If 0, router will not be the default router.■ If 1 – 255, the highest priority becomes the designated router.	<ul style="list-style-type: none">■ 0 – 255	1

**ip ospf interface
areaID**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Associates an interface with an OSPF area.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o i a

Important Considerations

- Set the area ID to the same value for all routers on the network segment because they are in the same area.
- 0.0.0.0 indicates the OSPF backbone area.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces that you want to associate with the area	<ul style="list-style-type: none">■ One or more valid IP interface index numbers■ all■ ? (for a list of selectable indexes)	all (factory default), or current value
Area ID	ID of area, in the form n.n.n.n (where 0 <= n <= 255) with which you want to associate the specified interfaces	Valid area ID	0.0.0.0 (factory default), or current value

ip ospf interface cost

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Assigns a cost to an OSPF interface.

Valid Minimum Abbreviation

ip o i c

Important Consideration

- The interface cost reflects the line speed of the port. Although the system calculates a default cost value based on the module media type, you can use this command to manually change the cost to a different value.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces to which you want to assign a cost	<ul style="list-style-type: none">■ One or more valid IP interface index numbers■ all■ ? (for a list of selectable indexes)	all (factory default), or current value
Cost	Cost that you want to assign to the specified interface (Higher values are slower ports.)	1 – 65535	Cost of slowest port (usually 1)

**ip ospf interface
delay*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Sets the OSPF interface transmit delay.

✓ 3500
✓ 9000
94003900
9300**Valid Minimum Abbreviation**`ip o i del`**Important Considerations**

- The system adds the value of the transmit delay to all link state advertisements (LSAs) that it sends out to the network. Set the transmit delay according to the link speed: use a longer transmit delay time for slower link speeds.
- The transmit delay must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces for which you want to set the transmit delay	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Transmit delay	Delay (in seconds) that you want to assign to the specified interface	1 – 65535 seconds	1 (factory default), or current value

ip ospf interface hello *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500
✓ 9000
9400

Sets the interface Hello interval.

Valid Minimum Abbreviation

ip o i he

3900
9300

Important Considerations

- Hello packets inform other routers that the sending router is still active on the network.
- If a router does not send Hello packets for a period of time specified by the dead interval, the router is considered inactive by its neighbors.
- The Hello packet interval must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces for which you want to set the Hello interval	<ul style="list-style-type: none">■ One or more valid IP interface index numbers■ all■ ? (for a list of selectable indexes)	all (factory default), or current value
Hello packet interval	Interval (in seconds) at which the interface transmits Hello packets	1 – 65535 seconds	10 (factory default), or current value

**ip ospf interface
retransmit*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Specifies the OSPF link state advertisement (LSA) retransmit interval for an interface.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ip o i r`**Options**

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces on which you want to set the LSA retransmit interval	<ul style="list-style-type: none"> ■ One or more valid IP interface index numbers ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
LSA retransmit time	Interval at which the specified interface retransmits LSAs	1 – 65535 seconds	5 (factory default), or current value

ip ospf interface dead

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Specifies the dead interval for an interface.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ip o i dea

3900

9300

Important Consideration

- Set the dead interval to the same value for all routers on the network.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces on which you want to set the dead interval	<ul style="list-style-type: none">■ One or more valid IP interface index numbers■ all■ ? (for a list of selectable indexes)	all (factory default), or current value
Dead interval	Maximum duration (in seconds) that neighbor routers wait for a Hello packet before they determine that the transmitting router is inactive	1 – 65535 seconds	40 (factory default), or current value

**ip ospf interface
password*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Sets password security for an OSPF interface.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation`ip o i pa`**Important Considerations**

- To remove a previously assigned password, set the password to `none`.
- The password must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of one or more interfaces for which you want to assign or remove a password	<ul style="list-style-type: none">■ One or more valid IP interface index numbers■ all■ ? (for a list of selectable indexes)	all (factory default), or current value
Password	Password for the specified interface The none option removes a previously assigned password.	<ul style="list-style-type: none">■ Up to eight ASCII characters	none (factory default), or current value

ip ospf linkStateData
databaseSummary

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Summarizes link state advertisements (LSAs) in the link state database.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o l d

Important Consideration

- To view link state database information, OSPF must be active (enabled).

Options

Prompt	Description	Possible Values	[Default]
Area ID	Area ID (in the form n.n.n.n where 0 <= n <= 255) that corresponds to the OSPF area for which you want to view LSA summary information	Valid area ID	0.0.0.0 (factory default), or current value
Area mask	Subnet mask of OSPF area for which you want to view LSA summary information	Valid area mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data Database Summary Display

Field	Description
Checksum summation	Total of all LSA checksums
External LSAs	Number of external link LSAs
LSA count	Number of LSAs
Network LSAs	Number of network link LSAs
Router LSAs	Number of router link LSAs
Summary LSAs	Number of summary link LSAs

ip ospf linkStateData router

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays router link state advertisements (LSAs) in the link state database.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o l r

Important Consideration

- To view link state database information, OSPF must be active (enabled).

Options

Prompt	Description	Possible Values	[Default]
Area ID	Area ID (in the form n.n.n.n where 0 <= n <= 255) that corresponds to the OSPF area for which you want to view router link state advertisement information	Valid area ID	0.0.0.0 (factory default), or current value
Area mask	Subnet mask of OSPF area for which you want to view router link state advertisement information	Valid area mask	0.0.0.0 (factory default), or current value
LSID	Link State ID: router ID of the originating router (in the form of an IP address)	Router ID	0.0.0.0 (factory default), or current value
LSID mask	Link State ID bit mask (Example: 255.0.0.0)	Link State ID bit mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data Router Display

Field	Description
Flags	<ul style="list-style-type: none"> ■ V — Router is the endpoint of an active virtual link that is using the area as a transmit area. ■ ASBR — Router is an autonomous system boundary router. ■ ABR — Router is an area border router.

Field	Description
Link Data	<ul style="list-style-type: none">■ PTP — MIB II index value for an unnumbered point-to-point interface.■ Transit Net — IP address of the router's interface■ Stub Net — Network IP address mask■ Virtual link — IP interface address of neighboring router
Link ID	<ul style="list-style-type: none">■ PTP — Neighboring router's router ID■ Transit Net — Address of designated router■ Stub Net — IP network/subnetwork number■ Virtual link — Neighboring router's router ID
Link Type	<ul style="list-style-type: none">■ PTP — Connection is point-to-point to another router.■ Transit Net — Connection is to a transit network (one that has more than one OSPF router on it).■ Stub Net — Connection is to a stub network.■ Virtual link — Connection is to a far-end router that is the endpoint of a virtual link.
LS Age	Time (in seconds) since LSA was originated
LS Seq	Sequence number of the LSA (used to detect older duplicate LSAs)
LSID	ID number of the router that originated the LSA
Metric	Cost of the link
Router ID	Originating router ID

ip ospf linkStateData network

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays network link state advertisements (LSAs) in the link state database.

Valid Minimum Abbreviation

ip o l n

Important Consideration

- To view link state database information, OSPF must be active (enabled).

Options

Prompt	Description	Possible Values	[Default]
Area ID	Area ID (in the form n.n.n.n where 0 <= n <= 255) that corresponds to the OSPF area for which you want to view network LSA information	Valid area ID	0.0.0.0 (factory default), or current value
Area Mask	Subnet mask of OSPF area for which you want to view network LSA information	Valid area mask	0.0.0.0 (factory default), or current value
LSID	Link State ID: interface address of the designated router	Valid IP address	0.0.0.0 (factory default), or current value
LSID mask	Link State ID bit mask (Example: 255.0.0.0)	Link State ID bit mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data Network Display

Field	Description
Attached routers	List of routers that are fully adjacent to the designated router (DR); also the DR
LS Age	Time (in seconds) since the LSA was originated
LS Seq	Sequence number of the LSA (used to detect older duplicate LSAs)
LSID	Interface address (in the form of an IP address) of the designated router
Network mask	IP address mask for the network
Router ID	Originating router ID

ip ospf linkStateData summary

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays summary link state advertisements (LSAs) in the link state database.

Valid Minimum Abbreviation

ip o l s

Important Consideration

- To view link state database information, OSPF must be active (enabled).

Options

Prompt	Description	Possible Values	[Default]
Area ID	Area ID (in the form n.n.n.n where 0 <= n <= 255) that corresponds to the OSPF area for which you want to view summary LSA information	Valid area ID	0.0.0.0 (factory default), or current value
Area mask	Subnet mask of the OSPF area for which you want to view summary LSA information	Valid area mask	0.0.0.0 (factory default), or current value
LSID	Link State ID: <ul style="list-style-type: none"> ■ For type 3 summary LSAs, this is the IP address of the destination network ■ For type 4 summary LSAs, this is the autonomous system boundary router's Router ID (in the form of an IP address) 	<ul style="list-style-type: none"> ■ For type 3 summary LSAs, a valid IP address ■ For type 4 summary LSAs, a valid router ID 	0.0.0.0 (factory default), or current value
LSID mask	Link State ID bit mask (Example: 255.0.0.0)	Link State ID bit mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data Summary Display

Field	Description
LS Age	Time (in seconds) since LSA was originated
LS Seq	Sequence number of the LSA (used to detect older duplicate LSAs)
LSID	<ul style="list-style-type: none">■ Type 3 — Destination network's IP address■ Type 4 — ASBR's OSPF router ID
Metric	Cost to reach the network
Network mask	<ul style="list-style-type: none">■ For Type 3 — destination network's IP address mask■ For Type 4 — Not used, must be 0 (--)
Router ID	Originating router ID

ip ospf linkStateData
external

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays external network link state advertisements (LSAs) in the link state database.

Valid Minimum Abbreviation

ip o l e

Important Consideration

- To view link state database information, OSPF must be active (enabled).

Options

Prompt	Description	Possible Values	[Default]
LSID	Link State ID (in the form of the destination network's IP address)	Valid IP address	0.0.0.0 (factory default), or current value
LSID mask	Link State ID bit mask (Example 255.0.0.0)	Link State ID bit mask	0.0.0.0 (factory default), or current value

Fields in the IP OSPF Link State Data External Display

Field	Description
Fwd Address	Forwarding address for data traffic to the advertised destination
LS Age	Time (in seconds) since LSA was originated
LS Seq	Sequence number of the LSA (used to detect older duplicate LSAs)
LSID	IP network number
Metric	Cost to reach advertised destination
Network Mask	IP address mask for the advertised destination
Router ID	Originating router ID
RouteTag	Not used by OSPF; these 32 bits may be used to communicate other information between boundary routers. Tag contents are defined by applications.
Type	<ul style="list-style-type: none">■ Type 1 — normal link state metric■ Type 2 — metric is larger than any local link state path

**ip ospf neighbors
display****For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays information about currently defined neighbors in an OSPF area.

✓ 3500
✓ 9000
9400**Valid Minimum Abbreviation**

ip o n d

3900
9300**Fields in the IP OSPF Neighbors Display**

Field	Description
Flags	Neighbor identification flags: <ul style="list-style-type: none"> ■ D — Dynamic neighbor ■ S — Static neighbor ■ BDR — Backup designated router ■ DR — Designated router <p>Example: [S, BDR] + [D, DR] is a static neighboring backup designated router and a dynamic neighboring designated router</p>
Indx	Interface index that corresponds to the interface to which a neighbor belongs
Neighbor Addr	Interface address of neighbor
Pri	Neighbor's OSPF router priority
ReqQ	Number of LSAs being requested from neighbor
Router ID	Neighbor's OSPF router ID
RxQ	Number of LSAs in local retransmit queue to the neighbor
State	Neighbor's adjacency: <ul style="list-style-type: none"> ■ Down — No recent data received from neighbor, connection is down. ■ Attempt — Only used on nonbroadcast networks. No recent data received from neighbor (will attempt to contact). ■ Init — Have recently seen Hello packet from neighbor; however, two-way communication has not been established. ■ Two-way — Bidirectional communication has been established. ■ ExStart — Taking initial step to create adjacency between neighboring routers. ■ Exchange — Database descriptions are being exchanged. ■ Loading — LSA databases are being exchanged. ■ Full — Neighboring routers are fully adjacent.
SumQ	Number of LSAs in LSA summary queue for the neighbor

ip ospf neighbors add

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

Adds a neighbor static IP address to an existing interface.

Valid Minimum Abbreviation

ip o n a

Important Consideration

- The system learns neighbor addresses dynamically on interfaces that support multicast routing. Define static neighbors only on nonmulticast interfaces.

Options

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the interface to which you want to add a neighbor	<ul style="list-style-type: none">■ Valid interface index number■ ? (for a list of selectable indexes)	First available (factory default), or current value
Static neighbor address	Address of neighbor that you want to define	Valid IP address on interface subnetwork	–

**ip ospf neighbors
remove**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes a static neighbor from an existing interface.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

`ip o n r`

Options

3900
9300

Prompt	Description	Possible Values	[Default]
IP interface	Index number of the interface from which you want to remove a neighbor	<ul style="list-style-type: none">Valid interface index number? (for a list of selectable indexes)	First available (factory default), or current value
Neighbor address	Address of neighbor that you want to remove	Valid IP address on interface subnetwork	–

ip ospf routerID *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Sets the OSPF router ID.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o r

Important Considerations

- The OSPF router ID identifies the router to other routers within an autonomous system. Three types of router identifiers are available; all three take the form of an IP address:
 - **Default** — A unique ID that the system generates and uses as the default router ID
 - **Interface** — The index of an IP interface on the router
 - **Address** — An ID that you define in the form of an IP address
- OSPF routing must be inactive (disabled) before you can add or modify an OSPF router ID. To set the OSPF mode to `disabled`, see “ip ospf interface mode” earlier in this chapter. After you modify the router ID, you can set the OSPF mode to `enabled` on the interface
- The router ID must be unique from all other router IDs and ip interfaces in the autonomous system for OSPF to operate correctly. Choose the `default` setting to ensure unique router IDs.
- The resulting prompt depends on the router ID type that you choose.

Options

Prompt	Description	Possible Values	[Default]
Router ID type	Type of router identifier that you want to define	<ul style="list-style-type: none"> ■ default ■ interface ■ address 	default (factory default), or current value
IP interface	<i>For interface router ID type only.</i> Index number of IP interface to use as router ID.	<ul style="list-style-type: none"> ■ Valid IP interface ■ ? (for a list of selectable indexes) 	First available (factory default), or current value
Router ID	<i>For address router ID type only.</i> Identifier that is assigned to router in the form of an IP address 0.0.0.0 and 255.255.255.255 are invalid and will be rejected	User-defined router ID	Unique router ID generated by the system (factory default), or current value

IP OSPF Router ID Example (Interface Type)

```
Current OSPF router id = 0.43.66.0 (default)
Enter router ID type {default,interface,address|?} [default]: interface
Select IP interface {1-3|?}: 1
```

IP OSPF Router ID Example (Address Type)

```
Current OSPF router id = 24.23.11.23 (address)
Enter router ID type {default,interface,address|?} [address]: address
Enter router ID [24.23.11.23]: 101.89.2.4
```


ip ospf partition display

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays OSPF memory allocation.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o pa d

Important Consideration

- See “ip ospf partition modify” later in this chapter for information on how OSPF memory allocation works and how to modify it.

3900
9300

Fields in the IP OSPF Partition Display

Field	Description
Current partition maximum size	OSPF memory partition upper limit as implemented at the last system reboot.
Configured partition maximum size	Last value that you entered, which will become the current partition maximum size after the next system reboot. <ul style="list-style-type: none">■ 0 means that OSPF has been set to use the system memory partition at the next reboot.■ 1 means that OSPF has been set to use the default memory allocation scheme, deriving its partition size from the maximum size of the IP routing table at the next reboot.■ Any other value that does not equal the current partition maximum size means that OSPF has been manually set to use a specific maximum partition size at the next reboot.
Allocated partition size	Module’s current working memory. OSPF dynamically allocates memory in 100,000-byte chunks, up to the current partition maximum size.
OSPF is using the system partition	The administrator used the ip ospf partition modify command to set a partition value of 0. The OSPF protocol is using the system memory partition instead of its own partition, and there is no specified OSPF memory limit.

ip ospf partition modify

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the maximum memory that OSPF can allocate.

✓ 3500
✓ 9000
9400

3900
9300

Valid Minimum Abbreviation

ip o pa m

Important Considerations

- There are three choices for memory allocation:
 - Have the system intelligently determine the maximum OSPF memory partition size (partition size = 1). This is the default.
 - Have OSPF be part of system memory, growing as needed and without limit (partition size = 0).
 - Configure the maximum OSPF memory partition size manually (partition size = 4096 - <maximum available memory>).
- You typically do not have to modify the OSPF memory allocation. However, if the softRestarts statistic shown by the `ip ospf statistics` option begins to climb, it means that OSPF is thrashing for memory and you must increase the maximum memory.



For a complete description of OSPF memory allocation, see the “OSPF Memory Partition” section in the OSPF chapter of the Implementation Guide.

- The partition size option that you enter takes effect after a system reboot.

Options

Prompt	Description	Possible Values	[Default]
New partition maximum size	Maximum memory size (in bytes) to allocate to OSPF system operations	<ul style="list-style-type: none"> ■ 4096 to <maximum available size> ■ 0 (to specify system memory partition) ■ 1 (to specify a size based on amount of memory and the maximum routing table size. On extended memory systems, this is 4,200,000.) 	1 (factory default), or current OSPF partition size

ip ospf
stubDefaultMetric
display

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the stub default metric value for an area border router.

Valid Minimum Abbreviation

ip o stu di

Important Considerations

- The stub default metric value determines if the router generates the default route into the stub areas of the network. This value applies to area border routers (ABRs) that have attached stub areas.



If a stub default metric is not defined, the router does not advertise a default route into the attached stub area.

- By default, the `stub default metric` is not defined.

Field in the IP OSPF Stub Default Metric Display

Field	Description
Stub default metric	Currently defined OSPF stub default metric

**ip ospf
stubDefaultMetric
define**

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines the stub default metric value for an OSPF area border router.

Valid Minimum Abbreviation

`ip o stu de`

Important Considerations

- The stub default metric value determines if the router generates the default route into the stub areas of the network. This value applies to area border routers (ABRs) that have attached stub areas.



If a stub default metric is not defined, the router does not advertise a default route into the attached stub area.

- By default, the `stub default metric` is not defined.

Options

Prompt	Description	Possible Values	[Default]
Stub default metric	Stub default metric value to define for the area border router. Higher numbers are slower.	1 – 65535	Current stub default metric

**ip ospf
stubDefaultMetric
remove**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Disables the stub default metric on an OSPF area border router.

Valid Minimum Abbreviation

`ip o stu r`

Important Considerations

- The system removes the current stub default metric value immediately after you enter the command.
- The stub default metric value determines if the router generates the default route into the stub areas of the network. This value applies to area border routers (ABRs) that have attached stub areas.



If a stub default metric is not defined, the router does not advertise a default route into the attached stub area.

- By default, the `stub default metric` is not defined.

ip ospf virtualLinks summary

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays summary information about a virtual link.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v su

Options

3900
9300

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to display summary information	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–

Fields in the IP OSPF Virtual Links Summary Display

Field	Description
Dead Intvl	Number of seconds before the area border router's neighbors declare it down, when they stop hearing the router's Hellos
Hello Intvl	Length of time (in seconds) between Hello packets
Indx	Index number of the virtual link
Password	Password for the virtual link
Rxmit Intvl	Length of time (in seconds) between link state advertisement retransmissions
Target Router	End-point area border router where the virtual link terminates
Transit Area	Common area that the virtual link uses to reach the target router
Xmit Delay	Estimated number of seconds that it takes to transmit a link state update packet over the virtual link

ip ospf virtualLinks detail

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays detailed information about a virtual link.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v det

Important Consideration

- This display also contains virtual link detail and neighbor information.

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to display detail information	<ul style="list-style-type: none">■ Index number of a currently defined virtual link■ all■ ? (for a list of selectable indexes)	–

Fields in the IP OSPF Virtual Links Detail Display

Field	Description
Dead Intvl	Number of seconds before the area border router’s neighbors declare it down, when they stop hearing the router’s Hellos
Hello Intvl	Length of time (in seconds) between Hello packets
Indx	Index number of the virtual link
Password	Password for the virtual link
Rxmit Intvl	Length of time (in seconds) between link state advertisement retransmissions
Target Router	End-point area border router where the virtual link terminates
Transit Area	Common area that the virtual link uses to reach the target router
Xmit Delay	Estimated number of seconds that it takes to transmit a link state update packet over the virtual link

Fields in the IP OSPF Virtual Links Detail Display

Field	Description
Cost	Cost of sending a packet over the virtual link, expressed in the link state metric
Indx	Index number of the virtual link
Local Address	Address of the local router
Remote Address	Address of the remote router
State	State of the virtual link

Fields in the IP OSPF Virtual Links Neighbor Display

Field	Description
Indx	Index number for the interface to which a neighbor belongs
ReqQ	Number of LSAs that are being requested from the neighbor
RxQ	Number of LSAs that are in the local retransmit queue to the neighbor
State	Neighbor's adjacency: <ul style="list-style-type: none"> ■ Down — No recent data received from neighbor, connection is down. ■ Attempt — Only used on nonbroadcast networks. No recent data received from neighbor (will attempt to contact). ■ Init — Have recently seen Hello packet from neighbor; however, two-way communication has not been established. ■ Two-way — Bidirectional communication has been established. ■ ExStart — Taking initial step to create adjacency between neighboring routers. ■ Exchange — Database descriptions are being exchanged. ■ Loading — LSA databases are being exchanged. ■ Full — Neighboring routers are fully adjacent.
SumQ	Number of LSAs in LSA summary queue for the neighbor

ip ospf virtualLinks statistics

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays statistics that are associated with virtual links.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v st

Options

3900
9300

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to display statistics	<ul style="list-style-type: none">Valid interface index numberall? (for a list of selectable indexes)	–

Fields in the IP OSPF Virtual Links Statistics Display

Field	Description
adjacencyDown	Number of times that OSPF adjacencies have gone down
adjacencyUp	Number of times that OSPF adjacencies have been formed
authError	Number of packets discarded due to OSPF authentication errors Interpretation: <ul style="list-style-type: none">A non-zero value is bad and means that packets from some OSPF routers are being discarded due to authentication errors. This statistic is incremented under the following circumstances: <ul style="list-style-type: none">If the OSPF packet authentication type is something other than simple password (that is, cryptographic authentication is not supported in the current implementation).If the OSPF packet contains a password but the interface does not have a password configured.If the OSPF packet has a simple password that does not match the password defined for the OSPF interface.
computeDR	Number of times that the designated router was computed
lsaXsumError	Number of LSA checksum errors that have been detected
mismatchAreaID	Number of interface area ID mismatches that have been detected
mismatchAreaType	Number of interface area type mismatches that have been detected

Field	Description
mismatchDead	<p>Number of router dead interval mismatches that were detected</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value is bad and means that some OSPF routers on the interface are configured with a different dead interval than this router. This prevents the router from becoming a neighbor with these other routers. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF Hello packet is received and the dead interval it defines is different from the dead interval configured on the OSPF interface.
mismatchHello	Number of Hello packet interval mismatches that have been detected
mismatchMask	Number of subnet mask mismatches that have been detected
packetXsumError	Number of packet checksum errors since the interface has come up
receiveDD	<p>Number of database description packets that were received from valid OSPF neighbors.</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value is OK. <p>Database description packets are sent when forming adjacencies with valid neighbors. A large number of receiveDD packets in a network whose configuration has not changed could indicate that adjacencies are being torn down and reestablished.</p> <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF database descriptor packet from a valid OSPF neighbor is received.
receivedUnknown	Number of unknown LSAs that have been received

Field	Description
receiveError	<p>Number of general receive errors.</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value indicates that OSPF packets are being dropped and that this could be causing routing problems. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF Hello packet is received and the packet length is too short. ■ When an OSPF Hello packet is received that has the same router ID as the router receiving the packet. ■ When an OSPF database descriptor packet is received and the packet length is too short. ■ When an OSPF link state request (LSR) packet is received and the packet length is too short. ■ When processing an LSR packet, if the area is not configured on the interface. ■ When an OSPF link state update (LSU) packet is received and the packet length is too short. ■ When processing an LSU packet, if there are more than 500 advertisements the packet is not processed. ■ When an OSPF link state acknowledgement (LSAck) packet is received and the packet length is too short. ■ When processing an LSAck packet, if the area described by the packet is not known by the router receiving the packet. ■ When processing any OSPF packet, if the packet length is less than the OSPF header length then it must have been truncated and the packet is dropped. ■ When an OSPF packet is received on an interface that is not running OSPF. ■ When an OSPF packet is received over a virtual link, but the virtual link is down or not configured. ■ When an OSPF packet is received (over a non-virtual link) from a source whose IP network does not match the IP network of the interface on which it was received. ■ When an OSPF packet is received on a Non-Broadcast Multiple Access network from an unknown neighbor. ■ When an OSPF packet is received whose version is not OSPF version 2.
receiveHello	Number of Hello packets that have been received
receiveLsAck	Number of LSA acknowledgments that have been received
receiveLSR	Number of LSA request packets that have been received

Field	Description
receiveLSU	Number of link state update packets that have been received
transmitDD	<p>Number of database description packets that were transmitted</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value is OK. <p>Database description packets are sent when forming adjacencies with valid neighbors. A large number in a network whose configuration has not changed could indicate that adjacencies are being torn down and re-established.</p> <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF database descriptor packet is transmitted.
transmitError	<p>Number of general transmit errors</p> <p>Interpretation:</p> <ul style="list-style-type: none"> ■ A non-zero value indicates that an OSPF packet could not be sent either out a particular interface, or to a particular destination. This could prevent OSPF from running properly within the autonomous system and lead to routing problems. <p>This statistic is incremented under the following circumstances:</p> <ul style="list-style-type: none"> ■ When an OSPF Hello, LSU, or LSAck is being sent as a multicast packet on a non-broadcast multiple access network.
transmitHello	Number of Hello packets that have been transmitted
transmitLSAck	Number of LSA acknowledgments that have been transmitted
transmitLSR	Number of LSA request packets that have been transmitted
transmitLSU	Number of link state update packets that have been transmitted

ip ospf virtualLinks
define

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Creates a new virtual link to a destination router.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v def

Important Considerations

- All areas of an OSPF routing domain must connect to the backbone area. In cases where an area border router does not have direct, physical access to the backbone, you must configure a virtual link to act as a logical link to the backbone area.
- You can define up to 32 virtual links per router.

Options

Prompt	Description	Possible Values	[Default]
Transit area	Area ID (in the form n.n.n.n where 0 <= n <= 255) through which the link is going	Currently defined area ID	–
Target router	ID of the target router, which is the router where the virtual link terminates	Valid IP address of OSPF area border router	–

Removes a virtual link.

9400

ip o v rem

3900

9300

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link that you want to remove	<ul style="list-style-type: none"> Index number of a currently defined virtual link all ? (for a list of selectable indexes) 	–

ip ospf virtualLinks
areaID

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Modifies the transit area that is associated with a virtual link.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v a

Options

3900
9300

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify a new area ID	<div><div>■</div>Index number of a currently defined virtual link</div> <div><div>■</div>all</div> <div><div>■</div>? (for a list of selectable indexes)</div>	–
Target area	Area ID (in the form n.n.n.n where 0 <= n <= 255) of the transit area through which the virtual link must pass to reach the target router	ID of a currently defined area	Current value

**ip ospf virtualLinks
router**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies the target router that is associated with a virtual link.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ip o v ro

Options

3900

9300

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify a new target router	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–
Target router	IP address of the new destination area border router where the virtual link terminates	Valid IP address of an OSPF area border router	Current value

ip ospf virtualLinks
delay

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Sets the virtual link transmit delay, in seconds.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ip o v del

Important Consideration

- The virtual link transmit delay must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify the transmit delay	<ul style="list-style-type: none">■ Index number of a currently defined virtual link■ all■ ? (for a list of selectable indexes)	–
Transmit delay	New virtual link transmit delay (in seconds)	1 – 65535 seconds	1 (factory default), or current value

ip ospf virtualLinks hello

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the virtual link Hello interval, in seconds.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o v he

Important Considerations

- Hello packets inform other routers that the sending router is still active on the network.
- If a router does not send Hello packets for a period of time specified by the dead interval, the router is considered inactive by its neighbors.
- The virtual link Hello interval must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify the Hello interval	<ul style="list-style-type: none"> ■ Index number of a currently defined virtual link ■ all ■ ? (for a list of selectable indexes) 	–
Hello packet interval	Interval (in seconds) at which the area border router transmits Hello packets	1 – 65535 seconds	10 (factory default), or current value

ip ospf virtualLinks
retransmit

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Sets the virtual link retransmit interval, in seconds.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ip o v ret

Options

- 3900
- 9300

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify the retransmit interval	<ul style="list-style-type: none">■ Index number of a currently defined virtual link■ all■ ? (for a list of selectable indexes)	–
LSA retransmit time	Interval (in seconds) at which the area border router retransmits LSAs over the virtual link	1 – 65535 seconds	50 (factory default), or current value

**ip ospf virtualLinks
dead*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Sets the virtual link dead interval, in seconds.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o v dea

Important Consideration

- Set the dead interval to the same value for all routers on the network.

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify the dead interval	<ul style="list-style-type: none">■ Index number of a currently defined virtual link■ all■ ? (for a list of selectable indexes)	–
Dead interval	Maximum duration (in seconds) that neighbor routers wait for a Hello packet before they determine that the transmitting router is inactive	1 – 65535 seconds	40 (factory default), or current value

ip ospf virtualLinks password

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets password security for a virtual link.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

ip o v p

Important Considerations

- Set the virtual link password to none to remove a previously assigned password.
- The password must be consistent throughout the autonomous system.

Options

Prompt	Description	Possible Values	[Default]
Virtual link	Index number of the virtual link for which you want to specify a password	<ul style="list-style-type: none">■ Valid IP interface index number■ all■ ? (for a list of selectable indexes)	–
Virtual link password	Password for the specified virtual link	Up to eight ASCII characters	none (factory default), or current value

3900
9300

Field	Description
Action	Action for the route (accept or reject)
Idx	Index number of the interface
Protocol	Protocol (for example, OSPF)
Route	Source network
Source	Source router
Type	Whether the policy is an import or export policy
Wt	Administrative weight (range of values: 1 through 16)

ip ospf policy detail

Displays summary and detailed information about OSPF routing policies.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o po det

Important Considerations

- This display contains the summary information plus three additional fields: interface, metric, and ASEType.
- Your system has one unified IP routing table. Routing policies allow you to control the flow of information among the network, the protocols, and the routing tables on your system.
- There are two classes of routing policies:
 - **Import policies** — Control which OSPF non-self-originated external routes are stored in the routing table. OSPF import policies control only what the local router uses. They do not affect the propagation of non-self-originated external routes to other routers.
 - **Export policies** — Used on OSPF boundary routers to control which self-originated external routing updates are placed in the link-state database for propagation over the network. In this way, export policies govern what other routers learn with regard to the local boundary router's self-originated information.
- The system tracks policies that you define in both OSPF and Routing Information Protocol (RIP), so the indexes that are assigned to your policies may have gaps. For example, if you have OSPF policies 1 and 2, and RIP policies 3 through 6, the next policy is 7.

Fields in the IP OSPF Policy Detail Display

Field	Description
Action	Action for the route (accept or reject)
ASETType	Type of external metric — Type 1 or Type 2 — specified in the AS external link advertisement. OSPF boundary routers use Type 1 as default. Only applicable to export policies.
Index	Index number of the policy
Interface	Origin interface (only applicable when specifying direct as Origin Protocol)
Metric	Adjustment to the cost metric of routes that match the policy
Protocol	Origin protocol (for export policies only). Can also specify a direct or static route.

Field	Description
Route	Route against which the policy is applied
Source	Source router (only applicable to export policies that do not specify <code>direct</code> as Origin Protocol)
Type	Whether the policy is an <code>import</code> or <code>export</code> policy
Weight	Administrative weight (range of values: 1 through 16)

ip ospf policy define

Defines import and export OSPF routing policies.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o po def

Important Considerations

- The system assigns an index number to each policy and takes into account all route policies, Routing Information Protocol (RIP) and OSPF, that are set on the system.
- There are certain conditions associated with import and export policies. See the “OSPF Routing Policies” section in the OSPF chapter of your product’s *Implementation Guide* for more information.
- Your system has one unified IP routing table. Routing policies allow you to control the flow of information among the network, the protocols, and the routing tables on your system.
- There are two classes of routing policies:
 - **Import policies** — Control which OSPF non-self-originated external routes are stored in the routing table. OSPF import policies control only what the local router uses. They do not affect the propagation of non-self-originated external routes to other routers.
 - **Export policies** — Used on OSPF boundary routers to control which self-originated external routing updates are placed in the link-state database for propagation over the network. In this way, export policies govern what other routers learn with regard to the local boundary router’s self-originated information.
- You can set up an IP RIP or OSPF import or export policy to accept or advertise the default route, as long as the default route exists in the routing table. When you define a policy, you are always prompted for the route subnet mask after the route address, even though you specify the wildcard route address of 0.0.0.0.

Specify a route subnet mask as follows:

- If you want the wildcard subnet mask for all routes, use the default subnet mask (0.0.0.0).
- If you want the default route (not all routes), use 255.255.255.255.
- For more information about IP routing policies, see the *Implementation Guide* for your system.

Options

Prompt	Description	Possible Values	[Default]
Policy type	Type of policy	<ul style="list-style-type: none"> ■ import ■ export 	import
Origin protocols	For export policies only. Defines from which protocol the route originated	<ul style="list-style-type: none"> ■ direct ■ sta (static) ■ rip 	sta, rip
Source address	Source router from which the route was learned. Not applicable to the following: <ul style="list-style-type: none"> ■ Import policies ■ Export policies that define <code>direct</code> as the Origin Protocol 	Any valid IP address	0.0.0.0 (all)
Route address	Route IP address. Not applicable to export policies that define <code>direct</code> as the Origin Protocol.	Any valid IP address	0.0.0.0 (all)
Route subnet mask	Subnet mask for the route (for example, 255.255.0.0). Not applicable to export policies that define <code>direct</code> as the Origin Protocol.	Any valid subnet mask	0.0.0.0 (all)
IP interfaces	Index number of the interface for which you want to define a routing policy. Only applicable when specifying <code>direct</code> as the origin protocol when defining an export policy.	<ul style="list-style-type: none"> ■ Valid interface index ■ all ■ ? (for a list of selectable indexes) 	all (factory default), or current value
Policy action	Accept or reject the route	<ul style="list-style-type: none"> ■ accept ■ reject 	accept
Metric adjustment	For <code>accept</code> conditions only, increases or decreases the converted route metric by the specified value. Options: <ul style="list-style-type: none"> + (add) - (subtract) * (multiply metric by value) / (divide metric by value) % (modulo, remainder of division operation as integer) 	0 – 65535 with or without options	0, which does not change the metric

Prompt	Description	Possible Values	[Default]
ASE type	Type of external metric that is used in the AS external advertisement (ASE), defined as: <ul style="list-style-type: none"> ■ Type 1 — External metric is directly comparable (without translation) to the link state metric. ■ Type 2 — External metric is larger than any link state path. 	<ul style="list-style-type: none"> ■ Type 1 ■ Type 2 	1
Administrative weight	Metric value for this policy. (Higher values have higher priority.)	1 – 16	1

OSPF Import Policy Conditions

Route (address/mask)	Action	Description
Specified route/mask	accept	Add specified non-self-originated external route with or without metric adjustments (+, -, *, /, %) to the routing table.
all (0.0.0.0)	accept	Add all non-self-originated external routes with or without metric adjustments (+, -, *, /, %) to the routing table.
Specified route/mask	reject	Do not add specified non-self-originated external route to the routing table.
all	reject	Do not add any external routes to the routing table; reject all non-self-originated external routes.

OSPF Export Policy Conditions

Protocol	Source Router	Route	Action	Description
RIP or static	Specified router or all routers	Specified route/mask	accept	Advertise in external LSAs specified RIP/static route from specified router with or without metric adjustments (+, -, *, /, %).
RIP or static	Specified router or all routers	all (0.0.0.0)	accept	Advertise in external LSAs all RIP/static routes from specified router with or without metric adjustments (+, -, *, /, %).
RIP or static	Specified router or all routers	Specified route/mask	reject	Do not advertise in external LSAs RIP/static routes from specified routers.
RIP or static	Specified router or all routers	all (0.0.0.0)	reject	Do not advertise in external LSAs any RIP/static route from specified routers.

Export Policy Conditions for Direct Routes

Protocol	Interface	Action	Description
Direct	Specified non-OSPF interface or All non-OSPF interfaces	accept	Advertise in external LSAs all direct routes off of specified interfaces.
Direct	Specified non-OSPF interface or All non-OSPF interfaces	reject	Do not specify in external LSAs any direct routes off of specified interfaces.

Example of Import Policy

```
Select menu option (ip/ospf/policy): define
Enter policy type (import,export) [import]: import
Enter route address [0.0.0.0]: 204.201.89.9
Enter route subnet mask [255.255.255.0]:
Enter policy action (accept,reject) [accept]: accept
Enter metric adjustment ([+,-,*,/,%,0-65535) [0]:
Enter administrative weight (1-16) [1]: 2
```

Example of Export Policy

```
Select menu option (ip/ospf/policy): define
Enter policy type (import,export) [import]: export
Enter origin protocols (dir,sta,rip|all|?) [dir,sta,rip]: sta
Enter source address [0.0.0.0]: 204.243.30.4
Enter route address [0.0.0.0]: 22.32.4.2
Enter route subnet mask [255.0.0.0]:
Enter policy action (accept,reject) [accept]: accept
Enter metric adjustment ([+,-,*,/,%,0-65535) [0]:
Enter ASE type (type1,type2) [type1]: 2
Enter administrative weight (1-16) [1]: 3
```

Example of Export Policy for a Directly Connected Interface

```
Select menu option (ip/ospf/policy): define
Enter policy type (import,export) [import]: export
Enter origin protocols (dir,sta,rip|all|?) [dir,sta,rip]: dir
Select IP interfaces (1|all|?) [1]:
Enter policy action (accept,reject) [accept]: accept
Enter metric adjustment ([+,-,*,/,%,0-65535) [0]: 3
Enter ASE type (type1,type2) [type1]: 2
Enter administrative weight (1-16) [1]: 4
```

ip ospf policy modify Modifies an existing OSPF routing policy.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ip o po m

Important Considerations

- The system assigns an index number to each policy and takes into account all route policies, Routing Information Protocol (RIP) and OSPF, that are set on the system.
- There are certain conditions associated with import and export policies. See the *Implementation Guide* for your system for more information.
- Your system has one unified IP routing table. Routing policies allow you to control the flow of information among the network, the protocols, and the routing tables on your system.
- There are two classes of routing policies:
 - **Import policies** — Control which OSPF non-self-originated external routes are stored in the routing table. OSPF import policies control only what the local router uses. They do not affect the propagation of non-self-originated external routes to other routers.
 - **Export policies** — Used on OSPF boundary routers to control which self-originated external routing updates are placed in the link-state database for propagation over the network. In this way, export policies govern what other routers learn with regard to the local boundary router's self-originated information.
- You can set up an IP RIP or OSPF import or export policy to accept or advertise the default route, as long as the default route exists in the routing table. When you define a policy, you are always prompted for the route subnet mask after the route address, even though you specify the wildcard route address of 0.0.0.0.

Specify a route subnet mask as follows:

- If you want the wildcard subnet mask for all routes, use the default subnet mask (0.0.0.0).
- If you want the default route (not all routes), enter 255.255.255.255.
- For more information about IP routing policies, see the *Implementation Guide* for your system.

Options

Prompt	Description	Possible Values	[Default]
Policy	Index number of the policy that you want to modify	<ul style="list-style-type: none"> Valid policy index number ? (for a list of selectable indexes) 	–
Origin protocols	For export policies only. Defines from which protocol the route originated	<ul style="list-style-type: none"> direct sta (static) rip 	Current value
Source address	Source router from which the route was learned. Not applicable to the following: <ul style="list-style-type: none"> Import policies Export policies that define <code>direct</code> as the Origin Protocol 	Any valid IP address	Current value
Route address	Route IP address. Not applicable to export policies that define <code>direct</code> as the Origin Protocol.	Any valid IP address	Current value
Route subnet mask	Subnet mask for the route (for example, 255.255.0.0). Not applicable to export policies that define <code>direct</code> as the Origin Protocol.	Any valid mask	Current value
IP interfaces	Index number of the interface for which you want to define a routing policy. Only applicable when you specify <code>direct</code> as the origin protocol when defining an export policy.	<ul style="list-style-type: none"> Valid IP interface index all ? (for a list of selectable indexes) 	Current value
Policy action	Accept or reject the route.	<ul style="list-style-type: none"> accept reject 	Current value

Prompt	Description	Possible Values	[Default]
Metric adjustment	For accept conditions only, increases or decreases the converted route metric by the specified value. Options: + (add) - (subtract) * (multiply metric by value) / (divide metric by value) % (modulo, remainder of division operation as integer)	0 – 16, with or without options	Current value
ASE type	Type of external metric used in the AS external advertisement (ASE), defined as: <ul style="list-style-type: none"> ■ Type 1 — External metric is directly comparable (without translation) to the link state metric. ■ Type 2 — External metric is larger than any link state path. 	<ul style="list-style-type: none"> ■ Type 1 ■ Type 2 	Current value
Administrative weight	Metric value for this policy. (Higher values have higher priority.)	1 – 16	Current value

OSPF Import Policy Conditions

Route (address/mask)	Action	Description
Specified route/mask	accept	Add specified non-self-originated external route with or without metric adjustments (+, -, *, /, %) to the routing table.
All (0.0.0.0)	accept	Add all non-self-originated external routes with or without metric adjustments (+, -, *, /, %) to the routing table.
Specified route/mask	reject	Do not add specified non-self-originated external route to the routing table.
All	reject	Do not add any external routes to the routing table; reject all non-self-originated external routes.

OSPF Export Policy Conditions

Protocol	Source Router	Route	Action	Description
RIP or static	Specified router or all routers	Specified route/mask	accept	Advertise in external LSAs specified RIP/static route from specified router with or without metric adjustments (+, -, *, /, %).
RIP or static	Specified router or all routers	all (0.0.0.0)	accept	Advertise in external LSAs all RIP/static routes from specified router with or without metric adjustments (+, -, *, /, %).
RIP or static	Specified router or all routers	Specified route/mask	reject	Do not advertise in external LSAs RIP/static routes from specified router(s).
RIP or static	Specified router or all routers	all (0.0.0.0)	reject	Do not advertise in external LSAs any RIP/static route from specified router(s).

Export Policy Conditions for Direct Routes

Protocol	Interface	Action	Description
Direct	Specified non-OSPF interface or All non-OSPF interfaces	accept	Advertise in external LSAs all direct routes off of specified interfaces.
Direct	Specified non-OSPF interface or All non-OSPF interfaces	reject	Do not specify in external LSAs any direct routes off of specified interfaces.

IP OSPF Policy Modify Example

```
Select menu option (ip/ospf/policy): modify
Select policy {1|?} [1]:
Enter origin protocols (dir,sta,rip|all|?) [rip]:
Enter source address [0.0.0.0]:
Enter route address [0.0.0.0]:
Enter policy action (accept,reject) [accept]:
Enter metric adjustment ([+,-,*,/,%]0-65535) [0]:
Enter administrative weight (1-16) [1]:
Enter ASE type (type1,type2) [type1]:
```

9300

- The system assigns an index number to each policy that you define. This index number takes into account all route policies that are set on the system, Routing Information Protocol (RIP) and OSPF, so the assigned index may be higher than you expect.
- When you remove a policy, the associated index number is available for future use.

Prompt	Description	Possible Values	[Default]
Policy index	Index number of the policy that you want to delete	<ul style="list-style-type: none"> Valid policy index number all ? (for a list of selectable indexes) 	–

ip ospf statistics

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays general OSPF statistics.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ip o sta

3900

9300

Fields in the IP OSPF Statistics Display

Field	Description
extLsaChanges	Number of external LSA changes that have been made to the database
LSAsReceived	Number of link state advertisements that have been received
LSAsTransmitted	Number of link state advertisements that have been transmitted
memoryFailures	Number of nonfatal memory-allocation failures
recvErrors	Number of general receive errors
routeUpdateErrors	Number of nonfatal routing table update failures
softRestarts	Number of OSPF router soft restarts due to insufficient memory resources (implies a fatal memory-allocation failure). To fix this problem, use ip ospf partition modify to change the OSPF memory partition, add memory, or reconfigure the network topology to generate smaller OSPF databases.
SPFComputations	Number of shortest-path-first computations that have been made

This chapter provides guidelines and other key information about how to use the Internet Packet eXchange (IPX) protocol routing commands to route packets from your system to an external destination.

The IPX protocol is a NetWare LAN communications protocol that moves data between servers and workstation programs running on various network nodes. IPX is a User Datagram Protocol (UDP) that is used for connectionless communications. IPX packets are encapsulated and carried by Ethernet packet and Token Ring frames.

To route packets using the IPX protocol, you:

- 1 Define an IPX routing interface
- 2 Decide which IPX routing and server options you want to use
- 3 Enable IPX forwarding.

An IPX routing interface defines the relationship between an IPX virtual LAN (VLAN) and the subnetworks in the IPX network. Each routing IPX VLAN interface is associated with a VLAN that supports IPX. The system has one interface defined for each subnetwork that is directly connected to it. You must first define a VLAN, as described in Chapter 14, before you define an associated IPX VLAN interface.




For more information about IPX, see the Implementation Guide for your system.



For the CoreBuilder® 9000, the commands in this chapter apply to Layer 3 switching modules only.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



ipx interface display

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

3900

9300

Displays information about the IPX parameters and IPX interfaces that are configured on the system.

Valid Minimum Abbreviation

ipx i di

Important Considerations

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.

Fields in the IPX Interface Display

Field	Description
Format	Frame encapsulation format.
Index	System-assigned index number for the interface.
IPX address	Unique 4-byte network address.
State	Status of the IPX interface. It indicates whether the interface is available for communications (up) or unavailable (down).
Ticks	Number that the system uses to calculate route time. (A tick is an estimate of how long a packet takes to reach the network segment.) There are 18.21 ticks in a second. The possible values are 1 – 65534 and are defined as: <ul style="list-style-type: none">■ 1 = FDDI■ 4 = Ethernet■ 10+ = Serial Links
VLAN index	Index number of the VLAN that is associated with the IPX interface. When the system prompts you for this option, the menu identifies the available VLAN indexes.

ipx interface define

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Defines an IPX interface.

Valid Minimum Abbreviation

`ipx i de`

Important Considerations

- An IPX interface defines the relationships among an IPX virtual LAN (VLAN), the IPX router, and the IPX network. The IPX router has one IPX interface defined for each network than is directly connected to it.
- When you define an interface, you define the interface’s IPX address, cost, format, and any associated IPX VLAN index.
- Before you define the IPX (routing) interface, you must specify a VLAN and select IPX, IPX-II, IPX-802.2, IPX 802.2 LLC, IPX-802.3, or IPX-802.2-SNAP as a protocol that the VLAN supports, as described in Chapter 14. (For routing, a VLAN can now support multiple protocols.)
- Unless your network has special requirements such as the need for redundant paths, assign a cost of 1 to each interface.
- The two Fiber Distributed Data Interface (FDDI) encapsulation formats correspond to the Ethernet 802.2 LLC and 802.3 SNAP encapsulation formats. If you select either of these Ethernet encapsulation formats, the corresponding FDDI encapsulation format is automatically selected for shared Ethernet and FDDI ports.

Options

Prompt	Description	Possible Values	[Default]
IPX network address	4-byte IPX address of the interface. The address must be unique within the network.	0x1 – 0xffffffffe	–
Ticks	Number that the system uses to calculate route time. (A tick is an estimate of how long a packet takes to reach the network segment.) There are 18.21 ticks in a second.	1 – 65534	1

Prompt	Description	Possible Values	[Default]
Frame format	Frame encapsulation format for the interface. IPX uses four Ethernet and two FDDI formats: Ethernet Type II, Novell 802.3 RAW, 802.2 LLC, and 802.3 SNAP. The FDDI formats are available with 802.2 and SNAP.	<div><div>■ Ethernet_II</div><div>■ 802.2</div><div>■ 802.2 LLC</div><div>■ RAW_802.3</div><div>■ SNAP</div><div>■ 802.3_SNAP</div></div>	–
VLAN Interface Index	Index number of the VLAN to associate with the IPX interface.	<div><div>■ A selectable VLAN interface</div><div>■ ? (to view a list of selectable indexes)</div></div>	–

IPX Interface Define Example

```
Select menu option: ipx interface define
Enter IPX Address (0x1-0xfffffffffe): 0x45468f30
Enter Ticks (1-65534) [1]:1
Enter Frame Format (Ethernet_II,802.2,Raw_802.3,SNAP): 802.2
Enter VLAN interface index {4|?} [4]: 4
```

✓ 3500

✓ 9000

9400

3900

9300

ipx interface modify

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Changes the characteristics of an existing IPX interface.

Valid Minimum Abbreviation

`ipx i m`

Important Considerations

- An IPX interface defines the relationships among an IPX virtual LAN (VLAN), the IPX router, and the IPX network. The IPX router has one IPX interface defined for each network that is directly connected to it.
- When you modify an interface, you can change the interface’s IPX address, ticks, format, and the associated IPX VLAN index.
- Unless your network has special requirements (for example, a need for redundant paths), do not change the cost value of 1 that is assigned by default to each interface.

Options

Prompt	Description	Possible Values	[Default]
Index	Number associated with the interface that you want to modify.	<ul style="list-style-type: none">■ One or more selectable IPX interfaces■ ? (to view a list of selectable interfaces)	1 (if only 1 interface)
IPX network address	4-byte IPX address of the interface. The address must be unique within the network.	0x1 – 0xfffffffffe	Current address
Ticks	Number that the system uses to calculate route ticks. (A tick is an estimate of how long a packet takes to reach the network segment.) There are 18.21 ticks in a second.	1 – 65534 where: <ul style="list-style-type: none">■ 1 = FDDI■ 4 = Ethernet■ 10+ = Serial Link	Current setting

Prompt	Description	Possible Values	[Default]
Frame format	Frame encapsulation format for the interface. IPX uses four Ethernet and two FDDI formats: Ethernet Type II, Novell 802.3 RAW, 802.2 LLC, and 802.3 SNAP. The FDDI formats are available with 802.2, SNAP, and 802.3/SNAP.	<ul style="list-style-type: none">■ Ethernet_II■ 802.2■ 802.2 LLC■ RAW_802.3■ SNAP■ 802.3_SNAP	Current format
VLAN interface index	Index number of the VLAN that is associated with the IPX interface.	<ul style="list-style-type: none">■ A selectable VLAN interface■ ? (to view a list of selectable indexes)	Current VLAN index

ipx interface remove

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

Removes an IPX interface if you no longer perform routing on the ports that are associated with the interface.

✓ 9000

Valid Minimum Abbreviation

9400

ipx i r

3900

Options

Prompt	Description	Possible Values	[Default]
Index	Index number for the interface that you want to remove	<div><div>■</div>One or more selectable IPX interface indexes</div> <div><div>■</div>? (to view a list of selectable indexes)</div>	1 (if only 1 interface)

ipx interface
SAPadvertising

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Controls whether the system advertises IPX services.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ipx i s

Options

Prompt	Description	Possible Values	[Default]
IPX SAP advertising state	Whether the system advertises IPX services	<div><div>■ enable</div><div>■ disable</div></div>	disable

**ipx interface
RIPadvertising**

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Controls whether the system advertises IPX routes.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

`ipx i r`

Options

Prompt	Description	Possible Values	[Default]
IPX RIP advertising state	Whether the system advertises IPX services	<ul style="list-style-type: none">■ enable■ disable	disable

ipx route display

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500
✓ 9000
9400

3900
9300

Displays the routing tables for the system. The routing tables include all configured routes.

Valid Minimum Abbreviation

ipx ro d

Important Considerations

- Your system maintains a table of routes to other IPX networks. You can:
 - Use the Routing Information Protocol (RIP) to exchange routing information automatically.
 - Make static entries in this table using the Administration Console.
- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.
- For a CoreBuilder 3500 system, the route table display shows the range for the routing table primary entries in the format $n - m$, where n is the current number of entries and m is the maximum number of primary entries.
- The maximum number of hops, or routers, that a packet can cross, is 16 (except NetBIOS packets, which can cross no more than 7 routers).

Options (3500 only)

Prompt	Description	Possible Values	[Default]
Start of address range	First address in a range for which you want to display routes	0x0 – 0xffffffff	0x0
End of address range	Last address in a range for which you want to display routes	0x0 – 0xffffffff	0xffffffff

Fields in the IPX Route Display

Field	Description
Address	Unique 4-byte network address of a segment in the system's routing table.
Age	Number of seconds that have elapsed since the last time the router sent a packet.
Hops	Number of hops, or the number of routers that must be crossed to reach the network segment.
Interface	System-assigned number for the interface.
Node	6-byte MAC address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.
Ticks	Number of ticks, which is an estimate of time in seconds, that the packet takes to reach the network segment. There are 18.21 ticks in a second.

ipx route secondary Displays any secondary routes that are available.

- ✓ 3500
- ✓ 9000
- 9400

- 3900
- 9300

Valid Minimum Abbreviation

ipx ro se

Important Considerations

- To see entries for any secondary routes, you must:
 - Establish alternate paths to the same IPX network.
 - Enable the IPX secondary route/server option. See “ipx secondary” at the end of this chapter.
- A secondary route entry can replace a primary route entry when the primary route is removed from the routing table for any reason (for example, if the route reaches its age limit).
- For a CoreBuilder 3500 system, the route table display shows the range for the routing table primary entries in the format $n - m$, where n is the current number of entries and m is the maximum number of primary entries.

Fields in the IPX Secondary Route Display

Field	Description
Address	Unique 4-byte network address of a segment in the system’s routing table.
Age	Number of seconds that have elapsed since the last time the router sent a packet.
Hops	Number of hops, or the number of routers that must be crossed to reach the network segment.
Interface	System-assigned number for the interface.
Node	6-byte MAC address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.
Ticks	Number of ticks, which is an estimate of time in seconds, that the packet takes to reach the network segment. There are 18.21 ticks in a second.

ipx route static

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

Defines a static route.

Valid Minimum Abbreviation

ipx ro st

Important Considerations

- Before you define static routes on the system, define at least one IPX interface. See “ipx interface define” earlier in this chapter for more details.
- Static routes remain in the routing table until you remove them or until you remove the corresponding interface.
- If an interface goes down, routes are temporarily removed from the routing table until the interface comes back up.
- Static routes take precedence over dynamically learned routes to the same destination. You can have a maximum of 32 static routes.

Options

Prompt	Description	Possible Values	[Default]
IPX network address	4-byte IPX address of the interface. The address must be unique within the network.	0x1 – 0xfffffffffe	–
Hops	Number of hops, or number of routers that must be crossed to reach the network segment.	1 – 15	1
Interface number	Interface number to associate with the route. Depends on number of configured IPX interfaces.	<ul style="list-style-type: none"> ■ A selectable IPX interface number ■ ? (for a list of selectable IPX interfaces) 	–
Node address	6-byte MAC address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	A node address in the format xx-xx-xx-xx-xx-xx	–

IPX Static Route Example

Select menu option: **ip route static**

Enter IPX address (0x1-0xfffffffffe): **0x44648f30**

Enter Hops (1-15): **1**

Enter interface number (1-32) [1]: **1**

Enter node address: **08-00-3e-21-14-78**

ipx route remove

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Deletes a route from the IPX routing table.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

ipx ro r

Important Considerations

- The route is immediately deleted. You are not prompted to confirm the deletion.
- All servers that depend upon this route are removed from the server table, including static servers.

Options

Prompt	Description	Possible Values	[Default]
IPX network address	4-byte IPX network address	0x1 – 0xfffffffffe	–

ipx route flush ***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Deletes all dynamically learned routes from the IPX routing table.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

`ipx ro f`

3900

9300

Important Considerations

- All learned routes are immediately deleted. You are not prompted to confirm the deletion.
- All dynamic servers that depend on these routes are removed from the server table.

ipx server display***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays the server table for the system to determine which servers are learned.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ipx ser d`**Important Considerations**

- Your system maintains a table of servers that reside on other IPX networks. You can:
 - Use the Service Advertising Protocol (SAP) to exchange server information automatically.
 - Make static entries in this server table.
- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.
- For a CoreBuilder 3500 system, the route table display shows the range for the routing table primary entries in the format $n - m$, where n is the current number of entries and m is the maximum number of primary entries.

Options (3500 only)

Prompt	Description	Possible Values	[Default]
Service type	<p>Number for the type of service that the server performs.</p> <p>Enter up to 6 hex characters. For example, 0x4 = file server</p> <p>For more details, consult your Novell documentation.</p> <p>Use quotation marks (") around any string with embedded spaces.</p> <p>Use double quotes (") to enter an empty string.</p>	<ul style="list-style-type: none"> ■ * ■ 0x1 –0xfffff 	*
Service name pattern	<p>Pattern for the service name.</p> <p>Use quotation marks (") around any string with embedded spaces.</p> <p>Use double quotes (") to enter an empty string.</p>	<ul style="list-style-type: none"> ■ * ■ Up to 79 alphanumeric characters 	*

Fields in the IPX Server Display

Field	Description
Age	Number of seconds that have elapsed since the last time a server in the table sent a packet.
Hops	Number of networks that must be crossed to reach the server. The maximum number is 15.
Interface	Index number of the interface.
Name	Name for the server that you define.
Network	4-byte IPX network address of the server.
Node	6-byte MAC address of the server that forwards packets to the segment.
Socket	2-byte socket address of the server that receives service requests.
Type	Type of service that the server provides. The IPX protocol defines various types of services. One common type is 0x4 , which is for a file server. For more information on IPX type values, consult your Novell documentation.

ipx server static***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Defines a static IPX server.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation**ipx ser st****Important Considerations**

- Static servers remain in the table until you remove them, until you remove the corresponding interface, or until you remove the route to the corresponding network address.
- A static server must have an IPX network address that corresponds to a configured interface or to a static route. If an interface goes down, any static servers on that interface are permanently removed from the server table until the interface comes back up.
- Static servers take precedence over dynamically learned servers to the same destination. You can have a maximum of 32 static servers.
- Before you define static servers on the system, first define at least one IPX interface. See “ipx interface define” earlier in this chapter for more details.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Interface index number for the server	<ul style="list-style-type: none"> ■ A selectable IPX interface index ■ ? (for a list of selectable IPX interfaces) 	–
Service type	Number for the type of service that the server performs	<ul style="list-style-type: none"> ■ * ■ 0x1 – 0xffff 	*
Service name	Service name of the server, up to 79 characters	<ul style="list-style-type: none"> ■ Any selectable service name ■ ? (for a list of selectable names) 	–
IPX network address	IPX network address of the server	0x0 – 0xfffffffffe	–
Socket value	Socket value of the server	0x0 – 0xffff	–
Node address	Node address of the server		–
Hops	Number of hops to the server	0 – 15	–

IPX Static Server Example

```
Enter Interface index {1|?} [1]: 1
Enter service type {0x1-0xFFFF}: 0x4
Enter service name {?}: gb201
Enter IPX address (0x0-0xffffffff): 0x8c14a228
Enter socket (0x0-0xffff): 0x8059
Enter node address : 00-00-2e-f3-56-02
Enter hops (0-15): 2
```

ipx server remove

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Deletes a server from the IPX server table.

Valid Minimum Abbreviation

ipx ser r

Important Consideration

■ The server is immediately deleted. You are not prompted to confirm the deletion.

Options

Prompt	Description	Possible Values	[Default]
Service name	Service name of the server	<div>■ A selectable service name</div> <div>■ ? (for a list of selectable names)</div>	
Service type	Number for the type of service that the server performs.	<div>■ *</div> <div>■ 0x1 – 0xffff</div>	*

ipx server flush

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Deletes all dynamically learned servers from the server table.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ipx ser f

3900

9300

Important Consideration

- All learned servers are immediately deleted. You are not prompted to confirm the deletion.

ipx server secondary

Displays any secondary servers that are available.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`ipx ser se`

Important Considerations

- To see entries for any secondary server, you must:
 - Establish alternate paths to the same IPX server.
 - Enable the IPX secondary route/server option. See “ipx secondary” at the end of the chapter.
- A secondary server entry can replace a primary server entry when the primary server is removed from the server table for any reason (for example, if the associated interface goes down, or the primary entry reaches its age limit).
- For a CoreBuilder 3500 system, the route table display shows the range for the routing table primary entries in the format $n - m$, where n is the current number of entries and m is the maximum number of primary entries.

Fields in the IPX Secondary Server Display

Field	Description
Age	Number of seconds that have elapsed since the last time a server in the table sent a packet.
Hops	Number of networks that must be crossed to reach the server. The maximum number is 15.
Interface	Index number of the interface.
Name	Name for the secondary server.
Network	4-byte IPX network address of the server.
Node	6-byte MAC address of the server that forwards packets to the segment.
Socket	2-byte socket address of the server that receives service requests.
Type	Type of service that the server provides. The IPX protocol defines various types of services. One type is 0x4, which is for a file server. For more information on IPX type values, consult your Novell documentation.

ipx forwarding

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Controls whether the system forwards or discards IPX packets.

Valid Minimum Abbreviation

ipx f

Important Considerations

- When you enable IPX forwarding, the system acts as a normal IPX router, forwarding IPX packets from one network to another when required.
- When you disable IPX forwarding, the system discards all IPX packets.

Options

Prompt	Description	Possible Values	[Default]
IPX forwarding state	Whether the system forwards or discards IPX packets	<div><div>■ disabled</div><div>■ enabled</div></div>	disabled (factory default), or current value

ipx rip mode

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Selects the Routing Information Protocol (RIP) mode that is appropriate for your network.

Valid Minimum Abbreviation

ipx ri m

Important Considerations

- RIP allows the exchange of routing information on a NetWare network. IPX routers use RIP to create and maintain their dynamic routing tables.
- The system has three RIP modes:
 - **Off** — The system processes no incoming RIP packets and generates no RIP packets of its own.
 - **Passive** — The system processes all incoming RIP packets and responds to RIP requests, but it does not broadcast periodic or triggered RIP updates.
 - **Active** — The system processes all incoming RIP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered RIP updates.

Options

Prompt	Description	Possible Values	[Default]
RIP mode	Whether the system processes RIP packets	<ul style="list-style-type: none">■ off■ passive■ active	disabled (factory default), or current value

ipx rip triggered

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Sets the RIP Triggered update mode, which dictates when the IPX protocol broadcasts newly learned routes.

Valid Minimum Abbreviation

ipx ri t

Important Considerations

- The system has two RIP triggered modes:
 - **Disabled** — Broadcasts IPX routes 3 seconds after learning them.
 - **Enabled** — Broadcasts IPX routes immediately after learning them.

Options

Prompt	Description	Possible Values	[Default]
Triggered update mode	Mode that determines when IPX broadcasts newly learned routes	<ul style="list-style-type: none">■ disabled■ enabled	enabled

ipx rip policy
summary

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Display a list of IPX RIP (Routing Information Protocol) policies.

Valid Minimum Abbreviation

`ipx ri p s`

Fields in an IPX RIP Policy Summary Display

Field	Description
Idx	Index number of the IPX RIP policy.
Origin	Source of the route to which this policy applies. If the policy type is set to Export, the possible values of this parameter are RIP or Static. This parameter is not applicable if the policy type is set to Import.
Type	Import (apply the policy to received routes) or Export (apply the policy to advertised routes).
Route	One or more IPX network addresses where this policy applies.
Interface	One or more IP interfaces on this router associated with the RIP policy.
Source	6-byte MAC address of the router that can forward packets to the network. A source node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.
Action	Whether this router accepts or rejects a route that matches the policy.
Metric	Value the system uses to increase or decrease a route metric. (This parameter is valid only if the Policy Action is set to Accept.)
Weight	Metric value of this policy.

ipx rip policy define

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

3900

9300

Define a RIP (Routing Information Protocol) policy.

Valid Minimum Abbreviation

ipx ri p d

Important Considerations

- Every router maintains a table of current routing information in a routing table.
- Routing protocols receive or advertise routes from the network.
- Routing Policies control the flow of routing information between the network, the protocols, and the routing table manager.

Prompt	Description	Possible Values	[Default]
Type	Type of the policy: Import (apply the policy to received routes) or Export (apply the policy to advertised routes).	<ul style="list-style-type: none">■ Import■ Export	Import
Route Origin	Origin of the route to which this policy applies. This parameter is valid only if the policy Type is set to Export.	<ul style="list-style-type: none">■ Dir■ Static■ RIP■ All	All
Route address	Route to which this policy applies.	<ul style="list-style-type: none">■ 0x1-0fffffffe■ All	All
IP interfaces	One or more IP interfaces on this router associated with the RIP policy.	One or more IP interface numbers	All
Source node address	6-byte MAC address of the router that can forward packets to the network. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	<ul style="list-style-type: none">■ A node address in the format xx-xx-xx-xx-xx-xx■ All	All
Policy action	Whether this router accepts or rejects a route that matches the policy.	<ul style="list-style-type: none">■ Accept■ Reject	Accept

Prompt	Description	Possible Values	[Default]
Metric adjustment	Increase or decrease a route metric by a value that you specify. Specify an integer and an operand (+,-,*,/,%) to adjust the metric value. This parameter is valid only if the Policy Action is set to Accept.	<ul style="list-style-type: none"> ■ 0-16 ■ + (add) ■ - (subtract) ■ * (multiply) ■ / (divide) ■ % (modulo - remainder of integer division) 	0 (does not change the metric)
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same route. A higher value takes precedence over a lower value.	1 – 16	1

IPX RIP Policy Define Example

```

Select menu option (ipx/rip/policy): define
Enter policy type (import,export) [import]:export
Enter route origin (dir,static,rip,all) [all]:rip
Enter route address (0x1-0xfffffffffe|all) [all]:all
Select IP interfaces (2|all?) [all]:
Enter the source node address [all]:
Enter the policy action (accept, reject) [accept]: accept
Enter the metric adjustment ([+,-,*,/]|0-16) [0]:
Enter the administrative weight (1-16) [1]:2

```

ipx rip policy modify

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

Modify an existing RIP (Routing Information Protocol) policy.

Valid Minimum Abbreviation

ipx ri p m

3900

9300

Important Considerations

- Every router maintains a table of current routing information in a routing table.
- Routing protocols receive or advertise routes from the network.
- Routing Policies control the flow of routing information between the network, the protocols, and the routing table manager.

Prompt	Description	Possible Values	[Default]
Policy	Index number of the policy you want to modify.	<ul style="list-style-type: none">■ 1■ ? (to view a list of selectable policies)	1 (if only one policy)
Route Origin	Origin of the route to which this policy applies. This parameter is valid only if the policy Type is set to Export.	<ul style="list-style-type: none">■ Static■ RIP■ All	All
Route address	IPX route to which this policy applies.	<ul style="list-style-type: none">■ 0x1-0xffffffffe■ All	All
IP interfaces	One or more IP interfaces on this router associated with the RIP policy.	One or more IP interface numbers	All
Source node address	6-byte MAC address of the router that can forward packets to the segment. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	<ul style="list-style-type: none">■ A node address in the format xx-xx-xx-xx-xx-xx■ All	All
Policy action	Whether this router accepts or rejects a route that matches the policy.	<ul style="list-style-type: none">■ Accept■ Reject	Accept

Prompt	Description	Possible Values	[Default]
Metric adjustment	Increase or decrease a route metric by a value that you specify. Specify an integer and an operand (+,-,*,/,%) to adjust the metric value. This parameter is valid only if the Policy Action is set to Accept.	<ul style="list-style-type: none"> ■ 0-16 ■ + (add) ■ - (subtract)\ ■ * (multiply) ■ / (divide) ■ % (modulo - remainder of integer division) 	0 (does not change the metric)
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same route. A higher value takes precedence over a lower value.	1 – 16	1

IPX RIP Policy Modify Example

```

Select menu option (ipx/rip/policy): modify
Select policy {1|?}: 1
Enter route origin (static,rip,all) [all]: rip
Enter route address (0x1-0xfffffffffe|all) [all]:
Select IP interfaces (2|all?) [all]:
Enter the source node address [all]:
Enter the policy action (accept, reject) [accept]:
Enter the metric adjustment ([+,-,*,/]|0-16) [0]:
Enter the administrative weight (1-16) [1]:

```

ipx rip policy remove

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

Remove an existing RIP (Routing Information Protocol) policy.

✓ 9000

Valid Minimum Abbreviation

9400

ipx ri p r

3900

Options

9300

Prompt	Description	Possible Values	Default
Policy	Index number of the policy you want to remove	<div><div>■</div> 1</div> <div><div>■</div> ? (to view a list of selectable policies)</div>	1 (if only one policy)

ipx sap mode *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Selects a Service Advertising Protocol (SAP) mode that is appropriate for your network.

Valid Minimum Abbreviation

`ipx sa m`

Important Considerations

- SAP provides routers and servers that contain SAP agents with a means of exchanging network service information.
- The system has three SAP modes:
 - **Off** — The system does not process any incoming SAP packets and does not generate any SAP packets of its own.
 - **Passive** — The system processes all incoming SAP packets and responds to SAP requests, but it does not broadcast periodic or triggered SAP updates.
 - **Active** — The system processes all incoming SAP packets, responds to explicit requests for routing information, and broadcasts periodic and triggered SAP updates.

Options

Prompt	Description	Possible Values	[Default]
SAP mode	Whether the system processes SAP packets	<ul style="list-style-type: none">■ off■ passive■ active	disabled (factory default), or current value

ipx sap triggered

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Sets the SAP Triggered Update mode, which dictates when the IPX protocol broadcasts newly learned SAP server addresses.

Valid Minimum Abbreviation

ipx sa t

Important Considerations

- The system has two SAP triggered modes:
 - **Disabled** — Broadcasts IPX SAP server addresses 3 seconds after learning them.
 - **Enabled** — Broadcasts IPX SAP server addresses immediately after learning them.

Options

Prompt	Description	Possible Values	[Default]
Triggered update mode	Setting for IPX SAP broadcast timing	<ul style="list-style-type: none">■ disabled■ enabled	enabled

ipx sap policy
summary

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Display a list of IPX SAP (Service Advertising Protocol) policies.

Valid Minimum Abbreviation

ipx sa p s

Fields in an IPX SAP Policy Summary Display

Field	Description
Idx	Index number of the IPX SAP policy.
Origin	Source of the service to which this policy applies. If the policy type is set to Export, the possible values of this parameter are SAP, Static, or All. This parameter is not applicable if the policy type is set to Import.
Type	Policy type. Import (apply the policy to received services) or Export (apply the policy to advertised services).
Name	Object name that assigned to the server.
Type	Service type, represented by a one-digit number. Refer to Novel documentation for a complete list of service types.
Network	IPX network address for the server, or All, which implies all routes.
Node	6-byte MAC address of the router that can forward packets to the network. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.
Action	Whether this router accepts or rejects a route that matches the policy.

ipx sap policy detail

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Display information about IPX SAP (Service Advertising Protocol) policies.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ipx sap p det

3900
9300

Fields in an IPX SAP Policy Detail Display

Field	Description
Idx	Index number of the IPX SAP policy.
Interface	Index number of the IP interface associated with this policy.
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same service. A higher value takes precedence over a lower value.

ipx sap policy define

- ✓ 3500
- ✓ 9000
- 9400

- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Define a SAP (Service Advertising Protocol) policy.

Valid Minimum Abbreviation

ipx sa p def

Important Considerations

- Every router maintains a table of current configured services in a service table.
- The SAP running on the router receives and advertises services from the network.
- Service policies control the services in the service table and those that the router advertises.
- Novell defines several different service types using specific numbers for the server advertising the service. You enter a Novell service type when you define a SAP policy. Some of the most common service types are:

0x0004	File Server
0x0005	Job Server
0x0007	Print Server
0x0009	Archive Server
0x000A	Job Queue
0x0047	Advertising Print Server
0x0098	NetWare Access Server

For a complete list of Novell service types, consult your Novell documentation.

Options

Prompt	Description	Possible Values	[Default]
Policy Type	Type of the policy: Import (apply the policy to received services) or Export (apply the policy to advertised services).	<ul style="list-style-type: none">■ Import■ Export	Import
Service Origin	Origin of the service to which this policy applies. This parameter is valid only if the policy Type is set to Export.	<ul style="list-style-type: none">■ Static■ SAP■ All	All

Prompt	Description	Possible Values	[Default]
Service Type	Number for the type of service that the server performs. Enter up to 6 hex characters. For example, 0x4 = file server For more details, consult your Novell documentation.	<ul style="list-style-type: none"> ■ 0x1 – 0xfffff ■ All 	All
Server Name	Name of the server providing the services.	<ul style="list-style-type: none"> ■ Server name ■ All 	All
IPX Address	IPX network address of the network where the server resides.	<ul style="list-style-type: none"> ■ 0x0 – 0xfffffffffe ■ All 	All
Node Address	6-byte MAC address of the router that can forward packets to the network. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	<ul style="list-style-type: none"> ■ A node address in the format xx-xx-xx-xx-xx-xx ■ All 	All
Interface Index	Index number of the IP interface associated with this policy.	<ul style="list-style-type: none"> ■ One or more interface numbers ■ All ■ ? (to view a list of selectable interfaces) 	All
Policy action	Whether this router accepts or rejects a service that matches the policy.	<ul style="list-style-type: none"> ■ Accept ■ Reject 	Accept
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same service. A higher value takes precedence over a lower value.	1 – 16	1

IPX SAP Policy Define Example

```
Select menu option (ipx/rip/policy): define
Enter policy type (import,export) [import]:
Enter service origin (static,sap,all) [all]:sap
Enter the service type (0x1-0x1ffff|all) [all]:0x0004
Enter the server name (?) [all]:
Enter the IPX address (0x0-0xffffffffe|all) [all]:
Enter the node address [all]:
Select interface index (2|all?) [all]:
Enter the policy action (accept, reject) [accept]: accept
Enter the administrative weight (1-16) [1]:2
```

ipx sap policy modify

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

3900

9300

Modify a SAP (Service Advertising Protocol) policy.

Valid Minimum Abbreviation

ipx sa p m

Important Considerations

- Every router maintains a table of current configured services in a service table.
- The SAP running on the router receives and advertises services from the network.
- Service policies control the services in the service table and those that the router advertises.
- Novell defines several different service types using specific numbers for the server advertising the service. You can change the Novell service type when you modify a SAP policy. Some of the most common service types are:

0x0004	File Server
0x0005	Job Server
0x0007	Print Server
0x0009	Archive Server
0x000A	Job Queue
0x0047	Advertising Print Server
0x0098	NetWare Access Server

For a complete list of Novell service types, consult your Novell documentation.

Options

Prompt	Description	Possible Values	[Default]
Policy	Index number of the policy you want to modify.	<ul style="list-style-type: none">■ 1■ ? (to view a list of selectable policies)	1 (if only one policy)
Service Origin	Origin of the service to which this policy applies. This parameter is valid only if the policy Type is set to Export.	<ul style="list-style-type: none">■ Static■ SAP■ All	All

Prompt	Description	Possible Values	[Default]
Service Type	Number for the type of service that the server performs. Enter up to 6 hex characters. For example, 0x4 = file server For more details, consult your Novell documentation.	<ul style="list-style-type: none"> ■ 0x1 – 0xfffff ■ All 	All
Server Name	Name of the server providing the services.	<ul style="list-style-type: none"> ■ Server name ■ All 	All
IPX Address	IPX network address of the network where the server resides.	<ul style="list-style-type: none"> ■ 0x0 – 0xfffffffffe ■ All 	All
Node Address	6-byte MAC address of the router that can forward packets to the network. A node address of all zeroes (00-00-00-00-00-00) means that the route is connected directly to the router.	<ul style="list-style-type: none"> ■ A node address in the format xx-xx-xx-xx-xx-x ■ All 	All
Interface Index	Index number of the IP interface associated with this policy.	<ul style="list-style-type: none"> ■ One or more interface numbers ■ All ■ ? (to view a list of selectable interfaces) 	All
Policy action	Whether this router accepts or rejects a service that matches the policy.	<ul style="list-style-type: none"> ■ Accept ■ Reject 	Accept
Weight	Metric value of this policy. This parameter specifies the order of precedence for policies that match the same service. A higher value takes precedence over a lower value.	1 – 16	1

IPX SAP Policy Modify Example

```
Select menu option (ipx/rip/policy): modify
Select policy {1|?}:1
Enter service origin (static,sap,all) [all]:sap
Enter the service type (0x1-0x1ffff|all) [all]:all
Enter the server name (?) [all]:
Enter the IPX address (0x0-0xffffffffe|all) [all]:
Enter the node address [all]:
Select interface index (2|all?) [all]:
Enter the policy action (accept, reject) [accept]: accept
Enter the administrative weight (1-16) [1]:2
```

ipx sap policy remove

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Remove an existing SAP (Service Advertising Protocol) policy.

Valid Minimum Abbreviation

ipx sa p r

Options

Prompt	Description	Possible Values	Default
Policy	Index number of the policy you want to remove	<ul style="list-style-type: none">1? (to view a list of selectable policies)	1 (if only one policy)

ipx output-delay***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***✓ **3500**✓ **9000****9400**

Sets the IPX output-delay option for RIP (Routing Information Protocol) and SAP (Service Advertising Protocol) packets. This option delays the updating of the RIP and SAP server information table.

Valid Minimum Abbreviation**ipx i o****3900****9300****Options**

Prompt	Description	Possible Values	[Default]
Output-delay mode	Whether you want to enable or disable the output-delay option	<ul style="list-style-type: none">■ enable■ disable	disable

ipx statistics
summary

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Displays IPX summary statistics.

Valid Minimum Abbreviation

ipx st su

Important Considerations

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.

Fields in the IPX Statistics Summary Display

Field	Description
Forwarded	Number of IPX packets that were forwarded
Fwd Received	Number of IPX packets that were received to be forwarded
Fwd Transmitted	Number of IPX forwarded packets that were successfully transmitted
Host Delivers	Number of IPX packets that were delivered to the IPX host's RIP and SAP applications
Host Dropped	Number of IPX packets to or from the IPX hosts's RIP and SAP applications that were dropped
Host Tx	Number of IPX packets from the IPX host's RIP and SAP applications that were successfully transmitted

ipx statistics rip**For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays IPX RIP (Routing Information Protocol) statistics.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ipx st r`**Important Considerations**

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.

Fields in the IPX RIP Statistics Display

Field	Description
RIP Dropped	Number of IPX RIP packets that have been dropped
RIP Entries	Number of routes in the routing table (including local routes)
Routes Aged	Number of times the system marked a route entry unreachable, because it did not receive an update for that entry during the timeout period
RIP Received	Number of IPX RIP packets that have been received
RIP Requests	Number of IPX RIP requests that have been processed
RIP Responses	Number of IPX RIP responses that have been processed
RIP Transmitted	Number of IPX RIP packets that have been transmitted
Metric Changed	Number of times the metric changed on a route entry

ipx statistics sap***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays IPX SAP (Service Advertising Protocol) statistics.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation`ipx st sa`**Important Considerations**

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered updates are enabled.
 - Secondary route/server option is enabled.

Fields in the IPX SAP Statistics Display

Field	Description
SAP Dropped	Number of IPX SAP packets that have been dropped
SAP Entries	Number of servers in the server table
Servers Aged	Number of times the system marked a server entry unreachable, because it did not receive an update for that entry during the timeout period
SAP GNS Requests	Number of IPX SAP Get Nearest Service Requests that have been processed
SAP GNS Responses	Number of IPX SAP Get Nearest Service Responses that have been received
SAP Received	Number of IPX SAP packets that have been received
SAP Requests	Number of IPX SAP Requests that have been processed
SAP Responses	Number of IPX SAP Responses that have been processed
SAP Transmitted	Number of IPX SAP packets that have been transmitted
Metric Changed	Number of times the metric changed on a server entry

ipx statistics forwarding

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays IPX forwarding statistics.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ipx st f

Important Considerations

- The first line in the output (the status line) indicates whether:
 - IPX forwarding is enabled.
 - RIP is active.
 - SAP is active.
 - RIP Triggered updates are enabled.
 - SAP Triggered Updates are enabled.
 - Secondary route/server option is enabled.

Fields in the IPX Forwarding Statistics Display

Field	Description
Addr Errors	Number of IPX packets that were dropped that due to IPX address errors in network layer header
Forwarded	Number of IPX packets that were forwarded
Fwd Discards	Number of IPX packets to be forwarded that could not be forwarded
Fwd Received	Number of IPX packets that were received to be forwarded
Fwd Transmitted	Number of IPX forwarded packets that were successfully transmitted
Hdr Errors	Number of IPX packets that were dropped due to IPX Network layer header errors
Hop Count Errors	Number of IPX packets that were dropped due to exceeded maximum transport control
Host Delivers	Number of IPX packets that were delivered to the IPX host's RIP and SAP applications
Host In Discards	Number of IPX packets that were received for the IPX host's RIP and SAP applications that were dropped
Host Rx	Number of IPX packets that were delivered to the IPX host's RIP and SAP applications
Host Tx	Number of IPX packets that were transmitted from the IPX host's RIP and SAP applications

Field	Description
Host Tx Discards	Number of IPX packets from the IPX host's RIP and SAP applications that were dropped on transmission
Host Tx Request	Number of IPX packets from the IPX host's RIP and SAP applications to be transmitted
NetBIOS Max Hops	Number of IPX NetBIOS packets that exceeded the transport control maximum
NetBIOS Rx	Number of IPX NetBIOS packets that were received
NetBIOS Tx	Number of IPX NetBIOS packets that were transmitted
No Routes	Number of IPX packets that were dropped because the IPX route is unknown
Total Received	Number of IPX packets that were received
Tx Discards	Number of IPX packets that were forwarded but not successfully transmitted
Tx MTU Exceeded	Number of IPX packets that were forwarded but dropped because the MTU was exceeded

ipx statistics interface

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

Displays IPX interface statistics.

✓ 9000

Valid Minimum Abbreviation

9400

ipx st i

3900

Fields in the IPX Interface Statistics Display

9300

Field	Description
Addr Errors	Number of IPX packets that were dropped due to IPX address errors in the network layer header
Forwarded	Number of IPX packets that were forwarded
Fwd Discards	Number of IPX packets to be forwarded that were not
Fwd Received	Number of IPX packets that were received to be forwarded
Fwd Transmitted	Number of IPX forwarded packets that were successfully transmitted
Hdr Errors	Number of IPX packets that were dropped due to IPX Network layer header errors
Hop Count Errors	Number of IPX packets that were dropped due to exceeded maximum transport control
Host In Discards	Number of IPX packets that were received for the IPX host's RIP and SAP applications that were dropped
Host Rx	Number of IPX packets that were received for the IPX host's RIP and SAP applications
Host Tx	Number of IPX packets that were transmitted from the IPX host's RIP and SAP applications
Host Tx Discards	Number of IPX packets from the IPX host's RIP and SAP applications that were dropped on transmission
Index	Index number that is assigned to the IPX interface
NetBIOS Max Hops	Number of IPX NetBIOS packets that exceeded the transport control maximum
NetBIOS Rx	Number of IPX NetBIOS packets that were received
NetBIOS Tx	Number of IPX NetBIOS packets that were transmitted
No Routes	Number of IPX packets that were dropped because the IPX route is unknown
Total Received	Number of IPX packets that were received
Tx Discards	Number of IPX packets that were forwarded but not successfully transmitted
Tx MTU Exceeded	Number of IPX packets that were forwarded but dropped because the MTU was exceeded

Field	Description
Routes Aged	Number of times the system marked a route entry unreachable, because it did not receive an update for that entry during the timeout period
Servers Aged	Number of times the system marked a server entry unreachable, because it did not receive an update for that entry during the timeout period
Rip Metric Changed	Number of times the metric changed on a route entry
Sap Metric Changed	Number of times the metric changed on a server entry

ipx
oddLengthPadding

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets the compatibility mode for older network interface cards (NICs). This mode enables an interface to pad IPX packets that have an odd number of bytes. (Older NICs discard IPX packets that have an odd number of bytes.)

Valid Minimum Abbreviation

ipx od

Important Considerations

- This feature supports 10 MB switching modules only.
- If you use this feature, be careful to select only those interfaces that require odd-length padding. Enabling this feature for every interface slows network performance.

Options

Prompt	Description	Possible Values	[Default]
Interface index	Index number of the interface for which you want to set the oddLengthPadding state	<ul style="list-style-type: none">■ A selectable IPX interface index■ ? (for a list of selectable indexes)	1 (if only 1)
IPX odd-length padding state	State for odd-length padding for the specified interface	<ul style="list-style-type: none">■ disabled■ enabled	disabled

✓ 3500
9000
9400

3900
9300

ipx NetBIOS Determines whether the system handles IPX Type 20 packet forwarding on a per-interface basis.

Valid Minimum Abbreviation

ipx n

Options

Prompt	Description	Possible Values	[Default]
Interface index	index number of the interface for which you want to set the NetBIOS forwarding state	<div><div>■ One or more selectable IPX interface indexes</div><div>■ all</div><div>■ ? (for a list of selectable indexes)</div></div>	1 (if only 1)
IPX NetBIOS forwarding state	State for NetBIOS forwarding for the specified interface	<div><div>■ disabled</div><div>■ enabled</div></div>	enabled (factory default), or current value

IPX NetBIOS Example (3500)

Select menu option (ipx): **netBIOS**
Select interface index(es) (1-6|all|?): **1**
Interface 1 - Enter state for NetBIOS packets (disabled,enabled) [enabled]: **disabled**

ipx secondary Determines whether the system enables secondary routes and servers.

✓ 3500
9000
9400

Valid Minimum Abbreviation

`ipx sec`

Important Considerations

- This option allows the system to learn about secondary routes and secondary servers.
- With this option, a secondary route/server entry can replace a primary route/server entry when the primary route/server is removed from the routing/server table for any reason (for example, if the associated interface goes down, or if the primary entry reaches its age limit).
- For this option to have any effect, you must establish alternate paths to the same IPX network or server.
- After you enable the IPX secondary route/server option, you can display entries for any secondary routes or servers. (See “ipx route secondary” and “ipx server secondary” earlier in this chapter.)

Options

Prompt	Description	Possible Values	[Default]
IPX secondary route/server state	How to handle secondary routes and servers	<ul style="list-style-type: none">■ disabled■ enabled	enabled (factory default), or current value

IPX Secondary Example

```
Select menu option (ipx): secondary
Enter secondary route/server state (disabled,enabled)
[disabled]: enabled
```


This chapter provides guidelines and other key information about commands that you can use to configure AppleTalk routing on your system. Configuring and managing AppleTalk routing involves these tasks:

- Administering AppleTalk interfaces
- Administering routes
- Administering the AARP cache
- Displaying the Zone Table
- Configuring forwarding
- Configuring checksum
- Enabling DDP Source Socket Verification
- Pinging an AppleTalk node
- Viewing AppleTalk statistics




For more information about administering AppleTalk routing on your network, see the Implementation Guide for your system.



For the CoreBuilder® 9000, the commands in this chapter apply only to Layer 3 switching modules.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



appletalk interface
summary

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays summary information for all AppleTalk interfaces.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap i su

Fields in the AppleTalk Interface Summary Display

Field	Description
Address	AppleTalk interface address, which is based on the network range and the network node (Example: 20301.7)
Index	Index number of the AppleTalk interface
Network range	Range of numbers that are assigned to the interface (Example: 20301 – 20310)
State	Status of the AppleTalk interface, which indicates whether the interface is available (enabled) or unavailable (down)
VLAN index	Index number of the virtual LAN (VLAN) that is associated with the AppleTalk interface

appletalk interface
detail

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays detailed information for all AppleTalk interfaces.

Valid Minimum Abbreviation

ap i det

Fields in the AppleTalk Interface Detail Display

Field	Description
Address	AppleTalk interface address, which is based on the network range and the network node. (Example: 20301.7)
Index	Index number of the AppleTalk interface
Network Range	Range of numbers that are assigned to the interface Example: (20301 – 20310)
Seed	Whether the interface is configured as a seed (y) or non-seed (n) interface
State	Status of the AppleTalk interface, that is, whether the interface is available (enabled) or unavailable (down)
VLAN index	Index number of the virtual LAN (VLAN) that is associated with the AppleTalk interface
Zone List	All zone names that are associated with the AppleTalk interface

appletalk interface
define

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Defines an AppleTalk interface.

Valid Minimum Abbreviation

ap i def

Important Considerations

- An AppleTalk interface defines the relationship between a virtual LAN (VLAN) and an AppleTalk network:
 - Every AppleTalk interface has one VLAN associated with it.
 - For routing purposes, you define a range of network numbers that are assigned to the AppleTalk interface. Example: 20301 – 20310
- You can configure the interface to be a seed or nonseed interface:
 - **Seed interface** — Initializes (“seeds”) the network with your configuration information. This information includes the network range and zone name list.
 - **Nonseed interface** — Listens for a seed router and then takes the zone and network range information from the first seed interface that it detects. After a nonseed interface obtains this information, it can participate in AppleTalk routing.
- Before you define the AppleTalk interface, you must define a VLAN and select AppleTalk as a protocol that the VLAN supports.
- Clients that have not been configured to use a particular zone use the default zone name.
- You can enter up to 16 zone names per interface.

Options

Prompt	Description	Possible Values	[Default]
Seed Interface	Whether an interface is configured as an AppleTalk seed (y) or non-seed interface (n).	<ul style="list-style-type: none">■ n (no)■ y (yes)	y (factory default), or current value
Start of network range	Start of the network range that is associated with the seed interface. <i>Seed interfaces only.</i>	1 – 65279	–
End of network range	End of the network range that is associated with the seed interface. <i>Seed interfaces only.</i>	1 – 65279	Value specified for start of network range, or current value

Prompt	Description	Possible Values	[Default]
Default zone name	User-defined default AppleTalk zone name. Clients that have not been configured to use a particular zone use the default zone name. <i>Seed interfaces only.</i>	Up to 32 ASCII characters	–
Zone name	AppleTalk zone that is associated with the interface. You are prompted to enter up to 15 additional zone names. <i>Seed interfaces only.</i>	<ul style="list-style-type: none"> ■ Up to 32 ASCII characters ■ q (to quit specifying zone names) 	–
VLAN interface index	Index number of the VLAN that you want to associate with the AppleTalk interface.	<ul style="list-style-type: none"> ■ Available valid VLAN index number ■ ? (for a list of available VLAN indexes) 	–

appletalk interface
modify

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Modifies an existing AppleTalk interface.

Valid Minimum Abbreviation

ap i m

Important Considerations

- An AppleTalk interface defines the relationship between a virtual LAN (VLAN) and an AppleTalk network:
 - Every AppleTalk interface has one VLAN associated with it.
 - For routing purposes, you define a range of network numbers that are assigned to the AppleTalk interface. Example: 20301 – 20310
- You can configure the interface to be a seed or nonseed interface:
 - **Seed interface** — Initializes (“seeds”) the network with your configuration information. This information includes the network range and zone name list.
 - **Nonseed interface** — Listens for a seed router and then takes the zone and network range information from the first seed interface that it detects. After a nonseed interface obtains this information, it can participate in AppleTalk routing.
- Before you define the AppleTalk interface, you must define a VLAN and select AppleTalk as a protocol that the VLAN supports.
- Clients that have not been configured to use a particular zone use the default zone name.
- You can enter up to 16 zone names per interface.

Options

Prompt	Description	Possible Values	[Default]
Interface	Index number of the AppleTalk interface that you want to modify	<ul style="list-style-type: none"> ■ AppleTalk interface index number ■ ? (for a list of selectable indexes) 	–
Seed Interface	Whether you want to configure the interface as an AppleTalk seed (y) or nonseed interface (n)	<ul style="list-style-type: none"> ■ n (no) ■ y (yes) 	Current value

Prompt	Description	Possible Values	[Default]
Start of network range	Start of the network range that is associated with the seed interface. <i>Seed interfaces only.</i>	1 – 65279	Current value
End of network range	End of the network range that is associated with the seed interface. <i>Seed interfaces only.</i>	1 – 65279	Current value
Default zone name	User-defined default AppleTalk zone name. Clients that have not been configured to use a particular zone use the default zone name. <i>Seed interfaces only.</i>	Up to 32 ASCII characters	Current value
Zone name	First AppleTalk zone that is associated with the interface. You are then prompted to enter up to 15 additional zone names. <i>Seed interfaces only.</i>	<ul style="list-style-type: none"> ■ Up to 32 ASCII characters ■ q (to quit specifying zone names and move on to the VLAN interface index prompt) 	Current value
VLAN interface index	Index number of the VLAN that you want to associate with the AppleTalk interface. When the system prompts you for a VLAN interface index, it indicates the available VLANs that you can associate with a new AppleTalk interface.	<ul style="list-style-type: none"> ■ Available valid VLAN index number ■ ? (for a list of selectable indexes) 	Current value
Interface down time	Number of minutes that you want to bring down the AppleTalk interface after you change zone information. This prompt appears only when you modify the zone information that is associated with the interface.	1 – 120 minutes	–

appletalk interface
remove

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Removes an existing AppleTalk interface.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap i r

Important Considerations

- You can specify a single interface, multiple AppleTalk interfaces, or all AppleTalk interfaces.
- If only one AppleTalk interface exists on the system, the interface is immediately removed after you enter this command.
- The system prompts you to select an interface number only if more than one AppleTalk interface exists on the system.

Options

Prompt	Description	Possible Values	[Default]
Interface	Index number of one or more interfaces that you want to remove	<ul style="list-style-type: none">■ One or more valid AppleTalk interface index numbers■ ? (for a list of selectable indexes)■ all	–

appletalk interface statistics

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays statistics for each AppleTalk interface. You can specify a single AppleTalk interface, multiple interfaces, or all interfaces. If you have multiple interfaces and you do not specify one of them, the system prompts you to specify the appropriate interface index number.

Valid Minimum Abbreviation

ap i st

Important Consideration

- The display includes statistics for the AppleTalk Address Resolution Protocol (ARP), Datagram Delivery Protocol (DDP), Routing Table Maintenance Protocol (RTMP), Zone Information Protocol (ZIP), Name Binding Protocol (NBP), and AppleTalk Echo Protocol (AEP).

Fields in the AppleTalk Interface Statistics Display

Field	Description
aarpInProbes	Number of ARP probes that have been received
aarpInReqs	Number of ARP requests that have been received
aarpInResp	Number of ARP responses that have been received
aarpOutProbes	Number of ARP probes that have been sent
aarpOutReqs	Number of ARP requests that have been sent
aarpOutResp	Number of ARP responses that have been sent
ddpForwRequests	Total number of packets for which an attempt was made to forward them to their final destination
ddpInChecksumErrors	Number of DDP datagrams that were dropped because of a checksum error
ddpInLocals	Number of DDP datagrams for which this entity was the final DDP destination
ddpInReceives	Total number of packets that have been received, including those with errors
ddpInTooLongs	Number of input DDP datagrams that have been dropped because they exceeded the maximum DDP datagram size
ddpInTooShorts	Number of input DDP datagrams that have been dropped because the received data length was less than the data length that was specified in the DDP header, or the received data length was less than the length of the expected DDP header
ddpNoProtoHandlers	Number of DDP datagrams without protocol handlers
echoInReplies	Number of echo replies that have been received

Field	Description
echoInRequests	Number of echo requests that have been received
echoOutReplies	Number of echo replies that have been sent
echoOutRequests	Number of echo requests that have been sent
nbpInBroadcastReqs	Number of NBP broadcast requests that have been received
nbpInErrors	Number of NBP packets that have been received and rejected for any error
nbpInForwardReqs	Number of NBP forward requests that have been received
nbpInLookupReqs	Number of NBP lookup requests that have been received
rtmpInDataPkts	Number of RTMP data packets that have been received
rtmpInRequestPkts	Number of RTMP request packets that have been received
rtmpOutDataPkts	Number of good RTMP data packets that have been sent
rtmpRouteDeletes	Number of times that RTMP has deleted a route that was aged out of the table
zipAddressInvalids	Number of times that this entity had to broadcast a ZIP GetNetInfo reply because the GetNetInfo request had an invalid address
zipInErrors	Number of ZIP packets that have been received and rejected for any error
zipInExReplies	Number of ZIP extended replies that have been received
zipInGniRequests	Number of ZIP GetNetInfo request packets that have been received
zipInZipQueries	Number of ZIP queries that have been received
zipInZipReplies	Number of ZIP replies that have been received
zipOutGniReplies	Number of ZIP GetNetInfo reply packets that have been sent
zipOutInvalids	Number of ZIP GetNetInfo replies that have been sent with the indication that the previous client zone name was invalid

appletalk route display

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Displays AppleTalk routes that are listed in the system's routing table.

Valid Minimum Abbreviation

ap r d

Important Consideration

- Your system maintains a table of local and remote routes to all reachable AppleTalk networks. The Routing Table Maintenance Protocol (RTMP) automatically generates the routing table. RTMP defines rules for:
 - Information that is contained within each routing table
 - Exchanging information between routers so that the routers can maintain their routing tables

Fields in the AppleTalk Route Display

Field	Description
Distance	Distance in hops to the destination network
Interface	Interface that is used to reach the destination network
Network Range	Range of numbers that identify a network
Next Hop	Next hop Internet router to which the packet must be sent
State	Status of each route. One of the following: <ul style="list-style-type: none">■ good■ suspect■ bad■ really bad

appletalk route flush

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Deletes all dynamically learned AppleTalk routes from the routing table.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

ap r f

3900

9300

Important Consideration

- The system deletes all dynamically learned AppleTalk routes immediately after you enter the command. You are not prompted to confirm the deletion.

appletalk aarp
display

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the AppleTalk Address Resolution Protocol (AARP) cache.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap a d

Fields in the AppleTalk AARP Display

Field	Description
AARP address	AppleTalk protocol address
Age (secs)	Age of the ARP entry (in seconds)
Interface	Index number of the interface on which the address was learned
MAC address	Hardware address that corresponds to the AppleTalk address

3900
9300

appletalk aarp
remove

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Removes an AppleTalk Address Resolution Protocol (AARP) cache entry.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap a r

Options

Prompt	Description	Possible Values	[Default]
AARP address	AARP address that you want to remove from the system's AARP cache	Any valid AARP address	–

appletalk aarp flush***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Deletes all AppleTalk Address Resolution Protocol (ARP) entries from the system's ARP cache.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation**ap a f**

3900

9300

Important Consideration

- The system deletes all ARP entries immediately after you enter the command. You are not prompted to confirm the deletion.

appletalk zone
display network

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the AppleTalk Zone table, indexed by network numbers.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap z d n

Important Considerations

- AppleTalk routers use the Zone Information Protocol (ZIP) to map network numbers to Zones.
- Each AppleTalk router maintains a Zone Information Table (ZIT), which lists the zone-to-network mapping information.

3900
9300

appletalk zone
display zone

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the AppleTalk Zone table indexed by zones.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap z d z

Important Considerations

- AppleTalk routers use the Zone Information Protocol (ZIP) to map network numbers to Zones.
- Each AppleTalk router maintains a Zone Information Table (ZIT), which lists the zone-to-network mapping information.

3900
9300

appletalk forwarding

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

Enables and disables AppleTalk Data Delivery Protocol (DDP) forwarding.

✓ 9000

Valid Minimum Abbreviation

9400

ap f

3900

Options

9300

Prompt	Description	Possible Values	[Default]
Forwarding state	Whether to enable or disable AppleTalk forwarding	<div><div>■</div> enabled</div> <div><div>■</div> disabled</div>	disabled (factory default), or current value

appletalk checksum

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

Enables Data Delivery Protocol (DDP) checksum error detection for the AppleTalk protocol.

Valid Minimum Abbreviation

ap c

Important Considerations

- The AppleTalk protocol uses checksums to detect errors in data transmissions. A *checksum* totals all data bytes and adds the sum to the checksum field of the data packet. The receiving station computes a verification checksum from the incoming data and compares the new checksum with the value that is sent with the data. If the values do not match, the transmission contains an error.
- Disabled is the preferred setting. Enabling the checksum generation or verification significantly impacts the router's performance.

Options

Prompt	Description	Possible Values	[Default]
Checksum generation state	Whether to enable or disable generation of checksums for AppleTalk packets	<ul style="list-style-type: none">■ enabled■ disabled	disabled (factory default), or current value
Checksum verification state	Whether to enable or disable verification of checksums for AppleTalk packets	<ul style="list-style-type: none">■ enabled■ disabled	disabled (factory default), or current value

appletalk
sourceSocket

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Enables and disables AppleTalk Data Delivery Protocol (DDP) source socket verification.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap so

3900
9300

Options

Prompt	Description	Possible Values	[Default]
source Socket	Whether to enable or disable source socket verification	<div><div>■</div> enabled</div> <div><div>■</div> disabled</div>	disabled (factory default), or current value

Prompt	Description	Possible Values	[Default]
Destination AARP address	AppleTalk node that you want to test for network connectivity	Valid AARP address	–

appletalk statistics
ddp

For CoreBuilder 9000: Applies to Layer 3 switching modules only.
Displays AppleTalk Datagram Delivery Protocol (DDP) statistics.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation
ap s d

3900
9300

Fields in the AppleTalk DDP Statistics Display

Field	Description
inBcastErrors	Number of dropped DDP datagrams for which the system was not their final destination and that were sent to the broadcast MAC address
inCsumErrors	Number of DDP datagrams that were dropped because of a checksum error
inDiscards	Number of DDP Datagrams that were discarded during routing
inForwards	Total number of packets that were forwarded, including those with errors
inLocals	Number of DDP datagrams for which an attempt was made to forward them to their final destination
inNoClients	Number of DDP datagrams that were dropped for unknown DDP types
inNoRoutes	Number of DDP datagrams that were dropped for unknown routes
inReceives	Total number of packets that were received, including those with errors
inShortDdps	Number of input DDP datagrams that were dropped because the system was not their final destination and their type was short DDP
inTooFars	Number of input datagrams that were dropped because the system was not their final destination and their hop count would exceed 15
inTooLongs	Number of input DDP datagrams that were dropped because they exceeded the maximum DDP datagram size
inTooShorts	Number of input DDP datagrams that were dropped because the received data length was less than the data length that was specified in the DDP header, or the received data length was less than the length of the expected DDP header
outLocals	Number of host-generated DDP datagrams

appletalk statistics rtmp

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays AppleTalk Routing Table Maintenance Protocol (RTMP) statistics.

Valid Minimum Abbreviation

ap s r

Fields in the AppleTalk RTMP Statistics Display

Field	Description
inDatas	Number of good RTMP data packets that were received
inOtherErrs	Number of RTMP packets that have been received and rejected for an error other than a version mismatch
inRequests	Number of good RTMP request packets that were received
inVersionErrs	Number of RTMP packets that have been received and rejected due to a version mismatch
outDatas	Number of RTMP data packets that were sent
outRequests	Number of RTMP request packets that were sent
routeDeletes	Number of times that RTMP deleted a route that was aged out of the table
routeEqChgs	Number of times that RTMP changed the Next Internet Router in a routing entry because the hop count that was advertised in a routing table was equal to the current hop count for a particular network
routeLessChgs	Number of times that RTMP changed the Next Internet Router in a routing entry because the hop count that was advertised in a routing table was less than the current hop count for a particular network
routeOverflows	Number of times that RTMP attempted to add a route to the RTMP table but failed because of lack of space

appletalk statistics zip

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays AppleTalk Zone Information Protocol (ZIP) statistics.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

ap s z

Fields in the AppleTalk ZIP Statistics Display

Field	Description
inErrors	Number of ZIP packets that have been received and rejected for any error
inExReplies	Number of ZIP extended replies that have been received
inGniReplies	Number of ZIP GetNetInfo reply packets that have been received
inGniRequests	Number of ZIP GetNetInfo request packets that have been received
inLocalZones	Number of ZIP GetLocalZones requests packets that have been received
inObsoletes	Number of ZIP Takedown or ZIP Bringup packets that have been received
inQueries	Number of ZIP queries that have been received
inReplies	Number of ZIP replies that have been received
inZoneCons	Number of times that a conflict has been detected between this system's zone information and another entity's zone information
inZoneInvs	Number of times that this system has received a ZIP GetNetInfo reply with the zone invalid bit set because the corresponding GetNetInfo request had an invalid zone name
inZoneLists	Number of ZIP GetZoneLists requests packets that have been received
outAddrInvs	Number of times that this system had to broadcast a ZIP GetNetInfo reply because the GetNetInfo request had an invalid address
outExReplies	Number of ZIP extended replies that have been sent
outGniReplies	Number of ZIP GetNetInfo reply packets that have been sent out of this port
outGniRequests	Number of ZIP GetNetInfo packets that have been sent
outLocalZones	Number of transmitted ZIP GetLocalZones reply packets
outQueries	Number of ZIP queries that have been sent
outReplies	Number of ZIP replies that have been sent
outZoneInvs	Number of times that this system has sent a ZIP GetNetInfo reply with the zone invalid bit set in response to a GetNetInfo request with an invalid zone name
outZoneLists	Number of transmitted ZIP GetZoneList reply packets

**appletalk statistics
nbp*****For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays AppleTalk Name Binding Protocol (NBP) statistics.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation**ap s n**

3900

9300


Fields in the AppleTalk NBP Statistics Display

Field	Description
inBcastReqs	Number of NBP Broadcast Requests that have been received
inErrors	Number of NBP packets that have been received and rejected for any error
inFwdReqs	Number of NBP Forward Requests that have been received
inLkupReplies	Number of NBP Lookup Replies that have been received
inLkupReqs	Number of NBP Lookup Requests that have been received



TRAFFIC POLICY

Chapter 22 Quality of Service (QoS) and RSVP



QUALITY OF SERVICE (QoS) AND RSVP

Quality of Service (QoS) and the *Resource Reservation Protocol (RSVP)* are advanced features that provide policy-based services. *Policy-based services* establish various grades of network services to accommodate the needs of different types of traffic (for example, multimedia, video, and file backups). QoS software relies on RSVP to provide admission control.

This chapter provides guidelines and other key information about how to configure QoS and RSVP in your system.

QoS and RSVP features include classifiers, controls, and RSVP parameters. Configure these features in the following order:

- 1 You first enter the command `qos`
- 2 to define how the system groups packets so that it can schedule them with the appropriate service level.
- 3 You then enter the command `qos control define` to assign rate limits and priorities to the packets that are associated with one or more of your classifiers. A classifier has no effect until you associate it with a control.

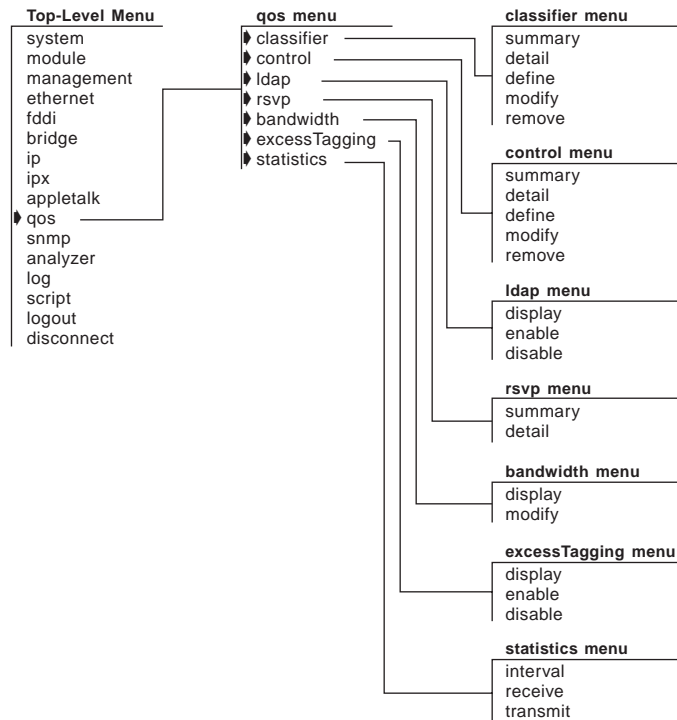
The system provides predefined classifiers and controls that are suitable for many configurations, or you can define your own classifiers, apply controls to the classifiers, and then decide whether to use RSVP. For more information about QoS and RSVP, see the *Implementation Guide* for your system.



For the CoreBuilder® 9000, the commands in this chapter apply only to Layer 3 switching modules.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



qos classifier
summary

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays summary information about the QoS classifiers on your system.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

q cl s

Fields in the QoS Classifier Summary Display

Field	Description
802.1p	For <i>nonflow</i> classifiers, IEEE 802.1p tag value
Cast	Cast type for the classifier: <ul style="list-style-type: none">■ <i>Flow</i> classifiers: unicast, multicast, or all■ <i>Nonflow</i> classifiers: unicast, multicast, broadcast, or all
Classifier	Number of the flow or nonflow classifier: <ul style="list-style-type: none">■ <i>Flow</i> classifiers in the range of 1 – 399 (Note: 20 and 23 are predefined.)■ <i>Nonflow</i> classifiers in the range of 400 – 498 (Note: 401 – 407, 420, 430, 440, 450, 460, 470, 480, and 490 are predefined, but you can modify or remove them.)
Control	Control number that you assign to the classifier
Name	Name that you assign to the classifier
Protocol	Protocol type, if applicable, that is associated with the classifier/control: <ul style="list-style-type: none">■ <i>Flow</i> classifiers: IP protocol type TCP, UDP, or all■ <i>Nonflow</i> classifiers: TCP, IP, IPX, AppleTalk, or any

qos classifier detail***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays detailed information about one or more QoS classifiers.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q cl det

Options

Prompt	Description	Possible Values	[Default]
Classifier number	Number of the classifier for which you want detail information	<ul style="list-style-type: none"> One or more numbers of configured classifiers all ? (for a list of selectable classifiers) 	–

Fields in the QoS Classifier Detail Display

Field	Description
802.1p	For nonflow classifiers, IEEE 802.1p tag value (any combination of priority tag values in the range 0 – 7)
Cast	The Cast type for the classifier: <ul style="list-style-type: none"> <i>Flow</i> classifiers: unicast, multicast, or all <i>Nonflow</i> classifiers: unicast, multicast, broadcast, or all
Classifier	Number of the flow or nonflow classifier: <ul style="list-style-type: none"> <i>Flow</i> classifiers in the range of 1 – 399 (Note: 20 and 23 are predefined.) <i>Nonflow</i> classifiers in the range of 400 – 498 (Note: 401 – 407, 420, 430, 440, 450, 460, 470, 480, and 490 are predefined, but you can modify or remove them.)

Field	Description
Classifier – Filters (flow classifiers only)	Filters (address and port patterns): <ul style="list-style-type: none">■ Source IP address■ Source IP address mask■ Destination IP address■ Destination IP address mask■ Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port range
Destination Port range (flow classifiers only)	Beginning and end of the TCP or UDP destination port range
Source Port range (flow classifiers only)	Beginning and end of the TCP or UDP source port range
Classifier – Installed Flows (if flows exist)	Actual flows seen on the system, with the following data: <ul style="list-style-type: none">■ Port■ Source IP address/source port■ Destination IP address /destination port■ Protocol type■ Number of flow cache misses
Control	Control number that you assign to the control
Name	Name that you assign to the classifier
Protocol	Protocol type, if applicable, that is associated with the classifier and control

qos classifier define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a flow or nonflow classifier.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

`q c1 def`

Important Considerations

- *Classifiers* define how the system groups packets so that it can schedule them with the appropriate service level. QoS supports flow and nonflow classifiers:
 - *Flow classifiers* apply to routed IP multicast and IP unicast packets. You can define up to 100 flow classifiers. Each filter (address and port pattern) in a flow classifier counts toward the limit.
 - *Nonflow classifiers* apply to bridged or routed traffic that is associated with a specific protocol (IP, TCP/IP, IPX, and AppleTalk) or to a custom protocol (Ethernet or Destination Service Access Point/Source Service Access Point (DSAP/SSAP). You can also use them to apply IEEE 802.1p tag values to forwarded frames. You can define up to 16 nonflow classifiers. All 16 nonflow classifiers are in use by default.
- The default classifier number is 499. You cannot remove or modify this default classifier. However, you can remove any of the predefined classifiers (for example, if you need another nonflow classifier). See “qos classifier remove” later in this chapter for more information.
- When you define a filter (address and port pattern) for a flow classifier, select a source and destination start and end port ranges that are as small as possible (for example, a single port). If the classifier applies to a wide range of Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports, you increase the amount of classified traffic on the system and consume valuable QoS resources.
- A classifier can have only one control applied to it.
- If you select `custom` when you define a nonflow classifier, you are prompted to select the protocol by Ethernet or DSAP/SSAP. After you select a protocol, you are prompted to provide the hexadecimal ranges.



Depending on the number of VLANs defined, you can define a maximum of 3 custom protocols that can have controls applied to them. This limitation does not apply to non-controlled custom protocols.

Options

Prompt	Description	Possible Values	[Default]
Classifier number	Number of the flow or nonflow classifier in the range of 1 – 498	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: 1 – 399 (except 20 and 23, which are predefined flow classifiers) ■ <i>Nonflow</i> classifiers: 400 – 498, (except 401 – 407, 420, 430, 440, 450, 460, 470, 480, and 490. 401 – 407 are predefined nonflow classifiers with applied controls and IEEE 802.1p tag values of 1 – 7.) 	–
Classifier name	Name that you assign to the classifier	<ul style="list-style-type: none"> ■ Unique name with up to 32 characters (Use " around any string with embedded spaces. Use " " to enter an empty string.) ■ ? (for a list of selection criteria) 	–
Cast type	Cast type for the flow or nonflow classifier	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: unicast, multicast, or all ■ <i>Nonflow</i> classifiers: unicast, multicast, broadcast, or all ■ ? (for a list of selectable cast types) 	–
Protocol type	IP or other protocol type, if applicable, that you want to associate with the flow or nonflow classifier	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: IP protocol type with TCP, UDP, or all ■ <i>Nonflow</i> classifiers: TCP, IP, IPX, AppleTalk, custom, or any ■ ? (for a list of selectable protocol types) 	–
Source IP address	For <i>flow</i> classifiers only, IP address of the source	Up to 255.255.255.255	0.0.0.0 (factory default, wildcard match)

Prompt	Description	Possible Values	[Default]
Source IP address mask	For <i>flow</i> classifiers only, source IP address mask, or how many portions of the IP address you want to match (Example: 255.255.255.0 matches the first three portions of the specified IP address.)	Up to four portions (255.255.255.255)	0.0.0.0 (factory default)
Destination IP address	For <i>flow</i> classifiers only, destination IP address	Up to 255.255.255.255	0.0.0.0 (factory default, wildcard match)
Destination IP address mask	For <i>flow</i> classifiers only, destination IP address mask, or how many portions of the address you want to match	Up to four portions (255.255.255.255)	0.0.0.0 (factory default, wildcard match)
Start and end of TCP or UDP source port range	For <i>flow</i> classifiers only, start and end of the TCP or UDP source port range. The start value determines the end value.	<ul style="list-style-type: none"> ■ 0 – 65535 (start) ■ 2049 – 65535 (end) See “QoS Classifier Define Example (Flow Classifier)”.	0 and 65535 (factory defaults)
Start and end of TCP or UDP destination port range	For <i>flow</i> classifiers only, start and end of the TCP or UDP destination port range. The start value determines the end value.	<ul style="list-style-type: none"> ■ 0 – 65535 (start) ■ 2049 – 65535 (end) See “QoS Classifier Define Example (Flow Classifier)”.	0 and 65535 (factory defaults)
Additional filter (address/port pattern)	For <i>flow</i> classifiers only, additional source, destination, and port information for this classifier	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	no (factory default)
Custom protocol type (custom nonflow classifiers only)	For <i>nonflow</i> classifiers with the custom protocol type	<ul style="list-style-type: none"> ■ Ethertype ■ DSAP/SSAP 	–

Prompt	Description	Possible Values	[Default]
Custom protocol hexadecimal value (custom nonflow classifiers only)	Hex values for <i>nonflow</i> classifiers with the protocol custom type	<ul style="list-style-type: none"> Ethertype hex value of 0x0 – 0xfffe DSAP hex value of 0x0 – 0xff Note: You cannot enter 0xaa - 0xaa SSAP hex value of 0x0 – 0xff Note: You cannot enter 0xaa - 0xaa 	0x0 – 0x0
802.1p tag	For <i>nonflow</i> classifiers only, IEEE 802.1p tag values	<ul style="list-style-type: none"> Any combination of priority tag values in the range of 0 – 7 all ? (for a list of possible values) 	–

Flow Classifier Procedure

To accept the default or current values that appear in brackets [], press Enter.

- i

1 Enter a classifier number in the range of from 1 through 399.

Flow classifiers 20 and 23 are predefined for FTP and Telnet.

2 Enter the classifier name (a unique name of up to 32 characters).

3 Enter a cast type.

For a flow classifier, the options are `unicast`, `multicast`, and `all`.

4 Enter the IP protocol type of `TCP`, `UDP`, or `all`.

5 Enter the source IP address. The default value is 0.0.0.0.

6 Enter the source IP address mask. The default value is 0.0.0.0.

7 Enter the destination IP address.

8 Enter the destination IP address mask.

9 Enter the start of the TCP or UDP source port range, in the range of from 0 through 65535. The default is 0.

10 Enter the end of the TCP or UDP source port range using a value of up to 65535.

The value that you enter for the start of the range determines the default for the end of the range. The end value must be greater than or equal to the start value.



To avoid severely affecting applications using the network, select a port range that is as small as possible (for example, a single port).

- 11** Enter the start of the TCP or UDP destination port range, in the range of from 0 through 65535. The default is 0.
- 12** Enter the end of the TCP or UDP destination port range using a value of up to 65535. The end value must be greater than or equal to the start value.
- 13** At the prompt, specify whether you want any other filters (address and port patterns) with this classifier (`yes` or `no`). The default is `no`.

If you specify `yes`, the system prompts you for additional information, beginning with the source IP address.



Flow classifiers classify traffic only at the network layer and therefore affect only traffic that is being routed from one subnet to another.

QoS Classifier Define Example (Flow Classifier)

```
Select menu option (qos/classifier): define
Enter classifier number (1-498): 26
Enter classifier name {?}: IPFilter1
Select cast type (unicast,multicast|all|?): all
Select IP protocol type (TCP,UDP|all|?): all
Enter source IP address [0.0.0.0]:168.20.30.0
Enter source IP address mask [255.255.0.0]:255.255.255.0
Enter destination IP address [0.0.0.0]:192.1.0.0
Enter IP address mask [255.255.255.0]:255.255.0.0
Enter start of UDP source port range (0-65535) [0]:0
Enter end of UDP source port range (0-65535) [65535]:65535
Enter start of UDP destination port range (0-65535) [0]:0
Enter end of UDP destination port range (0-65535)
[65535]:65535
Enter another filter (yes,no) [no]: n
```

Nonflow Classifier Procedure

To accept the default or existing values that appear in brackets [], press Return.

- 1 Enter a classifier number in the range of from 400 through 498.

Numbers 401 through 407 are predefined nonflow classifiers with applied controls; numbers 420, 430, 440, 450, 460, 470, 480, and 490 are predefined nonflow classifiers without controls. If you have not removed any of the predefined nonflow classifiers, you need to remove them before you can define another nonflow classifier. (With the default classifier, there is a limit of 16 predefined nonflow classifiers.)

- 2 Enter the classifier name (a unique name of up to 32 characters long).

- 3 Enter a cast type.

For a *nonflow* classifier, the options are *unicast*, *multicast*, *broadcast*, and *all*.

- 4 Enter one or any protocols.

The options are *TCP/IP*, *IP*, *IPX*, *Appletalk*, *any*, or *custom*.

- 5 If you choose *custom*, enter the protocol type (*ethernet* or *DSAP/SSAP*).

- For *ethernet* type enter the hexadecimal value.
- For *DSAP/SSAP* type, enter the DSAP and SSAP hexadecimal values.

- 6 Enter one or all IEEE 802.1p tags. Specify any combination of values in the range of from 0 through 7, or *all*.

QoS Classifier Define Example (Nonflow Classifier)

```
Select menu option (qos/classifier): define
Enter classifier number (1-498): 481
Enter classifier name {?}: AppleBcast
Select cast type (unicast,multicast,broadcast|all|?):
broadcast
Select protocols {TCP/IP,IP,IPX,Appletalk,any,custom|?}:
Appletalk
Select IEEE 802.1p tag(s) (0-7|all|?): all
```

qos classifier modify

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

3900

9300

Modifies a previously defined classifier.

Valid Minimum Abbreviation

q cl m

Important Consideration

- If the classifier that you want to modify is associated with a control, you *must* remove the control before you can modify the classifier. See “qos classifier remove” later in this chapter for more information.

Options

Prompt	Description	Possible Values	[Default]
Classifier number	Number of the flow or nonflow classifier that you want to modify. Existing classifiers are shown in braces.	<ul style="list-style-type: none">■ <i>Flow</i> classifiers: 1 – 399 (except 20 and 23, which are predefined flow classifiers)■ <i>Nonflow</i> classifiers: 400 – 498, (except 401 – 407, 420, 430, 440, 450, 460, 470, 480, 490. 401 – 407 are predefined nonflow classifiers with applied controls.)■ ? (for a list of selectable values)	–
Classifier name	Name of the classifier that you want to modify	<ul style="list-style-type: none">■ Unique name with up to 32 characters (Use " around any string with embedded spaces. Use " " to enter an empty string.)■ ? (for a list of selection criteria)	Current name
Cast type	Cast type for the flow or nonflow classifier	<ul style="list-style-type: none">■ <i>Flow</i> classifiers: unicast, multicast, or all■ <i>Nonflow</i> classifiers: unicast, multicast, broadcast, or all	Current cast type

Prompt	Description	Possible Values	[Default]
Protocol type	IP or other protocol type, if applicable, that is associated with the flow or nonflow classifier.	<ul style="list-style-type: none"> ■ <i>Flow</i> classifiers: IP protocol type with TCP, UDP, or all ■ <i>Nonflow</i> classifiers: TCP, IP, IPX, AppleTalk, any, or custom ■ ? (for a list of selectable values) 	Current protocol type
Source IP address	For <i>flow</i> classifiers only, IP address of the source.	Up to 255.255.255.255	0.0.0.0 (factory default, wildcard match), or current value
Source IP address mask	For <i>flow</i> classifiers only, source IP address mask, or how many portions of the IP address you want to match. (Example: 255.255.255.0 matches the first three portions of the specified IP address.)	Up to four portions (255.255.255.255)	0.0.0.0 (factory default, wildcard match), or current value
Destination IP address	For <i>flow</i> classifiers only, destination IP address.	Up to 255.255.255.255	0.0.0.0 (factory default), or current value
Destination IP address mask	For <i>flow</i> classifiers only, destination IP address mask, or how many portions of the IP address you want to match.	Up to four portions (255.255.255.255)	0.0.0.0 (factory default), or current value
Start and end of TCP or UDP source port range	For <i>flow</i> classifiers only, start and end of the TCP or UDP source port range. Specify as small a range as possible. The start value determines the end value.	0 – 65535	0 and 65535 (factory defaults), or current values

Prompt	Description	Possible Values	[Default]
Start and end of TCP or UDP destination port range	For <i>flow</i> classifiers only, start and end of the TCP or UDP destination port range. Specify as small a range as possible. The start value determines the end value.	0 – 65535	0 and 65535 (factory defaults), or current values
Additional filters (address/port patterns)	For <i>flow</i> classifiers only, additional source, destination, and port information for this classifier. Each set of information counts toward the classifier limit.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	no (factory default)
Custom protocol type (custom nonflow classifiers only)	For <i>nonflow</i> classifiers with the custom protocol type.	<ul style="list-style-type: none"> ■ Ethertype ■ DSAP/SSAP 	–
Custom protocol hexadecimal value (custom nonflow classifiers only)	Hex values for <i>nonflow</i> classifiers with the custom protocol type.	<ul style="list-style-type: none"> ■ Ethertype hex value of 0x0 – 0xffff ■ DSAP hex value of 0x0 – 0xff ■ SSAP hex value of 0x0 – 0xff 	0x0 – 0x0
802.1p tag	For <i>nonflow</i> classifiers only, the IEEE 802.1p tag value	<ul style="list-style-type: none"> ■ Any combination of priority tag values in the range of 0 – 7 ■ all ■ ? (for a list of selectable values) 	Current value, if any

Procedure (Flow Classifier)

- 1 Enter the number of the classifier that you want to modify. The current numbers are shown in braces { }.

- 2 To modify the name, enter the new name for the classifier.

The name that is associated with the classifier number that you specified is shown in brackets.

- 3** To modify the cast type, enter a new cast type.
For a flow classifier, the options are `unicast`, `multicast`, and `all`.
To accept the default or current value that appears in brackets, press Enter.
- 4** To modify the IP protocol type, enter another IP protocol type (`TCP`, `UDP`, or `all`).
- 5** To modify the current source IP address, enter a new source IP address.
- 6** To modify the current source IP address mask, enter a new source IP address mask.
- 7** To modify the current destination IP address, enter a new destination IP address.
- 8** To modify the current destination IP address mask, enter a new destination IP address mask.
- 9** To modify the TCP or UDP source port range, enter the new start of the TCP or UDP port range (in the range of from 0 through 65535).
Limit the source port range as much as possible.
- 10** Enter the new end of the TCP or UDP source port range (in the range of from 0 through 65535).
- 11** To modify the TCP or UDP destination port range, enter the new start of the TCP or UDP port range (in the range of from 0 through 65535).
- 12** Enter the new end of the TCP or UDP destination port range (in the range of from 0 through 65535).
Limit the destination port range as much as possible.
- 13** At the prompt, specify whether you want any other address and port patterns (filters) with this classifier: `yes` or `no`; the default is `no`.
If you specify `yes`, the system prompts you for additional filtering information, beginning with the source IP address.



If you have several existing address and port patterns, you must specify all of them again during the modification process. Any address and port patterns that you do not reenter are deleted.

Nonflow Classifier Procedure

- 1 To modify the cast type, enter a new cast type.
For a nonflow classifier, the options are unicast, multicast, broadcast, and all
- 2 To modify the associated protocols, enter another protocol.
The options are TCP/IP, IP, IPX, Appletalk, any, or custom.
- 3 If you choose custom, select the protocol type (ethernet or DSAP/SSAP).
 - For the ethernet type, enter the hexadecimal value
 - For the DSAP/SSAP type, enter the DSAP and SSAP hexadecimal values
- 4 To modify the handling of IEEE 802.1p tags, enter the appropriate tags using a value in the range of 0 through 7, or enter **all**

QoS Classifier Modify Example (Flow Classifier)

```
Select menu option (qos/classifier): modify
Enter classifier number
{20,23,26,401-407,420,430,440, 450, 460, 470,480,490|?}:26
Enter classifier name {?} [IPFilter1]:
Select cast type (unicast,multicast|all|?)
[unicast,multicast]:
Select IP protocol type (TCP,UDP|all|?) [TCP,UDP]:
Enter source IP address [168.20.30.0]:
Enter source IP address mask [255.255.0.0]:
Enter destination IP address [192.1.1.0]:
Enter destination IP address mask [255.255.255.0]:
Enter start of TCP source port range (0-65535) [0]:
Enter end of TCP source port range (0-65535) [65535]:
Enter start of TCP destination port range (0-65535) [0]:
Enter end of TCP destination port range (0-65535) [65535]:
Enter another filter (yes,no) [no]: n
```

qos classifier remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Removes a previously defined classifier.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q cl r

Important Considerations

- If the classifier that you want to remove is associated with a control, you *must* remove the control before you can remove the classifier. See “qos control remove” later in this chapter for more information.
- When you enter the command, specify the number that represents the classifier that you want to remove, or specify ? to view the selectable classifiers.

Options

Prompt	Description	Possible Values	[Default]
Classifier number	Number for the classifier that you want to remove	<ul style="list-style-type: none"> ■ Any selectable classifier number ■ ? (for a list of selectable classifiers) 	–

QoS Classifier Remove Example (3500)

Select menu option: **qos classifier remove**

Enter classifier number

{20,23,26,401-407,420,430,440,450,460,470,480,490|?}: **26**

qos control summary

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

3900

9300

Displays summary information about QoS controls.

Valid Minimum Abbreviation

q co s

Fields in the QoS Control Summary Display

Field	Description
802.1p Tag	For controls for nonflow classifiers, the IEEE 802.1p tag value (0 – 7).
Classifiers controlled	Classifiers that this control affects.
Control number	Number of the control.
Control name	Name of the control.
Excess loss eligible	For receivePort or aggregate rate limit types, whether excess packets are loss eligible.
Excess service	For receivePort or aggregate rate limit types, the service level for excess packets.
Loss eligible	Whether conforming packets are loss eligible. If a packet is loss eligible, it can be dropped if the transmit queue for which it is destined exceeds its threshold.
Service	Service level for the conforming packets.

qos control detail***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays detailed information about the QoS controls that you specify.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q co det

Options

Prompt	Description	Possible Values	[Default]
Control number	Number of the control for which you want detail information	<ul style="list-style-type: none"> ■ One or more configured controls ■ all ■ ? (for a list of selectable controls) 	–

Fields in the QoS Control Detail Display

Field	Description
802.1p tag	IEEE 802.1p priority tag value (0 – 7) that is applied to forwarded frames. Can be defined for both flow and nonflow classifiers.
Burst	Burst size in KBytes.
Classifiers controlled	Classifiers that this control affects.
Control (number)	Number of the control.
Control name	Name that you assign to the control.
End time	Control end time.
Excess loss eligible	For receivePort or aggregate rate limit types, whether excess packets are loss eligible.
Excess service	For receivePort or aggregate rate limit type, service level for excess packets.
Limit	Rate limit in KBytes/sec or percentage.
Loss eligible	Whether conforming packets are loss eligible. A loss-eligible packet can be dropped if the transmit queue for which it is destined is over its threshold.
Ports	Receive ports for which you want to enable the rate limit.
Rate limits control	Number of the control that the rate limit affects.
Service	Service level for the conforming packets (high, best, low, or drop).

Field	Description
Source Port range	Beginning and end of the source port range.
Start time	Control start time
TCP drop control	Whether TCP drop control filtering is enabled.
Time control type	Time control type (<i>specific</i> , <i>daily</i> , <i>weekdays</i> , and so forth).
Type	Rate limit type, <i>none</i> (no rate limit), <i>receivePort</i> , or <i>aggregate</i> .

qos control define *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Defines a control for one or more existing classifiers.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q co def

Important Considerations

- A *control* can assign multiple rate limit values and an IEEE 802.1p priority tag value to the packets that are associated with one or more classifiers.
- The system prompts you according to the rate limit type that you select. You can only use one rate limit type (*none*, *receivePort*, or *aggregate*) per control. For a type of *receivePort* or *aggregate*, you can specify multiple rate-limit values for groups of ports or individual ports. The *aggregate* rate limit type can only be applied to flow classifiers.
- *Loss-eligible packets* are conforming packets that are discarded instead of queued when transmit queues back up beyond a threshold. You can specify whether conforming packets (as well as nonconforming excess packets) are loss eligible when you define the control. Marking packets loss eligible is useful for an intelligent discard of traffic in a congestion situation. Nonconforming excess packets are packets that exceed the specified rate limit.
- With the QoS timer control, you can configure QoS control sessions with starting and ending times (similar to using a VCR).

Options

Prompt	Description	Possible Values	[Default]
Control number	Number of the control. Control numbers 1 – 4 are predefined controls.	<ul style="list-style-type: none"> ■ 5 – 50 ■ ? (for a list of selectable values) 	1 (factory default)

Prompt	Description	Possible Values	[Default]
Control name	Name that you assign to the control. Predefined names are as follows: <ul style="list-style-type: none"> ■ Default/Best Effort (for control 1) ■ Background (for control 2) ■ Business Critical (for control 3) ■ Controlled Load (for control 4) 	<ul style="list-style-type: none"> ■ Unique name with up to 32 characters (Use " around any string with embedded spaces. Use " " to enter an empty string.) ■ ? (for a list of selection criteria) 	Default/Best Effort
Rate limit type	Type of rate limit: <ul style="list-style-type: none"> ■ none (no rate limit) ■ receivePort (a rate limit on the specified ports) ■ aggregate (the bandwidth for all ports chosen for the associated classifier). For flow classifiers only. 	<ul style="list-style-type: none"> ■ none ■ receivePort ■ aggregate 	none (factory default)
Service level	Service level for the conforming packets (a transmit priority that corresponds to a transmit queue). Drop causes the system to drop all traffic on all ports that are associated with the classifier and control.	<ul style="list-style-type: none"> ■ For rate limit receivePort or aggregate: high, best (best effort), or low ■ For a rate limit of none: high, best, low, or drop 	best (factory default)
Loss eligible	Whether conforming packets are loss-eligible. A loss-eligible packet can be dropped if the transmit queue for which it is destined exceeds its threshold.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	no (factory default)
Excess packet service	For receivePort or aggregate rate limit types, the service level for excess packets (packets that exceed the rate limit).	<ul style="list-style-type: none"> ■ high ■ best ■ low ■ drop 	best (factory default)
Excess loss eligible	For receivePort or aggregate rate limit types, whether excess packets are loss-eligible.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	yes (factory default)

Prompt	Description	Possible Values	[Default]
How rate limit is expressed	For receivePort or aggregate rate limit types, in KBytes/sec or percentage.	<ul style="list-style-type: none"> ■ KBytes/sec ■ percentage 	KBytes/sec (factory default)
Rate limit value	For receivePort or aggregate rate limit types, in KBytes/sec or percentage. 0 makes all packets excess packets.	<ul style="list-style-type: none"> ■ 0 – 65434 KBytes/sec ■ 0 – 100 percent 	–
Burst size	For receivePort or aggregate rate limit types, the maximum amount of data in Kbytes that you can transmit at the line rate before the transmission is policed.	16 – 8192 KBytes	Determined by your specified rate limit
Bridge ports	Receive ports for which you want to enable the rate limit. If you specify a subset of ports, you can specify multiple rate limit values. On the CoreBuilder® 9000, the list of ports includes the front-panel ports and any enabled backplane ports.	<ul style="list-style-type: none"> ■ Any subset of selectable ports ■ all ■ ? (for a list of selectable ports) 	Selectable ports
802.1p tag	IEEE 802.1p priority tag value to apply to forwarded frames (for both flow and nonflow classifiers).	<ul style="list-style-type: none"> ■ 0 – 7 ■ none ■ ? (for a list of selectable values) 	none (factory default)
Apply another rate limit?	If you specified a subset of available ports, whether you want to define another rate limit for other ports.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n
TCP drop control enabled (flow classifiers only)	Whether one-way filtering is used so that drop packets establish a TCP connection.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n
Start and end times	Whether you want to set starting and ending times for a control.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n

Prompt	Description	Possible Values	[Default]
Input time type	Type of time control that you want to establish. See Table 7 for a complete listing of input time type options.	<div><div>■</div>specific</div> <div><div>■</div>daily</div> <div><div>■</div>dayoftheweek</div> <div><div>■</div>everydayofthe-week</div> <div><div>■</div>weekdays</div> <div><div>■</div>weekends</div> <div><div>■</div>everyweekdays</div> <div><div>■</div>everyweekends</div>	specific
Classifiers to be controlled	Classifiers for this control to affect. See "qos control summary" for a list of defined classifiers that are associated with controls.	Selectable classifiers (that is, those not already associated with a control)	–

Table 7 lists the options for the input time types. The key to the prompts are:

- mm/dd = month–day
- hh:mm = hour:minute

Table 7 Input Time Type Options

Input Time Type	Options
Specific (default)	Starting day (mm–dd) Starting time (hh:mm) Ending day (mm–dd) Ending time (hh:mm)
Daily	Starting day (mm–dd) Starting time (hh:mm) Ending day (mm–dd) Ending time (hh:mm)
	<div> <div>i</div> <div>The Ending day and time cannot exceed 24 hours from the Starting day and time.</div> </div>
Day of the week	Starting day (Monday = 1, Tuesday = 2, Wednesday = 3, Thursday = 4, Friday = 5, Saturday = 6, Sunday = 7) Starting time (hh:mm) Ending day (Monday = 1, Tuesday = 2, Wednesday = 3, Thursday = 4, Friday = 5, Saturday = 6, Sunday = 7) Ending time (hh:mm)
	<div> <div>i</div> <div>The Ending day and time cannot exceed 24 hours from the Starting day and time.</div> </div>
Every day of the week	Starting day (Monday = 1, Tuesday = 2, Wednesday = 3, Thursday = 4, Friday = 5, Saturday = 6, Sunday = 7) Starting time (hh:mm) Ending day (mm–dd) Ending time (hh:mm)
	<div> <div>i</div> <div>The Ending day and time cannot exceed 24 hours from the Starting day and time.</div> </div>
Weekdays	Starting time (hh:mm) Ending time (hh:mm)
Weekends	Starting time (hh:mm) Ending time (hh:mm)

Input Time Type	Options
Every weekday	Starting time (hh:mm) Ending time (hh:mm)
Every weekend	Starting time (hh:mm) Ending time (hh:mm)

Procedure

- 1 Enter a control number.

The valid range is 5 through 50, with the next available number as the default.

- 2 Enter a control name.

- 3 Enter the rate limit type: `none`, `receivePort`, or `aggregate`.

The default is `none`. To drop all conforming packets for a set of ports, use `receivePort` or `aggregate`, set the rate limit to 0, and specify the appropriate set of ports.



You can apply aggregate rate limits only to flow classifiers.

- 4 For the `receivePort` or `aggregate` limit type, enter the service level for conforming packets as `high`, `best`, or `low`.

For the `none` rate limit type, enter the service level for conforming packets as `high`, `best`, `low`, or `drop`.

The default is `best` (best effort).



If you use `drop`, the system drops all traffic on all ports for the classifier that is associated with the control. Ping packets are ICMP, not UDP/TCP, so they are not dropped.

- 5 Specify whether the conforming packets are loss eligible (`yes` or `no`).

The default is `no`.

- 6 If you have selected `receivePort` or `aggregate` for the rate limit type, you are prompted for the following information:

- a Enter the service level for excess packets (`high`, `best`, `low`, or `drop`).

The default is `best`.

- b Specify whether excess packets are loss eligible (`yes` or `no`). The default is `yes`.

- c Specify how the rate limit is expressed (percentage of port bandwidth or `KBytes/sec`. `KBytes/sec` is the default).

- d** If you specified `KBytes/sec` for the rate limit, enter the value for the rate limit in KBytes/sec (0 through 65434).

If you specify that you want a `percentage` for the rate limit, specify the percentage in the range of from 0 through 100 percent. These numbers are rounded to the nearest 16 KBytes/sec. A value of 0 makes all packets excess packets.

- e** Enter the burst size in KBytes (16 through 8192, with the default value depending on your specified rate limit). The *burst size* is the maximum amount of data that you can transmit at the line rate before the transmission is policed.
- f** Specify the receive ports for which you want to enable the rate limit (specific bridge ports or all bridge ports).

If you apply the rate to only one or a subset of the bridge ports, you are prompted to specify whether you want to define another rate limit for another set of bridge ports. If you specify `yes`, you are prompted to enter another rate limit and burst size for another set of ports. This sequence of prompting continues until you specify `n`, meaning that you do not want to define another rate limit for another set of ports.



If the receive port is the anchor port for a trunk, the rate limit applies to each port that is associated with the trunk. For example, a rate limit of 1000 KBytes on a three-port trunk means that each port in the trunk has the 1000-KByte limit.

- 7** Enter an IEEE 802.1p tag value in the range of from 0 through 7 or `none` (the default) to apply to forwarded frames.
- 8** Specify whether drop packets used to establish a TCP connection (`yes`, `no`). The default is `no`.
- 9** Set the start and end time for the control (`yes`, `no`). The default is `no`.
 - a** If you specified a start and end time, enter the time type.
Time type selections are variations on days of the week and weekends or it can be specific day (or range of days) and time. See Table 7 for a complete listing of input time type options.
 - b** Enter the starting day and/or time.
 - c** Enter the ending day and/or time.

10 Enter the classifiers that are subject to this control.

The system displays the available classifiers in parentheses. If you select `aggregate` as the rate limit type, or if you said yes to the drop TCP connection packets option, only flow classifiers appear in parentheses.

QoS Control Define Example (3500)

This example shows a control for a nonflow classifier. Because the control has a rate limit of `none`, the system does not prompt you for information that applies to the other rate limit types.

```
Select menu option (qos/control): define
Enter control number {5-50|?} [5]:
Enter control name {?}: definetest
Enter rate limit type (none,receivePort,aggregate) [none]:
Enter service for conforming packets (high,best,low,drop)
[best]:
Are conforming packets loss eligible (yes,no) [no]:
Select IEEE 802.1p tag to apply to forwarded frames.
Tag {0-7|none|?} [none]:
Drop packets used to establish a TCP connection (yes,no)
[no]:
Set start and end time for the control (yes,no) [no]: yes
Enter input time type
(specific,daily,dayoftheweek,everydayoftheweek,weekdays,
weekends,everyweekdays,everyweekends) [specific]: weekdays
Enter the Qos control starting time (hh:mm): 09:00
Enter the Qos control ending time (hh:mm): 17:00
Select classifiers which are subject to this control.
Enter classifiers (20,23,420,430,440,450,4...: 450
```

qos control modify *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

✓ 3500

✓ 9000

9400

3900

9300

Modifies the characteristics of a previously defined control (including controls 1 through 4, which the system provides by default).

Valid Minimum Abbreviation

q co m

Important Considerations

- The software prompts you according to the rate limit type that you select.
- If the existing control has a rate limit type of `receivePort` or `aggregate` with multiple rate limits, you can now change one rate limit without affecting the other defined rate limits.

Options

Prompt	Description	Possible Values	[Default]
Control number	Number of the control that you want to modify. Existing control numbers appear in braces. Control numbers 1-4 are predefined.	5 – 50	No default, but the next sequential number is automatically entered
Control name	Name of the control that you want to modify.	<ul style="list-style-type: none"> ■ Unique name with up to 32 characters (Use " around any string with embedded spaces. Use " " to enter an empty string.) ■ ? (for a list of selection criteria) 	Current name for specified control

Prompt	Description	Possible Values	[Default]
Rate limit type	Type of rate limit: <ul style="list-style-type: none"> ■ none (no rate limit) ■ receivePort (a rate limit on the specified ports) ■ aggregate (the bandwidth for all ports specified for the associated classifier) 	<ul style="list-style-type: none"> ■ none ■ receivePort ■ aggregate 	Current rate limit type
Service level	Service level for the conforming packets. Drop causes the system to drop all traffic on all ports that are associated with the classifier/control.	<ul style="list-style-type: none"> ■ For a rate limit of receivePort or aggregate: high, best (best effort), or low ■ For a rate limit of none: high, best, low, or drop 	Current service level
Loss eligible	Whether conforming packets are loss-eligible. A loss-eligible packet can be dropped if the transmit queue for which it is destined exceeds its threshold.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	Current value
Excess packets service	For receivePort or aggregate rate limit types, service level for excess packet.	<ul style="list-style-type: none"> ■ high ■ best ■ low ■ drop 	Current value
Excess loss eligible	For receivePort or aggregate rate limit types, whether excess packets are loss-eligible.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	Current value
How rate limit is expressed	For receivePort or aggregate rate limit types, format of the rate limit.	<ul style="list-style-type: none"> ■ KBytes/sec ■ percentage 	KBytes/sec
Rate limit value	For receivePort or aggregate rate limit types, number of Kbytes/sec or a percentage. 0 makes all packets excess packets.	<ul style="list-style-type: none"> ■ 0 – 65434 KBytes/sec ■ 0 – 100 percent 	–

Prompt	Description	Possible Values	[Default]
Burst size	For receivePort or aggregate rate limit types, maximum amount of data (in Kbytes) that you can transmit at the line rate before the transmission is policed.	16 – 8192 KBytes	Determined by your specified rate limit
Bridge ports	For receivePort or aggregate rate limit types, the receive ports for which you want to enable the rate limit.	<ul style="list-style-type: none"> ■ One or more selectable ports ■ all ■ ? (for a list of selectable ports) 	Current bridge ports
802.1p tag	IEEE 802.1p priority tag value that you want to apply to forwarded frames (for flow or nonflow classifiers).	<ul style="list-style-type: none"> ■ 0 – 7 ■ none ■ ? (for a list of selectable values) 	Current value
TCP drop control enabled (flow classifiers only)	Whether one-way filtering is used so that drop packets establish a TCP connection.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n
Start and end times	Whether you want to set starting and ending times for a control.	<ul style="list-style-type: none"> ■ y (yes) ■ n (no) 	n
Input time type	Type of time control that you want to establish. See Table 7 for a complete listing of input time type options.	<ul style="list-style-type: none"> ■ specific ■ daily ■ dayoftheweek ■ everydayoftheweek ■ weekdays ■ weekends ■ everyweekdays ■ everyweekends 	specific
Classifiers controlled	Classifiers that this control affects. See “qos control summary” for a list of defined classifiers associated with controls.	Selectable classifiers (that is, those not already associated with a control)	Current classifier or classifiers for the control

Procedure

- 1 Enter the control number that you want to modify. The existing controls are displayed in braces { }.
- 2 To modify the name, enter the new name for the classifier.
The name that is associated with the specified control number appears in brackets [].
- 3 Enter the rate limit type (for example, `none`, `receivePort`, or `aggregate`).
The available values depend on how the control was defined; the current limit appears in brackets.
- 4 For the `receivePort` or `aggregate` rate limits, enter the service level for conforming packets as `high`, `best`, or `low`.
For the `none` rate limit, enter the service level for conforming packets as `high`, `best`, `low`, or `drop`. If you use `drop`, the system drops all traffic on all ports for the classifier that is associated with the control. The current value appears in brackets.
- 5 Specify whether the conforming packets are loss eligible (`yes` or `no`).
- 6 If you have selected `receivePort` or `aggregate` for the rate limit type, you are prompted for the following information:
 - a Enter the service level for excess packets (`high`, `best`, `low`, or `drop`).
 - b Specify whether excess packets are loss eligible (`yes` or `no`). Your current value is the default.
 - c Specify whether you want to modify the existing rate limits (`yes` or `no`).
If you enter `no`, the system maintains the existing values for all associated rate limits. If you enter `yes`, specify how the first rate limit should be expressed (percentage of port bandwidth or KBytes/sec). KBytes/sec is the default. If the control has multiple per-port rate limits, you can change one rate limit without affecting the others.
 - d If you specified KBytes/sec for the first (or only) rate limit, enter the value for the rate limit in KBytes/sec (0 through 65434).
If you specified percentage for the rate limit, specify the percentage in the range of from 0 through 100 percent.
 - e Enter the burst size in KBytes (in the range of from 16 through 8192). The default value depends on your specified rate limit.

- f** Specify the bridge ports for which you want to enable the new rate limit (for example, 1-13, or all).

If you modify the rate limit and apply it to only one or a subset of the bridge ports, you are prompted to specify whether you want to modify or define another rate limit for another set of bridge ports. If you specify *yes*, you are prompted to enter another rate limit and burst size. This sequence of prompting continues until you specify *n*, meaning that you do not want to modify or define another rate limit for another set of ports. The rate limit applies only to those ports that you explicitly specified; any ports that you did not specify are not associated with your rate limit.

- 7** Select an IEEE 802.1p tag value in the range of from 0 through 7 or the value *none* to apply to forwarded frames.
- 8** Specify whether drop packets are used to establish a TCP connection (*yes*, *no*). The default is *no*.
- 9** Set the start and end time for the control (*yes*, *no*). The default is *no*.
 - a** If you specified a start and end time, enter the time type.

Time type selections are variations on days of the week and weekends or it can be specific day (or range of days) and time. See Table 7 for a complete listing of input time type options.
 - b** Enter the starting day and/or time.
 - c** Enter the ending day and/or time.
- 10** Enter the classifiers that are subject to this control. The system displays the associated classifiers in brackets. (If you select *aggregate* as the rate limit type, or select the drop packets use to establish a TCP connection option, the system displays only flow classifiers.)

QoS Control Modify Example (3500)

This example shows modifications to a predefined control (4) for a predefined classifier (405).

```
Select menu option: qos control modify
Enter control number {1-5}: 4
Enter control name {?} [Controlled Load]:
Interactive_Multimedia
Enter rate limit type (none,receivePort,aggregate) [none]:
receivePort
Enter service for conforming packets (high,best,low) [high]:
Are conforming packets loss eligible (yes,no) [no]:
Enter service for excess packets (high,best,low,drop) [low]:
drop
How should rate limit be expressed (percentage,KBytes/sec)
[KBytes/sec]:
Enter rate limit in KBytes/sec (0-65434): 2048
Enter burst size in KBytes (16-8192) [181]:
Select bridge ports (1-13|all|?) [1-13]:
Select IEEE 802.1p tag to apply to forwarded frames.
Enter IEEE 802.1p tag {0-7|none|?} [none]:
Drop packets used to establish a TCP connection (yes,no) [no]:
Do you want to modify/add the start and end time for the control (yes,no) [no]
y
Do you want to have any time control (yes,no) [no]: y
Enter input time type (specific,daily,dayoftheweek,everydayoftheweek,weekdays,
weekends,everyweekdays,everyweekends) [specific]:
Enter the Qos Control starting day (mm-dd): 06-02
Enter the Qos control starting time (hh:mm): 09:00
Enter the Qos Control ending day (mm-dd): 06-02
Enter the Qos control ending time (hh:mm): 17:00
Select classifiers which are subject to this control.
Enter classifiers (20,23,404-407,420,430,4...[404-407]): 405
```

qos control remove *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Removes a previously defined control.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q co r

Important Consideration

- When you remove a control, the associated classifiers are no longer controlled and no longer have a set rate limit, service level, or 802.1p tag.

Options

Prompt	Description	Possible Values	[Default]
Control number	Number for the control that you want to remove	<ul style="list-style-type: none"> ■ One or more selectable control numbers ■ ? (for a list of the selectable controls) 	–

QoS Control Remove Example (9000 Layer 3)

```
CB9000@slot2.1 [12-E/FEN-TX-L3] (qos/control): remove
Enter control number {2-5|?}: 5
```

qos ldap display

Displays Lightweight Directory Access Protocol (LDAP) status information.

✓ 3500

9000

9400

Valid Minimum Abbreviation

q 1 disp

Important Considerations

- When LDAP is enabled, displays server IP address and polling period.
- When LDAP is disabled, displays QOS, Resource Reservation Protocol (RSVP), and LDAP status.

Fields in the QoS LDAP Display

Field	Description
LDAP server address	The IP address of the LDAP server
Poll period	Selected poll period

qos ldap enable

Enables QoS parameter directory services which are located on the Lightweight Directory Access Protocol (LDAP) server.

✓ 3500
9000
9400

Valid Minimum Abbreviation

q l e

3900
9300

Important Considerations

- An LDAP server must be configured.
- Before you enable LDAP, the LDAP server must have a directory group configured with QoS parameters in an *ldif* file.
- Parameter changes for a specific group may affect more than one system. If you know that a change will affect more than one system, disable LDAP to test the change. After you are sure you want the change, you can then enable LDAP.

Options

Prompt	Description	Possible Values	[Default]
Enable	Connects your system to the LDAP server	–	Disabled
Poll period		600 – 2000	–
LDAP server address	The IP address of the LDAP server you have configured	–	–
LDAP group name	Name of an LDAP entry on the LDAP server that indexes other entries containing QoS classifier and control information.	–	Wildcard

qos ldap disable

Disables QoS parameter directory services, which are located on the Lightweight Directory Access Protocol (LDAP) server.

✓ 3500
9000
9400

Valid Minimum Abbreviation

q l disa

Important Considerations

- By default, LDAP is disabled.
- If LDAP is disabled, you do not receive automatic updates.

Options

Prompt	Description	Possible Values	[Default]
Disabled	Removes the connection to the LDAP server	–	Disabled

qos rsvp summary***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

Displays summary Resource Reservation Protocol (RSVP) information when RSVP is enabled.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation**q r s****Fields in the QoS RSVP Summary Display**

Field	Description
Excess loss eligible	Whether excess packets are loss-eligible.
Excess service	Service level for excess/policed traffic (<i>best</i> or <i>low</i>).
Per resv bandwidth	Largest reservation that RSVP attempts to install.
Policing option	When to drop excess packets. <i>Edge policing</i> causes excess packets to be dropped only at the edge (that is, when the traffic has not yet passed through any network device that has already performed policing for that flow). Options are <i>edge</i> , <i>always</i> , or <i>never</i> .
Total resv bandwidth	Admission control policy. RSVP begins to refuse reservations when the requested bandwidth on an output link exceeds the total reservable bandwidth.

qos rsvp detail **For CoreBuilder 9000: Applies to Layer 3 switching modules only.**

Displays detailed RSVP information when RSVP is enabled.

Valid Minimum Abbreviation

q r de

Important Consideration

- If no flows are installed on the system or on a Layer 3 module, the command displays only the summary information.

Options

Prompt	Description	Possible Values	[Default]
Level of RSVP information (when flows are installed)	If RSVP flows are available to report, the amount of RSVP information you want	<ul style="list-style-type: none"> ■ all ■ session ■ IP 	–

Fields in the QoS RSVP Detail Display

Field	Description
Excess loss eligible	Whether excess packets are loss-eligible.
Excess service	Service level for excess/policed traffic (best or low).
Per resv bandwidth	Largest reservation that RSVP attempts to install.
Policing option	When to drop excess packets. <i>Edge policing</i> causes excess packets to be dropped only at the edge (that is, when the traffic has not yet passed through any network device that has already performed policing for that flow).
Session	Session numbers, destination IP addresses and ports, protocols, number of senders, receivers, and RSVP reservations.
Session – receiver and session reservation	Port numbers, an RSVP style (ST) of fixed filter (FF), shared explicit (SE), or wildcard filter (WF), next hop addresses, LIH values, TTD values, bandwidth values, burst values, and filters.
Session – sender	Port numbers, source IP addresses, previous hop addresses, Logical Interface Handle (LIH) values, Time To Die (TTD) values, bandwidth values, burst size values, and output ports.
Session – installed flows	Actual flow that was installed on the system (shown in the last portion of the output).
Total resv bandwidth	Admission control policy. RSVP begins to refuse reservations when the requested bandwidth on an output link exceeds the total reservable bandwidth.

qos rsvp enable *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Enables RSVP on the system RSVP settings that you specify.

✓ 3500

✓ 9000

9400

3900

9300

Valid Minimum Abbreviation

q r e

Important Considerations

- By default, RSVP is disabled.
- In general, when you enable RSVP, use the default settings.
- You are allowing RSVP to reserve this amount of bandwidth in the system. You can oversubscribe (over 100) and specify a value of up to 200.

Options

Prompt	Description	Possible Values	[Default]
Maximum total reservable bandwidth	Admission control policy. RSVP begins to refuse reservations when the requested bandwidth on an output link exceeds the total reservable bandwidth.	0 – 200 percent	50 (factory default)
Maximum per-reservation bandwidth	Largest reservation that RSVP attempts to install.	0 – 100 percent	50 (factory default)
Policing option	When to drop excess packets. <i>Edge policing</i> drops excess packets only at the edge (that is, when traffic has not yet passed through any network device that has already performed policing for that flow). <ul style="list-style-type: none"> ■ With <i>edge</i>, the system polices the flow when RSVP requests it. ■ With <i>always</i>, the system polices the flow regardless of whether RSVP requests it. ■ With <i>never</i>, the system never polices the flow even if RSVP requests it. 	<ul style="list-style-type: none"> ■ edge ■ always ■ never 	edge (factory default)

Prompt	Description	Possible Values	[Default]
Service level for excess /policed traffic	Service level for excess/policed traffic. Low is recommended. This setting applies to the excess traffic with the reserved bandwidth (that is, which queue it should be placed in).	<div><div>■</div> best</div> <div><div>■</div> low</div>	low (factory default)
Excess Loss Eligible	Whether excess packets are loss-eligible	<div><div>■</div> yes</div> <div><div>■</div> no</div>	no (factory default)

Procedure

- 1
- Enter the maximum total reservable bandwidth, using a percentage of the output link (a value of from 0 through 200, with 50 as the default).
- 2
- Enter the maximum per-reservation bandwidth, using a percentage of the output link (a value of from 0 through 100, with 50 as the default).
- 3
- Enter the policing option (edge, always, or never, with edge as the default).
- 4
- Enter the service level for excess/policed traffic (best or low, with low as the default).
- 5
- Specify whether excess packets are loss eligible (yes or no, with no as the default).

qos rsvp disable *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Disables RSVP on the system.

✓ 3500

✓ 9000

9400

Valid Minimum Abbreviation

q r di

3900

9300

Important Considerations

- By default, RSVP is disabled.
- This command does not verify that RSVP has been disabled.

qos bandwidth display

- ✓ 3500
- ✓ 9000
- 9400
- 3900
- 9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays the link bandwidth as the ratio of bandwidth that is allocated to high priority traffic versus best effort traffic. Link bandwidth is the total link bandwidth less the bandwidth that RSVP and network control traffic use.

Valid Minimum Abbreviation

q b d

Important Consideration

- By default, 75 percent of bandwidth is allocated to high-priority traffic.

**qos bandwidth
modify**

✓ 3500
✓ 9000
9400

3900
9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Sets how to weigh the high priority and best effort transmit queues, and sets RSVP bandwidth for the control queue. Low priority packets do not have bandwidth explicitly allocated.

Valid Minimum Abbreviation

q b m

Important Considerations

- When you enter the command, the system prompts you to enter the percentage of bandwidth to use for high-priority traffic on the output link.
- The value 75 specifies that three high-priority packets are transmitted for each best effort packet.
- The value 50 sets equal priority for high priority and best effort packets.
- The value 100 is strict prioritization; it allows best effort packets to be sent only when no high priority packets need to be sent.

Options

Prompt	Description	Possible Values	[Default]
Percentage of bandwidth	Percentage of bandwidth that you want to be used for high-priority traffic on the output link	0 – 100 percent	75

qos excessTagging
display

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Displays status information about whether excess packets are tagged with a special IEEE 802.1p tag value.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

q e disp

3900
9300

**qos excessTagging
enable**

✓ 3500

✓ 9000

9400

3900

9300

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Tags or retags excess packets with a special 802.1p tag value. This special value refers to any packets that are marked as excess that you want to tag.

Valid Minimum Abbreviation

q e e

Important Considerations

- Excess tagging is disabled by default.
- When you enter this command, you are prompted to enter an IEEE 802.1p tag value for excess packets in the range of 0 through 7, with 0 as the default. For example, if you specify 1, excess packets become background traffic.

Options

Prompt	Description	Possible Values	[Default]
IEEE 802.1p tag value	Tag value to use to tag or retag excess packets	0 – 7	0

qos excessTagging
disable

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

Disables the tagging of excess packets with a special 802.1p tag value.

✓ 3500
✓ 9000
9400

Valid Minimum Abbreviation

q e disa

Important Consideration

- Excess tagging is disabled by default.

3900
9300

qos statistics interval *For CoreBuilder 9000: Applies to Layer 3 switching modules only.*

Sets a sampling interval for gathering QoS statistics.

- ✓ 3500
- ✓ 9000
- 9400

Valid Minimum Abbreviation

q s i

- 3900
- 9300

Important Considerations

- The default interval is 5 seconds.
- When you enter this command, the system prompts you to enter the appropriate interval. The existing value appears in brackets.
- A nonzero value shows the byte or packet-count-per-interval period. A zero value shows byte or packet counters.

Options

Prompt	Description	Possible Values	[Default]
Interval	Interval, in seconds, during which you want to gather QoS statistics	0 – 60 seconds	5 (factory default), or current value

qos statistics receive ***For CoreBuilder 9000: Applies to Layer 3 switching modules only.***

✓ 3500
✓ 9000
9400

Displays QoS receive statistics.

Valid Minimum Abbreviation

q s r

3900
9300

Important Considerations

- The system displays the statistics at the interval that you specified. The default interval is 5 seconds.
- The receive statistics shows the effect of the traffic control services that you configured.

Options

Prompt	Description	Possible Values	[Default]
Bridge ports	Port numbers whose receive statistics you want to display. On the CoreBuilder® 9000, the list of ports includes the front-panel ports and any enabled backplane ports.	<ul style="list-style-type: none">■ One or more port numbers■ all■ ? (for a list of selectable ports)	–

Fields in the QoS Receive Statistics Display

Field	Description
droppedPackets	Number of packets that were dropped when they were received
droppedPacketsPeak	Highest number of packets that were dropped on receipt up to this point
flowExcess	Number of flow classifier bytes that are excess
flowExcessPeak	Highest number of flow excess bytes that have been received up to this point
flowReserved	Number of conforming flow classifier bytes that have been received
flowReservedPeak	Highest number of flow classifier bytes that have been received up to this point
nonFlowExcess	Number of nonflow classifier bytes that have been received that are excess
nonFlowExcessPeak	Highest number of nonflow excess bytes that have been received up to this point

Field	Description
nonFlowReserved	Number of conforming non-flow classifier bytes that have been received
nonFlowResvPeak	Peak count: The highest number of conforming nonflow classifier bytes that have been received up to this point
port	If you display statistics for multiple ports, the port number that is associated with the statistics

qos statistics transmit

For CoreBuilder 9000: Applies to Layer 3 switching modules only.

✓ 3500

✓ 9000

9400

3900

9300

Displays QoS transmit statistics.

Valid Minimum Abbreviation

q s t

Important Considerations

- The transmit statistics help you track bandwidth utilization and packet loss by physical port and queue (reserved, high, best, and low).
- The RSVP and network control packets go out on the reserved queue.
- When you mark any packet (conforming or excess) as loss eligible, the packet is dropped if the transmit queue for which it is destined is over its threshold. A packet that is marked loss-eligible falls into one of the two highLoss statistic categories:
 - If the transmit queue is not over its threshold, the packet is sent and counted as a highLossSent packet.
 - If the transmit queue is over its threshold, it is dropped and counted as a highLossDropped packet.
- If you do *not* mark a packet as loss-eligible, it falls into one of the three lowLoss statistics.
 - If the queue is not over the threshold, it is counted as a lowLossSent.
 - If the queue is over its threshold, it is counted as lowLossDelayed.
 - If the queue is full, it is counted as lowLossDropped.
- *Loss-eligible packets* are conforming packets that are discarded instead of queued when transmit queues back up beyond a threshold. You can specify whether conforming packets (as well as nonconforming excess packets) are loss-eligible when you define a control. Marking packets loss-eligible is useful to enable intelligent discard of traffic in a congestion situation. When the system is congested, you can decide which traffic can be discarded and mark that traffic as loss eligible.

Options

Prompt	Description	Possible Values	[Default]
Bridge ports	Port numbers of ports for which you want to display transmit statistics. On the CoreBuilder® 9000, the list of ports includes the front-panel ports and any enabled backplane ports.	<ul style="list-style-type: none"> ■ One or more port numbers ■ all ■ ? (for a list of selectable ports) 	–
Queues	Transmit queues (types of service) whose statistics you want to display.	<ul style="list-style-type: none"> ■ reserved ■ high ■ best ■ low ■ all ■ ? (for a list of selectable values) 	–

Fields in the QoS Transmit Statistics Display


Field	Description
highLossDropped	Number of loss-eligible packets that were discarded and were over the threshold
highLossDroppedPeak	Current highest count of loss-eligible packets that were discarded and were over the threshold
highLossSent	Number of loss-eligible packets that were sent and were under the threshold (at low latency)
highLossSentPeak	Current highest count of loss-eligible packets that were sent and were under the threshold
lowLossDelayed	Number of non-loss-eligible packets that were sent and over the threshold (that is, the transmit queues were backing up but not overflowing)
lowLossDelayedPeak	Current highest count of non-loss-eligible packets that were sent and were over the threshold
lowLossDropped	Number of packets that were discarded because they exceeded the length of the transmit queue
lowLossDroppedPeak	Current highest count of packets that were discarded because they exceeded the length of the transmit queue
lowLossSent	Number of non-loss-eligible packets that were sent and were under the threshold (at low latency)



MONITORING

Chapter 23 Event Log

Chapter 24 Roving Analysis



EVENT LOG

This chapter provides guidelines and other key information about how to administer event logs in your system, including the following tasks:

- Display the event log configuration
- Configure the output devices
- Configure the services

Use event logging to capture different types of log messages from various services (applications) and send them to the Administration Console. The log messages display real-time information about the state of the system or a specific service, and can help you diagnose site-specific problems.



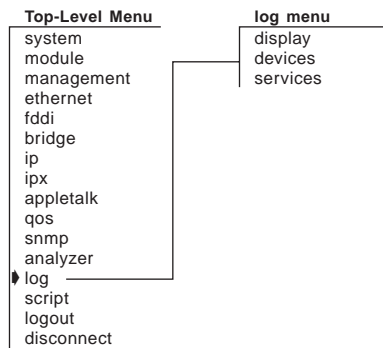
On CoreBuilder® 9000 systems, event logging is controlled entirely through the Enterprise Management Engine (EME), not through the Administration Consoles of individual modules as described here. See the CoreBuilder 9000 Enterprise Management Engine User Guide for information on how to keep logs of switch events.



For more information about implementing event logging on your network, see the CoreBuilder 3500 Implementation Guide.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



log display Displays the current log settings.

✓ 3500
9000
9400

Valid Minimum Abbreviation

log di

Important Consideration

- The CoreBuilder 3500 by default enables logging to the serial port session and disables logging to any Telnet or modem session. However, you can toggle the current logging state on the CoreBuilder 3500 from serial port to Telnet or modem by entering Ctrl+L.

Fields in the Log Display

Field	Description
consoleOut	Administration Console output device. You can enable or disable the Console to display event log messages for each severity level.
Logging message	Whether logging to this console session is enabled or disabled.
Supported Event Log Services	
AppleTalk	Appletalk log service. Enabled or disabled for each severity level.
IPX	IPX log service. Enabled or disabled for each severity level.
System	System log service. Enabled or disabled for each severity level.
Severity Levels	
Config	Configuration changes.
Error	Application-specific error. Default: enabled
Info	Severity level of changes in the state of the system that are not caused by events at any other severity level
Warning	Nonfatal problem. Default: enabled

✓ 3500
9000
9400

3900
9300

log devices Configures severity levels for event logging on the Administration Console.

Valid Minimum Abbreviation

log de

Important Considerations

- You can set the console to log events for one or more of the four severity levels.
- To specify multiple severity levels, separate the levels with a comma (for example, warning,config).

Options

Prompt	Description	Possible Values	[Default]
Levels for console	Event logging severity level for console output	<ul style="list-style-type: none">■ error■ warning■ config■ info■ all■ ? (for a list of valid severity levels)	–
Selected levels	Whether selected event logging for console output is enabled or disabled	<ul style="list-style-type: none">■ y (yes)■ n (no)	y

Log Devices Examples (3500)

```
Select menu option (log): devices
Select levels for console (error,warning,config,info|all|?): ?

Selectable values

error,warning,config,info

Select levels for console (error,warning,config,info|all|?): all
Enable the selected levels (n,y) [y]: y
```

To disable the config and info severity levels:

```
Select menu option (log): devices
```

```
Select levels for console (error,warning,config,info|all|?): config,info
```

```
Enable the selected levels (n,y) [y]: n
```

The display now indicates that the error and warning severity levels remain enabled and the config and info levels are disabled.

✓ 3500
9000
9400

3900
9300

log services

Enables the logging of messages that pertain to the following services:

- System level
- AppleTalk
- IPX

Valid Minimum Abbreviation

log s

Important Considerations

- For a specific service or all services, you can configure up to four severity levels.
- Use a comma to separate multiple service names and severity levels (for example, `system, appletalk` and `error, warning`).

Options

Prompt	Description	Possible Values	[Default]
Services	Services to configure	<ul style="list-style-type: none">■ system■ ipx■ appletalk■ all■ ? (for a list of valid services to configure)	–
Levels	Severity levels to enable	<ul style="list-style-type: none">■ error■ warning■ config■ info■ all■ ? (for a list of valid severity levels to enable)	–
Selected services/levels	Whether the selected services and severity levels are enabled or disabled	<ul style="list-style-type: none">■ y (yes)■ n (no)	y

Log Services Examples

To enable all severity levels for the AppleTalk service:

```
Select menu option (log): services
```

```
Select services (system,ipx,appletalk|all|?): ?
```

```
Selectable values
```

```
system,ipx,appletalk
```

```
Select services (system,ipx,appletalk|all|?): appletalk
```

```
Select levels (error,warning,config,info|all|?): all
```

```
Enable the selected services/levels (n,y) [y]: y
```

To show that all severity levels are enabled for the AppleTalk service, enter

```
log display
```

To disable the warning and info severity levels for the AppleTalk service, follow this example:

```
Select menu option (log): services
```

```
Select services (system,ipx,appletalk|all|?): appletalk
```

```
Select levels (error,warning,config,info|all|?): warning,info
```

```
Enable the selected services/levels (n,y) [y]: n
```

To show that the AppleTalk service is associated with only the error and config severity levels, enter **log display**

ROVING ANALYSIS

This chapter provides guidelines and other key information about how to set up roving analysis in your system, including the following tasks:

- Display roving analysis configuration
- Add and remove analyzer
- Start and stop monitoring

Roving analysis is the mirroring of traffic on one port to another port of the same media type.

- The port being monitored is called the *monitor port*.
- The port that receives the mirrored traffic is called the *analyzer port*.

The analyzer port typically has a network analyzer or RMON *probe* attached through which you can watch the network traffic.

Use roving analysis to monitor Fast Ethernet, Gigabit Ethernet, or Fiber Distributed Data Interface (FDDI) port traffic for network management and troubleshooting purposes. You use the Administration Console to choose any network segment that is attached to a system and monitor its activity.

You can monitor a designated roving analysis port to:

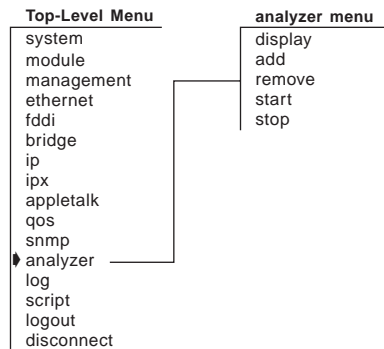
- Analyze traffic loads on each segment so that you can continually optimize your network loads by moving network segments
- Troubleshoot network problems (for example, to find out why a particular segment has so much traffic)



For more information about implementing roving analysis on your network, see the Implementation Guide for your system.

Menu Structure

The commands that you can use depend on the system that you have, your level of access, and the types of modules and other hardware options that are configured for your system. The following diagram shows the complete list of commands for all systems. See the checklist at the beginning of each command description in this chapter for whether your system supports the command.



analyzer display

Displays the roving analysis configuration, showing which ports are designated as analyzer ports and which bridge ports are currently being monitored.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

an d

Fields in the Analyzer Display

Field	Description
Ports configured as analyzer ports	List of analyzer ports on the system, including the port number and MAC address. These are the ports that can accept traffic that is mirrored from a monitored port. Analyzer ports are typically connected to a network analyzer or probe. There may be multiple analyzer ports defined on the switch.
Port	Analyzer port number
Type	Media type and Port Speed (FDDI, Fast Ethernet, or Gigabit Ethernet)
Address	MAC address of the analyzer port
Ports being monitored	List of ports that the system is monitoring. Includes the MAC address of the analyzer port to which the monitored port traffic will be forwarded.
Port	Monitored port number
Type	Media type and Port Speed (FDDI, Fast Ethernet, or Gigabit Ethernet)
Analyzer Address	MAC address of the analyzer port to which the monitored port traffic will be forwarded and to which your network analyzer or probe is attached. There may be multiple analyzer ports defined on the switch.

Analyzer Display Example (3500)

```
Select menu option (analyzer): display
Ports configured as analyzer ports:
      Port      Type      Address
      8      Fast Ethernet  00-80-3e-2b-42-08

Ports being monitored:
      Port      Type      Analyzer Address
      12      Fast Ethernet  00-80-3e-2b-42-08
```

analyzer add Defines a bridge port to serve as a dedicated analyzer port.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

an a

Important Considerations

- On CoreBuilder® 3500 and CoreBuilder 9000 systems, you can connect as many as 16 network analyzers to a system. On other platforms, you can connect one network analyzer. For more accurate analysis, attach the analyzer to a dedicated port instead of through a repeater.
- After a port is selected to serve as an analyzer port, it cannot receive or transmit any other data. Instead, it receives only the data from the ports to be monitored. If you have enabled the Spanning Tree Protocol (STP) on the port, STP is automatically disabled.
- If the physical port configuration changes in the system (that is, if you remove or rearrange modules), the MAC address of the analyzer port remains fixed. If you replace the module with the analyzer port with a module of a different media type, the roving analysis port (RAP) configuration for that port is cleared.
- When you configure a port that is part of a virtual LAN (VLAN) as an analyzer port, a warning is displayed because adding the port removes the port from all VLANs. When the port is restored (when you remove the analyzer port), it becomes a member of the default VLAN.
- If the probe is attached to a 10 Mbps Ethernet analyzer port and the roving analysis port (RAP) is monitoring a 100 Mbps Ethernet port with a sustained traffic rate greater than 10 Mbps, the analyzer may not see all of the frames.
- After you enter a bridge port number, the system displays the MAC address of the analyzer port. Record this information for setting up the port that you want to monitor.
- On the CoreBuilder 9000, the port to which the analyzer is attached and the port you wish to monitor must be on the same blade.
- Trunked ports and resilient link ports can not be configured as analyzer ports.

analyzer remove

Restores the port to be a regular bridge port. Restores the Spanning Tree state to its state before the port was configured as an analyzer port.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

an r

Important Considerations

- Use this command when you no longer need the bridge port for the analyzer.
- The analyzer port can not be removed if it still has monitor ports.
- The port becomes a member of the default virtual LAN (VLAN) when it is restored (when you remove it as an analyzer port).
- The port will not be automatically restored to any VLAN it might have been a member of before it was configured as an analyzer port — you must do this yourself.

Options

Prompt	Description	Possible Values	[Default]
Bridge port	Number of the bridge port to which the analyzer is attached n varies by platform. Only active analyzer port numbers are displayed.	<ul style="list-style-type: none"> ■ 1 – n ■ ? (for a list of available bridge ports) 	–

Analyzer Remove Example (3500)

Select menu option (analyzer): **remove**

Select bridge port {2,7|?}: **7**

analyzer start Starts port monitoring activity on the selected bridge port.

✓ 3500

✓ 9000

✓ 9400

✓ 3900

✓ 9300

Valid Minimum Abbreviation

an sta

Important Considerations

- You must already have an analyzer port configured. First designate a bridge port to serve as the analyzer port and connect the analyzer to that port. See “analyzer add” earlier in this chapter for details.
- On the CoreBuilder 9000, the analyzer port and the monitor port must be on the same module.
- The MAC address of the analyzer port is displayed when you configure that port, and when you display the roving analysis configurations on the system to which the analyzer is attached.
- The media type of the analyzer port must match the media type of the port being monitored. Fast Ethernet and Gigabit Ethernet are the same media type.
- You can use a Fast Ethernet (10 Mbps) port to monitor a Gigabit Ethernet (100 Mbps) port, but a warning message will be printed. If the sustained traffic load is greater than 10 Mbps, the analyzer on the slower port may not see all the frames on the faster port.
- When you successfully configure a bridge port to be monitored, all the data that the monitored port receives and transmits is copied to the selected analyzer port.
- Once a port is selected to serve as a monitor port, the RMON data that it can record is limited to the RMON groups (statistics, history, alarm, event, protocolDir, and probeConfig) that do not require hardware sampling.
- If you replace the module that the monitored port resides on with a module of a different media type, the roving analysis port (RAP) configuration for the monitored port is reset.

Options

Prompt	Description	Possible Values	[Default]
Bridge port	Number of the bridge port to be monitored n varies by platform.	■ 1 – n ■ ? (for a list of available bridge ports)	–
Target analyzer port address	MAC address of the port to which the analyzer is attached	A valid MAC address of an analyzer port	–

Analyzer Start Example (9000 1000BASE-SX module)

```
CB9000@slot 3.1 [9-GEN-SX-L2] (analyzer): start  
Select bridge port {1-8,10-12|?}: 1  
Enter the target analyzer port address: 00-20-9c-0d-e1-2a
```

analyzer stop Stops port monitoring activity on the selected bridge port.

- ✓ 3500
- ✓ 9000
- ✓ 9400

- ✓ 3900
- ✓ 9300

Valid Minimum Abbreviation

an sto

Important Consideration


- Port data is no longer copied and forwarded to the selected analyzer port from the port that you specify. See “analyzer start” earlier in this chapter for details.

Options

Prompt	Description	Possible Values	[Default]
Bridge port	Number of the bridge port that is being monitored n varies by platform.	<ul style="list-style-type: none">■ 1 – n■ ? (for a list of available bridge ports)	–


Analyzer Stop Example (3500)


Select menu option (analyzer): **stop**
Select bridge port {3,4|?}: **3**



REFERENCE

Appendix A Technical Support





TECHNICAL SUPPORT

3Com provides easy access to technical support information through a variety of services. This appendix describes these services.

Information contained in this appendix is correct at time of publication. For the most recent information, 3Com recommends that you access the 3Com Corporation World Wide Web site.

Online Technical Services

3Com offers worldwide product support 24 hours a day, 7 days a week, through the following online systems:

- World Wide Web site
- 3Com Knowledgebase Web Services
- 3Com FTP site
- 3Com Bulletin Board Service (3Com BBS)
- 3Com FactsSM Automated Fax Service

World Wide Web Site

To access the latest networking information on the 3Com Corporation World Wide Web site, enter this URL into your Internet browser:

<http://www.3com.com/>

This service provides access to online support information such as technical documentation and software, as well as support options that range from technical education to maintenance and professional services.

3Com Knowledgebase Web Services

This interactive tool contains technical product information compiled by 3Com expert technical engineers around the globe. Located on the World Wide Web at **<http://knowledgebase.3com.com>**, this service gives all 3Com customers and partners complementary, round-the-clock access to technical information on most 3Com products.

3Com FTP Site Download drivers, patches, software, and MIBs across the Internet from the 3Com public FTP site. This service is available 24 hours a day, 7 days a week.

To connect to the 3Com FTP site, enter the following information into your FTP client:

- Hostname: **ftp.3com.com**
- Username: **anonymous**
- Password: **<your Internet e-mail address>**



You do not need a user name and password with Web browser software such as Netscape Navigator and Internet Explorer.

3Com Bulletin Board Service

The 3Com BBS contains patches, software, and drivers for 3Com products. This service is available through analog modem or digital modem (ISDN) 24 hours a day, 7 days a week.

Access by Analog Modem

To reach the service by modem, set your modem to 8 data bits, no parity, and 1 stop bit. Call the telephone number nearest you:

Country	Data Rate	Telephone Number
Australia	Up to 14,400 bps	61 2 9955 2073
Brazil	Up to 28,800 bps	55 11 5181 9666
France	Up to 14,400 bps	33 1 6986 6954
Germany	Up to 28,800 bps	4989 62732 188
Hong Kong	Up to 14,400 bps	852 2537 5601
Italy	Up to 14,400 bps	39 2 27300680
Japan	Up to 14,400 bps	81 3 5977 7977
Mexico	Up to 28,800 bps	52 5 520 7835
P.R. of China	Up to 14,400 bps	86 10 684 92351
Taiwan, R.O.C.	Up to 14,400 bps	886 2 377 5840
U.K.	Up to 28,800 bps	44 1442 438278
U.S.A.	Up to 53,333 bps	1 847 262 6000

Access by Digital Modem

ISDN users can dial in to the 3Com BBS using a digital modem for fast access up to 64 Kbps. To access the 3Com BBS using ISDN, call the following number:

1 847 262 6000

3Com Facts Automated Fax Service

The 3Com Facts automated fax service provides technical articles, diagrams, and troubleshooting instructions on 3Com products 24 hours a day, 7 days a week.

Call 3Com Facts using your Touch-Tone telephone:

1 408 727 7021

Support from Your Network Supplier

If you require additional assistance, contact your network supplier. Many suppliers are authorized 3Com service partners who are qualified to provide a variety of services, including network planning, installation, hardware maintenance, application training, and support services.

When you contact your network supplier for assistance, have the following information ready:

- Product model name, part number, and serial number
- A list of system hardware and software, including revision levels
- Diagnostic error messages
- Details about recent configuration changes, if applicable

If you are unable to contact your network supplier, see the following section on how to contact 3Com.

Support from 3Com

If you are unable to obtain assistance from the 3Com online technical resources or from your network supplier, 3Com offers technical telephone support services. To find out more about your support options, call the 3Com technical telephone support phone number at the location nearest you.

When you contact 3Com for assistance, have the following information ready:

- Product model name, part number, and serial number
- A list of system hardware and software, including revision levels
- Diagnostic error messages
- Details about recent configuration changes, if applicable

Here is a list of worldwide technical telephone support numbers:

Country	Telephone Number	Country	Telephone Number
Asia, Pacific Rim			
Australia	1 800 678 515	P.R. of China	10800 61 00137 or
Hong Kong	800 933 486		021 6350 1590
India	+61 2 9937 5085	Singapore	800 6161 463
Indonesia	001 800 61 009	S. Korea	
Japan	0031 61 6439	From anywhere in S. Korea:	00798 611 2230
Malaysia	1800 801 777	From Seoul:	(0)2 3455 6455
New Zealand	0800 446 398	Taiwan, R.O.C.	0080 611 261
Pakistan	+61 2 9937 5085	Thailand	001 800 611 2000
Philippines	1235 61 266 2602		
Europe			
From anywhere in Europe, call: +31 (0)30 6029900 phone			
+31 (0)30 6029999 fax			
Europe, South Africa, and Middle East			
From the following countries, you may use the toll-free numbers:			
Austria	0800 297468	Netherlands	0800 0227788
Belgium	0800 71429	Norway	800 11376
Denmark	800 17309	Poland	00800 3111206
Finland	0800 113153	Portugal	0800 831416
France	0800 917959	South Africa	0800 995014
Germany	0800 1821502	Spain	900 983125
Hungary	00800 12813	Sweden	020 795482
Ireland	1800 553117	Switzerland	0800 55 3072
Israel	1800 9453794	U.K.	0800 966197
Italy	1678 79489		
Latin America			
Argentina	AT&T +800 666 5065	Mexico	01 800 CARE (01 800 2273)
Brazil	0800 13 3266	Peru	AT&T +800 666 5065
Chile	1230 020 0645	Puerto Rico	800 666 5065
Colombia	98012 2127	Venezuela	AT&T +800 666 5065
North America			
	1 800 NET 3Com		
	(1 800 638 3266)		
	Enterprise Customers:		
	1 800 876-3266		

Returning Products for Repair

Before you send a product directly to 3Com for repair, you must first obtain an authorization number. Products sent to 3Com without authorization numbers will be returned to the sender unopened, at the sender's expense.

To obtain an authorization number, call or fax:

Country	Telephone Number	Fax Number
Asia, Pacific Rim	+ 65 543 6500	+ 65 543 6348
Europe, South Africa, and Middle East	+ 31 30 6029900	+ 31 30 6029999
Latin America	1 408 326 2927	1 408 326 3355
From the following countries, you may call the toll-free numbers; select option 2 and then option 2:		
Austria	0800 297468	
Belgium	0800 71429	
Denmark	800 17309	
Finland	0800 113153	
France	0800 917959	
Germany	0800 1821502	
Hungary	00800 12813	
Ireland	1800553117	
Israel	1800 9453794	
Italy	1678 79489	
Netherlands	0800 0227788	
Norway	800 11376	
Poland	00800 3111206	
Portugal	0800 831416	
South Africa	0800 995014	
Spain	900 983125	
Sweden	020 795482	
Switzerland	0800 55 3072	
U.K.	0800 966197	
U.S.A. and Canada	1 800 NET 3Com (1 800 638 3266)	1 408 326 7120 (not toll-free)
	Enterprise Customers: 1 800 876 3266	

INDEX

Symbols

? character 473, 478

Numbers

3C number 69, 129
3Com bulletin board service (3Com BBS) 768
3Com Knowledgebase Web Services 767
3Com URL 767
3ComFacts 769
802.3_RAW packets 256

A

AARP (AppleTalk Address Resolution Protocol) 674 to 676
access levels 37, 38
 and passwords 75
addModify (snmp trap) 195
address group
 adding port addresses 391
address threshold 252
address/port patterns for QoS classifiers 696, 701, 703
addresses
 adding static 294
 for SNMP trap reporting 194
addressThresholdEvent 257
Administration Console 29 to 42
 password access 75
administration console
 of an ATM switch 30
advancedPing 179, 475, 476
advancedTraceRoute 184
 packet size 480
 ttl option 480
 wait option 480
advertise RIP mode 451, 453
AEP (AppleTalk Echo Protocol) 682
aggregate rate limit 710, 712
 for flow classifiers 715
aggregated links 299
aging time 258

allClosed mode for VLANs
 and Ignore STP mode 365
 displaying 338
 selecting 364
allOpen mode for VLANs
 displaying 338
 selecting 364
analyzer port
 MAC address 758
anchor ports
 rate limits affecting 716
 trunking 313
AppleTalk
 AARP (AppleTalk Address Resolution Protocol) 674 to 676
 AEP (AppleTalk Echo Protocol) 682
 checksums 680
 DDP statistics 683
 forwarding 679
 interfaces 663 to 670
 NBP statistics 686
 ping 682
 removing interfaces 669
 routes 672, 673
 RTMP statistics 684
 source socket verification 681
 ZIP statistics 685
 zones 677, 678
applying controls to classifiers 713
areas 531 to 537
ARP (Address Resolution Protocol)
 cache 171
 deleting cache entries 170
 deleting dynamic cache entries 172
 displaying cache 168
 flushing all entries 171
 flushing dynamic entries 172
 remove 170
 static cache entry 169
ARP cache 428 to 433
ASCII-based editor
 and scripts 124
ATM switch 30
autonegotiation, Ethernet ports 212

B

backplane ports, interface module 31
 backup
 saving NV data 107
 bandwidth, QoS
 displaying 733
 modifying 734
 bandwidth, RSVP 725, 728
 baseline, setting current 133
 baud rate
 serial port 94, 97
 baud setting 95
 best service level 711
 blocking, ignoring STP 365
 BOOTP (Boot Protocol)
 as UDP service 442
 hop count 442
 relay threshold 446
 bridge ports
 adding MAC addresses 294
 defining VLANs 345, 352
 deleting VLANs 363
 listing MAC addresses 293
 modifying VLANs 355, 360
 VLAN summary 339, 342
 bridge-wide parameters, allOpen or allClosed VLAN mode 364
 bulletin board service 768
 burst size, QoS control 712
 burst, advancedPing option 476

C

cast types for QoS classifiers 695, 718
 changing VLANs 355, 360
 channels
 management and data 31
 chassis
 management architecture 33
 power management 33
 checksums, AppleTalk 680
 Class of Service 267
 Class of Service (CoS) 267
 classifiers, QoS
 applying controls to 713
 default 694
 defining 694
 displaying detail information 692
 displaying summary information 691
 guidelines for using 689
 modifying 701
 parameters for defining 695, 718

 predefined flow and nonflow 691
 removing 706
 specifying address/port patterns 696, 701, 703
 specifying IP addresses 695, 702
 command strings
 entering abbreviated 41
 entering values 41
 quick 34
 commands 150
 system menu
 for baselining statistics 90
 for managing NV data 108
 community strings 192
 configuration tasks 34
 conforming packets
 service levels 711
 console access 73
 control packets 733, 734
 controls, QoS
 associating with classifiers 713
 burst size 712
 defining 710
 displaying detail information 708
 displaying summary information 707
 modifying 718, 721
 names 711
 parameters for defining 710
 removing 724
 service levels 711
 specifying IEEE 802.1p tags 712
 conventions
 notice icons 23
 text 24
 CoreBuilder 3500 system
 and network monitoring 755
 CoreBuilder 9000
 management features 33
 system management overview 30
 cost
 IP RIP mode 456
 Spanning Tree settings 254

D

DAS (dual attach station) pairs
 trunks and 307
 data channels
 management 31
 date
 displaying 137
 DDP (Datagram Delivery Protocol) 683
 defaults
 control service level (best) 715

- IP RIP mode (learn) 453
- OSPF route metric 538 to 540
- QoS classifier 694
- route for IP 421
- screen height 76
- Spanning Tree Protocol 261
- ttl value for advancedTraceRoute 480
- ttl value for traceRoute 182
- UDP port number for advancedTraceRoute 480
- UDP port number for traceRoute 182
- defining
 - QoS controls 710, 712, 715
 - VLANs 310, 331
- deleting
 - links 336
 - trunks 318
 - VLANs 363
- designated root 253
- destination address
 - for SNMP trap reporting 194
- destination IP address for QoS classifiers 702
- destination IP address masks 702
- detail
 - trunks 329
- detail information
 - trunks 305
 - VLANs 341
- details, AppleTalk interface 664
- Diagnostics status 69
- disabled RIP mode 451, 453
- disabling
 - excess packet tagging 737
 - RSVP 732
- displaying
 - QoS bandwidth 733
 - QoS classifier detail 692
 - QoS classifier summary 691
 - QoS control detail information 708
 - QoS excess packet tagging 735
 - QoS summary information 691, 707
 - RSVP detail information 729
 - summary RSVP information 728
- displaying TCMP state 301, 304, 328
- DNS (Domain Name System) servers 436 to 440
- documentation
 - comments 25
- drop service level 711
- duplex mode, Ethernet ports 212, 213
- DVMRP (Distance-Vector Multicast Routing Protocol) 507
- dynamic versus static VLAN origin 345, 352

E

- edge policing option 725, 728
- editor for scripts
 - EMACS 124
 - vi 124
- EME (Enterprise Management Engine)
 - console 30
 - overview 33
- enabled RIP mode 451, 453
- enabling
 - excess packet tagging 736
 - RSVP 730
- enabling and disabling Ethernet ports 220
- errors
 - routing interface 406
- Ethernet
 - address
 - and restoring NV data 110
 - and roving analysis 757
 - autonegotiation 212
 - enabling and disabling ports 220
 - fragmenting packets 255
 - menu options 203
 - PACE Access 217
 - PACE Interactive Access 218
 - port duplex mode 212, 213
 - port flow control 215
 - port labels 219
 - port monitoring 221, 222
 - port numbering 204, 207
 - port speed 212, 213
 - port state 220
 - statistics 204, 208
- event log 747, 752
- examples
 - defining QoS classifiers 699
 - defining QoS controls 717
 - defining VLANs (Layer 2 devices) 354
 - defining VLANs (Layer 3 devices) 350, 351, 358, 359
 - modifying QoS classifiers 705
 - modifying QoS controls 722
 - modifying VLANs (Layer 2 devices) 362
 - modifying VLANs (Layer 3 devices) 359, 362
 - of a script 125
 - removing VLANs 363
 - setting Ignore STP mode 365
 - trunk changes 316
 - trunk definitions 311
- excess packet tagging, QoS 735, 737
- excess packets
 - treatment with RSVP 728, 730
- extended diagnostics version number 69

F

- fax service (3ComFacts) 769
 - FDDI (Fiber Distributed Data Interface)
 - fragmenting packets 255
 - port label 241
 - FDDI MAC
 - condition report 237
 - LLC Service, enabling 239
 - FDDI station
 - and SRFs 224, 228
 - FDDI_Snap packets 256
 - feedback on documentation 25
 - File Transfer Protocol (FTP) 87, 89, 107
 - filter id 371
 - filters for QoS flow classifiers
 - defining 696
 - modifying 701
 - flow classifiers
 - cast types 695, 701, 702
 - defining 694
 - predefined 691
 - protocol types 695, 701, 702
 - removing 706
 - using aggregate rate limit 715
 - flow control
 - defining for Gigabit trunk 310, 315
 - displaying for trunks 305, 329
 - flow control, Ethernet ports 215
 - flush
 - for management ip routes 164
 - snmp trap 197
 - flushing
 - learned IP routes 424
 - SNMP trap addresses 197
 - forwarding
 - AppleTalk 679
 - IPX 629
-

G

- gateway IP address 421
 - Gigabit Ethernet
 - trunks 318
 - guidelines
 - for using QoS 689
 - GVRP (GARP VLAN Registration Protocol)
 - displaying status 338
 - using 345, 352
-

H

- hardware revision numbers 69
 - high service level 711
 - hop count 442
-

I

- ICMP statistics 187, 483
- ID, VLAN 341
- IEEE 802.1p priority tagging
 - for excess packets 735, 736
 - for nonflow classifiers 697, 703, 705
 - for QoS controls 712
- IEEE 802.1Q tagging 349
- IGMP (Internet Group Management Protocol)
 - query mode 524, 525
 - snooping mode 524, 525
- IGMP snooping, Layer 2 devices 270
- IGMP snooping, Layer 3 devices 524
- Ignore STP mode
 - selecting 365
- ignoring blocking for VLANs 365
- in-band-management 149
- index 356
 - VLAN interface 360, 395
- interface module
 - backplane ports 31
- interfaces
 - IP 449
 - OSPF 541 to 556
- interfaces, AppleTalk
 - define 665
 - detail display 664
 - modify 667
- interval, QoS statistics 738
- IP (Internet Protocol)
 - address masks for QoS classifiers 695
 - addresses 395, 421
 - loading software 89
 - QoS classifiers 702
 - addresses and restoring NV data 110
 - advancedPing 475
 - advancedTraceRoute 480
 - ARP cache 428 to 433
 - defining routes 422
 - DNS 436
 - enabling or disabling routing 450
 - interfaces
 - displaying 608, 610
 - removing 613, 649
 - statistics 418
 - summary information 398

- overlapped interfaces 447 to 449
- overview 149
- ping functions 473, 478
- RIP mode 451
- routes 450
- statistics 482
 - ICMP 483
 - UDP 483
- traceRoute functions 478
- UDP Helper 442, 447
- IP multicast
 - cache 518
 - DVMRP metric 507, 510
 - hop count 517
 - IGMP 524, 525
 - prune messages 518, 519
 - routing table 517
 - TTL threshold 510
 - tunnels 511, 513, 517, 519, 527
- IP multicast filtering
 - IGMP snooping 524, 525
- IP multicast routing
 - DVMRP 507
 - IGMP 525
 - routeDisplay 517
- IP protocol types
 - modifying 704
- IP routes
 - flushing 424
 - interface status 421
- IP routing
 - enabling or disabling 450
- IPX
 - forwarding
 - enabling or disabling 629
 - statistics 653
 - interfaces
 - statistics 655
 - RIP mode
 - setting 630
 - statistics 651
 - triggered updates 631
 - RIP policy
 - define 633
 - summary 632, 637
 - routes
 - defining static 618
 - flushing learned routes 621
 - removing 620
 - SAP (Service Advertisement Protocol) mode
 - statistics 652
 - triggered updates 639

- SAP mode 638
- SAP policy
 - define 642
 - detail 641
 - modify 645
 - remove 648
 - summary 640
- static servers 622, 624

K

- KBytes/sec rate limit 712

L

- labels, Ethernet ports 219
- LANs
 - virtual 337
- Layer 2 devices
 - defining VLANs 352
 - modifying VLANs 360
- Layer 3 addresses
 - for VLANs 347
 - modifying 355
- Layer 3 devices
 - defining VLANs 346
 - modifying VLANs 355
- learn RIP mode 451, 453
- learned routes, IP 424
- learning state 264
- LER (Link Error Rate)
 - alarm value 242
- lerCutoff
 - and lerAlarm value 243
- levels, service 711
- limits
 - for QoS classifiers 694
 - QoS rate 710, 712, 715
- link aggregation 299
- link state database, OSPF 557 to 563
- links
 - removing 336
 - resources 336
- listening state 264
- LLC (Logical Link Control)
 - service description 239
- LMA (Local Management Application), ATM
 - Switch 31
- log, event 747
- logout 126
- low service level 711

M

- MAC (Media Access Control) addresses
 - adding 294
 - displaying 293
- MAC type for trunk 310, 315
- management
 - and naming the system 101, 136
 - configuring system access 190
 - displaying detailed information 153
 - displaying summary information 151
 - SNMP community strings 192
 - Transcend Network Control Services 30
 - Web Management applications 30
- management data channels 31
- management ip
 - advancedPing 179
 - advancedTraceRoute 184
 - displaying statistics 186
 - ping 177
 - statistics 176
 - tracing a route destination 182
- management ip arp
 - defining a static cache entry 169
 - displaying cache 168
 - flushing all entries from cache 171
 - flushing dynamic entries 172
 - removing cache entries 170
- management ip interface
 - defining the IP address 157
 - displaying summary information 156
 - modifying 158
 - removing 159
- management ip rip
 - displaying RIP information 173
- management ip route
 - default 165
 - defining a static route 162
 - deleting default 166
 - displaying the routing table 160
 - finding in table 167
 - flushing learned routes 164
 - noDefault 166
 - removing an existing route 163
 - searching the routing table 167
- masks
 - source and destination IP address 695, 702
 - subnet 395
- maximum per-reservation bandwidth 730
- maximum total reservable bandwidth 730
- memory partition, OSPF 569, 570
- memory size 69
- menu structure 150
- menus
 - and command strings 40
 - entering abbreviated command strings 41
 - entering values 41
 - navigating 42
 - selecting options 40
- MIBs 768
- MLAN channel 31
- mode, operating
 - defining for Ethernet trunk 310, 315
 - displaying for trunks 305, 329
- modem
 - external, configuring 99, 100
- modes, VLAN
 - definition 340, 343
 - displaying 338
 - selecting allOpen or allClosed 364
 - selecting Ignore STP 365
- modifying
 - QoS bandwidth 734
 - QoS classifiers 701
 - QoS controls 718
 - VLANs (Layer 2 devices) 360
 - VLANs (Layer 3 devices) 355, 358
- module
 - diagnostic messages 129
 - displaying date 137
- module status information 69
- monitoring
 - ports, Ethernet 221, 222
- MultiPoint Link Aggregation (MPLA) 321
 - mode 324
 - Peer Switch Interface State 322

N

- name server, DNS 436 to 440
- names
 - for QoS classifiers 695, 718
 - for QoS controls 710
 - trunk 316
 - VLAN 349
- navigating menus 42
- NBP (Name Binding Protocol) 686
- neighbor notification
 - and LLC Service 239
- neighbors, OSPF 564 to 566
- network supplier support 769
- network troubleshooting 755
- none, for rate limit 710, 712, 715
- nonflow classifiers
 - cast types 695, 701, 702
 - defining 694

- predefined 691
- protocol types 695, 701, 702
- removing 706
- specifying IEEE 802.1p tags 697, 703, 705
- numbering
 - ports, Ethernet 204, 207
- numbers
 - for QoS classifiers 695, 718
 - for QoS controls 710
- NV data
 - restoring 110

O

- online technical services 767
- origin, VLAN 340, 343
- OSPF (Open Shortest Path First)
 - areas 531 to 537
 - default route metric 538 to 540
 - interfaces 541 to 556
 - link state database 557 to 563
 - memory partition 569, 570
 - neighbors 564 to 566
 - router ID 567
 - routing policies 590 to 602
 - soft restarts 570
 - statistics 603
 - stub default metrics 571 to 573
 - virtual links 574 to 589
- out-of-band
 - management 149
- overlapped IP interfaces 447 to 449

P

- PACE Access, Ethernet 217
- PACE Interactive Access, Ethernet 218
- packet filter
 - displaying contents 372, 373, 374, 376, 377, 379, 382, 384
 - filter id 371
 - processing paths 382
- packet size
 - advancedPing 475
 - advancedTraceRoute 480
- packets
 - tagging of excess 736, 737
- password
 - access levels 35
 - configuring 75
 - IP RIP-2 interface 459
- percentage rate limit 712
- per-reservation bandwidth 725, 728

- ping 177
 - advanced ping example 181
 - example 178
- ping command
 - possible responses 473
- pings, AppleTalk 682
- policing options, RSVP 725, 728
- policy
 - IPX RIP
 - define 633
 - modify 635
 - summary 632, 637
 - IPX SAP
 - define 642
 - detail 641
 - modify 645
 - remove 648
 - summary 640
- policy-based services 689
- port
 - label 241
- port group
 - adding ports 391
- port number
 - setting the traceRoute 182, 480
- port ranges for QoS flow classifiers 696, 701, 703
- port speed 95
 - terminal port, setting the 93, 96
- port state, Ethernet 220
- ports
 - autonegotiation, Ethernet 212
 - defining for VLANs 346
 - defining in trunks 310, 315
 - duplex mode, Ethernet 212, 213
 - enabling and disabling, Ethernet 220
 - flow control, Ethernet 215
 - labels, Ethernet 219
 - maximum number in group 391
 - monitoring, Ethernet 221, 222
 - numbering, Ethernet 204, 207
 - PACE Access, Ethernet 217
 - PACE Interactive Access, Ethernet 218
 - receive ports for controls 712, 716
 - speed, Ethernet 212, 213
 - speed, setting 95
 - state, Ethernet 220
 - statistics, Ethernet 204, 208
 - tagging 349
- predefined QoS classifiers 691
- priority tags
 - excess packets 736
 - nonflow classifiers 697, 703, 705
 - QoS controls 712

prioritization 267
 probe
 RMON 755
 procedures
 defining controls 715
 defining flow classifiers 697
 defining nonflow classifiers 700
 defining RSVP 731
 defining VLANs (Layer 2 devices) 354
 defining VLANs (Layer 3 devices) 349
 modifying VLANs (Layer 2 devices) 361
 modifying VLANs (Layer 3 devices) 358
 protocol types
 for QoS classifiers 695, 718
 modifying for VLANs 355
 modifying QoS classifier 704
 selecting for VLANs 345, 347, 352
 prune messages
 IP multicast 519

Q

QoS (Quality of Service) bandwidth
 displaying 733
 modifying 734
 QoS (Quality of Service) classifiers
 defining 694
 displaying detail information 692
 displaying summary information 691
 example of defining 699
 example of modifying 705
 guidelines for using 689
 modifying 701
 removing 706
 QoS (Quality of Service) controls
 applying to classifiers 713
 defining 710
 displaying detail information 708
 displaying summary information 707
 example of defining 717
 example of modifying 722
 modifying 718, 721
 removing 724
 service levels 711
 specifying rate limits 710, 712, 715
 QoS (Quality of Service) excess packet tagging
 disabling 737
 displaying 735
 enabling 736
 QoS (Quality of Service) statistics
 interval 738
 receive 739
 transmit 741

quiet
 advancedPing option 475

R

rate limits, QoS control 710, 712, 715
 modifying one or more 718
 using with trunks 716
 reboots
 trunks and 307
 receive ports
 rate limit 712
 specifying for trunks 716
 receive statistics, QoS 739
 receivePort rate limit 710, 712, 715
 relay threshold
 BOOTP 446
 remote access 73
 removing
 IP interfaces 612, 613, 649
 links 336
 QoS classifiers 706
 QoS controls 724
 trunks 318
 VLANs 363
 reserved packets 733, 734
 returning products for repair 771
 RIP (Routing Information Protocol)
 display 173
 management statistics 176
 mode example 175
 modes 173, 174
 RIP mode
 IP
 interface information 451
 IPX
 setting 630
 statistics 651
 triggered updates 631
 RIP policy
 define
 IPX 633
 modify 635
 summary
 IPX 632, 637
 RIP-2 password 459
 rlogin
 and rebooting the system 123
 router ID, OSPF 567
 routes
 adding default 165
 AppleTalk 672, 673
 defining static IP 422

- deleting default 166
- finding in table 167
- flushing from the routing table 164
- IPX
 - displaying in routing table 615
 - flushing all learned 621
 - removing 620
 - SAPadvertising 613
- types of 421
- routing policies, OSPF 590 to 602
- roving analysis
 - and Spanning Tree 758
- RSVP (Resource Reservation Protocol)
 - definition of 689
 - disabling 732
 - displaying detail information 729
 - displaying summary information 728, 729
 - enabling 730
 - policing options 725, 728
 - procedure for defining 731
 - session information 729
 - treatment of excess packets 728, 730
- RTMP (Routing Table Maintenance Protocol) 684

S

- SAP (Service Advertisement Protocol) mode
 - statistics 652
 - triggered updates 639
- SAP mode
 - IPX 638
- SAP policy
 - define
 - IPX 642
 - detail
 - IPX 641
 - modify
 - IPX 645
 - remove
 - IPX 648
 - summary
 - IPX 640
- script 124
- scripts for the Administration Console
 - examples 125
 - script command 124
- serial number 69, 129
- serial port (modem)
 - setting baud rate 95
- server information 116
- servers
 - defining static IPX 624
 - displaying static IPX 622
 - table for 622
- service levels
 - conforming packets 711
 - default 715
 - RSVP 728, 730
- services for event logging 752
- Simple Network Time Protocol (SNTP) 116 to 122
- size, burst 712
- SMT (Station Management)
 - lerAlarm value 242
 - lerCutoff value 243
- snapshot feature 86
- sniffer 755
- SNMP (Simple Network Management Protocol)
 - agent 189
 - community strings 192
 - display 189
 - displaying configurations 191
 - trap reporting
 - flushing addresses 197
- SNMP trap
 - addressThresholdEvent 257
- soft restarts 570
- software
 - backup NV data 107
 - build date and time 69, 129
 - version 69
- source address 480
 - advancedPing option 476
 - advancedTraceRoute option 480
 - traceRoute option 182
- source IP address for QoS classifiers 695, 702, 710, 718
- source IP address mask 695, 702
- source socket verification
 - AppleTalk 681
- speed, Ethernet ports 212, 213
- split horizon 457
- SRF (Status Report Frames)
 - and FDDI stations 224, 228
 - and lerAlarm 242
- state
 - STP mode (VLANs) 365
- state of IP interface 395
- static routes
 - defining for IP 422
 - defining for IPX 618
- static servers
 - defining IPX 624
- statistics
 - DDP (Datagram Delivery Protocol) 683
 - displaying IP, UDP, and ICMP 186
 - Ethernet 204, 208
 - general IP 482

- ICMP (Internet Control Message Protocol) 187, 483
- IP interface 418
- IPX forwarding 653
- IPX interface 655
- IPX RIP 651
- IPX SAP 652
- NBP (Name Binding Protocol) 686
- OSPF (Open Shortest Path First) 603
- OSPF soft restart 570
- QoS (Quality of Service) interval for 738
- QoS receive 739
- QoS transmit 741
- RTMP (Routing Table Maintenance Protocol) 684
- trunk 305, 329
- UDP (User Datagram Protocol) 187, 483
- VLAN (virtual LAN) 341
- ZIP (Zone Information Protocol) 685
- statistics, AppleTalk protocol 683 to 686
- STP (Spanning Tree Protocol)
 - stpMode 365
- stub default metrics, OSPF 571 to 573
- subnet masks
 - defining 404, 406
 - displaying 399
 - for VLANs 349
- summary information
 - trunk 301, 304, 328
 - VLAN 338
- system baseline display 90
- system baseline set 91
- system console access 73
- system console webAccess 71
- system diagErrLog 115
- system ID 69
- system information
 - displaying 69, 129
- system name
 - displaying 69
 - setting 101, 135, 136
- system reboot 123
- system serial port 93, 96
- system snmp define 117
- system snmp display 116
- system snmp modify 118
- system snmp pollInterval 121
- system snmp remove 119
- system snmp state 120
- system snmp timezone 104, 106, 121
- system snmp tolerance 122
- system up time 69

T

- T_Opr 232
- tagging, VLAN
 - defining 345, 352
 - displaying 341
 - modifying 355, 360
 - specifying 347, 352
- tags, priority
 - for controls 712
 - for excess packets 736
 - for nonflow classifiers 697, 703, 705
- TCMP (Trunk Control Message Protocol)
 - displaying state 301, 304, 305, 328, 329
- technical support
 - 3Com Knowledgebase Web Services 767
 - 3Com URL 767
 - bulletin board service 768
 - fax service 769
 - network suppliers 769
 - product repair 771
- telnet
 - rebooting the system 123
- terminal port
 - port speed 93, 96
- terminal speed 93
- terminalSpeed command (system serialPort) 93, 96
- terminate a Telnet session 126
- TFTP (Trivial File Transfer Protocol) 87
- time
 - displaying module 137
- time in service 69
- tolerance threshold 122
- total reservable bandwidth 725, 728
- traceRoute
 - port number 182, 480
 - source address 182
 - using 478
- traceroute, IP multicast 528
- transmit statistics, QoS 741
- trap reporting
 - adding 195
 - flushing addresses 197
 - modifying 195
- T-Req 232
- triggered updates
 - RIP 631
 - SAP 639
- trunk groups
 - supported 314
- trunking
 - and VLANs 346
 - definition 299
 - overview 299

- trunks
 - defining 310, 315
 - definition 299
 - detail information 305, 329
 - maximum ports 316
 - names 316
 - removing 318
 - resources 318
 - sample definition 311
 - summary information 304, 328
- trusted IP clients 77 to 82
- ttl (time to live)
 - advancedTraceRoute 480
 - example 185
- type of module 69

U

- UDP (User Datagram Protocol) Helper
 - overlapped IP interfaces 447 to 449
 - port and IP forwarding addresses 444
- UDP Helper
 - BOOTP 442
- UDP port number
 - advancedTraceRoute 480
 - traceRoute 182
- UDP statistics 187, 483
- unspecified protocol type 348
- updates
 - RIP triggered 631
 - SAP triggered 639
- URL 767
- user configuration information 116

V

- values 252
 - default 41
 - entering in command strings 41
- vi editor 124
- VID (VLAN ID) 341
 - range 347, 352
- virtual links, OSPF 574 to 589
- VLAN interface 356
- VLAN interface index 356, 360
 - specifying for Ignore STP mode 365
 - used to delete VLANs 363
- VLANs
 - bridge VLAN commands
 - modify 316
- VLANs (virtual LANs) 337
 - defining for Layer 2 devices 352
 - defining for Layer 3 devices 345

- detail information and statistics 341
 - displaying summary information 338
 - errors 406
 - interface index 401, 607
 - modifying (Layer 2 devices) 360
 - modifying (Layer 3 devices) 355, 358
 - removing 363
 - setting allOpen or allClosed mode 364
 - setting Ignore STP mode 365
 - trunking 346
- VRRP (Virtual Router Redundancy Protocol)
 - defining 492
 - enabling or disabling 499
 - introduction 485
 - modifying 495
 - removing 495, 498

W

- wait
 - advancedPing option 475
 - advancedTraceRoute option 480
- Web Management
 - access 71, 72
 - applications 30
- World Wide Web (WWW) 767

Z

- ZIP (Zone Information Protocol) 685
- zones 677, 678

