

314720-D Rev 00
May 2004

4655 Great America Parkway
Santa Clara, CA 95054

Configuring IP Routing Operations

Passport 8000 Series Software Release 3.7



NORTEL
NETWORKS™

Copyright © 2004 Nortel Networks

All rights reserved. May 2004.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is proprietary to Nortel Networks Inc.

The software described in this document is furnished under a license agreement and may be used only in accordance with the terms of that license. The software license agreement is included in this document.

Trademarks

Nortel Networks, the Nortel Networks logo, the Globemark, Unified Networks, Passport, and BayStack are trademarks of Nortel Networks.

Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation.

The asterisk after a name denotes a trademarked item.

Restricted rights legend

Use, duplication, or disclosure by the United States Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013.

Notwithstanding any other license agreement that may pertain to, or accompany the delivery of, this computer software, the rights of the United States Government regarding its use, reproduction, and disclosure are as set forth in the Commercial Computer Software-Restricted Rights clause at FAR 52.227-19.

Statement of conditions

In the interest of improving internal design, operational function, and/or reliability, Nortel Networks Inc. reserves the right to make changes to the products described in this document without notice.

Nortel Networks Inc. does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Portions of the code in this software product may be Copyright © 1988, Regents of the University of California. All rights reserved. Redistribution and use in source and binary forms of such portions are permitted, provided that the above copyright notice and this paragraph are duplicated in all such forms and that any documentation, advertising materials, and other materials related to such distribution and use acknowledge that such portions of the software were developed by the University of California, Berkeley. The name of the University may not be used to endorse or promote products derived from such portions of the software without specific prior written permission.

SUCH PORTIONS OF THE SOFTWARE ARE PROVIDED “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

In addition, the program and information contained herein are licensed only pursuant to a license agreement that contains restrictions on use and disclosure (that may incorporate by reference certain limitations and notices imposed by third parties).

Nortel Networks Inc. software license agreement

This Software License Agreement (“License Agreement”) is between you, the end-user (“Customer”) and Nortel Networks Corporation and its subsidiaries and affiliates (“Nortel Networks”). PLEASE READ THE FOLLOWING CAREFULLY. YOU MUST ACCEPT THESE LICENSE TERMS IN ORDER TO DOWNLOAD AND/OR USE THE SOFTWARE. USE OF THE SOFTWARE CONSTITUTES YOUR ACCEPTANCE OF THIS LICENSE AGREEMENT. If you do not accept these terms and conditions, return the Software, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

“Software” is owned or licensed by Nortel Networks, its parent or one of its subsidiaries or affiliates, and is copyrighted and licensed, not sold. Software consists of machine-readable instructions, its components, data, audio-visual content (such as images, text, recordings or pictures) and related licensed materials including all whole or partial copies. Nortel Networks grants you a license to use the Software only in the country where you acquired the Software. You obtain no rights other than those granted to you under this License Agreement. You are responsible for the selection of the Software and for the installation of, use of, and results obtained from the Software.

1. Licensed Use of Software. Nortel Networks grants Customer a nonexclusive license to use a copy of the Software on only one machine at any one time or to the extent of the activation or authorized usage level, whichever is applicable. To the extent Software is furnished for use with designated hardware or Customer furnished equipment (“CFE”), Customer is granted a nonexclusive license to use Software only on such hardware or CFE, as applicable. Software contains trade secrets and Customer agrees to treat Software as confidential information using the same care and discretion Customer uses with its own similar information that it does not wish to disclose, publish or disseminate. Customer will ensure that anyone who uses the Software does so only in compliance with the terms of this Agreement. Customer shall not a) use, copy, modify, transfer or distribute the Software except as expressly authorized; b) reverse assemble, reverse compile, reverse engineer or otherwise translate the Software; c) create derivative works or modifications unless expressly authorized; or d) sublicense, rent or lease the Software. Licensors of intellectual property to Nortel Networks are beneficiaries of this provision. Upon termination or breach of the license by Customer or in the event designated hardware or CFE is no longer in use, Customer will promptly return the Software to Nortel Networks or certify its destruction. Nortel Networks may audit by remote polling or other reasonable means to determine Customer’s Software activation or usage levels. If suppliers of third party software included in Software require Nortel Networks to include additional or different terms, Customer agrees to abide by such terms provided by Nortel Networks with respect to such third party software.

2. Warranty. Except as may be otherwise expressly agreed to in writing between Nortel Networks and Customer, Software is provided “AS IS” without any warranties (conditions) of any kind. NORTEL NETWORKS DISCLAIMS ALL WARRANTIES (CONDITIONS) FOR THE SOFTWARE, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OF NON-INFRINGEMENT. Nortel Networks is not obligated to provide support of any kind for the Software. Some jurisdictions do not allow exclusion of implied warranties, and, in such event, the above exclusions may not apply.

3. Limitation of Remedies. IN NO EVENT SHALL NORTEL NETWORKS OR ITS AGENTS OR SUPPLIERS BE LIABLE FOR ANY OF THE FOLLOWING: a) DAMAGES BASED ON ANY THIRD PARTY CLAIM; b) LOSS OF, OR DAMAGE TO, CUSTOMER’S RECORDS, FILES OR DATA; OR c) DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS OR SAVINGS), WHETHER IN CONTRACT, TORT OR OTHERWISE (INCLUDING NEGLIGENCE) ARISING OUT OF YOUR USE OF THE SOFTWARE, EVEN IF NORTEL NETWORKS, ITS AGENTS OR SUPPLIERS HAVE BEEN ADVISED OF THEIR POSSIBILITY. The forgoing limitations of remedies also apply to any developer and/or supplier of the Software. Such developer and/or supplier is an intended beneficiary of this Section. Some jurisdictions do not allow these limitations or exclusions and, in such event, they may not apply.

4. General

- a. If Customer is the United States Government, the following paragraph shall apply: All Nortel Networks Software available under this License Agreement is commercial computer software and commercial computer software documentation and, in the event Software is licensed for or on behalf of the United States

Government, the respective rights to the software and software documentation are governed by Nortel Networks standard commercial license in accordance with U.S. Federal Regulations at 48 C.F.R. Sections 12.212 (for non-DoD entities) and 48 C.F.R. 227.7202 (for DoD entities).

- b. Customer may terminate the license at any time. Nortel Networks may terminate the license if Customer fails to comply with the terms and conditions of this license. In either event, upon termination, Customer must either return the Software to Nortel Networks or certify its destruction.
- c. Customer is responsible for payment of any taxes, including personal property taxes, resulting from Customer's use of the Software. Customer agrees to comply with all applicable laws including all applicable export and import laws and regulations.
- d. Neither party may bring an action, regardless of form, more than two years after the cause of the action arose.
- e. The terms and conditions of this License Agreement form the complete and exclusive agreement between Customer and Nortel Networks.
- f. This License Agreement is governed by the laws of the country in which Customer acquires the Software. If the Software is acquired in the United States, then this License Agreement is governed by the laws of the state of New York.

Contents

Preface	27
Before you begin	27
Text conventions	28
Hard-copy technical manuals	29
How to get help	30
 Chapter 1	
IP routing concepts	31
IP addressing	32
Subnet addressing	33
Supernet addressing and CIDR	35
Types of IP routing	36
Virtual routing between VLANs	37
Router ports	38
Static routes	39
Static IP for Management port	39
Domain Name Server	40
Implementation for DNS	40
Black hole static routes	40
IP enhancements and policies	41
Equal Cost MultiPath (ECMP)	42
Alternate route	42
Route filtering/IP policies	44
Accept policies/in filters	45
Redistribution filters	46
Announce policies/out filters	47
Route filtering stages	48
Prefix list	50

Defining route policies	50
Configuration sequence	50
Per-port routing control	51
PPPoE VLANs	51
IP connectivity protocols	53
Address Resolution Protocol (ARP)	53
Enabling ARP traffic	54
Proxy ARP	55
Flushing router tables	56
UDP broadcast forwarding	56
Reverse Address Resolution Protocol (RARP)	57
Virtual Router Redundancy Protocol (VRRP)	58
VRRP Fast Hello Timers	60
RIP and OSPF	61
Routing Information Protocol (RIP)	62
Open Shortest Path First (OSPF) Protocol	64
Overview	65
Benefits	65
OSPF routing algorithm	66
Autonomous system and areas	67
Backbone area	67
Stub area	68
Not so stubby area (NSSA)	68
Neighbors	69
Neighbors on NBMA networks	69
Neighbor adjacencies	70
NBMA adjacencies	70
OSPF routers	70
Router types	71
OSPF interfaces	72
Broadcast interface	72
Non-broadcast multiaccess interface	73
Passive interface	77
OSPF and IP	77
OSPF packets	78

Link state advertisements	79
AS external routes	80
OSPF virtual links	80
Specifying ASBRs	81
Metric Speed	82
Circuitless IP	82
HA-CPU/Layer 3 CPU Redundancy	84
OSPF	85
RIP	85
Prefix Lists and Route Policy	86
VRRP	86
Route Discovery	86
DHCP Relay	87
UDP Forwarding	87
IP Filters	87
RSMLT	88
SMLT/RSMLT operation in L3 environments	88
Failure scenarios	89
Router R1 failure:	89
Router R1 recovery	89
Designing and configuring an RSMLT network	91
Chapter 2	
IP routing configuration examples	93
ARP configuration examples	93
Adding a static ARP entry to a brouter port	94
Adding a static ARP entry to a VLAN	94
Deleting a static ARP entry	95
Changing the default ARP aging time	95
RIP configuration examples	96
RIP send modes	97
Configuring send mode parameters	98
Configuring receive mode parameters	98
RIP configuration tasks	99
Configuration example — Base configuration	99

Displaying configuration files	102
Configuration example — Configuring RIPv2	104
Configuration example — Spanning tree in Passport 8000 routed networks . . .	106
Configuration example - Supplying a Default Route	108
Displaying configuration files	114
Configuration example - Using RIP accept policies	117
Displaying configuration files	120
Configuration example - Using RIP announce policies	121
OSPF configuration examples	122
Configuration example — OSPF interface types	123
Configuring a circuitless IP interface	123
Configuring an IP OSPF interface	125
Configuration example — Equal Cost Multi Path	127
Configuration example — OSPF security mechanisms	130
Simple Password Mechanism	130
Message Digest 5	131
Configuration example — Diagnosing OSPF neighbor state problems	134
Displaying the current state of all OSPF neighbors	135
INIT State problems	136
EXSTART/EXCHANGE Problems	137
Configuration example — OSPF network types	138
Configuration example — OSPF area types	139
Normal area	142
Stub area	143
NSSA	147
Configuration example — OSPF ABR	151
Configuration examples — OSPF ASBR configurations	154
Distributing OSPF routes to RIP and RIP to OSPF using AS-external-LSA Type 1 metrics	155
Distributing an Internet default route to OSPF using AS-external-LSA Type 2 metrics	159
Viewing advertised AS_External LSAs	161
Configuration example — Controlling NSSA external routes advertised	162
Displaying configuration files	166
Configuration example — Multi-area complex	169
Displaying configuration files	171

VRRP configuration examples	177
VRRP configuration example—Normal operation	179
Configuring R1	180
Configuring R2	183
Viewing the VRRP status	186
Displaying VLAN configuration files	187
VRRP configuration example—VRRP operation with SMLT	189
Configuring R1 for VRRP and SMLT	190
Configuring R2 for VRRP and SMLT	193
Displaying VLAN configuration files	196
 Chapter 3	
Configuring IP routing using Device Manager	199
 Router interface types	200
Assigning an IP address on a brouter port	201
Assigning an IP address to a virtual routing port	203
Enabling or disabling per-port routing	205
Globally enabling IP routing features	206
Enabling IP forwarding globally	206
Enabling ECMP globally	209
Enabling alternative routes globally	210
Alternative routes overview	210
Globally enabling alternative routes	211
IP router management	211
Configuring a router's IP protocol stack	211
Viewing IP addresses and their associated router interfaces	213
Viewing and managing the system routing table	214
IP static route table overview	217
Creating IP static routes	218
Creating a static default route	220
Creating a Black hole static route	221
Deleting a static route	222
Configuring IP route preferences	223
Flushing routing tables	224
Flushing by VLAN	224

Flushing by port	225
Configuring circuitless IP	228
Configuring a circuitless IP interface	228
Enabling OSPF on a circuitless IP interface	229
Deleting a circuitless IP interface	231
Configuring ICMP router discovery	231
Enabling ICMP router discovery globally	232
Viewing the ICMP router discovery table	232
Configuring router discovery on a VLAN	234
Configuring router discovery on a port	236

Chapter 4

Configuring IP routing using the CLI 239

Roadmap of IP commands	240
IP routing commands	245
Configuring global parameters	245
Configuring alternative routes	247
Configuring IP forwarding	248
Configuring IP routes	248
Configuring IP route preferences	249
Showing IP route preference information	250
Configuring route discovery	251
Configuring IP route policies	253
Configuring IP static routes	262
Creating Layer 3 static routes	268
Creating a black hole static route	269
Configuring an IP mroute interface	269
Configuring an IP mroute static-source-group	270
Show IP commands	270
Showing IP forwarding status	271
Showing IP interfaces	271
Showing IP route discovery status	272
Showing IP route table information	272
Showing IP static-route information	273
Enabling or disabling per-port routing	274

Configuring Ethernet IP commands	275
Configuring Ethernet IP addresses	276
Creating a brouter port	276
Configuring a directed broadcast on a port	277
Showing routing IP information	278
Configuring route discovery on a port	279
Showing ICMP router discovery information for all interfaces	281
Showing ICMP router discovery information for all VLANs	281
Showing ICMP router discovery information for all ports	282
VLAN IP commands	283
Configuring a VLAN	284
Configuring a directed-broadcast on a VLAN	285
Configuring route discovery on a VLAN	285
Showing VLAN information	287
Configuring circuitless IP	288
Configuring circuitless IP on an interface	288
Showing circuitless IP information	290
 Chapter 5	
Configuring ARP using Device Manager	291
 Enabling or disabling ARP on the routing interface	291
Enabling or disabling ARP on the brouter port	292
Viewing and managing ARP	293
Creating static ARP entries	294
Configuring Proxy ARP	296
ARP Threshold	297
 Chapter 6	
Configuring ARP using the CLI	299
 Roadmap of IP commands	300
Configuring ARP on a port	301
Configuring an ARP proxy on a port	302
Showing ARP port information	302
Configuring ARP on a VLAN	303
Configuring an ARP proxy on a VLAN	304

Showing ARP VLAN information	305
Configuring IP ARP	306
Configuring ARP static entries	307
Configuring ARP Threshold	311
Showing ARP information	313

Chapter 7

Configuring RIP using Device Manager..... 315

Configuration prerequisites	316
Enabling RIP globally	316
Enabling and configuring RIP on a brouter port	318
Enabling and configuring RIP on a VLAN	321
Viewing RIP protocol statistics	323
Configuring RIP interface parameters	324
Configuring Advanced featured on a RIP interface	326

Chapter 8

Configuring RIP using the CLI 329

Roadmap of IP commands	330
Configuring RIP global parameters	332
Configuring RIP parameters on an interface	336
Showing RIP global configuration information	339
Showing information on a RIP interface	340
Configuring RIP on a port	341
Showing RIP information on a port	344
Setting RIP parameters for a VLAN	345
Showing RIP information for VLANs	349

Chapter 9

Configuring OSPF using Device Manager..... 351

Viewing general OSPF routing information	352
Enabling or disabling OSPF on a router	355
Manually initiating a SPF run	356
Configuring OSPF interfaces	357
Viewing OSPF interface information	358

Creating an OSPF interface	360
Changing an OSPF interface type	362
Configuring OSPF NBMA interfaces	363
Adding NBMA neighbors	364
Viewing OSPF neighbor information	365
Managing an OSPF brouter port interface	366
Assigning an IP address to a brouter port interface	367
Configuring OSPF on a brouter port interface	368
Managing an OSPF VLAN interface	372
Assigning an IP address to a VLAN interface	372
Configuring OSPF on a VLAN interface	373
Managing OSPF areas information	375
Viewing OSPF areas information	376
Creating a stub area or NSSAs	377
Creating a virtual link	378
Managing an automatic virtual link	379
Configuring a manual virtual link	379
Viewing virtual links on neighboring devices	381
Managing router hosts	382
Specifying ASBRs	384
Configuring metric speed	385
Configuring global default metric speed	385
Managing metrics with the peer layer interface	385
Viewing stub area metrics	387
Viewing advertisements in the Link State Database	388
Viewing characteristics in the Ext. Link State database	390
Inserting OSPF area aggregate ranges	391
Configuring an OSPF redistribute policy	394
 Chapter 10	
Configuring OSPF using the CLI	397
 Roadmap of IP commands	398
Configuring OSPF global parameters	403
Configuring OSPF host route parameters	404
Configuring an OSPF interface	407

Configuring OSPF areas	410
Configuring OSPF area ranges	412
Configuring OSPF area virtual interface	412
Configuring OSPF neighbors	414
Show OSPF commands	415
Showing OSPF areas	415
Showing OSPF ASE link state advertisements	416
Showing OSPF default metric information	417
Showing OSPF host route configuration	417
Showing OSPF interface statistics	417
Showing OSPF information	418
Showing OSPF interface information	419
Showing OSPF interface timer settings	420
Showing the OSPF link state database table	420
Showing OSPF neighbors	423
Showing OSPF range statistics	423
Configuring port-based OSPF parameters	424
Showing OSPF port statistics	427
Showing OSPF errors on a port	428
Showing OSPF configuration settings on a port	428
Showing basic OSPF information on a port	429
Showing extended OSPF information	430
Configuring OSPF parameters for a VLAN	430
Showing OSPF parameters configured for VLANs	433

Chapter 11

Configuring VRRP using Device Manager..... 435

Configuration prerequisites	436
VRRP and Split-MLT	437
Configuring VRRP for the interface	437
Configuring VRRP secondary features	440
Configuring VRRP on a port	442
Configuring VRRP on a VLAN (or brouter port)	444
Configuring Fast Advertisement Interval on a Port	447
Configuring Fast Advertisement Interval on a VLAN	447

Chapter 12	
Configuring IP VRRP using the CLI	449
Roadmap of IP commands	450
Configuring VRRP on a port	450
Showing VRRP port information	453
Configuring VRRP on a VLAN	454
Showing vlan info vrrp extended command	457
Showing VRRP interface information	458
Dependencies and rules	459
Chapter 13	
Configuring IP policies using Device Manager	461
Route Policy configuration prerequisites	462
Configuring the prefix list	462
Creating and editing the As-Path-List	465
Creating and editing a Community List	467
Creating and editing a route policy	469
Applying routing policies	476
Configuring an OSPF accept policy	477
Configuring an OSPF redistribute policy	479
Configuring inbound/outbound filtering policies on a RIP interface	483
Deleting inbound/outbound filtering policies on a RIP interface	484
Configuring inbound/outbound filtering policies on a DVMRP interface	485
Chapter 14	
Configuring IP Policies using the CLI	489
Roadmap of IP commands	490
IP policy commands	493
Configuring prefix-lists	494
Configuring route policies	496
Configuring a policy for accepting external routes from a router	501
Applying OSPF accept policy changes	503
Configuring OSPF redistribute policies	504
Applying configuration changes to OSPF redistribute policies	506
Showing IP policies	507

Showing prefix lists used by route policies	507
Showing information about route policies	508
Showing information about OSPF accept policies	509
Showing information about OSPF route redistribute policies	510

Chapter 15

Configuring RSMLT using Device Manager and the CLI 511

Configuring RSMLT on a VLAN using Device Manager	511
Viewing and editing RSMLT local information	513
Viewing and editing RSMLT peer information	514
Configuring RSMLT on a VLAN using the CLI	516
Showing IP RSMLT information	519

Index. 521

Figures

Figure 1	Network and host boundaries in IP address classes	33
Figure 2	Class C address supernet	35
Figure 3	IP routing between VLANs	37
Figure 4	Route filtering for unicast routing protocols	45
Figure 5	Route filtering stages	48
Figure 6	Route filtering logic	49
Figure 7	PPPoE and IP configuration	52
Figure 8	Proxy ARP operation	55
Figure 9	Virtual Router Redundancy Protocol configuration	59
Figure 10	Hop count or metric in RIP	62
Figure 11	NBMA subnet	73
Figure 12	NBMA subnet configuration example	76
Figure 13	Virtual link between ABRs through a transit area	81
Figure 14	Routers with IBGP connections	83
Figure 15	SMLT and RSMLT in L3 environments	90
Figure 16	Configuration example—base configuration	100
Figure 17	Configuration example—configuring RIPv2	104
Figure 18	Single spanning tree group	106
Figure 19	Multiple spanning tree groups	107
Figure 20	Supplying a default route	108
Figure 21	RIP accept policy	118
Figure 22	CLIP interface	123
Figure 23	OSPF example	125
Figure 24	ECMP example	128
Figure 25	show ip route info	129
Figure 26	MD5 authentication example	132
Figure 27	show ip ospf neighbors	135
Figure 28	Configuring OSPF network type example	138
Figure 29	Normal Area example	142

Figure 30	Stub Area example	143
Figure 31	Configuration example — Stub Area	144
Figure 32	NSSA Area example	148
Figure 33	Configuration example — NSSA Area	148
Figure 34	OSPF ABR example	151
Figure 35	show ip OSPF area	153
Figure 36	show ip OSPF info	154
Figure 37	OSPF routes: OSPF/RIP and RIP/OSPF	155
Figure 38	OSPF routes: OSPF/RIP and RIP/OSPF	159
Figure 39	show ip ospf ase command	162
Figure 40	Controlling external routes advertised	163
Figure 41	Multi-area complex configuration example	169
Figure 42	VRRP example	179
Figure 43	show ip vrrp info command for R1	186
Figure 44	show ip vrrp info command for R2	187
Figure 45	VRRP example with SMLT	189
Figure 46	Port dialog box—IP Address tab	202
Figure 47	Port, Insert IP Address dialog box	202
Figure 48	VLAN dialog box — Basic tab	203
Figure 49	IP, VLAN dialog box—IP Address tab	204
Figure 50	IP, VLAN, Insert IP Address dialog box	204
Figure 51	IP dialog box—Globals tab	207
Figure 52	IP dialog box—Addresses tab	213
Figure 53	IP dialog box—Routes tab	215
Figure 54	IP dialog box—Static Routes tab	218
Figure 55	IP, Insert Static Routes dialog box	219
Figure 56	The Static Routes tab	222
Figure 57	IP dialog box—RoutePref tab	224
Figure 58	VLAN dialog box—Advanced tab	225
Figure 59	Port dialog box—Interface tab	226
Figure 60	IP dialog box—Circuitless IP tab	228
Figure 61	IP, Insert Circuitless dialog box	229
Figure 62	OspfCircuitless dialog box	230
Figure 63	IP dialog box—Router Discovery tab	233
Figure 64	IP, VLAN—Router Discovery tab	235

Figure 65	Port dialog box—Router Discover tab	237
Figure 66	config ip info command output	247
Figure 67	config ip route preference command	250
Figure 68	show ip route preference command output	251
Figure 69	config ip route-policy <policy name> seq <seq number> command	259
Figure 70	config ip route-policy <policy name> seq <seq number> info command	260
Figure 71	show ip route-policy info command	261
Figure 72	show ip route-policy info ? command	261
Figure 73	config ip static-route info command output	263
Figure 74	creating an L3 static route	268
Figure 75	creating a black hole static route	269
Figure 76	show ip forwarding command output	271
Figure 77	show ip interface command output	272
Figure 78	show ip route info command output	273
Figure 79	show ip static-route info command output	274
Figure 80	show ports info brouter-port command output	277
Figure 81	config ethernet ip info command output	277
Figure 82	show ports info ip command output	278
Figure 83	Route-discovery configuration examples	280
Figure 84	show config verbose command output	281
Figure 85	show vlan info route-discovery command output	282
Figure 86	show port info route-discovery command output	283
Figure 87	config vlan <vid> ip info command output	284
Figure 88	show vlan info ip command output	287
Figure 89	config ip circuitless-ip-int info command output	289
Figure 90	show ip circuitless-ip-int info command output	290
Figure 91	Port dialog box—ARP tab	292
Figure 92	IP dialog box—ARP tab	294
Figure 93	IP, Insert ARP dialog box	295
Figure 94	IP, VLAN dialog box—ARP tab	297
Figure 95	IP, Global tab	298
Figure 96	config ethernet <ports> ip arp-response info command output	301
Figure 97	config ethernet <ports> ip arp-response info command output	302
Figure 98	show ports info arp command (partial output)	303
Figure 99	config vlan <vid> ip arp-response info command output	304

Figure 100	config vlan <aid> ip proxy info command output	305
Figure 101	show vlan info arp command	306
Figure 102	config ip arp info command (partial output)	308
Figure 103	ARP Threshold	313
Figure 104	show ip arp info command output	314
Figure 105	RIP dialog box—Globals tab	317
Figure 106	Port dialog box—RIP tab	319
Figure 107	IP, VLAN dialog box—IP Address tab	321
Figure 108	IP, VLAN dialog box—RIP tab	322
Figure 109	RIP dialog box—Status tab	323
Figure 110	RIP dialog box—Interface tab	325
Figure 111	RIP dialog box—Interface Advance tab	326
Figure 112	config ip rip info command output	333
Figure 113	config ip rip interface command	338
Figure 114	show ip rip <i>info</i> command output	339
Figure 115	show ip rip interface command output	340
Figure 116	config ethernet ip rip info command output	343
Figure 117	show ports info rip command (partial output)	345
Figure 118	config vlan ip rip info command output	348
Figure 119	show vlan info rip command output	349
Figure 120	OSPF dialog box—General tab	353
Figure 121	Force SPF run dialog box	356
Figure 122	OSPF dialog box—Interfaces tab	358
Figure 123	OSPF Insert Interfaces dialog box	361
Figure 124	Neighbors tab—NBMA manually-configured neighbors	362
Figure 125	OSPF dialog box—Neighbors tab	364
Figure 126	OSPF, Insert Neighbors dialog box	365
Figure 127	Port dialog box — IP Address tab	367
Figure 128	Port, Insert IP Address dialog box	367
Figure 129	Port dialog box — OSPF tab	369
Figure 130	VLAN dialog box—Basic tab	372
Figure 131	IP, VLAN dialog box—IP Address tab	373
Figure 132	IP, VLAN dialog box—Insert IP Address dialog box	373
Figure 133	IP, VLAN dialog box—OSPF tab	374
Figure 134	OSPF dialog box—Areas tab	376

Figure 135	OSPF dialog box—Areas tab	378
Figure 136	OSPF dialog box—Virtual If tab	380
Figure 137	OSPF, Insert Virtual If dialog box	380
Figure 138	OSPF dialog box—Virtual Neighbor tab	382
Figure 139	OSPF dialog box—Hosts tab	383
Figure 140	OSPF dialog box—If Metrics tab	386
Figure 141	OSPF dialog box—Stub Area Metrics tab	387
Figure 142	OSPF dialog box—Link State Database tab	389
Figure 143	OSPF dialog box—Ext. Link State DB tab	390
Figure 144	OSPF dialog box—Area Aggregate tab	392
Figure 145	OSPF, Insert Area Aggregate dialog box	392
Figure 146	OSPF dialog box—Redistribute tab	394
Figure 147	OSPF, Insert OSPF Redistribute dialog box	395
Figure 148	config ip ospf info command output	404
Figure 149	config ip ospf interface info command output	410
Figure 150	config ip ospf area info command output	411
Figure 151	show ip ospf area command output	416
Figure 152	show ip ospf ase command output	416
Figure 153	show ip ospf default-metric command output	417
Figure 154	show ip ospf ifstats command output	418
Figure 155	show ip ospf info command output	418
Figure 156	show ip ospf interface command output	419
Figure 157	show ip ospf int-timers command output	420
Figure 158	show ip ospf lsdb command output	421
Figure 159	show ip ospf lsdb detail command output	422
Figure 160	show ospf neighbors command output	423
Figure 161	show ip ospf stats command output	424
Figure 162	config ethernet ip ospf info command output	427
Figure 163	show ports error ospf command output	428
Figure 164	show ports info ospf command (partial output)	429
Figure 165	show ports stats ospf main command output	429
Figure 166	show ports stats interface extended command output	430
Figure 167	config vlan ip ospf info command output	432
Figure 168	show vlan info ospf command output	433
Figure 169	VRRP dialog box—Interface tab	438

Figure 170	VRRP dialog box—Secondary Feature tab	440
Figure 171	Port dialog box—VRRP tab	442
Figure 172	Port, Insert VRRP dialog box	443
Figure 173	VLAN dialog box—Basic tab	445
Figure 174	IP, VLAN dialog box—VRRP tab	446
Figure 175	IP, VLAN, Insert VRRP dialog box	446
Figure 176	config ethernet ports ip vrrp info command output	453
Figure 177	show ip vrrp info command output	454
Figure 178	config vlan ip vrrp info command output	456
Figure 179	show vlan info vrrp extended command output	457
Figure 180	show ip vrrp info command output	458
Figure 181	Policy dialog box—Prefix List tab	463
Figure 182	Policy, Insert Prefix List dialog box	464
Figure 183	Policy dialog box—As Path List tab	465
Figure 184	Policy, Insert As Path List dialog box	466
Figure 185	Policy dialog box—Community List tab	467
Figure 186	Policy, Insert Community List dialog box	468
Figure 187	Policy dialog box—Route Policy tab	471
Figure 188	Policy, Insert Route Policy dialog box	472
Figure 189	Policy dialog box—Applying Policy tab	476
Figure 190	Policy dialog box—OSPF Accept tab	478
Figure 191	Policy, Insert OSPF Accept dialog box	478
Figure 192	Policy dialog box—OSPF Redistribute tab	480
Figure 193	Policy, Insert OSPF Redistribute dialog box	481
Figure 194	Policy dialog box—RIP In/Out Policy tab	483
Figure 195	Policy dialog box—DVMRP In/Out Policy tab	486
Figure 196	config ip prefix-list command	495
Figure 197	config ip prefix-list <name> info command	495
Figure 198	config ip route-policy <policy name> seq <seq number> command	500
Figure 199	config ip route-policy <policy name> seq <seq number> info command	501
Figure 200	config ip ospf accept adv-rtr command	503
Figure 201	config ip ospf redistribute direct syntax command	506
Figure 202	show ip prefix-list command	508
Figure 203	show ip route-policy info command	509
Figure 204	show ip ospf accept info command output	510

Figure 205	show ip ospf redistribute command output	510
Figure 206	VLAN dialog box—Basic tab	512
Figure 207	IP, VLAN2 dialog box—RSMLT tab	512
Figure 208	RSMLT dialog box—Local tab	513
Figure 209	RSMLT dialog box—Peer tab	515
Figure 210	config vlan <vid> ip rsmlt info command output	518
Figure 211	show ip rsmlt info local/peer command output	520

Tables

Table 1	IP addresses	32
Table 2	Subnet masks for Class B and Class C IP addresses	34
Table 3	Router types in an OSPF network	71
Table 4	RIP send modes	97
Table 5	Neighbor states	134
Table 6	OSPF network types	139
Table 7	Port, Insert IP Address dialog box fields	203
Table 8	IP dialog box—Globals tab fields	208
Table 9	Addresses tab fields	213
Table 10	Routes tab fields	215
Table 11	IP dialog box, Static Routes tab fields	220
Table 12	RoutePref tab dialog box fields	224
Table 13	IP dialog box, Circuitless IP tab fields	229
Table 14	IP dialog box—Router Discovery tab fields	233
Table 15	IP, VLAN—Router Discovery tab fields	235
Table 16	Port dialog box—Router Discovery tab fields	237
Table 17	Port dialog box—ARP tab fields	292
Table 18	IP dialog box—ARP tab fields	294
Table 19	Globals tab fields	317
Table 20	Port dialog box—RIP tab fields	319
Table 21	RIP dialog box—Status tab fields	323
Table 22	RIP dialog box—Interface tab fields	325
Table 23	RIP dialog box—Interface Advance tab fields	327
Table 24	RIP supply and listen settings and switch action	343
Table 25	General tab fields	354
Table 26	OSPF dialog box—Interfaces tab fields	358
Table 27	Neighbors tab fields	366
Table 28	OSPF tab fields	370
Table 29	Areas tab fields	376

Table 30	OSPF dialog box—Virtual If tab fields	381
Table 31	OSPF dialog box—Virtual Neighbor tab fields	382
Table 32	Host tab fields	384
Table 33	If Metrics tab fields	387
Table 34	Stub Area Metrics tab fields	388
Table 35	Link State Database tab fields	389
Table 36	Ext. Link State DB tab fields	391
Table 37	Area Aggregate tab fields	393
Table 38	OSPF, Insert OSPF Redistribute dialog box fields	395
Table 39	Interface tab fields	438
Table 40	Secondary Feature tab fields	441
Table 41	Port, Insert VRRP dialog box fields	443
Table 42	Policy, Insert Prefix List dialog box fields	464
Table 43	Policy, Insert As Path List dialog box fields	466
Table 44	Policy, Insert Community List dialog box fields	468
Table 45	Protocol Route Policy table	469
Table 46	Policy, Insert Route Policy dialog box fields	472
Table 47	Policy, Applying Policy dialog box fields	477
Table 48	Policy, Insert OSPF Accepts dialog box fields	479
Table 49	Policy, Insert OSPF Redistribute dialog box fields	481
Table 50	Policy, RIP In/Out Policy dialog box fields	484
Table 51	Policy, DVMRP In/Out Policy dialog box fields	486
Table 52	RSMLT dialog box—Local tab fields	514
Table 53	RSMLT dialog box—Peer tab fields	516

Preface

This guide provides instructions for using the Command Line Interface (CLI) and the Device Manager graphical user interface (GUI) to perform general network management operations on Passport 8000 switches.

For details about how to perform various IP routing tasks, with step-by-step procedures using the CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

For more information about using Passport 8000 Series switches, refer to the Related Publications section of the release notes that accompany this release.

Before you begin

This guide is intended for network administrators with the following background:

- Basic knowledge of networks, Ethernet bridging, and IP routing
- Familiarity with networking concepts and terminology
- Basic knowledge of network topologies
- Experience with windowing systems or graphical user interfaces (GUIs)

Text conventions

This guide uses the following text conventions:

angle brackets (< >)	Indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when entering the command. Example: If the command syntax is <code>ping <ip_address></code>, you enter <code>ping 192.32.10.12</code>
bold Courier text	Indicates command names and options and text that you need to enter. Example: Use the dinfo command. Example: Enter show ip {alerts routes}.
braces ({})	Indicate required elements in syntax descriptions where there is more than one option. You must choose only one of the options. Do not type the braces when entering the command. Example: If the command syntax is <code>show ip {alerts routes}</code>, you must enter either <code>show ip alerts</code> or <code>show ip routes</code>, but not both.
brackets ([])	Indicate optional elements in syntax descriptions. Do not type the brackets when entering the command. Example: If the command syntax is <code>show ip interfaces [-alerts]</code>, you can enter either <code>show ip interfaces</code> or <code>show ip interfaces -alerts</code>.
ellipsis points (. . .)	Indicate that you repeat the last element of the command as needed. Example: If the command syntax is <code>ethernet/2/1 [<parameter> <value>] . . .</code>, you enter <code>ethernet/2/1</code> and as many parameter-value pairs as needed.

<i>italic text</i>	Indicates new terms, book titles, and variables in command syntax descriptions. Where a variable is two or more words, the words are connected by an underscore. Example: If the command syntax is <code>show at <valid_route></code>, <code>valid_route</code> is one variable and you substitute one value for it.
plain Courier text	Indicates command syntax and system output, for example, prompts and system messages. Example: <code>Set Trap Monitor Filters</code>
separator (>)	Shows menu paths. Example: <code>Protocols > IP</code> identifies the IP command on the Protocols menu.
vertical line ()	Separates choices for command keywords and arguments. Enter only one of the choices. Do not type the vertical line when entering the command. Example: If the command syntax is <code>show ip {alerts routes}</code>, you enter either <code>show ip alerts</code> or <code>show ip routes</code>, but not both.

Hard-copy technical manuals

You can print selected technical manuals and release notes free, directly from the Internet. Go to the www.nortelnetworks.com/documentation URL. Find the product for which you need documentation. Then locate the specific category and model or version for your hardware or software product. Use Adobe* Acrobat Reader* to open the manuals and release notes, search for the sections you need, and print them on most standard printers. Go to Adobe Systems at the www.adobe.com URL to download a free copy of the Adobe Acrobat Reader.

A list of related publications for this manual can be found in the release notes that came with your software.

How to get help

If you purchased a service contract for your Nortel Networks product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased a Nortel Networks service program, contact Nortel Networks Technical Support. To obtain contact information online, go to the www.nortelnetworks.com/cgi-bin/comments/comments.cgi URL, then click on Technical Support.

From the Technical Support page, you can open a Customer Service Request online or find the telephone number for the nearest Technical Solutions Center. If you are not connected to the Internet, you can call 1-800-4NORTEL (1-800-466-7835) to learn the telephone number for the nearest Technical Solutions Center.

An Express Routing Code (ERC) is available for many Nortel Networks products and services. When you use an ERC, your call is routed to a technical support person who specializes in supporting that product or service. To locate an ERC for your product or service, go to the <http://www.nortelnetworks.com/help/contact/erc/index.html> URL.

Chapter 1

IP routing concepts

The router management features covered in this documentation apply regardless of which routing protocols are used and include router IP configuration, IP route table management, ARP configuration, ARP table management, BootP/DHCP relay configuration, and VRRP configuration. You should be familiar with the basics of routing and IP addresses.



Note: See [Chapter 2, “IP routing configuration examples,” on page 93](#), for configuration examples, including commands, for most of the concepts described in this chapter.

This chapter includes the following topics:

Topic	Page
IP addressing	32
Types of IP routing	36
Static routes	39
Static IP for Management port	39
IP enhancements and policies	41
PPPoE VLANs	51
IP connectivity protocols	53
RIP and OSPF	61
Routing Information Protocol (RIP)	62
Open Shortest Path First (OSPF) Protocol	64
Circuitless IP	82

IP addressing

An IP version 4 address consists of 32 bits expressed in a “dotted-decimal” format (x.x.x.x). The IP version 4 address space is divided into “classes,” with classes A, B, and C reserved for unicast addresses and accounting for 87.5 percent of the 32-bit IP address space. Class D is reserved for multicast addressing. [Table 1](#) lists the breakdown of IP address space by address range and mask.

Table 1 IP addresses


Class	Address range	Mask	Number of addresses
A	1.0.0.0 - 126.0.0.0	255.0.0.0	126
B	128.0.0.0 - 191.0.0.0	255.255.0.0	127 * 255
C	192.0.0.0 - 223.0.0.0	255.255.255.0	31 * 255 * 255
D	224.0.0.0 - 239.0.0.0		

To express an IP address in dotted-decimal notation, you convert each octet of the IP address to a decimal number and separate the numbers by decimal points. For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted-decimal notation as 128.32.10.167.

Each IP address class, when expressed in binary, has a different boundary point between the network and host portions of the address as illustrated in [Figure 1](#). The network portion is a network number field from 8 through 24 bits. The remaining 8 through 24 bits identify a specific host on the network.

This section includes the following topics:

- [“Subnet addressing,”](#) next
- [“Supernet addressing and CIDR”](#) on page 35

Figure 1 Network and host boundaries in IP address classes

Subnet addressing

The concept of subnetworks (or subnets) extends the IP addressing scheme by allowing an organization to use one IP address range for multiple networks. Subnets are two or more physical networks that share a common network-identification field (the network portion of the 32-bit IP address).

You create a subnet address by increasing the network portion to include a subnet address, thus decreasing the host portion of the IP address. For example, in the address 128.32.10.0, the network portion is 128.32, while the subnet is found in the first octet of the host portion (10). A subnet mask is applied to the IP address and identifies the network and host portions of the address.

Table 2 illustrates how subnet masks used with Class B and Class C addresses can create differing numbers of subnets and hosts. This example includes using the zero subnet, which is permitted on an Passport 8000 switch.

Table 2 Subnet masks for Class B and Class C IP addresses

Number of bits	Subnet mask	Number of subnets (recommended)	Number of hosts per subnet
Class B			
2	255.255.192.0	2	16,382
3	255.255.224.0	6	8,190
4	255.255.240.0	14	4,094
5	255.255.248.0	30	2,046
6	255.255.252.0	62	1,022
7	255.255.254.0	126	510
8	255.255.255.0	254	254
9	255.255.255.128	510	126
10	255.255.255.192	1,022	62
11	255.255.255.224	2,046	30
12	255.255.255.240	4,094	14
13	255.255.255.248	8,190	6
14	255.255.255.252	16,382	2
Class C			
1	255.255.255.128	0	126
2	255.255.255.192	2	62
3	255.255.255.224	6	30
4	255.255.255.240	14	14
5	255.255.255.248	30	6
6	255.255.255.252	62	2


Variable-length subnet masking (VLSM) allows you to divide your intranet into pieces that match your requirements. Routing will be based on the longest subnet mask/network that matches. RIPv2 and OSPF are routing protocols that support VLSM.

Supernet addressing and CIDR

A supernet is a group of networks identified by contiguous network addresses. IP service providers can assign customers blocks of contiguous addresses to define supernets as needed. Supernetting allows you to address an entire block of Class C addresses and avoid using large routing tables to track the addresses.

Each supernet has a unique supernet address that consists of the upper bits shared by all of the addresses in the contiguous block. For example, consider the Class C addresses shown in [Figure 2](#). By adding the mask 255.255.128.0 to IP address 192.32.128.0, you aggregate the addresses 192.32.128.0 through 192.32.255.255 and 128 Class C addresses use a single routing advertisement. In the bottom half of [Figure 2](#), you use 192.32.0.0/17 to aggregate the 128 addresses (192.32.0.0/24 to 192.32.127.0/24).

Figure 2 Class C address supernet



9577EA

Another example is the block of addresses 192.32.0.0 to 192.32.7.0. The supernet address for this block is 11000000 00100000 00000, with the 21 upper bits shared by the 32-bit addresses.

A complete supernet address consists of an *address/mask* pair:

- The *address* is the first 32-bit IP address in the contiguous block. In this example, the address is 11000000 00100000 00000000 00000000 (192.32.0.0 in dotted-decimal notation).
- The *mask* is a 32-bit string containing a set bit for each bit position in the supernet part of the address. The mask for the supernet address in this example is 11111111 11111111 1111000 00000000 (255.255.248.0 in dotted-decimal notation).

The complete supernet address in this example is 192.32.0.0/21.

The supernet address is also referred to as the classless interdomain routing (CIDR) address. Although “classful” prohibits using an address mask with the IP address, CIDR allows you to create networks of various sizes using the address mask. Although VLSM also allows you to divide up your address space, the division is not seen outside your network. With CIDR, your addresses are used by routers outside your network.

Types of IP routing

When routing on a VLAN, an IP address is assigned to the VLAN and is not associated with any particular physical port. Brouter ports are VLANs that route IP packets and bridge nonroutable traffic in a single port VLAN.

This section includes the following topics:


- [“Virtual routing between VLANs,”](#) next
- [“Brouter ports”](#) on page 38

Virtual routing between VLANs

Passport 8000 switches support wire-speed IP routing between VLANs. As shown in Figure 3, although VLAN 1 and VLAN 2 are on the same switch, for traffic to flow from VLAN 1 to VLAN 2, the traffic must be routed.

When you configure routing on a VLAN, you assign an IP address to the VLAN, which acts as a “virtual router interface” address for the VLAN (it is called a virtual router interface because it is not associated with any particular port). The VLAN IP address can be reached through any of the VLAN ports, and frames are routed from the VLAN through the gateway’s IP address. Routed traffic can be forwarded to another VLAN within the switch.

Figure 3 IP routing between VLANs



11040fa

When Spanning Tree Protocol is enabled in a VLAN, the spanning tree convergence must be stable before the routing protocol begins. This requirement can lead to an additional delay in the forwarding of IP traffic.

Because a given port can belong to multiple VLANs (some of which are configured for routing on the switch and some of which are not), there is no longer a one-to-one correspondence between the physical port and the router interface.

As with any IP address, virtual router interface addresses are also used for device management. For SNMP or TELNET management, you can use any virtual router interface address to access the switch as long as routing is enabled on the VLAN.

For more information about:	See:
Using Device Manager to configure Virtual routing	Chapter 3, “Configuring IP routing using Device Manager,” on page 199
Using the CLI to configure Virtual routing	Chapter 4, “Configuring IP routing using the CLI,” on page 239
Virtual routing configuration examples	Chapter 2, “IP routing configuration examples,” on page 93

Brouter ports

The Passport 8000 switch also supports the concept of brouter ports. A brouter port is a single-port VLAN that can route IP packets as well as bridge all nonroutable traffic. The difference between a brouter port and a standard IP protocol-based VLAN configured to do routing is that the routing interface of the brouter port is not subject to the spanning tree state of the port. A brouter port can be in the blocking state for nonroutable traffic and still be able to route IP traffic. This feature removes any interruptions caused by Spanning Tree Protocol recalculations in routed traffic.

A brouter port is actually a one-port VLAN; therefore, each brouter port decreases the number of available VLANs by one and uses one VLAN ID.

For more information about:	See:
Using Device Manager to configure brouter ports	Chapter 3, “Configuring IP routing using Device Manager,” on page 199
Using the CLI to configure brouter ports	Chapter 4, “Configuring IP routing using the CLI,” on page 239
Brouter ports configuration examples	Chapter 2, “IP routing configuration examples,” on page 93

Static routes

Static routes allow you to create routes to a destination IP address manually (see also, [“Static IP for Management port” on page 39](#)).

You can use a static default route to specify a route to all networks for which there are no explicit routes in the Forwarding Information Base or the routing table. This route is by definition a route with the prefix length of zero [RFC 1812]. The Passport 8000 switch can be configured with any route via the IP static routing table.



Note: To create a default static route, the destination address and subnet mask must be set to 0.0.0.0.

Static routes can also be configured with a next hop that is not directly connected, but that hop must be reachable. Otherwise, the static route will not be enabled.

For more information about:	See:
Using Device Manager to configure static routes	Chapter 3, “Configuring IP routing using Device Manager,” on page 199
Using the CLI to configure static routes	Chapter 4, “Configuring IP routing using the CLI,” on page 239
Static routes configuration examples	Chapter 2, “IP routing configuration examples,” on page 93

Static IP for Management port

The network management port is assigned a default IP address if the boot.cfg file not present at boot time. If boot.cfg is present, then the management IP address is taken from that file. At start the system searches for the boot.config file and if it is present then it assigns the IP address for Management Port.

However if at the start the boot.config file is not present in flash, then it sends a boot request for management port to be assigned. If assigned the result is so. Otherwise, IP address for the bootp server assigns the IP address on Mac address of the OOB interface.

Domain Name Server

Where the applications communicate with other machines like telnet, ping and ssh, they use the IP address of the machine to identify the machine. However, instead of the IP address the machine is identified by a hostname. The hostname is translated to the IP address.

The Domain Name Server (DNS) enables the user to use machine names instead of machine IP addresses with applications that need to communicate with other machines. The host name is translated to the IP address, then the DNS servers configured will be queried for mappings.

Implementation for DNS

DNS is enhanced to query the configured DNS server for mapping from hostname to IP address. Up to 3 different DNS servers can be configured. There is a search for the match for the local hosts file, and if it is not found the DNS server is queried for mappings.

Black hole static routes

A black hole static route is a route with an invalid next-hop, such that the data packets destined to this network will be dropped by the switch (see also, [“Static routes” on page 39](#)).

While aggregating or injecting routes to other routers, the router itself may not have a path to the aggregated destination. In such cases, the result is a “black hole” and a routing loop. To avoid such loops, you can configure a black hole static route to the destination it is advertising.

You can configure a preference value for a black hole route. However, you need to configure that preference value appropriately, so that when you wish the black hole route to be used, it gets elected as the best route.

Before adding a black hole static route, a check is performed to ensure that there is no other static route to that identical destination in an enabled state. If such a route exists, you will not be allowed to add the black hole route and an error message will display.

If there is a black hole route enabled, you will also not be allowed to add another static route to that destination. You will need to delete or disable the black hole route prior to adding a regular static route to that destination.

For more information about:	See:
Using Device Manager to configure black hole static routes	Chapter 3, “Configuring IP routing using Device Manager,” on page 199
Using the CLI to configure black hole static routes	Chapter 4, “Configuring IP routing using the CLI,” on page 239
Static routes configuration examples	Chapter 2, “IP routing configuration examples,” on page 93

IP enhancements and policies

In the Passport 8000 switch software, the behavior of IP route policies has been restructured to accommodate the following new scalability requirements:

- [“Equal Cost MultiPath \(ECMP\),” next](#)
- [“Alternate route” on page 42](#)
- [“Route filtering/IP policies” on page 44](#)
- [“Prefix list” on page 50](#)
- [“Defining route policies” on page 50](#)
- [“Configuration sequence” on page 50](#)
- [“Per-port routing control” on page 51](#)

Equal Cost MultiPath (ECMP)

The Equal Cost MultiPath (ECMP) feature allows routers to determine up to four equal cost paths to the same destination prefix. The multiple paths can be used for load sharing of traffic and allows faster convergence to other active paths in case of network failure. By maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP traffic.

The ECMP feature supports and complements the following protocols and route types:

- Open Shortest Path First (OSPF)
- Routing Information Protocol (RIP)
- Static route
- Default route

For more information about:	See:
Using Device Manager to configure ECMP	Chapter 3, “Configuring IP routing using Device Manager,” on page 199
Using the CLI to configure ECMP	Chapter 4, “Configuring IP routing using the CLI,” on page 239
ECMP configuration examples	Chapter 2, “IP routing configuration examples,” on page 93

Alternate route

Routers can learn several routes to a given destination network through several protocols. In the Passport 8000 switch software, if the alternate route feature is enabled, it stores all of these alternate routes sorted in order of network mask/cost/route preference. The “best” or first listed in this list is the best route, which is used by the hardware. The rest of the routes are referred to as alternate routes.

To avoid traffic interruption, alternate routes can be enabled globally to replace best routes with an alternate route if the best route becomes unavailable. The alternate route concept is applied between routing protocols, for example if an OSPF route becomes unavailable and an alternate RIP route is available it will be immediately activated without waiting for an update interval to expire.

The internal routing table manager records the route changes for protocols. It maintains separate tables of static (user-configured) and dynamic (protocol-learned) routes and, in the Passport 8000 switch software, you can configure preferences that determine the precedence given to one type of route over another.

In the event of learning a route with the same network mask and cost values from multiple sources (protocols), route preferences are taken into consideration to select the best route to be added to the forwarding database. Up to four other route(s) per destination are held available as an alternative route.

You can set route preferences for static routes and routing protocols. When you are configuring a static route on the Passport 8000 switch, you can specify a preference for the route. To modify the preference for a static route, disable the route before you edit the configuration, and then re-enable the route.



Note: Changing route preferences is a process-oriented operation that can affect system performance and network reachability while performing the procedures. Therefore, Nortel Networks recommends that if you want to change preferences for static routes or routing protocols, you should do so when configuring routes or before enabling routing protocols.

On an Passport 8000 switch, default preferences are assigned to all standard routing protocols. You can modify the default preference for a protocol to lend it higher or lower priority compared to other protocols. When you change the preference for a route, if all best routes remain best routes, only the local route tables are changed. However, if changing the protocol preference causes best routes to no longer be best routes, neighboring route tables may be affected.

In addition, you can modify the preference value for dynamic routes through route filtering/IP policies, and this value will override the global preference for the protocol. This alternative mechanism allows you to change the behavior of specific routes to have a different preference rather than acquiring the global


protocol preference. For a static route, you can specify an individual route preference that will override the global static route preference. The preference value can be anything between 0 and 255, with 0 reserved for local routes and 255 representing an unreachable route.

For more information about:	See:
Using Device Manager to configure alternate routes	Chapter 3, “Configuring IP routing using Device Manager,” on page 199
Using the CLI to configure alternate routes	Chapter 4, “Configuring IP routing using the CLI,” on page 239
Alternate route configuration examples	Chapter 2, “IP routing configuration examples,” on page 93

Route filtering/IP policies

When IP traffic is routed by the Passport 8000 switch a number of filters can be applied which manage accept, redistribute, and announce policies for unicast routing table information. The filtering process relies on the IP prefix lists in the common routing table manager infrastructure. Filters apply in different ways to different unicast routing protocols.

[Figure 4](#) shows how filters are applied to BGP, RIP, and OSPF protocol.

Figure 4 Route filtering for unicast routing protocols

11041fa

This section includes the following topics:

- [“Accept policies/in filters,”](#) next
- [“Redistribution filters”](#) on page 46
- [“Announce policies/out filters”](#) on page 47
- [“Route filtering stages”](#) on page 48

Accept policies/in filters

Accept policies or in filters are applied to incoming traffic to determine whether or not to add the route to the routing table. Accept policies/in filters are applied in different ways to different protocols, as follows:

- RIP and BGP — filters are applied to all incoming route information
- OSPF — filters are applied only to external route information. Internal routing information is not filtered because otherwise, other routers in the OSPF domain might have inconsistent databases that could affect the router's view of the network topology.

In a network with multiple routing protocols the network administrator can prefer specific routes from RIP instead of from OSPF. The network prefix is a commonly used match criteria for accept policies/in filters.

For more information about:	See:
Using Device Manager to configure accept policies	Chapter 13, "Configuring IP policies using Device Manager," on page 461
Using the CLI to configure accept policies	Chapter 14, "Configuring IP Policies using the CLI," on page 489
Accept policy configuration examples	"Configuration example - Using RIP accept policies" on page 117

Redistribution filters

Redistribution filters notify changes in the route table to the routing protocol (within the device). In earlier releases of the Passport 8000 switch software, redistribution was handled through announce policies. Announce policies should be strictly applied to Link-State Advertisements (LSAs), RIP updates, or BGP NLRI to their respective domains. With redistribution filters, providing you do not breach the protocol rules, you can choose not to advertise everything that is in the protocol database, or you can summarize or suppress route information. On the Passport 8000 switch, by default, no external routes are leaked to protocols that have not been configured.

For more information about:	See:
Using Device Manager to configure redistribution policies	Chapter 13, "Configuring IP policies using Device Manager," on page 461
Using the CLI to configure redistribution policies	Chapter 14, "Configuring IP Policies using the CLI," on page 489
Redistribution policy configuration examples	Chapter 2, "IP routing configuration examples," on page 93

Announce policies/out filters

Announce policies or out filters are applied to outgoing advertisements to neighbors/peers in the protocol domain, to determine whether or not to announce specific route information. Out filtering applies to RIP updates and BGP NLRI updates.


In contrast, out filtering is not applied to OSPF information because OSPF routing information must always be consistent across the domain. To restrict the flow of external route information in the OSPF protocol database, you can apply redistribution filters instead of out filters.

For more information about:	See:
Using Device Manager to configure announce policies	Chapter 13, "Configuring IP policies using Device Manager," on page 461
Using the CLI to configure announce policies	Chapter 14, "Configuring IP Policies using the CLI," on page 489
Announce policy configuration examples	Chapter 2, "IP routing configuration examples," on page 93

Route filtering stages

Figure 5 shows the three distinct filter stages that are applied to IP traffic.

Figure 5 Route filtering stages



These stages are:

1 Filter stage 1

Filter stage 1 is the accept policy/in filter that is applied to incoming traffic to detect changes in the dynamic (protocol-learned) routing information, which are then submitted to the routing table.


2 Filter stage 2

Filter stage 2 is the redistribution filter that is applied to the entries in the routing table to the protocol during leaking process.

3 Filter stage 3

Filter stage 3 is the announce policy/out filter that is applied to outgoing traffic within a protocol domain.

Figure 6 shows the logical process for route filtering on the Passport 8000 switch.

Figure 6 Route filtering logic

10533EB

Prefix list

In previous releases of Passport 8000 switch software, you defined lists per routing protocol to which you wanted to apply a policy or policies. The new IP enhancements and policies allow you to create one or more IP prefix lists and apply this list to any IP route policy.



Note: When you configure a prefix list for a route policy, be sure to add the prefix as “a.b.c.d/32.” You must enter the full 32-bit mask in order to exact a full match of a specific IP address.

For more information about:	See:
Using Device Manager to configure prefix lists	Chapter 13, “Configuring IP policies using Device Manager,” on page 461
Using the CLI to configure prefix lists	Chapter 14, “Configuring IP Policies using the CLI,” on page 489
Prefix list configuration examples	Chapter 2, “IP routing configuration examples,” on page 93

Defining route policies

As IP route policies are no longer tied to a specific protocol, you can define an IP route policy and its attributes globally, and then apply them individually to interfaces and protocols.

Configuration sequence

Using Device Manager, configure route filtering/IP policies in the following three stages:

- 1

Create prefix lists in Device Manager in the IP Routing > Policy > Prefix List tab.
- An IP prefix list is a list of IP networks with masks and a name for reference. Using the MaskLenFrom and MaskLenTo parameters, you can define the range in which this prefix list will be applied to networks.

- 2 Configure IP route policies in IP Routing > Policies > Route Policies tab.

The route policy defines the matching criteria and the actions taken if the policy matches. Prefix lists are used as an input for route policies.

- 3 Apply IP policies to IP interfaces as in- or out-filters by routing protocol.

Per-port routing control

You can enable or disable routing capabilities on specified switch ports, even when the port is part of a routed VLAN. For example, when you disable IP routing on a specific port, the IP traffic ingressing that port is not routed to any other interface on the switch.

You can use this feature as a security measure to prevent non-trusted VLAN ports from injecting IP traffic that is destined to be routed by the switch.

For more information about:	See:
Using Device Manager to configure per-port routing	“Assigning an IP address to a virtual routing port” on page 203
Using the CLI to configure per-port routing	“Enabling or disabling per-port routing” on page 274

PPPoE VLANs

Point-to-Point Protocol over Ethernet (PPPoE) allows you to connect multiple computers on an Ethernet to a remote site through *common customer premises equipment*¹. You can use PPPoE to allow multiple users (for example, an office environment, or a building with many users) to share a common line connection to the Internet.

PPPoE combines the Point-to-Point protocol, commonly used in dial-up connections, with the Ethernet protocol, which supports multiple users in a local area network. The PPP protocol information is encapsulated within an Ethernet frame (see RFC 2516: Point-to-Point Protocol over Ethernet).

1 A telephone company term used to indicate a modem and similar devices.


The Passport 8600 switch allows you to configure PPPoE VLANs using Protocol-based VLANs. The protocol types used by Passport 8600 switch to classify PPPoE packets within the VLANs are:

- 0x8863 (Discovery Stage)
- 0x8864 (PPP Session Stage)

In the example shown in [Figure 7](#), VLAN 1 is a PPPoE VLAN that transports PPPoE traffic to the Internet Service Provider (ISP) network. The traffic to the ISP is bridged. IP traffic can also be routed to the Local Area Network (LAN) using other types of VLANs (for example, port-based VLANs, IP protocol-based VLANs, or IP subnet-based VLANs).

For more information about configuring PPPoE VLANs, see [Configuring Layer 2 Operations: VLANs, Spanning Tree, MultiLink Trunking](#).

Figure 7 PPPoE and IP configuration



11011FA

IP connectivity protocols

This section describes the various protocols that are used for enhanced and resilient IP connectivity.

This section includes the following topics:

- [“Address Resolution Protocol \(ARP\)” next](#)
- [“UDP broadcast forwarding” on page 56](#)
- [“Reverse Address Resolution Protocol \(RARP\)” on page 57](#)
- [“Virtual Router Redundancy Protocol \(VRRP\)” on page 58](#)
- [“VRRP Fast Hello Timers” on page 60](#)

Address Resolution Protocol (ARP)

Network stations using the IP protocol need both a physical address and an IP address to transmit a packet. In situations where the station knows only the network host’s IP address, the Address Resolution Protocol (ARP) enables the network station to determine the physical address of the network host by binding a 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts.

If a network station wants to send a packet to a host but knows only the host’s IP address, the network station uses ARP to determine the host’s physical address as follows:

- 1 The network station broadcasts a special packet, called an ARP request, that asks the host at the specified IP address to respond with its physical address.
- 2 All network hosts receive the broadcast request.
- 3 Only the specified host responds with its hardware address.
- 4 The network station then maps the host’s IP address to its physical address and saves the results in an address-resolution cache for future use.
- 5 The network station’s ARP table displays the associations of the known MAC address to IP address.

Static ARP entries can be created, and individual ARP entries can be deleted.

This section includes the following topics:

- [“Enabling ARP traffic,”](#) next
- [“Proxy ARP”](#) on page 55
- [“Flushing router tables”](#) on page 56

Enabling ARP traffic

The Passport 8000 switch accepts and processes ARP traffic, Spanning Tree BPDUs and Topology Discovery Protocol (TDP) packets on *port-based* VLANs with the default port action set to DROP. To permit ARP traffic, you must use the command line interface to do the following:

- Configure a user-defined protocol-based VLAN for ARP EtherType (byprotocol usrDefined 0x0806)
- Set the ports with a default port action of DROP

You then need to add these ports to the VLAN as static members. Finally, set the port Default VLAN ID to the correct port-based VLAN where the ARPs will be processed.



Note: It is not necessary for you to make any configuration changes for the BPDUs and TDP packets.

The ARP configuration sequence is demonstrated in the following example:

- 1 To create a user-defined protocol-based VLAN with ethertype 0x0806 (specific to the ARP protocol), enter:

```
vlan 4000 create byprotocol 1 usrDefined 2054 name 'ARP'
```
- 2 To remove all ports from this user-defined protocol-based VLAN, type:

```
vlan 4000 ports remove 1/1-1/48,4/1-4/8 member portmember
```
- 3 To add all the ports with the default port action set to DROP for this protocol-based VLAN, enter:

```
vlan 4000 ports add 1/26,1/32 member portmember  
vlan 4000 ports add 1/26,1/32 member static
```


Only one user-defined protocol-based VLAN for ARP is allowed per STG. If the ports with the default port action set to DROP are in different STGs, you need to create additional user-defined protocol-based VLANs. Note that this procedure is effective ONLY with port based VLANs.

Proxy ARP

Proxy ARP allows a network station to respond to an ARP request from a locally attached host or end station for a remote destination. It does so by sending an ARP response back to the local host with its own MAC address of the network station interface for the subnet on which the ARP request was received. The reply is generated only if the switch has an active route to the destination network.

Figure 8 is an example of proxy ARP operation. In this example, host C with mask 24 appears to be locally attached to host B with mask 16, so host B sends an ARP request for host C. However, the Passport 8000 switch is between the two hosts. To enable communication between the two hosts, the Passport 8000 switch would respond to the ARP request with host C's IP address but with its own MAC address.

Figure 8 Proxy ARP operation



11012fa

Flushing router tables

For administrative and/or troubleshooting purposes, it is sometimes necessary to flush the routing tables. Device Manager enables you to flush routing tables either by VLAN or by port. In a VLAN context, all entries associated with the VLAN are flushed. In a port context, all entries associated with the port are flushed.

UDP broadcast forwarding

Some network applications such as the NetBIOS name service rely on a User Datagram Protocol (UDP) broadcast to request a service or locate a server for an application. If a host is on a network, subnet segment, or VLAN that does not include a server for the service, UDP broadcasts are by default not forwarded to the server located on a different network segment or VLAN. Resolve this problem by forwarding the broadcasts to the server through physical or virtual router interfaces.

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface out to other router IP interfaces as a rebroadcast or to a configured IP address.

- If the address is that of a server, the packet is sent as a unicast packet to this address.
- If the address is that of an interface on the router, the frame is rebroadcast.

To follow the basic steps for setting up UDP broadcast forwarding:

- 1** Enter protocols into a table.
- 2** Create policies (protocol/server pairs).
- 3** Assemble these policies into lists or profiles.
- 4** Apply the list to the appropriate interfaces.

When a UDP broadcast is received on a router interface, it must meet the following criteria if it is to be considered for forwarding:

- Must be a MAC-level broadcast
- Must be an IP limited broadcast

- Must be for the specified UDP protocol
- Must have a TTL value of at least 2

For each ingress interface and protocol, the policy specifies how the UDP broadcast is retransmitted: to a unicast host address or to a broadcast address.

Reverse Address Resolution Protocol (RARP)

Certain devices use the Reverse Address Resolution Protocol (RARP) to obtain an IP address from an RARP server. MAC address information for the port is broadcast on all ports associated with an IP protocol-based or port-based VLAN. To enable a device to request an IP address from a RARP server outside its IP VLAN a RARP protocol-based VLAN must be created.

RARP has the format of an Address Resolution Protocol (ARP) frame but its own Ethernet type (8035). So RARP can be removed from the IP protocol-based VLAN definition and treated as a separate protocol thus creating the concept of a RARP protocol-based VLAN.

A typical network topology provides desktop switches in wiring closets with one or more trunk ports extending to one or more data center switches where attached servers provide file, print, and other services. Using RARP functionality, all ports in a network requiring access to an RARP server could be defined as potential members of an RARP protocol-based VLAN. All tagged ports and data center RARP servers must be defined as static or permanent members of the RARP VLAN. Therefore, a desktop host would broadcast an RARP request to all other members of the RARP VLAN. In normal operation, these members would include only the requesting port, tagged ports, and data center RARP server ports. Because all other ports are potential members of this VLAN and RARP is only transmitted at bootup, all other port VLAN memberships would have expired. With this feature, one or more centrally located RARP servers could extend RARP services across traditional VLAN boundaries to reach desktops globally.

Virtual Router Redundancy Protocol (VRRP)


Because end stations are often configured with a static default gateway IP address, a loss of the default gateway router causes a loss of connectivity to the remote networks.

The Virtual Router Redundancy Protocol (VRRP), (RFC 2338) is designed to eliminate the single point of failure that can occur when the single static default gateway router for an end station is lost. VRRP introduces the concept of a virtual IP address (transparent to users) shared between two or more routers connecting the common subnet to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides a dynamic default gateway redundancy in the event of failover.

The VRRP router controlling the IP address(es) associated with a virtual router is called the primary router and forwards packets to these IP addresses. The election process provides a dynamic transition of forwarding responsibility if the primary router becomes unavailable.

In the configuration example shown in [Figure 9 on page 59](#), the first three hosts install a default route to R1 (virtual router 1) IP address and the other three hosts install a default route to R2 (virtual router 2) IP address.

This configuration not only has the effect of load sharing the outgoing traffic, but it also provides full redundancy. If either router fails, the other router assumes responsibility for both addresses.

Figure 9 Virtual Router Redundancy Protocol configuration

11013fa

The Passport 8000 switch supports 255 VRRP interfaces per switch. VRRP uses the following terms:

- VRRP router — a router running the VRRP protocol
- Virtual router — an abstract object acting as the default router for one or more hosts, consisting of a virtual router ID and a set of addresses
- IP address owner — the VRRP router that has virtual router IP addresses as real interface addresses (This router is the one that responds to packets sent to this IP address.)
- Primary IP address — an IP address selected from the real addresses and used as the source address of packets sent from the router interface (The virtual primary router sends VRRP advertisements using this IP address as the source.)
- Virtual primary router — the router assuming responsibility for forwarding packets sent to the IP address associated with the virtual router and answering ARP requests for these IP addresses
- Virtual primary router backup — the virtual router that becomes the primary router should the current primary router fail

When a VRRP router is initialized, if it is the IP address owner, its priority is 255 and it sends a VRRP advertisement. The VRRP router also broadcasts an ARP request containing the virtual router MAC address for each IP address associated with the virtual router. The VRRP router then transitions to the controlling state.

In the controlling state, the VRRP router functions as the forwarding router for the IP addresses associated with the virtual router. It responds to ARP requests for these IP addresses, forwards packets with a destination MAC address equal to the virtual router MAC address, and accepts only packets addressed to IP addresses associated with the virtual router if it is the IP address owner. If the priority is not 255, the router transitions to the backup state to ensure that all layer 2 switches in the down path relearn the new origin of the VRRP MAC addresses.

In the backup state, a VRRP router monitors the availability and state of the primary router. It does not respond to ARP requests and must discard packets with a MAC address equal to the virtual router MAC address. It does not accept packets addressed to IP addresses associated with the virtual router. If a shutdown occurs, it transitions back to the initialize state. If the primary router goes down, the backup router sends the VRRP advertisement and ARP request described in the preceding paragraph and transitions to the controlling state.

If an advertisement timer fires, the router sends an advertisement. If an advertisement is received with a 0 priority, the router sends an advertisement. If the priority is greater than the local priority or if it is the same as the local priority and the primary IP address of the sender is greater than the local primary IP address, the router transitions to the backup state. Otherwise, it discards the advertisement. If a shutdown occurs, the primary router sends a VRRP advertisement with a priority of 0 and transitions to the initialize state.

VRRP Fast Hello Timers

The current implementation of VRRP allows you to set the advertisement time interval (in seconds) between sending advertisement messages. This allows for faster network convergence with standardized VRRP failover. However, losing connections to servers for more than a second can result in missing critical failures. Customer network uptime in many cases requires faster network convergence which means detecting network problems within hundreds of milliseconds.

To achieve these requirements two new enhancements are introduced, Fast Advertisement Enable and the Fast Advertisement Interval.

Fast Advertisement Enable acts like a toggle switch for the Advertisement Interval and the Fast Advertisement Interval. When Fast Advertisement Enable is enabled, the Fast Advertisement Interval is used instead of the Advertisement Interval.

The Fast Advertisement Interval is similar to the current Advertisement Interval parameter except for the unit of measure and the range. The Fast Advertisement Interval is expressed in milliseconds and the range is from 200 to 1000 milliseconds. This unit of measure must also be in multiples of 200 milliseconds, otherwise an error is displayed.



Note: When the Fast Advertisement Interval is enabled, VRRP will only communicate other Passport 8600 modules with the same settings.

RIP and OSPF

The Passport 8000 switch supports wire-speed IP routing of frames using one of the following dynamic IP routing protocols:

- RIP version 1 (RFC 1058)
- RIP version 2 (RFC 1723)
- OSPF version 2 (RFC 2178)

Unlike static IP routing, where a manual entry must be made in the routing table to specify a routing path, dynamic IP routing uses a “learning” approach to determine the paths and routes to other routers. There are two basic types of routing algorithm: distance vector and link state. Routing Information Protocol (RIP) is a distance vector protocol and Open Shortest Path First (OSPF) Protocol is a link state protocol.


Routing Information Protocol (RIP)

In routed environments, routers communicate with one another to track available routes. Routers can learn about available routes dynamically using the Routing Information Protocol (RIP). The Passport 8000 switch software implements standard RIP for exchanging TCP/IP route information with other routers.

RIP uses broadcast User Datagram Protocol (UDP) data packets to exchange routing information. Each router “advertises” routing information by sending a routing information update every 30 seconds. If a router does not receive an update from another router within 90 seconds, it marks the routes served by the “nonupdating” router as being unusable. If no update is received within 240 seconds, the router removes all routing table entries for the “nonupdating” router.

RIP is known as a distance vector protocol. The vector is the network number and next hop, and the distance is the cost associated with the network number. RIP identifies network reachability based on cost, and cost is defined as hop count. One hop is considered to be the distance from one router to the next. This cost or hop count is known as the *metric* (Figure 10).

Figure 10 Hop count or metric in RIP



11014fa

RIP version 1 was distributed in the early years of the Internet and advertised default class address without subnet masking. RIP version 2 advertises more explicitly, based on the subnet mask.

The Passport 8000 switch supports RIP version 2, which advertises routing table updates using multicast instead of broadcasting. RIP version 2 supports variable length subnet masks (VLSM) and triggered updates of routers.

A directly connected network has a metric of zero. An unreachable network has a metric of 16. Therefore, the highest metric between any two networks can be 15 hops or 15 routers.

For more information about:	See:
Using Device Manager to configure RIP	Chapter 7, "Configuring RIP using Device Manager," on page 315
Using the CLI to configure RIP	Chapter 8, "Configuring RIP using the CLI," on page 329
RIP configuration examples	Chapter 2, "IP routing configuration examples," on page 93

Open Shortest Path First (OSPF) Protocol

Open Shortest Path First (OSPF) Protocol is an Interior Gateway Protocol (IGP) that distributes routing information between routers belonging to a single *autonomous system* (AS). Intended for use in large networks, OSPF is a link-state protocol which supports IP subnetting, TOS-based routing, and the tagging of externally-derived routing information.

This section includes the following topics:

- [“Overview,” next](#)
- [“Benefits” on page 65](#)
- [“OSPF routing algorithm” on page 66](#)
- [“Autonomous system and areas” on page 67](#)
- [“Neighbors” on page 69](#)
- [“OSPF routers” on page 70](#)
- [“Router types” on page 71](#)
- [“OSPF interfaces” on page 72](#)
- [“OSPF and IP” on page 77](#)
- [“OSPF packets” on page 78](#)
- [“Link state advertisements” on page 79](#)
- [“AS external routes” on page 80](#)
- [“OSPF virtual links” on page 80](#)
- [“Specifying ASBRs” on page 81](#)
- [“Metric Speed” on page 82](#)

For more information about:	See:
Using Device Manager to configure OSPF	Chapter 9, “Configuring OSPF using Device Manager
Using the CLI to configure OSPF	Chapter 10, “Configuring OSPF using the CLI,” on page 397
OSPF configuration examples	Chapter 2, “IP routing configuration examples,” on page 93

Overview

In an OSPF network, each router maintains a *link-state database* that describes the topology of the autonomous system (AS). The database contains the *local state* for each router in the AS, including the router's usable interfaces and reachable neighbors. Each router periodically checks for changes in its local state and shares any changes detected by flooding *link-state advertisements* (LSAs) throughout the AS. Routers synchronize their topological databases based on the sharing of information from LSAs.

From the topological database, each router constructs a *shortest-path tree*, with itself as the root. The shortest-path tree gives the optimal route to each destination in the AS. Routing information from outside the AS appears on the tree as leaves.

OSPF routes IP traffic based solely on the destination IP address and subnet mask, and IP Type of Service (TOS) contained in the IP packet header.

Benefits

In large networks OSPF offers the following benefits:

- Fast convergence

In the event of topological changes, OSPF recalculates routes quickly.

- Minimal routing protocol traffic

Unlike distance vector routing protocols such as RIP, OSPF generates a minimum of routing protocol traffic.

- Load sharing

OSPF provides support for equal-cost multipath routing. If several equal-cost routes to a destination exist, traffic is distributed equally among them.

- Type of Service

Separate routes can be calculated for each IP Type of Service.

OSPF routing algorithm

A separate copy of the OSPF routing algorithm runs in each area. Routers which are connected to multiple areas run multiple copies of the algorithm. The sequence of processes governed by the routing algorithm is as follows:

- 1** When a router starts, it initializes the OSPF data structures and then waits for indications from lower-level protocols that its interfaces are functional.
- 2** A router then uses the Hello Protocol to discover neighbors. On point-to-point and broadcast networks the router dynamically detects its neighbors by sending hello packets to the multicast address AllSPFRouters. On non-broadcast multiaccess networks, some configuration information is required in order to discover neighbors.
- 3** On all multiaccess networks (broadcast or non-broadcast), the Hello Protocol also elects a DR for the network.
- 4** The router attempts to form adjacencies with some of its neighbors. On multiaccess networks, the DR determines which routers become adjacent. This behavior does not occur if a router is configured as a passive interface, because passive interfaces do not form adjacencies.
- 5** Adjacent neighbors synchronize their topological databases.
- 6** The router periodically advertises its link-state, and also does so when its local state changes. LSAs include information about adjacencies enabling quick detection of dead routers on the network.
- 7** LSAs are flooded throughout the area, ensuring that all routers in an area have exactly the same topological database.
- 8** From this database each router calculates a shortest-path tree, with itself as root. This shortest-path tree in turn yields a routing table for the protocol.

Autonomous system and areas

The AS can be subdivided into areas that group together contiguous networks, routers connected to these networks, and attached hosts. Each area has its own topological database which is invisible from outside the area. Routers within an area know nothing of the detailed topology of other areas. Subdividing the AS into areas significantly reduces the amount of routing protocol traffic as compared to treating the entire AS as a single link-state domain.

You can attach a router to more than one area, which allows you to maintain a separate topological database for each connected area. Two routers within the same area maintain an identical topological database for that area. Each area is assigned a unique area ID and the area ID 0.0.0.0 is reserved for the backbone area.

Packets are routed in the AS based on their source and destination addresses. If the source and destination of a packet reside in the same area intra-area routing is used. If the source and destination of a packet reside in different areas inter-area routing is used. Intra-area routing protects the area from bad routing information because no routing information obtained from outside the area can be used. Inter-area routing must pass through the backbone area which is described in the following section.

This section includes the following topics:

- [“Backbone area,”](#) next
- [“Stub area”](#) on page 68
- [“Not so stubby area \(NSSA\)”](#) on page 68

Backbone area

The backbone area consists of the following network types:

- Networks and attached routers that are not contained in any other area
- Routers that belong to multiple areas

The backbone is usually contiguous but you can create a non-contiguous area by configuring virtual links.

Virtual links can be configured between any two backbone routers that have an interface to a common non-backbone area. Virtual links belong to the backbone and use intra-area routing only. Virtual links are described on [page 80](#).

The backbone is responsible for distributing routing information between areas. The topology of the backbone area is invisible to other areas, while it knows nothing of the topology of those areas.

In inter-area routing, a packet travels along three contiguous paths in a point-to-multipoint configuration, as follows:

- 1 An intra-area path from the source to an *area border router* (ABR)
- 2 A backbone path between the source and destination areas
- 3 Another intra-area path to the destination.

The OSPF routing algorithm finds the set of such paths that has the smallest cost. The topology of the backbone dictates the backbone paths used between areas. Inter-area paths are selected by examining the routing table summaries for each connected ABR. The OSPF behavior has been modified according to OSPF standards so that OSPF routes cannot be learned through an area border router (ABR) unless it is connected to the backbone or through a virtual link.

Stub area

A stub area is configured at the edge of the OSPF routing domain and has only one ABR. A stub area does not receive LSAs for routes outside its area, reducing the size of its link-state database. A packet destined outside the stub area is routed to the ABR, which examines it before forwarding the packet to its destination. The network behind a passive interface is treated as a stub area, and does not form adjacencies. It is advertised into the OSPF area as an internal route.

Not so stubby area (NSSA)

A not so stubby area prevents the flooding of external LSAs into the area by replacing them with a default route. An NSSA can import small stub (non-OSPF) routing domains into OSPF. Like stub areas, NSSAs are at the edge of an OSPF routing domain. Non-OSPF routing domains are attached to the NSSAs, forming NSSA transit areas. Accessing the addressing scheme of small stub domains permits the NSSA border router to also perform manual aggregation.

Neighbors

In an OSPF network, any two routers that have an interface to the same network are *neighbors*. Routers use the *Hello Protocol* to discover their neighbors and maintain neighbor relationships. On a broadcast or point-to-point network, the Hello Protocol dynamically discovers neighbors. On a non-broadcast multiaccess network (NBMA), you must manually configure neighbors for the network.

The Hello Protocol provides bi-directional communication between neighbors. Periodically OSPF routers send out hello packets over all interfaces. Included in these hello packets is the following information:

- The router's priority
- The router's Hello Timer and Dead Timer values
- A list of routers that have sent this router hello packets on this interface
- The router's choice for *designated router* (DR) and *backup designated router* (BDR)

Bidirectional communication is determined when one router discovers itself listed in its neighbor's hello packet.

This section includes the following topics:

- [“Neighbors on NBMA networks,”](#) next
- [“Neighbor adjacencies”](#) on page 70
- [“NBMA adjacencies”](#) on page 70

Neighbors on NBMA networks

NBMA interfaces whose router priority is a positive, non-zero value are eligible to become DR for the NBMA network and are configured with a list of all attached routers. The neighbors list includes each neighbor's IP address and router priority. In an NBMA network, any router with a priority other than zero is eligible to become the DR for the NBMA network. You must manually configure the IP address, mask, and router priority of neighbors on routers that are eligible to become the DR or BDR for the network.

Logging messages indicate when an OSPF neighbor state change occurs. This log message indicates the previous state and the new state of the OSPF neighbor. The log message generated for system traps also indicates the previous state and the current state of the OSPF neighbor.

Neighbor adjacencies

Neighbors may form an *adjacency* for the purpose of exchanging routing information. When two routers form an adjacency, they go through a *database exchange* process to synchronize their topological databases. When their databases are synchronized, the routers are said to be fully adjacent. Bandwidth is conserved because, from this point on, only routing change information is passed between the adjacent routers.

All routers connected by a point-to-point network or a virtual link always form an adjacency. All routers on a broadcast or NBMA multiaccess network form an adjacency with the DR and the BDR.

NBMA adjacencies

In an NBMA network, before a DR is elected, the router sends hello packets only to those neighbors eligible to become DR. The NBMA DR only forms adjacencies with its configured neighbors, and drops all packets coming from other sources. The neighbor configuration also tells the router the expected hello behavior for each neighbor.



Note: If a router receives a hello packet from a neighbor with a different priority than what is configured, the router will automatically change the configured priority to match the dynamically learned priority.

OSPF routers

To limit the amount of routing protocol traffic, the Hello Protocol elects a designated router (DR) and a backup designated router (BDR) on each multiaccess network. Instead of neighboring routers forming adjacencies and swapping link-state information with each other (which on a large network can mean a lot of routing protocol traffic), all routers on the network form adjacencies with the DR and the BDR *only* and send link-state information to them. The DR redistributes this information to every other adjacent router.

When operating in backup mode, the BDR receives link-state information from all routers on the network and listens for acknowledgements. Should the DR fail, the BDR can transition quickly to the role of DR because its routing tables are up-to-date.

Router types

Routers in an OSPF network can take on different roles depending on how they are configured. [Table 3](#) describes the router types you can configure in an OSPF network.

Table 3 Router types in an OSPF network

Router Type	Description
AS boundary router (ASBR)	A router attached at the edge of an OSPF network is called an AS boundary router (ASBR). An ASBR generally has one or more interfaces that run an inter-domain routing protocol such as BGP. In addition, any router distributing static routes or RIP routes into OSPF is considered an ASBR. The ASBR forwards external routes into the OSPF domain. In this way, routers inside the OSPF network learn about destinations outside their domain.
Area border router (ABR)	A router attached to two or more areas inside an OSPF network is considered an area border router (ABR). ABRs play an important role in OSPF networks by condensing the amount of OSPF information that is disseminated.
Internal router (IR)	A router that has interfaces only within a single area inside an OSPF network is considered an internal router (IR). Unlike ABRs, IRs have topological information only about the area in which they are contained.
Designated router (DR)	In a broadcast or NBMA network a single router is elected to be the designated router (DR) for that network. A DR assumes the responsibility of making sure all routers on the network are synchronized with one another and also advertises that network to the rest of the AS.
Backup designated router (BDR)	A backup designated router (BDR) is elected in addition to the designated router (DR) and, in the event of failure of the DR, will assume its role quickly.

OSPF interfaces

An OSPF interface, or link, is configured on an IP interface. In the Passport 8000 switch, an IP interface can be either a single link (brouter port) or a logical interface configured on a VLAN (multiple ports). The state information associated with the interface is obtained from the underlying lower level protocols and the routing protocol itself.

On an Passport 8000 switch, OSPF interfaces are designated as one of the following types:

- broadcast (active)
- non-broadcast multiaccess (NBMA)
- passive



Note: When an OSPF interface is enabled, you cannot change its interface type. You must first disable the interface. You can then change its type and re-enable it. If it is an NMBA interface, you must also first delete its manually-configured neighbors.

This section includes the following topics:

- [“Broadcast interface,” next](#)
- [“Non-broadcast multiaccess interface” on page 73](#)
- [“Passive interface” on page 77](#)

Broadcast interface

Broadcast interfaces support many attached routers and can address a single physical message to all attached broadcast routers (sent to AllSPFRouters and AllDRouters).

Broadcast interfaces discover neighboring routers dynamically using the OSPF Hello Protocol. Each pair of routers on a broadcast network, such as an Ethernet, communicate directly.

Non-broadcast multiaccess interface

Non-broadcast multiaccess (NBMA) interfaces support many routers, but cannot broadcast.


In contrast to a broadcast network where some OSPF protocol packets are multicast (sent to AllSPFRouters and AllDRouters), OSPF packets on an NBMA interface are replicated and sent to each neighboring router, in turn, as unicast. NBMA networks drop all OSPF packets with destination address AllSPFRouters and AllDRouters.

An example of an NBMA network is an ATM subnet that supports a mesh of PVCs containing each pair of routers.

Figure 11 shows an example of four routers attached to an NBMA subnet where each router is connected to every other router via an ATM permanent virtual circuit.

A single IP subnet is assigned to the NBMA segment and each router is assigned an IP address within the subnet.

Figure 11 NBMA subnet



11015fa

Designated router parameters

OSPF treats an NBMA network much like it treats a broadcast network. Since many routers are attached to the network, a designated router (DR) is elected to generate the network's link-state advertisements.

Because the NBMA network does not broadcast, you must manually configure neighbors for each router eligible to become DR (those whose router priority for the network is a positive, non-zero value). You must also configure a PollInterval for the network.

NBMA neighbors list and priorities

NBMA interfaces whose router priority is a positive, non-zero value are eligible to become DR for the NBMA network and are configured with a list of all attached routers, or neighbors. This neighbors list includes each neighbor's IP address and router priority.

This information is used both during and after the DR election process. When an interface to a non-broadcast network with a non-zero priority comes up, and before the Hello Protocol elects a DR, the router sends hello packets only to those neighbors eligible to become DR (or those whose router priority is a positive, non-zero value). Once a DR is elected, it only forms adjacencies with its configured neighbors, and drops all packets from other sources. This neighbor configuration also tells the router the expected hello behavior of each neighbor.



Note: If a router eligible to become DR receives a hello packet from a neighbor showing a different priority than what is already configured for this neighbor, the DR changes the configured priority to match the dynamically-learned priority.

NBMA PollInterval

An NBMA interface is also configured with a PollInterval. The PollInterval designates the interval at which hello packets are sent to inactive neighboring routers. Hello packets are typically sent at the HelloInterval, for example every 10 seconds. If a neighboring router becomes inactive, or if hello packets have not been received for the established RouterDeadInterval, hello packets are sent at the specified PollInterval, for example, every 120 seconds.

Sending hello packets

You must configure a neighbors list for the DR to allow an NBMA network to send hello packets. If the router is eligible to become a DR (if its router priority is a positive, non-zero value), it periodically sends hello packets to all neighbors that are also eligible. The effect of this is that any two eligible routers are always exchanging hello packets, which is necessary for the correct DR election. You can minimize the number of hello packets sent by minimizing the number of eligible routers on a non-broadcast network.

When the DR is elected, it begins sending hello packets to all manually configured neighbors, synchronizing their link-state databases, establishing itself as DR, and identifying the BDR.

If a router is not eligible to become DR, it periodically sends hello packets to both the DR and the BDR. It also sends a hello packet in reply to a hello packet received from any eligible neighbor (other than the current DR and BDR). This process establishes an initial bidirectional relationship with any potential DR.

When sending hello packets periodically to any neighbor, the interval between hello packets is determined by the neighbor's state. If the neighbor is in the Down state, hello packets are sent at the designated PollInterval, for example every 120 seconds. Otherwise, hello packets are sent at the designated HelloInterval, for example every 10 seconds.


Forming adjacencies

In an NBMA network, as in a broadcast network, all routers become adjacent to the DR and the BDR. The adjacencies are formed after the router priorities are assigned, the neighbors are configured, and the network DR is elected.

[Figure 12 on page 76](#) shows an NBMA subnet example with router priorities and manually configured neighbors.

Because R1 and R2 have a router priority of 0, they are not eligible to become the DR. Also, R1 and R2 do not require configuration of a neighbors list; neighbors are discovered dynamically through the Hello Protocol.

R3 and R4 both have a positive, non-zero priority and are eligible to become the DR. Neighbor lists must be manually configured on R3 and R4.

Figure 12 NBMA subnet configuration example

11016fa

To create the NBMA configuration example shown in [Figure 12](#):

- 1 Configure the following for each router:
 - NBMA interface type
 - PollInterval value
 - Router priority
- 2 Configure R1, R2, and R4 as neighbors on R3.
- 3 Configure routers R1, R2, and R3 as neighbors on R4.
- 4 Bring up all routers at the same time.
- 5 R3 and R4 send each other a hello packet to elect a DR.

- 6 The Hello Protocol elects R3 as the DR, and R4 as the BDR.
- 7 R3 (DR) and R4 (BDR) send hello packets to all other routers on the NBMA subnet, synchronizing their link-state databases, and establishing themselves as DR and BDR.
- 8 R1 and R2 reply to R3 and R4.
- 9 R3 and R4 each form three adjacencies (one with each router on the NBMA subnet).
- 10 R1 and R2 each form two adjacencies (one with the DR and one with the BDR).

Passive interface

The objective of the passive interface is to enable an interface to advertise into an OSPF domain while limiting its adjacencies.

By changing the interface's type value to passive, it is advertised into the OSPF domain as an internal stub network with the following behaviors:

- does not send hello packets into the OSPF domain
- does not receive hello packets from the OSPF domain
- does not form adjacencies in the OSPF domain

With the passive interface feature, the interface requires only a new interface type value to allow it to be advertised as an OSPF internal route. Without the passive interface feature, to advertise a network into OSPF and not form OSPF adjacencies, it must be configured as a non-OSPF interface and the local network must be redistributed as an AS-external-LSA.

OSPF and IP

OSPF runs “on top of” IP, which means that an OSPF packet is sent with an IP data packet header. The protocol field in the IP header is set to 89 which identifies it as OSPF, distinguishing it from other packets that use an IP header.

A destination in an OSPF route advertisement is expressed as an IP address and a variable-length mask. Taken together, the address and the mask indicate the range of destinations to which the advertisement applies.

The ability to specify a range of networks allows OSPF to send one summary advertisement that represents multiple destinations. For example, a summary advertisement for the destination 128.185.0.0 with a mask of 255.255.0.0 describes a single route to destinations 128.185.0.0 to 128.185.255.255.

OSPF packets

All OSPF packets start with a 24 octet header that contain information about the OSPF version, the packet type and length, the ID of the router transmitting the packet, and the ID of the OSPF area from which the packet is sent. An OSPF packet will be one of the following types:

- Hello packets

Hello packets are transmitted between neighbors and are never forwarded. The Hello Protocol requires routers to send hello packets to neighbors at pre-defined hello intervals. If hello packets are not received by a neighbor router within the specified dead interval, the neighbor router will declare the other router dead.

- Database description (DD) packets

DD packets are exchanged when a link is first established between neighboring routers which synchronize their link state databases.

- Link state request packets

Link state request packets describe one or more link state advertisements that a router is requesting from its neighbor. Routers send link state requests if the information received in DD packets from a neighbor is not consistent with its own link state database.

- Link state update packets

Link state update packets contain one or more link state advertisements, and are sent following a change in network conditions.

- Link state acknowledgement packets

Link state acknowledgement packets are sent to acknowledge receipt of link state updates, containing the headers of the link state advertisements that were received.

Link state advertisements

OSPF does not require each router to send its entire routing table to its neighbors. Instead, each OSPF router floods only link-state change information in the form of link-state advertisements (LSAs) throughout the area or AS. LSAs in OSPF are one of the following five types:

- Router links advertisement

A router links advertisement is flooded only within the area and contains information about neighbor routers and the LANs to which the router is attached. A backbone router can flood router link advertisements within the backbone area.

- Network links advertisement

A network links advertisement is generated by a DR on a LAN, listing all routers on that LAN and flooding only within the area. A backbone DR can flood network links advertisements within the backbone area.

- Network summary link advertisement

A network summary link advertisement is flooded into an area by an ABR that describes networks that are reachable outside the area. An ABR attached to two areas will generate a different network summary link advertisement for each of these areas. ABRs also generate area summary link advertisements containing information about destinations within an area, which are flooded to the backbone area.

- ASBR summary link advertisement

An ASBR summary link advertisement describes the cost of the path to an ASBR from the router generating the advertisement.

- AS external link advertisement

An AS external link advertisement is sent by an ASBR to describe the cost of the path to a destination outside the AS from the ASBR generating the advertisement. This information is flooded to all routers in the AS.

AS external routes

OSPF considers the following routes to be *AS external (ASE)* routes:

- A route to a destination outside the AS
- A static route
- A default route
- A route derived by RIP
- A directly connected network not running OSPF

OSPF virtual links


On an OSPF network, an Passport 8000 switch which is acting as an ABR must be connected directly to the backbone. If no physical connection is available, a virtual link can be established which you can configure automatically or manually.

An automatic virtual link can provide redundancy support for critical network connections. Automatic virtual linking creates virtual paths for vital traffic paths in your OSPF network. In the event of a connection failure on the network, such as when an interface cable providing connection to the backbone (either directly or indirectly) becomes disconnected from the switch, the virtual link is available to maintain connectivity.

Specifying automatic virtual linking ensures that a link will be created via another router. When you specify automatic virtual linking, it is always ready to create a virtual link. If automatic virtual linking uses more resources than you want to expend, creating a manual virtual link may be the better solution. This approach lets you conserve resources while having specific control of where virtual links are placed in your OSPF configuration.

Figure 13 shows how to configure a virtual link between the ABR in area 2.2.2.2 and the ABR in area 0.0.0.0.

Figure 13 Virtual link between ABRs through a transit area



11017fa

To configure a virtual link between the ABRs in Area 1 and Area 3, you define Area 2 as the transit area between the other two areas, and identify R2 as the neighbor router through which R2 must send information to reach the backbone via R1.

Specifying ASBRs

ASBRs advertise non-OSPF routes into OSPF domains so that they can be passed along throughout the OSPF routing domain. A router can function as an ASBR if one or more of its interfaces is connected to a non-OSPF network (for example, RIP, BGP, or EGP).

To conserve resources, you may want to limit the number of ASBRs in your network or to specifically control which routers perform as ASBRs to control traffic flow.

Metric Speed

For OSPF, the “best” path to a destination is the path that offers the least-cost metric delay. In OSPF, cost metrics are configurable, allowing you to specify preferred paths. You can configure metric speed globally or for specific ports and interfaces on your network. In addition, you can control redistribution options between non-OSPF interfaces and OSPF interfaces.

Default metric speeds are assigned for different port types, such as 10Mb/s or 100Mb/s ports. On a Passport 8000 switch, you can specify a new metric speed for an IP interface. An IP interface can be a brouter port or a VLAN.




Note: On the Passport 8000 switch when you enable a port for OSPF routing, the default metric in the port window in Device Manager is “0.” A value of “0” (zero) means that the port will use the default metrics for port types that are specified on the OSPF general window.

Circuitless IP

Circuitless IP (CLIP) is a virtual (or loop back) interface that is not associated with any *physical* port. You can use the CLIP interface to provide uninterrupted connectivity to your switch *as long as there is an actual path to reach the device*.

For example, as shown in [Figure 14 on page 83](#), a physical point-to-point link exists between R1 and R2 along with the associated addresses (195.39.1.1/30 and 195.39.1.2/30). Note also that an IBGP session exists between two additional addresses 195.39.128.1/30 (CLIP 1) and 195.39.281.2/30 (CLIP 2).

CLIP 1 and CLIP 2 represent the virtual CLIP addresses that are configured between R1 and R2. These virtual interfaces are not associated with the physical link or hardware interface. This allows the IBGP session to continue as long as there is a path between R1 and R2. An IGP (such as OSPF) is used to route addresses corresponding to the CLIP addresses. After all the CLIP addresses are learned by the routers in the AS, the IBGP is established and routes can be exchanged.

Figure 14 Routers with IBGP connections

The CLIP interface is treated as any other IP interface. The network associated with the CLIP is treated as a local network attached to the device. This route always exists and the circuit is always up because there is no physical attachment.

Routes are advertised to other routers in the domain either as external routes using the route-redistribution process or when you enable OSPF in a passive mode to advertise an OSPF internal route. You can configure the OSPF protocol only on the circuitless IP interface.

When you create a CLIP interface, the system software programs a local route with the CPU as destID. All packets that are destined to the CLIP interface address are processed by the CPU. Any other packets with destination addresses associated with this network (but not to the interface address) are treated as if they are from any unknown host.

For more information about:	See:
Using Device Manager to configure circuitless IP	Chapter 3, "Configuring IP routing using Device Manager," on page 199
Using the CLI to configure circuitless IP	Chapter 4, "Configuring IP routing using the CLI," on page 239
Circuitless IP configuration examples	Chapter 2, "IP routing configuration examples," on page 93

HA-CPU/Layer 3 CPU Redundancy

HA-CPU/Layer 3 CPU Redundancy provides continuous operation of Passport 8600 features when there is a CPU failover. The HA-CPU takes over in HA CPU mode from the Master CPU that has failed, so that operations and processes are not interrupted.

To access the HA CPU in root node, use the following command:

```
config bootconfig flag
```

To enable or disable HA CPU mode, use the following commands:

```
ha cpu true
```

Or,

```
ha cpu false
```

Once the HA-CPU flag is changed, the router must be rebooted.

Both the Master CPU and the HA-CPU must run the same version of the software, because hitless upgrade is not supported.

HA-CPU/Layer 3 CPU Redundancy provides redundancy for the following:

- OSPFv2, including MD5
- RIPv1, RIPv2
- Prefix lists and Route policies
- ECMP/Alternate Routes
- VRRP
- IRDP (ICMP Route Discovery Protocol)
- IEEE 802.3ad
- IEEE 802.1x
- DHCP Relay Agent
- UDP forwarding

- IP Filters



Note: NOTE: PCAP is now supported in HA, if you reboot the secondary CPU, and enable PCAP.

OSPF

HA-CPU/Layer 3 Redundancy avoids disruption of network traffic when a Master CPU that is running OSPF fails over. It maintains an exact copy of the OSPF instance of the Master CPU on the HA-CPU. When the HA-CPU comes up, all OSPF information on the Master CPU is Table Synchronized and all OSPF events are Event Synchronized to the HA-CPU. When there is a Master CPU fail-over, the OSPF instance on HA-CPU resumes without affecting router traffic and OSPF neighbors.

During HA-CPU to Master CPU transition, it may take up to 3 seconds for the New Master CPU to transmit OSPF packets. Therefore, router dead intervals of 5 seconds or higher are recommended.

RIP

HA-CPU/ Layer 3 Redundancy for RIP allows for CPU fail-over without altering any existing RIP routing information in the entire network and without disrupting network traffic.

HA-CPU/ Layer 3 Redundancy for RIP synchronizes all RIP routing and interface information from the Master CPU to the HA-CPU so that the RIP instances on both CPUs are exactly the same. The RIP Route TTL, however, may have different values between the Master CPU and the HA-CPU because of the nature of the synchronization process.

At fail-over transition, the new Master CPU may miss some RIP Update packets if they were in the receiving queue of the old Master CPU or if the Event Synchronization message did not reach or did not finish processing in the old HA-CPU. This information may be recovered by the RIP Protocol.

Prefix Lists and Route Policy

The Route Policy and Prefix List related configurations are synchronized to the HA-CPU from the Master CPU so that the routes announced or accepted from one protocol domain to another are not affected when Master CPU fails over. Table Synchronization synchronizes the configuration to the HA-CPU when it comes up, and any events triggered in the Master are notified to the HA-CPU by Event synchronization.

VRRP

VRRP provides layer 3 redundancy by protocol perspective. When the VRRP master router fails, the backup virtual router takes a period of time to function as the VRRP master. Layer 3 CPU redundancy of VRRP prevents disruption of IP routing and forwarding operations, and protects networks with Virtual Routers from interruption.

Layer 3 CPU redundancy of VRRP also provides protection and faster fail-over than protocol failover. VRRP statistics are not synchronized from the primary CPU to the secondary CPU. In HA-CPU mode, VRRP Fast advertisements is not supported.

Route Discovery

Layer 3 redundancy for Route Discovery synchronizes the Route Discovery configuration from the master CPU to the HA-CPU so that the Router Discovery advertisements are sent to the hosts without any delay when the Master CPU fails over. Layer 3 redundancy does this using Table Synchronization and Event Synchronization.

Router Solicitation messages are not synchronized from the Master CPU to the HA-CPU. If the Master CPU received the Solicitation message from the host and the Master CPU fails over, then the Router Advertisements from the new Master are sent only if the timer expires or it receives one more Solicitation message.

DHCP Relay

Layer 3 Redundancy for DHCP Relay synchronizes the DHCP relay configuration of the master CPU to the relay configuration of the HA-CPU so that DHCP requests are forwarded to the DHCP server without any delay when the master CPU fails. Any event triggered in the master CPU is then synchronized to the HA-CPU by Event Synchronization.

DHCP related packets received in the Master CPU are not synchronized to the HA-CPU. If a request is received in the master CPU, but it switches over to the HA-CPU before forwarding the request, the information is lost when the HA-CPU takes over. However, the client continues to send the DHCP Discover packet again until it receives the DHCP Offer packet. The DHCP statistics are not synchronized to the HA-CPU.

UDP Forwarding

Table Synchronization and Event Synchronization allow UDP Forwarding configurations to be notified to the HA-CPU. The UDP broadcasts are forwarded to the respective server without any delay once the Master CPU fails over.

UDP broadcast packets received in the Master CPU are not synchronized to the HA-CPU. If the Master received a packet and failover occurs before it forwards it, the forwarding will not occur until the client sends one more broadcast.

IP Filters

IP Traffic Filter Redundancy protects network filter activities when there is a Master CPU fail-over. It uses existing HA-CPU Table Synchronization (Table Sync) and Event Synchronization (Event Sync) mechanisms to synchronize all IP Traffic Filter information between the Master CPU and the HA-CPU. This allows the Master CPU and HA-CPU to have exactly the same IP Filter information.

Since IP Traffic Filter Counter information is retrieved directly from system hardware, only the Master CPU can provide correct Filter Counter information, not the HA-CPU.

RSMLT

In many cases, core network convergence-time is dependent on the length of time a routing protocols requires to successfully convergence. Depending on the specific routing protocol, this convergence time can cause network interruptions ranging from seconds to minutes.

The Nortel Networks RSMLT feature allows rapid failover for core topologies by providing an *active-active* router concept to core SMLT networks.

Supported scenarios are: SMLT triangles, squares and SMLT full mesh topologies, with routing enabled on the core VLANs.

Routing protocols can be any of the following protocol types: IP Unicast Static Routes, RIP1, RIP2, OSPF, BGP and IPX RIP.

In the case of core router failures RSMLT takes care of the packet forwarding, thus eliminating dropped packets during the routing protocol convergence.

SMLT/RSMLT operation in L3 environments

[Figure 15 on page 90](#) shows a typical redundant network example with user aggregation, core, and server access layers. To minimize the creation of many IP subnets, one VLAN (VLAN 1, IP subnet A) spans all wiring closets.

SMLT provides the loop-free topology and enables all links to be forwarding for VLAN 1, IP Subnet A.

The aggregation layer switches are configured with routing enabled and provide an active-active default gateway functionality through RSMLT.

In this case routers R1 and R2 are forwarding traffic for IP subnet A. RSMLT provides both router failover and link failover. For example, if the SMLT link in between R2 and R4 are broken, the traffic will failover to R1 as well.

For IP subnet A, VRRP with a Backup-Master could provide the same functionality as RSMLT, as long as no additional router is connected to IP subnet A.

RSMLT provides superior router redundancy in core networks (IP subnet B), where OSPF is used for the routing protocol. Routers R1 and R2 are providing router backup for each other, not only for the edge IP subnet A, but also for the core IP subnet B. Similarly routers R3 and R4 are providing router redundancy for IP subnet C and also for core IP subnet B.

Failure scenarios

Please refer to [Figure 15 on page 90](#) for the following failure scenarios.

Router R1 failure:

For example, R3 and R4 are using both R1 as their next hop to reach IP subnet A. Even though R4 sends the packets to R2, they will be routed directly at R2 into subnet A. R3 sends its packets towards R1 and they are also sent directly into subnet A. When R1 fails, all packets will be directed to R2, with the help of SMLT. R2 still routes for R2 and R1. After OSPF convergences, the routing tables in R3 and R4 change their next hop to R2 in order to reach IP subnet A. The network administrator can choose to set the hold-up timer (i.e., for the amount of time R2 will route for R1 in a failure case) for a time period greater than the routing protocol convergence, or set it as indefinite (i.e., the pair always routes for each other).

In the application where RSMLT is used at the edge instead of VRRP, the hold-up timer value of indefinite is recommended.


Router R1 recovery

When R1 reboots after a failure, it becomes active as a VLAN bridge first. Packets destined to R1 are switched, using the bridging forwarding table, to R2 for as long as the hold down timer is configured. Those packets are routed at R2 for R1. Similar to VRRP, the hold down timer value needs to be greater than what the routing protocol requires to converge its tables.

When the hold down time expires and the routing tables have converged, R1 starts routing packets for itself and also for R2. Therefore, it does not matter which one of the two routers is used as the next hop from R3 and R4 to reach IP subnet A.

If single-homed IP subnets are configured on R1 or R2, it is recommended to add another routed VLAN to the ISTs with lower routing protocol metrics as a traversal VLAN/subnet in order to avoid unnecessary ICMP redirect generation messages. This recommendation is also applicable to VRRP implementations.

Figure 15 SMLT and RSMLT in L3 environments



Designing and configuring an RSMLT network

Because RSMLT is based on SMLT, all SMLT configuration rules apply. In addition, RSMLT is enabled on the SMLT aggregation switches on a per VLAN basis. The VLAN has to be a member of SMLT links and the IST trunk.

The VLAN also must be routable (IP address configured) and on all four routers (as shown in [Figure 15 on page 90](#)) an Interior Routing Protocol (IGP) such as OSPF has to be configured, although it is independent from RSMLT.

There are no changes to any IGP state machines and any routing protocol, even static routes, can be used with RSMLT.

RSMLT pair switches provide backup for each other. As long as one of the two routers of an IST pair is active, traffic forwarding is available for both next hops R1/R2 and R3/R4.

Chapter 2

IP routing configuration examples

This chapter provides configuration examples for common IP routing tasks and includes the CLI commands you use to create the example configuration.



Note: For a complete description of the CLI commands you can use to configure specific IP Routing tasks, including those shown in this chapter, see the appropriate CLI chapter in this guide (refer to [“Contents,” on page 5](#)).

This chapter includes the following topics:

Topic	Page
ARP configuration examples	93
RIP configuration examples	96
OSPF configuration examples	122
VRRP configuration examples	177

ARP configuration examples

The Passport 8600 switch provides the following Address Resolution Protocol (ARP) features:

- Default ARP aging
- Enabling of Proxy ARP
- Static ARP entries.

To communicate with devices that do not respond to ARP requests, you can configure a static ARP entry on the Passport 8600 switch. Alternatively, if you do not want to age out an *existing* ARP entry, you can configure a static ARP entry on the Passport 8600 switch (a static ARP entry maps the device's IP address to its MAC address).

When you configure a static ARP entry on the Passport 8600 switch, you assign both the IP address and the MAC address to the physical port, including the VLAN number if the physical port is associated with a VLAN.

This section includes the following topics:

- [“Adding a static ARP entry to a brouter port,”](#) next
- [“Adding a static ARP entry to a VLAN”](#) on page 94
- [“Deleting a static ARP entry”](#) on page 95
- [“Changing the default ARP aging time”](#) on page 95

Adding a static ARP entry to a brouter port

To add a static ARP entry to a brouter port, use the following command:

```
Passport-8610:5# config ip arp add ports <value> ip <value>
mac <value>
```

Where:

- add ports *value* is the slot/port number of the brouter port
- ip *value* is the IP address of the interface.
- mac *value* is the MAC address.

Example:

```
Passport-8610:5# config ip arp add ports 1/46 ip 172.2.2.13
mac 00:00:98:22:33:44
```

Adding a static ARP entry to a VLAN

To add a static ARP entry to a VLAN, use the following command:

```
Passport-8610:5# config ip arp add ports <value> ip <value>
mac <value> [vlan <value>]
```

Where:

- *add ports value* is the slot/port number of the brouter port.
- *ip value* is the IP address of the interface.
- *mac value* is the MAC address.
- *vlan value* is the VLAN number (if the physical port is associated with a VLAN).

Example:

```
Passport-8610:5# config ip arp add ports 1/48 ip 10.1.1.23  
mac 00:00:11:43:54:23 vlan 10
```

Deleting a static ARP entry

To delete a static entry, use the following command:

```
Passport-8610:5# config ip arp delete <ipaddr>
```

Where:

- *ipaddr* is the IP address of the static entry.

Example:

```
Passport-8610:5# config ip arp delete 172.2.2.13
```

Changing the default ARP aging time

The default ARP aging time value is set for 360 minutes. To change this value, use the following command:

```
Passport-8610:5# config ip arp aging <minutes>
```

Where:

- *minutes* is the arp lifetime in minutes in the range 1 and 32767 (the default value is 360 minutes).

Example:

```
Passport-8610:5# config ip arp aging 180
```

RIP configuration examples

Routing Information Protocol (RIP) is an interior gateway protocol (IGP), which is one of a class of algorithms known as distance vector algorithms. The hop count, or distance, is used as a metric to determine the best path to a remote network or host. The hop count cannot exceed 15 hops (assuming a cost of one hop for each network). RIP uses User Datagram Protocol (UDP) data packets to exchange routing information.

RIP sends routing information updates every 30 seconds. The updates contain information about known networks and the distances (hop count) associated with each. For RIPv1, no mask information is exchanged; the natural mask is always applied by the router receiving the update. Mask information is always included for RIPv2.

If information about a network is not received for 90 seconds, the metric associated with the network is raised to infinity (the metric is set for 16), and the network then becomes unreachable. If information about a network is not received for 180 seconds (six update intervals), it is removed from the routing table. These default timers can be changed by configuring the RIP Interface Timeout Timer parameter and Holddown Timer parameters.

This section provides examples of the common RIP configuration tasks and includes the CLI commands used to create the configuration.

The following topics are included:

- [“RIP send modes,”](#) next
- [“RIP configuration tasks”](#) on page 99
- [“Configuration example — Base configuration”](#) on page 99
- [“Configuration example — Configuring RIPv2”](#) on page 104
- [“Configuration example — Spanning tree in Passport 8000 routed networks”](#) on page 106
- [“Configuration example - Supplying a Default Route”](#) on page 108
- [“Configuration example - Using RIP accept policies”](#) on page 117
- [“Configuration example - Using RIP announce policies”](#) on page 121

RIP send modes

Table 4 describes the four RIP *send* modes that are supported on the Passport 8600 switch. You can configure RIP send modes on all router interfaces.

Table 4 RIP send modes

Send mode:	Description	Result
rip1comp	This mode is used to broadcast RIP-2 updates using RFC 1058 route consumption rules. This mode is the default value on the Passport 8600 switch.	- Destination MAC is a broadcast, ff-ff-ff-ff-ff-ff - Destination IP is a broadcast for the network (for example, 192.1.2.255) - RIP Update is formed as a RIP-2 update, including network mask - RIP version = 2
rip1	This mode is used to broadcast RIP updates that are compliant with RFC 1058.	- Destination MAC is a broadcast, ff-ff-ff-ff-ff-ff - Destination IP is a broadcast for the network (for example, 192.1.2.255) - RIP Update is formed as a RIP-1 update, no network mask included - RIP version = 1
rip2	This mode is used to broadcast multicast RIP-2 updates.	- Destination MAC is a multicast, 01-00-5e-00-00-09 - Destination IP is the RIP-2 Multicast address, 224.0.0.9 - RIP Update is formed as a RIP-2 update including network mask - RIP version = 2
nosend	No RIP updates are sent on the interface.	None

You can choose any of three options for receiving RIP updates:

- rip1OrRip2 — accepts RIPv1 or RIPv2 updates
- rip1 — accepts RIPv1 updates only
- rip2 — accepts RIPv2 updates only

Configuring send mode parameters

To configure your switch send mode parameters at the IP interface level, use the following command:

```
Passport-8610:5# config ip rip interface <ipaddr>  
send-mode <mode>
```

Where:

- *ipaddr* is the IP address of the RIP interface.
- *mode* indicates that you must enter a send mode value:
{*notsend*|*rip1*|*rip1comp*|*rip2*}.

Example:

```
Passport-8610:5# config ip rip interface 10.1.1.9  
send-mode rip2
```

Configuring receive mode parameters

To configure your switch receive mode parameters at the IP interface level, use the following command:

```
Passport-8610:5# config ip rip interface <ipaddr>  
receive-mode <mode>
```

Where:

- *ipaddr* is the IP address of the RIP interface.
- *mode* indicates that you must enter a receive mode value:
{*rip1*|*rip2*|*rip1orrip2*}.

Example:

```
Passport-8610:5# config ip rip interface 10.1.1.9  
receive-mode rip2
```

RIP configuration tasks

You can configure RIP on a VLAN or on a brouter port. If you configure RIP on a VLAN, the following tasks are required:

- Configure VLANs, add ports and STG group
- Enable RIP
- Disable supply RIP updates, if required
- Disable listen for RIP updates, if required

RIP Split Horizon¹ is enabled, by default. If you set the Poison parameter to true, Poison Reverse is enabled.

- Enable Default Route Supply if a default route exists in the route table

Default Route listen can be enabled to add a default route to the route table if advertised from another router.

- Add in or out Route Policy
- Enable Triggered Updates, if required
- Cost of the link. Enter a value of 1 to 15 where 1 is default.


Configuration example — Base configuration

As shown in [Figure 16 on page 100](#), Passport 8600 switch (R1) is configured between a Business Policy Switch 2000 (BPS1) and the edge of the Network core. Two VLANs (VLAN 2 and VLAN 3) are associated with BPS1.

For this example, R1 is configured as follows:

- R1 is using IP Subnet VLANs to provide routing between VLAN 2 and VLAN 3 on port 1/48, that is connected to BPS1.
- Core port (2/7) is configured as a brouter port with RIP.

1 If Split Horizon is invoked, IP routes learned from an immediate neighbor are not advertised back to the neighbor. If Poison Reverse is enabled, the RIP updates sent to a neighbor from whom a route is learned are "poisoned" with a metric of 16. Therefore, the receiver neighbor ignores this route because the metric 16 indicates infinite hops in the network. These mechanisms are used to prevent routing loops.

Figure 16 Configuration example—base configuration

11018fa

The following section provides step-by-step procedures that show how to configure R1 for this example.

Configuring R1

1 Configure tagging on port 1/48:

The following command configures tagging on port 1/48. Note that tagging is required to support multiple VLANs on the same interface.

```
Passport-8610:5# config ether 1/48 perform-tagging enable
```

2 Configure R1 for VLAN 2 access:

- a The following command creates VLAN = 2 using Spanning Tree Group = 1 and the VLAN type for the IP Subnet. If you are using another STG group, create the new STG group first, then add port 1/48 to the new STG group:

```
Passport-8610:5# config vlan 2 create byipsubnet 1
10.1.20.0/24
```

- b** The following commands configure port 1/48 as a static member for VLAN 2 and remove all other potential members:

```
Passport-8610:5# config vlan 2 ports remove
1/1-1/48,2/1-2/8 member portmember
Passport-8610:5# config vlan 2 ports add 1/48 member
portmember
Passport-8610:5# config vlan 2 ports add 1/48 member
static
```

- c** The following command adds the IP address of 10.1.20.2/24 to IP Subnet VLAN 2:

```
Passport-8610:5# config vlan 2 ip create 10.1.20.2/24
```

- d** The following commands enable RIP for VLAN 2 and disable RIP supply and listen. Note that RIP supply and listen are not required because there is no Router attached to VLAN 2:

```
Passport-8610:5# config vlan 2 ip rip enable
Passport-8610:5# config vlan 2 ip rip supply disable
Passport-8610:5# config vlan 2 ip rip listen disable
```

3 Configure R1 for VLAN 3 access:

- a** The following command creates VLAN = 3 using Spanning Tree Group = 1 and VLAN type of IP Subnet. If using another STG group, create the new STG group first, then add port 1/48 to the new STG group:

```
Passport-8610:5# config vlan 3 create byipsubnet 1
10.1.30.0/24
```

- b** The following commands configure port 1/48 as a static member for VLAN 3 and remove all other potential members:

```
Passport-8610:5# config vlan 3 ports remove
1/1-1/48,2/1-2/8 member portmember
Passport-8610:5# config vlan 3 ports add 1/48 member
portmember
Passport-8610:5# config vlan 3 ports add 1/48 member
static
```

- c** The following command adds the IP address of 10.1.30.2/24 to IP Subnet VLAN 3:

```
Passport-8610:5# config vlan 3 ip create 10.1.30.2/24
```

- d** The following commands enable RIP for VLAN 3 and disable RIP supply and listen. Note that RIP supply and listen are not required because there is no Router attached to VLAN 3:

```
Passport-8610:5# config vlan 3 ip rip enable
Passport-8610:5# config vlan 3 ip rip supply disable
Passport-8610:5# config vlan 3 ip rip listen disable
```

4 Configure brouter port 2/7 on R1:

- a** The following command adds the IP address of 10.1.1.1/30 to port 2/7 using brouter VLAN = 2090:

```
Passport-8610:5# config ethernet 2/7 ip create
10.1.1.1/30 2090
```

- b** The following command enables RIP on this interface:

```
Passport-8610:5# config ethernet 2/7 ip rip enable
```

5 Enable RIP globally:

```
Passport-8610:5# config ip rip enable
```

Displaying configuration files

You can use the following show command to display the configuration commands and parameters used to create the topology shown in [Figure 16 on page 100](#):

```
Passport-8610:5# show config
```



Note: You can copy and paste the command outputs shown here to update your configuration files.

Configuration file for R1

```
# PORT CONFIGURATION - PHASE I
#
ethernet 1/48 perform-tagging enable

# VLAN CONFIGURATION
#
vlan 1 ip igmp mrdisc mrdisc-enable disable
vlan 2 create byipsubnet 1 10.1.20.0/255.255.255.0
vlan 2 ports remove 1/1-1/47,2/1-2/8,3/1-3/8 member portmember
vlan 2 ports add 1/48 member portmember
vlan 2 ports add 1/48 member static
```

```
vlan 2 ports remove 1/1-1/47,2/1-2/8,3/1-3/8 member portmember
vlan 2 ip create 10.1.20.2/255.255.255.0 mac_offset 0
vlan 2 ip rip enable
vlan 2 ip rip listen disable
vlan 2 ip rip supply disable
vlan 3 create byipsubnet 1 10.1.30.0/255.255.255.0
vlan 3 ports remove 1/1-1/47,2/1-2/8 member portmember
vlan 3 ports add 1/48,3/1-3/8 member portmember
vlan 3 ports add 1/48 member static
vlan 3 ports remove 1/1-1/47,2/1-2/8 member portmember
vlan 3 ip create 10.1.30.2/255.255.255.0 mac_offset 1
vlan 3 ip rip enable
vlan 3 ip rip listen disable
vlan 3 ip rip supply disable
#
# PORT CONFIGURATION - PHASE II
#
ethernet 2/7 ip create 10.1.1.1/255.255.255.252 2090 mac_offset 2
ethernet 2/7 ip rip enable
#
# IP ROUTE POLICY CONFIGURATION
#
ip rip enable
```

Configuration example — Configuring RIPv2


When RIP is enabled on a VLAN or router port, the default settings are:

- Send Mode: rip1compatible
- Receive Mode: rip1orRip2

Depending on your configuration requirements, you may want to configure the Passport 8600 switch to only operate in RIPv1 mode or RIPv2 mode.

This configuration example (see [Figure 17](#)) shows how to configure R1 to only operate in RIPv2 mode.

Figure 17 Configuration example— configuring RIPv2



11019fa

The following section provides step-by-step procedures that show how to configure R1 to add RIP version 2 to VLAN 2, VLAN 3, and the router port.

Configuring R1

1 Configure RIPv2 on VLAN 2:

The following commands enable RIPv2 mode on the IP address used for VLAN 2.

```
Passport-8610:5# config ip rip interface 10.1.20.2  
send-mode rip2  
Passport-8610:5# config ip rip interface 10.1.20.2  
receive-mode rip2
```

2 Configure RIPv2 on VLAN 3:

The following commands enable RIPv2 mode on the IP address used for VLAN 3.

```
Passport-8610:5# config ip rip interface 10.1.30.2  
send-mode rip2  
Passport-8610:5# config ip rip interface 10.1.30.2  
receive-mode rip2
```

3 Configure RIPv2 on the brouter port:

The following commands enable RIPv2 mode on the IP address used for the brouter port.

```
Passport-8610:5# config ip rip interface 10.1.1.1  
send-mode rip2  
Passport-8610:5# config ip rip interface 10.1.1.1  
receive-mode rip2
```


Configuration example — Spanning tree in Passport 8000 routed networks

In the previous configuration example (see [“Configuration example — Configuring RIPv2” on page 104](#)), a brouter port is used to connect to the network core.

A brouter port is a single-port VLAN that can route IP packets as well as bridge all non-routable traffic. The difference between a brouter port and a standard IP protocol-based VLAN (that is configured for routing), is the brouter port’s routing interface is not affected by the port’s spanning tree state. Therefore, when you use a brouter port, the spanning tree protocol is eliminated from the backbone network.

If VLAN connectivity is required in the core to support non-IP protocols, be careful that the spanning tree does not cause blocked ports. Blocked ports can occur if you are using a *single* Spanning Tree Group (STG) instance, with multiple VLANs ([Figure 18](#)).

Figure 18 Single spanning tree group




11020fa

You can prevent blocked ports from occurring by configuring *multiple* STGs (Figure 19). The multiple STGs can be used to eliminate loops at Layer 2, while still permitting both Layer 2 and Layer 3 connectivity between devices.

If you are using VLANs in the core network, adhere to the following configuration rules:

- A VLAN can exist in only one STG.
- Use only one STG on Access ports.
- Use multiple STGs on Trunk ports.

Figure 19 Multiple spanning tree groups




Configuration example - Supplying a Default Route

In the configuration example shown in [Figure 20](#), Passport 8600 switch (R1) is configured to add a default route that is directed to the Network Core and advertised to R2 and R3.

For this example:

- Router ports are used for all core links
- Port-based VLANs are configured for local networks.

Figure 20 Supplying a default route



11022fa

The following sections provide step-by-step procedures that show how to configure R1, R2, and R3, for the example configuration shown in [Figure 20](#).

Configuring R1

This section describes how to configure R1 for the configuration example shown in [Figure 20 on page 108](#). To configure R1, use the following commands:

1 Configure R1 for access to VLAN 2:

- a The following command creates VLAN = 2 using Spanning Tree Group = 1. If you are using another STG group, create the new STG group first, then add port 1/48 to the new STG group:

```
Passport-8610:5# config vlan 2 create byport 1
```

- b The following command adds the access port 1/48 to VLAN 2.

```
Passport-8610:5# config vlan 2 ports add 1/48
```

- c The following command adds IP address 10.1.20.2/24 to VLAN 2:

```
Passport-8610:5# config vlan 2 ip create 10.1.20.2/24
```

- d The following commands enable RIP for VLAN 2 and disable RIP supply and listen. Unless there is an external router attached to VLAN 2, there is no need to supply or listen for RIP updates.

```
Passport-8610:5# config vlan 2 ip rip enable
Passport-8610:5# config vlan 2 ip rip supply disable
Passport-8610:5# config vlan 2 ip rip listen disable
```

2 Configure brouter port 2/7 on R1:

- a The following command adds IP address 10.1.1.1/30 to port 2/7, using brouter VLAN = 2090:

```
Passport-8610:5# config ethernet 2/7 ip create
10.1.1.1/30 2090
```

- b The following commands enable RIP and advertise a default route out this interface. **Note:** A RIP Out-Policy for this interface is required to advertise the local default route, which will be described later in this procedure:

```
Passport-8610:5# config ethernet 2/7 ip rip
default-supply enable
Passport-8610:5# config ethernet 2/7 ip rip enable
```

3 Configure brouter port 2/8 on R1:

- a** The following command adds IP address 10.1.1.5/30 to port 2/8, using brouter VLAN = 2091:

```
Passport-8610:5# config ethernet 2/8 ip create
10.1.1.5/30 2091
```

- b** The following commands enable RIP and advertise a default route out this interface. **Note:** A RIP Out-Policy for this interface is required to advertise the local default route, which will be described later in this procedure:

```
Passport-8610:5# config ethernet 2/8 ip rip
default-supply enable
Passport-8610:5# config ethernet 2/8 ip rip enable
```

4 Configure brouter port 2/1 on R1:

- a** The following command adds IP address 10.1.1.9/30 to port 2/1, using brouter VLAN = 2092:

```
Passport-8610:5# config ethernet 2/1 ip create
10.1.1.9/30 2092
```

- b** The following commands enable RIP for this interface. For this interface, we will also disable RIP supply and listen:

```
Passport-8610:5# config ethernet 2/1 ip rip enable
Passport-8610:5# config ethernet 2/1 ip rip listen
disable
Passport-8610:5# config ethernet 2/1 ip rip supply
disable
```

5 Create the default route:

The following command creates the static default route:

```
Passport-8610:5# config ip static-route create
0.0.0.0/0 next-hop 10.1.1.10 cost 1
```

6 Enable RIP globally:

The following command globally enables RIP:

```
Passport-8610:5# config ip rip enable
```

7 Create an IP prefix:

The following command adds a prefix list named “default” with the default route address. This address will be used for the Route Policy in Step 8:

```
Passport-8610:5# config ip prefix-list "default"  
add-prefix 0.0.0.0/0
```

8 Create a route policy:

The following commands create a route policy named “default_route” with a match for the IP Prefix list created in Step 7:

```
Passport-8610:5# config ip route-policy  
"default_route" seq 1 create  
Passport-8610:5# config ip route-policy  
"default_route" seq 1 enable  
Passport-8610:5# config ip route-policy  
"default_route" seq 1 action permit  
Passport-8610:5# config ip route-policy  
"default_route" seq 1 match-network "default"
```

9 RIP policy configuration:

The following commands add the route policy created in Step 8 to the two core links to R2 and R3, from R1:

```
Passport-8610:5# config ip rip interface 10.1.1.1  
out-policy "default_route"  
Passport-8610:5# config ip rip interface 10.1.1.5  
out-policy "default_route"
```

Configuring R2

This section describes how to configure R2 for the configuration example shown in [Figure 20 on page 108](#). To configure R2, use the following commands:

1 Configure R2 for access to VLAN 2:

- a** The following command creates VLAN = 2 using Spanning Tree Group = 1. If you are using another STG group, create the new STG group first, then add port 1/48 to the new STG group:

```
Passport-8610:5# config vlan 2 create byport 1
```

- b** The following command adds the access port 1/48 to VLAN 2.

```
Passport-8610:5# config vlan 2 ports add 1/48
```

- c** The following command adds IP address 10.1.30.2/24 to VLAN 2:

```
Passport-8610:5# config vlan 2 ip create 10.1.30.2/24
```

- d** The following commands enable RIP for VLAN 2 and disable RIP supply and listen.

```
Passport-8610:5# config vlan 2 ip rip enable
Passport-8610:5# config vlan 2 ip rip supply disable
Passport-8610:5# config vlan 2 ip rip listen disable
```

2 Configure brouter port 2/7 on R2:

- a** The following command adds IP address 10.1.1.2/30 to port 2/7, using brouter VLAN = 2090:

```
Passport-8610:5# config ethernet 2/7 ip create
10.1.1.2/30 2090
```

- b** The following commands enable RIP and default route listen for this interface:

```
Passport-8610:5# config ethernet 2/7 ip rip
default-listen enable
Passport-8610:5# config ethernet 2/7 ip rip enable
```

3 Configure brouter port 2/1 on R2:

- a** The following command adds IP address 10.1.1.13/30 to port 2/1, using brouter VLAN = 2091:

```
Passport-8610:5# config ethernet 2/1 ip create
10.1.1.13/30 2091
```

- b** The following commands enable RIP and default route listen for this interface:

```
Passport-8610:5# config ethernet 2/1 ip rip
default-listen enable
Passport-8610:5# config ethernet 2/1 ip rip enable
```

4 Enable RIP globally:

The following command globally enables RIP:

```
Passport-8610:5# config ip rip enable
```


Configuring R3

This section describes how to configure R3 for the configuration example shown in [Figure 20 on page 108](#). To configure R3, use the following commands:

1 Configure R3 for access to VLAN 2:

- a** The following command creates VLAN = 2 using Spanning Tree Group = 1. If you are using another STG group, create the new STG group first, then add port 1/48 to the new STG group:

```
Passport-8610:5# config vlan 2 create byport 1
```

- b** The following command adds the access port 1/48 to VLAN 2.

```
Passport-8610:5# config vlan 2 ports add 1/48
```

- c** The following command adds IP address 10.1.20.2/24 to VLAN 2:

```
Passport-8610:5# config vlan 2 ip create 10.1.20.2/24
```

- d** The following commands enable RIP for VLAN 2 and disable RIP supply and listen.

```
Passport-8610:5# config vlan 2 ip rip enable
Passport-8610:5# config vlan 2 ip rip supply disable
Passport-8610:5# config vlan 2 ip rip listen disable
```

2 Configure brouter port 2/8 on R3:

- a** The following command adds IP address 10.1.1.6/30 to port 2/8, using brouter VLAN = 2090:

```
Passport-8610:5# config ethernet 2/8 ip create
10.1.1.6/30 2090
```

- b** The following commands enable RIP and default route listen for this interface:

```
Passport-8610:5# config ethernet 2/8 ip rip
default-listen enable
Passport-8610:5# config ethernet 2/8 ip rip enable
```

3 Configure brouter port 2/1 on R3:

- a** The following command adds IP address 10.1.1.14/30 to port 2/1, using brouter VLAN = 2091:

```
Passport-8610:5# config ethernet 2/1 ip create
10.1.1.14/30 2091
```

- b** The following commands enable RIP and default route supply and listen for this interface:

```
Passport-8610:5# config ethernet 2/1 ip rip
default-listen enable
Passport-8610:5# config ethernet 2/1 ip rip
default-supply enable
Passport-8610:5# config ethernet 2/1 ip rip enable
```

4 Enable RIP globally:

The following command globally enables RIP:

```
Passport-8610:5# config ip rip enable
```

Displaying configuration files

You can use the following show command to display the configuration commands and parameters used to create the topology shown in [Figure 20 on page 108](#):

```
Passport-8610:5# show config
```



Note: You can copy and paste the command outputs shown here to update your configuration files.

Configuration file for R1

```
#
# VLAN CONFIGURATION
#
vlan 2 create byport 1
vlan 2 ports remove 1/1-1/47,2/1-2/8,3/1-3/8 member portmember
vlan 2 ports add 1/48 member portmember
vlan 2 ip create 10.1.20.2/255.255.255.0 mac_offset 0
vlan 2 ip rip enable
vlan 2 ip rip listen disable
vlan 2 ip rip supply disable
#
# PORT CONFIGURATION - PHASE II
#
ethernet 2/1 ip create 10.1.1.9/255.255.255.252 2092 mac_offset 3
ethernet 2/1 ip ospf metric 0
ethernet 2/1 ip rip enable
ethernet 2/1 ip rip listen disable
ethernet 2/1 ip rip supply disable
```

```
ethernet 2/7 ip create 10.1.1.1/255.255.255.252 2090 mac_offset 1
ethernet 2/7 ip rip enable
ethernet 2/7 ip rip default-supply enable
ethernet 2/8 ip create 10.1.1.5/255.255.255.252 2091 mac_offset 2
ethernet 2/8 ip rip enable
ethernet 2/8 ip rip default-supply enable
#
# IP PREFIX LIST CONFIGURATION
#
ip prefix-list "default" add-prefix 0.0.0.0/0 maskLenFrom 0
maskLenTo 0
#
# IP ROUTE POLICY CONFIGURATION
#
ip route-policy "default_route" seq 1 create
ip route-policy "default_route" seq 1 enable
ip route-policy "default_route" seq 1 action permit
ip route-policy "default_route" seq 1 match-network "default"
ip route-policy "default_route" seq 1 set-metric-type type2
ip route-policy "default_route" seq 1 set-nssa-pbit enable
#
ip static-route create 0.0.0.0/0.0.0.0 next-hop 10.1.1.10 cost 1
preference 5
ip rip enable
ip rip interface 10.1.1.1 send-mode rip2
ip rip interface 10.1.1.5 send-mode rip2
ip rip interface 10.1.20.2 send-mode rip2
#
# RIP POLICY CONFIGURATION
#
ip rip interface 10.1.1.1 out-policy "default_route"
ip rip interface 10.1.1.5 out-policy "default_route"
```

Configuration file for R2

```
#
# VLAN CONFIGURATION
#
vlan 2 create byport 1
vlan 2 ports remove 1/1-1/47,2/1-2/8,3/1-3/8 member portmember
vlan 2 ports add 1/48 member portmember
vlan 2 ip create 10.1.30.2/255.255.255.0 mac_offset 0
vlan 2 ip rip enable
vlan 2 ip rip listen disable
vlan 2 ip rip supply disable
#
# PORT CONFIGURATION - PHASE II
#
```

```
ethernet 2/1 ip create 10.1.1.13/255.255.255.252 2091 mac_offset 2
ethernet 2/1 ip ospf metric 0
ethernet 2/1 ip rip enable
ethernet 2/1 ip rip default-listen enable
ethernet 2/1 ip rip default-supply enable
ethernet 2/7 ip create 10.1.1.2/255.255.255.252 2090 mac_offset 1
ethernet 2/7 ip ospf metric 0
ethernet 2/7 ip rip enable
ethernet 2/7 ip rip default-listen enable
ethernet 2/7 ip rip default-supply enable
ethernet 2/8 state disable
# IP ROUTE POLICY CONFIGURATION
#
ip rip enable
```

Configuration file for R3

```
#
# VLAN CONFIGURATION
#
vlan 2 create byport 1
vlan 2 ip create 10.1.40.2/255.255.255.0 mac_offset 0
vlan 2 ip rip enable
vlan 2 ip rip listen disable
vlan 2 ip rip supply disable
#
# PORT CONFIGURATION - PHASE II
#
ethernet 2/1 ip create 10.1.1.14/255.255.255.252 2091 mac_offset 2
ethernet 2/1 ip ospf metric 0
ethernet 2/1 ip rip enable
ethernet 2/1 ip rip default-listen enable
ethernet 2/1 ip rip default-supply enable
ethernet 2/2 state disable
ethernet 2/8 ip create 10.1.1.6/255.255.255.252 2090 mac_offset 1
ethernet 2/8 ip ospf metric 0
ethernet 2/8 ip rip enable
ethernet 2/8 ip rip default-listen enable
ethernet 2/8 ip rip default-supply enable
# IP ROUTE POLICY CONFIGURATION
#
ip rip enable
```

Configuration example - Using RIP accept policies

You can use RIP Accept policies on the Passport 8600 switch to selectively accept routes from RIP updates. If no policies are defined, the default behavior is applied which adds all learned routes to the route table.

RIP Accept policies can be used to:

- Listen to RIP updates only from certain gateways.
- Listen only for specific networks.
- Assign a specific mask to be included with a network in the routing table (such as a network summary).

In the configuration example shown in [Figure 21 on page 118](#), Passport 8600 switch (R1) is configured with a RIP Accept policy, which creates a single route directed to R3 for all networks configured on it. The accept policy accepts any network from 10.1.240.0 to 10.1.255.0, and creates a single entry in the routing table on R1.

You can calculate a summary route, by comparing the common bits in the address range to derive the summary address.

For example, if the range of IP addresses is from 10.1.240.0 to 10.1.255.0:

- 1 Determine the *third* octet of the first address:

10.1.**240**.0 = 1111 0000

- 2 Determine the *third* octet of the ending address:

10.1.**255**.0 = 1111 1111

- 3 Extract the common bits:


240 = ~~1111~~ 0000

255 = ~~1111~~ 1111

1111 = 20 bit mask

Therefore, the network address to use for this example is 10.1.240.0/[20](#).

Figure 21 RIP accept policy



11023fa

The following section provides step-by-step procedures that show how to configure R1 for the example configuration shown in [Figure 21](#).

Configuring R1

1 Configure the IP prefix list on R1:

The following command creates a prefix list named Prefix_1 with an IP range from 10.1.240.0 to 10.1.255.0

```
Passport-8610:5# config ip prefix-list "Prefix_1"  
add-prefix 10.1.240.0/20 maskLenFrom 20 maskLenTo 32
```

2 Configure the route policy:

The following commands configure a route policy named "rip_pol_1" with match criteria using the IP Prefix configured in Step 1. This has the effect of injecting one route of 10.1.240.0/20 into the route table.

```
Passport-8610:5# config ip route-policy "rip_pol_1" seq 1
create
Passport-8610:5# config ip route-policy "rip_pol_1" seq 1
enable
Passport-8610:5# config ip route-policy "rip_pol_1" seq 1
action permit
Passport-8610:5# config ip route-policy "rip_pol_1" seq 1
match-network "Prefix_1"
Passport-8610:5# config ip route-policy "rip_pol_1" seq 1
set-injectlist "Prefix_1"
```

3 Add the Route Policy to the appropriate RIP interfaces:

The following commands add the Route Policy created in Step 2 to both RIP core ports.

```
Passport-8610:5# config ip rip interface 10.1.1.1
in-policy "rip_pol_1"
Passport-8610:5# config ip rip interface 10.1.1.5
in-policy "rip_pol_1"
```

Displaying configuration files

You can use the following show command to display the configuration commands and parameters used to create the topology shown in [Figure 21 on page 118](#):

```
Passport-8610:5# show config
```



Note: You can copy and paste the command outputs shown here to update your configuration files.

Configuration file for R1

```
#
# IP PREFIX LIST CONFIGURATION
#
ip prefix-list "Prefix_1" add-prefix 10.1.240.0/20 maskLenFrom 20
maskLenTo 32
#
# IP ROUTE POLICY CONFIGURATION
#
ip route-policy "rip_pol_1" seq 1 create
ip route-policy "rip_pol_1" seq 1 enable
ip route-policy "rip_pol_1" seq 1 action permit
ip route-policy "rip_pol_1" seq 1 match-network "Prefix_1"
ip route-policy "rip_pol_1" seq 1 set-injectlist "Prefix_1"
ip route-policy "rip_pol_1" seq 1 set-metric-type type2
ip route-policy "rip_pol_1" seq 1 set-nssa-pbit enable
#
ip rip enable
#
# RIP POLICY CONFIGURATION
#
ip rip interface 10.1.1.1 in-policy "rip_pol_1"
ip rip interface 10.1.1.5 in-policy "rip_pol_1"
```


Configuration example - Using RIP announce policies

In the previous configuration example (see [“Configuration example - Using RIP accept policies” on page 117](#)), a RIP Accept policy is used on R1 to insert a single route into its route table for all networks from R3. Instead of using an Accept Policy on R1, you could use a RIP Announce Policy on R3 to announce a single route to both R1 and R2 for its local network range.

To configure the RIP Announce Policy on R3 (refer to [Figure 21 on page 118](#)), use the following configuration steps.

Configuring R3

- 1 Configure the IP prefix list on R3:

The following command creates a prefix list named Prefix_1 with IP address 10.1.240.0.

```
Passport-8610:5# config ip prefix-list "Prefix_1"  
add-prefix 10.1.240.0/20
```

- 2 Configure the route policy:

The following commands configure a route policy named "Policy_Rip" with match criteria using the IP Prefix configured in Step 1.

```
Passport-8610:5# config ip route-policy "rip_pol_1" seq 1  
create  
Passport-8610:5# config ip route-policy "rip_pol_1" seq 1  
enable  
Passport-8610:5# config ip route-policy "rip_pol_1" seq 1  
action permit  
Passport-8610:5# config ip route-policy "rip_pol_1" seq 1  
set-injectlist "Prefix_1"
```

- 3 Add the Route Policy to the appropriate RIP interfaces:

The following commands add the Route Policy created in Step 2 to both RIP core ports.

```
Passport-8610:5# config ip rip interface 10.1.1.14  
out-policy "Policy_Rip"  
Passport-8610:5# config ip rip interface 10.1.1.6  
out-policy "Policy_Rip"
```

OSPF configuration examples

The Open Shortest Path First (OSPF) protocol is a link-state protocol designed as a standards-based Internal Gateway Protocol (IGP) for interconnecting users and networks. OSPF maintains a link-state database of interface, link, router, and network status to calculate the shortest path to every network element. The Passport 8600 switch uses the link-state database to build a routing table. This calculation is based on Dijkstra's algorithm¹ model of calculating the shortest path from one point to another.

The Passport 8600 switch supports the following OSPF standards:

- RFC 2328 (OSPF version 2)
- RFC 1850 (OSPF Management Information Base)
- RFC 2178 (OSPF MD5 cryptographic authentication)

This section provides examples of the common OSPF configuration tasks and includes the CLI commands used to create the configuration.

The following topics are included:

- [“Configuration example — OSPF interface types,” next](#)
- [“Configuration example — Equal Cost Multi Path” on page 127](#)
- [“Configuration example — OSPF security mechanisms” on page 130](#)
- [“Configuration example — Diagnosing OSPF neighbor state problems” on page 134](#)
- [“Configuration example — OSPF network types” on page 138](#)
- [“Configuration example — OSPF area types” on page 139](#)
- [“Configuration example — OSPF ABR” on page 151](#)
- [“Configuration examples — OSPF ASBR configurations” on page 154](#)
- [“Configuration example — Controlling NSSA external routes advertised” on page 162](#)
- [“Configuration example — Multi-area complex” on page 169](#)

1 Dijkstra's algorithm, named after its discoverer, E.W. Dijkstra, solves the problem of finding the shortest path from a point in a graph (the source) to a destination. This calculation is used to determine the best path to any network based on the total path cost. All paths to a given network are determined and the cost calculated, however, only the best path will be used populate the routing table.

Configuration example — OSPF interface types

This section describes configuration examples for two OSPF interface types:

- “[Configuring a circuitless IP interface](#),” next
- “[Configuring an IP OSPF interface](#)” on page 125.


Configuring a circuitless IP interface

A circuitless IP (CLIP) address, sometimes referred to as a loopback address, is an IP address that is not tied to any specific interface. Because the CLIP is not tied to a physical port or VLAN, the CLIP state is always active.

Nortel Networks recommends that you use the CLIP address for the OSPF Router-ID. By doing so, the OSPF Router-ID is always active, regardless of the port state (up/down).

The CLIP interface is treated as any other IP interface and the network associated with the CLIP address is treated as a local network attached to the device. This route always exists and the circuit is always up because there is no physical attachment ([Figure 22](#)).

Figure 22 CLIP interface



11024fa

The following sections provide step-by-step procedures that show how to configure R1 and R2 for this example.

Configuring R1

This section describes how to configure CLIP on R1 and use it for the OSPF Router-ID. To configure CLIP, use the following commands:

1 Define a CLIP address on R1:

The following commands create a circuitless IP address 195.39.128.1/32 which is used for the OSPF Router-ID (where X is the CLIP ID and can be any instance from 1-32). The CLIP IP address typically uses a 32-bit mask.

```
Passport-8610:5# config ip circuitless-ip-int x create  
195.39.128.1/32  
Passport-8610:5# config ip circuitless-ip-int x ospf  
enable
```

2 Enable OSPF on the CLIP:

The following commands enable OSPF and sets the OSPF Router-ID using the CLIP address created in Step 1.

```
Passport-8610:5# config ip ospf admin-state enable  
Passport-8610:5# config ip ospf router-id 195.39.128.1/32  
Passport-8610:5# config ip ospf enable
```

Configuring R2

This section describes how to configure CLIP on R2 and use it for the OSPF Router-ID. To configure CLIP, use the following commands:

1 Define a CLIP address on R2:

The following commands create a circuitless IP address 195.39.128.2/32 which is used for the OSPF Router-ID (where y is the CLIP ID and can be any instance from 1-32). The CLIP IP address typically uses a 32-bit mask.

```
Passport-8610:5# config ip circuitless-ip-int y create  
195.39.128.2/32  
Passport-8610:5# config ip circuitless-ip-int y ospf  
enable
```

2 Enable OSPF on the CLIP:


The following commands enable OSPF and set the OSPF Router-ID using the CLIP address created in Step 1.

```
Passport-8610:5# config ip ospf admin-state enable
Passport-8610:5# config ip ospf router-id 195.39.128.2/32
Passport-8610:5# config ip ospf enable
```

Configuring an IP OSPF interface

You can configure an IP OSPF interface at a brouter port interface level or at a VLAN (port or IP-Subnet) level ([Figure 23](#)).

Figure 23 OSPF example



The following steps show how to configure OSPF on brouter port 2/1 and VLAN 2, as shown in [Figure 23](#).

1 Configure OSPF interface — brouter port

The following commands configure port 2/1 as a brouter port with VLAN ID 2134 and enable OSPF on this interface:

```
Passport-8610:5# config ethernet 2/1 ip create
10.1.1.21/30
Passport-8610:5# config ethernet 2/1 ip ospf enable
```

2 Configure OSPF interface — VLAN

The following commands create port-based VLAN 2 under STG 1 with OSPF.

```
Passport-8610:5# config VLAN 2 create byport 1
Passport-8610:5# config VLAN 2 ports add 1/2
Passport-8610:5# config VLAN 2 ip create 172.3.1.1/24
Passport-8610:5# config VLAN 2 ip ospf enable
```

3 Configure a CLIP address

The following commands create a CLIP address, which is used for the OSPF Router-ID:

```
Passport-8610:5# config ip circuitless-ip-int 1 create
1.1.1.1/32
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

4 Enable OSPF globally

The following commands enable OSPF and assign the CLIP address created in Step 3 as the OSPF Router-ID:

```
Passport-8610:5# config ip ospf router-id 1.1.1.1
Passport-8610:5# config ip ospf enable
```

Configuration example — Equal Cost Multi Path

Equal Cost Multi Path (ECMP) is an OSPF feature for load balancing routed IP traffic across (up to four) equal cost paths.

Some benefits you can gain with ECMP are:

- You do not need to rerun Dijkstra; if the main path fails, the other ECMP path(s) automatically take the load.
- Loadsharing implies better use of network facilities.
- The traffic distribution algorithm is identical to the MultiLink Trunk (MLT) algorithm for IP datagrams:
 - $\text{MOD}(\text{DestIP}(X)[5:0] \text{ XOR } \text{SrcIP}(Y)[5:0], \text{\#of active links})$
 - XOR the last 6 bits of the source and destination IP address, divide by the number of links, and take the remainder:

Example:

Assuming 192.1.1.3 sends to 192.1.1.4

3 = 00:00:00:00:11

4 = 00:00:00:01:00

XOR = 00:01:11 = 7

Divide by the number of ECMP ports (assume 4 for this example):


$7/4 = 3$

The remainder is 3, therefore this stream lines up with the fourth port of the 4-port ECMP group.

In the configuration example shown in [Figure 24](#), the following commands enable two ECMP paths for R1:

```
Passport-8610:5# config ip ecmp enable
Passport-8610:5# config ip ecmp-max-path 2
```

Figure 24 ECMP example



After you configure ECMP, you can verify the ECMP paths in the routing table. To display the routing table, use the following show command:

```
show ip route info
```

[Figure 25 on page 129](#) shows sample output for the `show ip route info` command.

As shown in [Figure 25](#), the paths shown with the letter “E” in the TYPE column are *designated* equal cost paths. In this example, you can see two routes to IP address 10.1.40.0, and two routes to IP address 10.1.30.0.

Figure 25 show ip route info

```
show ip route info
```

Response from R1:

Ip Route									
DST	MASK	NEXT	COST	VLAN	PORT	PROT	AGE	TYPE	PRF
1.1.1.1	255.255.255.255	1.1.1.1	1	0	-/-	LOC	0	DB	0
1.1.40.0	255.255.255.255	10.1.1.2	12	2190	2/7	OSPF	0	IBE	20
1.1.40.0	255.255.255.255	10.1.1.6	12	2191	2/8	OSPF	0	IBE	20
1.1.1.3	255.255.255.255	10.1.1.2	11	2190	2/7	OSPF	0	IB	20
1.1.1.4	255.255.255.255	10.1.1.6	11	2191	2/8	OSPF	0	IB	20
2.1.1.0	255.255.255.252	2.1.1.1	1	3999	-/-	LOC	0	DB	0
10.1.1.0	255.255.255.252	10.1.1.1	1	-	2/7	LOC	0	DB	0
10.1.1.4	255.255.255.252	10.1.1.5	1	-	2/8	LOC	0	DB	0
10.1.1.8	255.255.255.252	10.1.1.6	2	2191	2/8	OSPF	0	IB	20
10.1.1.12	255.255.255.252	10.1.1.2	2	2190	2/7	OSPF	0	IB	20
10.1.20.0	255.255.255.0	10.1.20.2	1	2	-/-	LOC	0	DB	0
10.1.30.0	255.255.255.0	10.1.1.2	11	2190	2/7	OSPF	0	IBE	20
10.1.30.0	255.255.255.0	10.1.1.6	11	2191	2/8	OSPF	0	IBE	20

13 out of 11 Total Num of Dest Networks,13 Total Num of Route Entries displayed.

TYPE Legend:

I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route,
E=Ecmp Route, U=Unresolved Route, N=Not in HW

Configuration example — OSPF security mechanisms

The Passport 8600 implementation of OSPF includes security mechanisms to prevent the OSPF routing domain from being attacked by unauthorized routers.

These security mechanisms are there to prevent a malicious person from joining an OSPF domain and advertising false information in its OSPF LSAs. Likewise, it prevents a misconfigured router from joining an OSPF domain.

There are two security mechanisms:

- [“Simple Password Mechanism,”](#) next
- [“Message Digest 5”](#) on page 131

Simple Password Mechanism

The Simple Password security mechanism is a text-simple password mechanism; only routers that contain the same authentication id in their LSA headers can communicate with each other.

Nortel Networks does not recommend that you use this security mechanism because the password is stored in plain text and can be read from the configuration file or from the LSA packet.

To configure simple password, use the following commands.

Configuring brouter Ports:

Use the following commands to configure brouter ports:

```
Passport-8610:5# config ethernet x/y ip ospf
authentication-type simple
Passport-8610:5# config ethernet x/y ip ospf
authentication-key <string>
```

Where:

x = slot number

y = port number

Configuring VLAN ports:

Use the following commands to configure VLAN ports:

```
Passport-8610:5# config vlan x ip ospf authentication-type  
simple  
Passport-8610:5# config vlan x ip ospf authentication-key  
<string>
```

Where:

x = VLAN number

Message Digest 5

Nortel Networks recommends that you use Message Digest 5 (MD5) for OSPF security because it provides standards based (RFC 1321) authentication, using 128-bit encryption. When you use MD5 for OSPF security, it is almost impossible for a malicious user to compute or extrapolate the decrypting codes from the OSPF packets.

Basically, each OSPF packet has a message digest appended to it, which needs to be matched between sending and receiving routers. The message digest is calculated on either side, based on the MD5 Key and any padding, then compared for a match. If the message digest does not meet the match criteria, the packet is rejected.

MD5 authentication configuration steps:

To configure MD5, complete the following steps:

1 Create a MD5 key and key-id

The following command configures the MD5 key and key-id:

```
Passport-8610:5# config ip ospf interface <ipaddr>  
add-message-digest-key <md5-key-id> md5-key <value>
```

Where:

- *ipaddr* is the IP address of the OSPF interface to be secured.
- *md5-key-id* is a numeric integer in the range 1 and 255.
- *md5-key value* is an alphanumeric password of up to 16 bytes {string length 0..16}

2 Set the authentication type to message-digest.

The following command configures the authentication type to message-digest:

```
Passport-8610:5# config ip ospf interface <ipaddr>
authentication-type message-digest
```

Where:


ipaddr is the IP address of the OSPF interface to be secured.

auth-type selects the authentication type {none|simple|message-digest}

Configuration example — MD5 authentication:

In the configuration example shown in [Figure 26](#), MD5 authentication is configured between Passport 8600 switches R1 and R2.

Figure 26 MD5 authentication example



The following sections provide step-by-step procedures that show how to configure R1 and R2 for this example.

Configuring R1

This section describes how to configure MD5 authentication on R1, using the following commands:

- Configure MD5 authentication on R1:

The following commands enable MD5 authentication for OSPF interface 10.1.1.1 using key “qw sdf89.”

```
Passport-8610:5# config ip ospf interface 10.1.1.1
add-message-digest-key 1 md5-key qw sdf89
Passport-8610:5# config ip ospf interface 10.1.1.1
authentication-type message-digest
```

Configuring R2

This section describes how to configure MD5 authentication on R2, using the following commands:

- Configure MD5 authentication on R2:

The following commands enable MD5 authentication for OSPF interface 10.1.1.2 using key “qw sdf89.”

```
Passport-8610:5# config ip ospf interface 10.1.1.2
add-message-digest-key 1 md5-key qw sdf89
Passport-8610:5# config ip ospf interface 10.1.1.2
authentication-type message-digest
```

Configuration example — Diagnosing OSPF neighbor state problems

At initial startup, routers transmit Hello packets in an attempt to find other OSPF routers to form adjacencies with. Once the Hello packets are received, the routers perform an initialization process, which causes the routers to transition through various states before the adjacency is established.

[Table 5](#) describes the various states a router can be in when forming an adjacency.

Table 5 Neighbor states

Step	State	Description
1	Down	Indicates that a neighbor has been configured manually, but the router has not received any information from the other router. This state can occur only on NBMA interfaces.
2	Attempt	On an NBMA interface, this state occurs when the router attempts to send unicast hellos to any configured interfaces.
3	Init	The router has received a general Hello packet (without its Router ID) from another router.
4	2-Way	The router received a Hello directed to it from another router. (The Hello contains its Router ID).
5	ExStart	Indicates the start of the Master/Slave election process.
6	Exchange	Indicates the Link State Database is exchanged
7	Loading	Indicates the processing state of the LSDB for input into the routing table. The router may request LSA for missing or corrupt routes.
8	Full	Indicates the normal full adjacency state.

This section describe some of the problems that can be encountered during the routers startup process. The following topics are included:

- [“Displaying the current state of all OSPF neighbors](#)
- [“INIT State problems” on page 136](#)
- [“EXSTART/EXCHANGE Problems” on page 137](#)

Displaying the current state of all OSPF neighbors


You can view status of all the OSPF neighbors and their current adjacency state to determine if problems occurred during the router’s initial startup sequence.

To view the current state of all OSPF neighbors and their current state of adjacency, use the following command:

```
Passport-8610:5# show ip ospf neighbors
```

Figure 27 shows sample output for the `show ip ospf neighbors` command.

Figure 27 show ip OSPF neighbors



When problems with OSPF occur, they most often occur during the initial startup, when the router cannot form adjacencies with other routers and the state is stuck in the INIT or EXSTART/EXCHANGE state.

INIT State problems

A router may be stuck in INIT state and not form an adjacency. There are several possible causes for this type of problem:

Authentication mismatch or configuration problem

There could be a mismatch in authentication keys or both sides are not configured for authentication.

To determine if this is causing the problem, issue the “trace Level 6 2” command, which allows you to see the OSPF packets that are received:

```
Passport-8610:5# trace level 6 2
Passport-8610:5# trace screen on
```

The example below shows the error received when there is an authentication failure:

```
[03/24/03 15:55:07:216] tMainTask OSPF: os_recv.c : 710 :
verify_ospf_packet: authType mismatch ipa= 10.1.1.18
```

Access Lists implemented on routers

Ensure that the path is not reachable due to Access Lists implemented on routers:

- Ensure the multicast address of 224.0.0.5 is able to traverse the link.
- If multicast traffic is being blocked for some reason, you may have to configure the Passport 8600 switch for OSPF NBMA, instead of Broadcast.

Inverse ARP misconfigured

When forming an adjacency over an ATM link, both routers must be able to support Inverse ARP, which maps the IP address to a PVC.

Passport 8600 switches do this automatically; however, if the Passport 8600 switch is connecting to another router, ensure that Inverse ARP is enabled on the other router. If Inverse ARP is not supported then it may be necessary to configure a static ARP entry.

EXSTART/EXCHANGE Problems

Although both routers may recognize each other and have moved beyond 2-way, the routers could be stuck in the EXSTART/EXCHANGE state (see [Table 5 on page 134](#)).

This type of problem is usually caused by a mismatch in MTU sizes between the routers. For example, one router could be set for a high MTU size and the other router's default value is a smaller value. Depending on the size of the LSDB, the router with the smaller value may not be able to process the larger packets and thus be stuck in EXSTART/EXCHANGE state. To avoid this problem, ensure that the MTU size value for both routers match.

This problem is usually encountered during interoperations in networks with other vendor devices. You can use the Trace Level 6 2 command to help troubleshoot this type of problem (refer to [“Authentication mismatch or configuration problem” on page 136](#)).



Note: All Passport 8600 switches (Software Release 3.2.0.0 and higher), automatically check for OSPF MTU mismatches.

In the Passport 8600 Software Release 3.2.0.0 and higher, the supported MTU size for OSPF is 1500 bytes, by default. Incoming OSPF DBD packets are dropped if their MTU size is greater than 1500 bytes. To allow the Passport 8600 switch to accept OSPF DBD packets with a different MTU size, enable `mtu-ignore` using the following command:

```
Passport-8610:5# config ip ospf interface <ipaddr>
mtu-ignore <enable|disable>
```

Where:

- *ipaddr* is the ip address of the OSPF interface.
- *enable|disable* enables or disables the feature.

(Note: the default value for `mtu-ignore` is `disable`.)


When `mtu-ignore` is set to `enable`, the MTU Check on the incoming OSPF DBD packet is not performed. The Passport 8600 switch cannot process packets sent on ATM links larger than 1950 bytes.

Configuration example — OSPF network types

OSPF network types were created to allow OSPF-neighboring between routers, over different types of network infrastructures. This allows you to configure each interface to support the various network types.

In the example configuration shown in [Figure 28](#), VLAN 2 on Passport 8600 switch R1 is configured for OSPF with the interface type field value set as passive. Because VLAN 2 is set as passive, OSPF Hello messages are Not sent on this segment, although R1 continues to advertise this interface to the remaining OSPF network.

Figure 28 Configuring OSPF network type example



The following step shows how to configure OSPF on VLAN 2.

➤ Configure OSPF interface — VLAN

The following commands create port-based VLAN 2 under STG 1 with OSPF and sets the interface type as passive:

```
Passport-8610:5# config vlan 2 create byport 1
Passport-8610:5# config vlan 2 ports add 1/2
Passport-8610:5# config vlan 2 ip create 172.3.1.1/24
Passport-8610:5# config vlan 2 ip ospf interface-type
passive
Passport-8610:5# config vlan 2 ip ospf enable
```

Table 6 describes the OSPF network interface types supported by the Passport 8600 switch.

Table 6 OSPF network types

Network interface type	Description
Broadcast	Automatically discovers every OSPF router on the network by sending OSPF Hello's to the multicast group AllSPFRouters (224.0.0.5). Neighboring is automatic and requires no configuration. This interface type is typically used in an Ethernet, ATM, or for certain Frame Relay environments.
Non-Broadcast-Multi-Access (NBMA)	The OSPF NBMA network type was used to correctly model network environments that do not have native Layer 2 broadcast/multicast capabilities, such as Frame Relay and X.25. The OSPF Hello's are unicasted to manually configured neighbors.
Passive	Allows interface network to be included in OSPF without generating LSAs or forming adjacencies. Typically used on an access network, or on an interface that is used for BGP peering. This also limits the amount of CPU cycles required to process Dijkstra.

Use the following command to configure an OSPF network type:

```
Passport-8610:5# config ethernet x/y ip ospf interface-type  
( {broadcast | nbma | passive} )
```

Where:

x = slot number

y = port number

Configuration example — OSPF area types

This section examines how to configure the Passport 8600 in OSPF networks that have more than one area.

In large networks with many routers and networks, the link state database (LSDB) and routing table can become very large. Large route tables and LSDBs consume memory. The processing of link-state advertisements results in more CPU cycles required to make forwarding decisions. To help reduce these undesired effects, an OSPF network can be divided into sub-domains called areas.



Note: An area is made up of a number of OSPF routers that have the same area identification.

By dividing a network into multiple areas, a separate LSDB, consisting of router LSAs and network LSAs are maintained for each area. Each router within an area maintains an LSDB only for the area to which it belongs. For example, the area router-LSAs and network-LSAs are not flooded beyond the area borders.

Therefore, the impact of a topology change is localized to the area to which it occurs. The only exception to this is for the area border routers, which must maintain a LSDB for each area to which they belong. Changes in topology are advertised to the rest of the network by the area border routers by advertising Summary-LSAs.

Area's are identified by a 32-bit Area ID, expressed in IP address format such as 0.0.0.0 for 0. Area 0 is also known as the backbone area and is responsible for distributing routing information to all other areas.

If multiple areas are used, they should all be attached to the backbone via an Area Border Router (ABR), which connects area 0.0.0.0 to the non-backbone area(s). If an area cannot be physically directly connected via an ABR to area 0, you will need to configure a Virtual Link to logically connect the area to the backbone area.

Three types of areas are supported by the Passport 8600 switch:

- [“Normal area,”](#) next
- [“Stub area” on page 143](#)
- [“NSSA” on page 147](#)


Normal area

A Normal Area is a collection of routers that use the same Area-ID that calculates inter-area and external routes through the use of the following Link-State Advertisements (LSAs):

- Summary-LSAs
- ASBR-summary-LSAs
- AS-external-LSAs

As shown in [Figure 29](#), a Normal Area supports Area Border Routers (ABRs) and Autonomous System Border Routers (ASBRs).

Figure 29 Normal Area example



11026fa

Configuring the ABR:

There are no configuration parameters for configuring a Passport 8600 switch like ABR. The switch automatically becomes an ABR when you configure more than one area on the switch (refer to [“Configuration example — OSPF ABR” on page 151](#)).

Configuring the ASBR:

You can configure the Passport 8600 switch as an OSPF ASBR, as follows:

- Distribute all OSPF routes to BGP or RIP.
- Distribute RIP, BGP, Direct, or static routes to OSPF

To configure a Passport 8600 as an ASBR, use the following command:

```
Passport-8610:5# config ip ospf as-boundary-router enable
```


For more information, refer to [“Configuration examples — OSPF ASBR configurations” on page 154](#))

Stub area

Stub Areas do not receive advertisements for external routes (AS-external LSAs, type 5) from an ABR, which reduces the size of the link state database. Instead, routing to external destinations from within a Stub Area is based on the default route that is originated by the Stub Area ABR.

As shown in [Figure 30](#), a Stub Area has only one ABR. All packets that are destined to be forwarded outside the Stub Area are routed to the Stub Area’s border exit point, where the packets are first examined by the ABR and then forwarded to a destination.

Figure 30 Stub Area example



11027fa

Configuring a Stub Area:

- Stub Areas do not support ASBRs.
- Stub Areas cannot support virtual links, without AS-external LSA support.


To configure an OSPF area as a Stub Area or NSSA, use the following commands:

```
Passport-8610:5# config ip ospf area <area IP address> stub
<true|false>
Passport-8610:5# config ip ospf area <area IP address> nssa
<true|false>
```


Configuration example — Stub Area

In the configuration example shown in [Figure 31](#), Passport 8600 switch R1 is configured in Stub Area 2, and R2 is configured as a Stub ABR for Area 2.

Figure 31 Configuration example — Stub Area



Legend



11028fa



Note: AS-external LSAs are not flooded into a Stub Areas. Instead, only one default route to external destinations is distributed into the Stub Area by the Stub ABR router.

The following sections provide step-by-step procedures that show how to configure R1 and R2 for this example.

Configuring R1

To configure R1, use the following commands:

1 Configure the OSPF interface on R1:

The following commands configure port 2/6 as a brouter port and enable OSPF on this interface.

```
Passport-8610:5# config ethernet 2/6 ip create
10.1.1.18/30 2090
Passport-8610:5# config ethernet 2/6 ip ospf enable
```

2 Configure VLAN 2 on R1:

The following commands create VLAN = 2 and enable OSPF for this interface.

```
Passport-8610:5# config vlan 2 create byport 1
Passport-8610:5# config vlan 2 ports add 1/48
Passport-8610:5# config vlan 2 ip create
172.3.3.1/255.255.255.0 vlan 2 ip ospf enable
```

3 Create a CLIP address for R1:

The following commands create a circuitless IP address which will be used for the OSPF Router-ID.

```
Passport-8610:5# config ip circuitless-ip-int 1 create
1.1.1.5/255.255.255.255
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

4 Enable OSPF on R1:

The following commands configure R1 as Stub Area 2, assign the Circuitless IP (created in Step 3) as the OSPF Router-ID, and adds the OSPF interfaces to Area 2.

```
Passport-8610:5# config ip ospf router-id 1.1.1.5
Passport-8610:5# config ip ospf enable
Passport-8610:5# config ip ospf area 0.0.0.2 create
Passport-8610:5# config ip ospf area 0.0.0.2 stub true
Passport-8610:5# config ip ospf interface 10.1.1.18 area
0.0.0.2
Passport-8610:5# config ip ospf interface 1.1.1.5 area
0.0.0.2
Passport-8610:5# config ip ospf interface 172.3.3.1 area
0.0.0.2
```

Configuring R2

To configure R2, use the following commands:

1 Configure the OSPF interface on R2:

The following commands configure port 2/6 as a brouter port and enable OSPF on this interface.

```
Passport-8610:5# config ethernet 2/6 ip create
10.1.1.17/30 2090
Passport-8610:5# config ethernet 2/6 ip ospf enable
```

2 Configure the second OSPF interface on R2:

The following commands configure port 2/1 as a brouter port and enable OSPF on this interface.

```
Passport-8610:5# config ethernet 2/1 ip create
10.1.1.22/30 2090
Passport-8610:5# config ethernet 2/1 ip ospf enable
```

3 Create a CLIP address for R2:

The following commands create a circuitless IP address which will be used for the OSPF Router-ID.

```
Passport-8610:5# config ip circuitless-ip-int 1 create
1.1.1.2/255.255.255.255
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

4 Enable OSPF on R2:

The following commands configure R2 as a Stub ABR. Note that, by default, OSPF interface 10.1.1.22 is placed into OSPF area 0.0.0.0. As one additional sub area of 0.0.0.2 is added to the configuration, R2 automatically becomes a Stub ABR.


```
Passport-8610:5# config ip ospf router-id 1.1.1.2
Passport-8610:5# config ip ospf enable
Passport-8610:5# config ip ospf area 0.0.0.2 create
Passport-8610:5# config ip ospf area 0.0.0.2 stub true
Passport-8610:5# config ip ospf interface 10.1.1.17 area
0.0.0.2
```

NSSA

Similar to Stub Areas, the Not So Stubby Areas (NSSAs) can also prevent the flooding of AS-External Link State advertisements into the NSSA Area by replacing them with a default route. However, NSSA Areas can also import small Stub (non-OSPF) routing domains into OSPF. This allows the NSSA Area to import external routes, such as RIP routes, and then advertise these routes throughout the network.

As shown in [Figure 32 on page 148](#), external routing information is imported into NSSA Areas by using Type-7 LSAs. These LSAs are translated at the NSSA Area boundary into LSA Type-5. The N/P bit in the Type-7 LSA Options field indicates whether the Type-7 LSA should be translated. Only those LSAs with the N/P-bit set are translated.

Figure 32 NSSA Area example




11029fa

Configuration example — NSSA Area

In the configuration example shown in [Figure 33](#), Passport 8600 switch R1 is configured as an NSSA ASBR router.

Figure 33 Configuration example — NSSA Area



11030fa

The following section provides a step-by-step procedure that shows how to configure R1 as in this example:

Configuring R1

To configure R1, use the following commands:

1 Configure the RIP interface on R1:

The following commands configure port 1/48 as a brouter port and enable RIP on this interface.

```
Passport-8610:5# config ethernet 1/48 ip create
20.1.1.2/30 2091
Passport-8610:5# config ethernet 1/48 ip RIP enable
```

2 Enable RIP globally and configure the RIPv2 interface:

The following commands globally enable RIP and configure the RIP interface on R1 for RIPv2.

```
Passport-8610:5# config ip rip enable
Passport-8610:5# config ip rip interface 20.1.1.2
send-mode rip2
Passport-8610:5# config ip rip interface 20.1.1.2
receive-mode rip2
```

3 Configure the OSPF interface on R1:

The following commands configure port 2/6 as a brouter port and enable OSPF on this interface.

```
Passport-8610:5# config ethernet 2/6 ip create
10.1.1.18/30 2090
Passport-8610:5# config ethernet 2/6 ip ospf enable
```

4 Create a CLIP address for R1:

The following commands create a circuitless IP address which will be used for the OSPF Router-ID.

```
Passport-8610:5# config ip circuitless-ip-int 1 create
1.1.1.5/255.255.255.255
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

5 Enable OSPF on R1:

The following commands configure R1 as an ASBR, assign the CLIP address (created in Step 4) as the OSPF Router-ID, creates OSPF NSSA Area 2, and adds the OSPF interfaces to Area 2.

```
Passport-8610:5# config ip ospf as-boundary-router
enable
Passport-8610:5# config ip ospf router-id 1.1.1.5
Passport-8610:5# config ip ospf enable
Passport-8610:5# config ip ospf area 0.0.0.2 create
Passport-8610:5# config ip ospf area 0.0.0.2 nssa true
Passport-8610:5# config ip ospf interface 10.1.1.18
area 0.0.0.2
Passport-8610:5# config ip ospf interface 1.1.1.5
area 0.0.0.2
```

6 Configure a route policy to distribute Direct and OSPF to RIP:

The following commands create a route policy named "Rip_Dist" that distributes directly-connected and OSPF routes into RIP.

```
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
create
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
enable
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
action permit
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
match-protocol local|ospf
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
set-metric-type type1
```

7 Apply a route policy to RIP Out-Policy:

The following command applies the "Rip_Dist" route policy, created in Step 6, to the RIP Out-Policy.

```
Passport-8610:5# config ip rip interface 20.1.1.2
out-policy "Rip_Dist"
```

8 Configure OSPF route distribution:

The following commands configure OSPF route distribution to distribute RIP routes as AS-external-LSA Type 1.


```
Passport-8610:5# config ip ospf redistribute rip create
Passport-8610:5# config ip ospf redistribute rip
metric-type type1
Passport-8610:5# config ip ospf redistribute rip enable
```

Configuration example — OSPF ABR

Configuration of an OSPF ABR is an automatic process on the Passport 8600 switch; no user intervention is required to complete the process. For example, when you configure more than one area, the Passport 8600 is automatically configured as an OSPF ABR.

In the configuration example shown in [Figure 34](#), Passport 8600 switch R1 is automatically configured as an OSPF ABR after it was configured with an OSPF interface for Area 0.0.0.2 and Area 0.0.0.0.

Figure 34 OSPF ABR example



Legend



Passport 8600 switch

11033fa

The following section provides a step-by-step procedure that shows how to configure R1 for this example.

Configuring R1

This section describes how to configure R1 for Area 0.0.0.2 and Area 0.0.0.0, which automatically configures R1 as an OSPF ABR.

To configure R1, use the following commands:

1 Configure an OSPF interface port 2/6:

The following commands configure port 2/6 as a brouter port and enable OSPF on this interface.

```
Passport-8610:5# config ethernet 2/6 ip create
10.1.1.17/30 2090
Passport-8610:5# config ethernet 2/6 ip ospf enable
```

2 Configure an OSPF interface port 2/1:

The following commands configure port 2/1 as a brouter port and enable OSPF on this interface.

```
Passport-8610:5# config ethernet 2/1 ip create
10.1.1.22/30 2090
Passport-8610:5# config ethernet 2/1 ip ospf enable
```

3 Create Circuitless IP

The following commands create a circuitless IP address which will be used for the OSPF Router-ID.

```
Passport-8610:5# config ip circuitless-ip-int 1 create
1.1.1.2/255.255.255.255
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```


4 Enable OSPF

The following commands configure R1 as an ABR. Note that, by default, OSPF interface 10.1.1.22 is placed into OSPF area 0.0.0.0. Because one additional area of 0.0.0.2 is added to the configuration, R1 automatically becomes an ABR.

```
Passport-8610:5# config ip ospf router-id 1.1.1.2
Passport-8610:5# config ip ospf enable
Passport-8610:5# config ip ospf area 0.0.0.2 create
Passport-8610:5# config ip ospf interface 10.1.1.17 area
0.0.0.2
```


Showing the created areas

To display the created areas, use the following command:

```
Passport-8610:5# show ip ospf area
```

Figure 35 shows sample output for the show ip ospf area command.

Figure 35 show ip OSPF area



Displaying the ABR status

To display the ABR status, use the following command:

```
Passport-8610:5# show ip ospf info
```

Figure 36 shows sample output for the `show ip ospf info` command.

Figure 36 show ip OSPF info

```

Passport-8610:6# show ip ospf info
=====
                        Ospf General
=====
    RouterId: 1.1.1.2
    AdminStat: enabled
    VersionNumber: 2
    AreaBdrRtrStatus: true ←
    ASBdrRtrStatus: false
    ExternLsaCount: 1
    ExternLsaChecksumSum: 29660 (0x73dc)
    TOSSupport: 0
    OriginateNewLsas: 270
    RxNewLsas: 1047
    TrapEnable: false
    AutoVirtLinkEnable: false
    SpfHoldDownTime: 10

```

Configuration examples — OSPF ASBR configurations

This section describes ASBR configuration examples and includes the CLI commands you can use to recreate the configurations: You can configure an OSPF ASBR on the Passport 8600 switch to:

- Distribute all OSPF routes to BGP or RIP.
- Distribute RIP, BGP, Direct, or static routes to OSPF


This section includes the following topics:

- [“Distributing OSPF routes to RIP and RIP to OSPF using AS-external-LSA Type 1 metrics,” next](#)
- [“Distributing an Internet default route to OSPF using AS-external-LSA Type 2 metrics” on page 159](#)
- [“Viewing advertised AS_External LSAs” on page 161](#)

Distributing OSPF routes to RIP and RIP to OSPF using AS-external-LSA Type 1 metrics

The configuration example shown in [Figure 37](#), shows a Passport 8600 switch (R1) configured as an ASBR between an OSPF network and a RIPv2 network. In this example, R1 distributes all OSPF routes to the RIP network, and all RIP routes to the OSPF network.

Figure 37 OSPF routes: OSPF/RIP and RIP/OSPF



11034fa

The following sections provide step-by-step procedures that show how to configure R1 for this example.

You can configure R1 as follows:

- “Configuring R1 to distribute all OSPF routes to RIP,” next
- “Configuring R1 to distribute a default route only to RIP” on page 158

Configuring R1 to distribute all OSPF routes to RIP

To configure R1 to distribute all OSPF routes to RIP, complete the following steps:

1 Configure RIP:

a Configure the RIP interface on R1:

Use the following *two* commands to configure port 1/31 as a brouter port and enable RIP on this interface.

```
Passport-8610:5# config ethernet 1/31 ip create
10.1.1.41/30 2136
Passport-8610:5# config ethernet 1/31 ip rip enable
```

b Configure the RIP interface for RIPv2 mode only:

The following commands enable RIP and configure the RIP interface for RIPv2 mode only.

```
Passport-8610:5# config ip rip enable
Passport-8610:5# config ip rip interface 10.1.1.41
send-mode rip2
Passport-8610:5# config ip rip interface 10.1.1.41
receive-mode rip2
```

2 Configure the OSPF interface:

Use the following *two* commands to configure port 2/7 as a brouter port and enable OSPF on this interface.

```
Passport-8610:5# config ethernet 2/7 ip create
10.1.1.2/30 2134
Passport-8610:5# config ethernet 2/7 ip ospf enable
```

3 Assign a circuitless IP address on R1:

Use the following commands to create a CLIP address, which will be used for the OSPF Router-ID.

```
Passport-8610:5# config ip circuitless-ip-int 1 create
1.1.1.3/32
```

4 Assign R1 as the ASBR:

Use the following commands to configure R1 as an ASBR and assign the CLIP address (created in Step 3) as the OSPF Router-ID.

```
Passport-8610:5# config ip ospf as-boundary-router enable
Passport-8610:5# config ip ospf router-id 1.1.1.3
Passport-8610:5# config ip ospf enable
```

5 Configure OSPF route distribution:

Use the following commands to configure OSPF route distribution to import RIP into OSPF. The Passport 8600 switch (R1) distributes the RIP routes as AS-external-LSA (LSA Type 5), using external metric type 1.

```
Passport-8610:5# config ip ospf redistribute rip create
Passport-8610:5# config ip ospf redistribute rip metric
10
Passport-8610:5# config ip ospf redistribute rip
metric-type type1
Passport-8610:5# config ip ospf redistribute rip enable
```

6 Configure a route policy:

A route policy is required for OSPF to RIP route redistribution. After you create the route policy, you must apply it to the RIP interface.

Use the following commands to create a route policy named “allow,” which will distribute both local interfaces and OSPF.

```
Passport-8610:5# config ip route-policy “allow” seq 1
create
Passport-8610:5# config ip route-policy “allow” seq 1
enable
Passport-8610:5# config ip route-policy “allow” seq 1
action permit
Passport-8610:5# config ip route-policy “allow” seq 1
match-protocol local|ospf
```

7 Apply the route policy to RIP Out-Policy:

Use the following command to apply the route policy created in Step 6 to RIP interface 10.1.1.41.

```
Passport-8610:5# config ip rip interface 10.1.1.41
out-policy “allow”
```

Configuring R1 to distribute a default route only to RIP

The configuration steps described in the previous section distributes *all* OSPF routes to RIP. However, there may be times when it may be more advantageous for you to distribute only a default route to RIP. The following configuration steps describe how to distribute only a default route to RIP instead of all OSPF routes to RIP.

To configure R1 to distribute a default route only to RIP, complete the following steps:

- 1 Configure an IP Prefix list with a default route:

The following command creates an IP Prefix list named “default” with IP address 0.0.0.0.

```
Passport-8610:5# config ip prefix-list "default"  
add-prefix 0.0.0.0/0
```

- 2 Configure a route policy:

The following commands create a route policy named “Policy_Default,” which distributes the IP Prefix list created in Step 1. Note that “ospf” is selected for the match-protocol value. This causes the default route to only be advertised via RIP if OSPF is up.

```
Passport-8610:5# config ip route-policy "Policy_Default"  
seq 1 create  
Passport-8610:5# config ip route-policy "Policy_Default"  
seq 1 enable  
Passport-8610:5# config ip route-policy "Policy_Default"  
seq 1 action permit  
Passport-8610:5# config ip route-policy "Policy_Default"  
seq 1 match-protocol ospf  
Passport-8610:5# config ip route-policy "Policy_Default"  
seq 1 set-injectlist "default"  
Passport-8610:5# config ip route-policy "Policy_Default"  
seq 1 set-metric-type type1
```

3 Apply the route policy to the RIP Out-Policy:


The following command applies the route policy created in Step 2 to RIP interface 10.1.1.41.

```
Passport-8610:5# config ip rip interface 10.1.1.41
out-policy "Policy_Default"
```

Distributing an Internet default route to OSPF using AS-external-LSA Type 2 metrics

The configuration example shown in Figure 38, shows a Passport 8600 switch (R1) configured as an ASBR between an OSPF network and the Internet. For this example, R1 is configured to distribute a default route for Internet traffic.

Figure 38 OSPF routes: OSPF/RIP and RIP/OSPF



Legend



Passport 8600 switch

11035fa

To configure R1 to distribute a default route for Internet traffic, complete the following steps:

1 Configure the OSPF interface:

The following command configures port 2/1 as a brouter port and enables OSPF on this interface.

```
Passport-8610:5# config ethernet 2/1 ip create 10.1.1.21/
30 2134
Passport-8610:5# config ethernet 2/1 ip ospf enable
```

2 Assign a CLIP address for R1:

The following command assigns a circuitless IP address to R1, which is used for both the OSPF Router-ID and the BGP identifier.

```
Passport-8610:5# config ip circuitless-ip-int 1 create
1.1.1.1/32
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

3 Enable OSPF:

The following commands configure R1 as an ASBR and assign the CLIP address created in Step 2 as the OSPF Router-ID

```
Passport-8610:5# config ip ospf as-boundary-router enable
Passport-8610:5# config ip ospf router-id 1.1.1.1/32
Passport-8610:5# config ip ospf enable
```

4 Configure the BGP interface:

The following commands configure the BGP interface on R1 and establish R1 as a BGP peer.

```
Passport-8610:5# config ethernet 1/2 ip create
131.168.1.2/30
Passport-8610:5# config ip bgp local-as 65500
Passport-8610:5# config ip bgp enable
Passport-8610:5# config ip bgp neighbor 131.168.1.1
create
Passport-8610:5# config ip bgp neighbor 131.168.1.1
remote-as 65503
Passport-8610:5# config ip bgp neighbor 131.168.1.1
admin-state enable
```

5 Configure a prefix list with the default route:

The following command adds a prefix list with the default route, which will be used in the next step (Step 6).

```
Passport-8610:5# config ip prefix-list "default_prefix"
add-prefix 0.0.0.0/0
```


6 Configure a route policy to distribute for default route distribution:

The following commands create a Route Policy named “Default-Route” and adds the Prefix List created in Step 5. Note that the external metric value is set for Type 2.

```
Passport-8610:5# config ip route-policy "Default_Route"
seq 1 create
Passport-8610:5# config ip route-policy "Default_Route"
seq 1 enable
Passport-8610:5# config ip route-policy "Default_Route"
seq 1 action-permit
Passport-8610:5# config ip route-policy "Default_Route"
seq 1 set-injectlist "Default_Prefix"
Passport-8610:5# config ip route-policy "Default_Route"
seq 1 set-metric 100
Passport-8610:5# config ip route-policy "Default_Route"
seq 1 set-metric-type type2
```

7 Configure OSPF route distribution:

The following commands enable BGP route importation into OSPF, but distribute only a Default Route. (For more information about distributing OSPF routes into BGP, see [Configuring BGP Services](#).)

```
Passport-8610:5# config ip ospf redistribute bgp create
Passport-8610:5# config ip ospf redistribute bgp metric 1
Passport-8610:5# config ip ospf redistribute bgp
route-policy "Default_Route"
Passport-8610:5# config ip ospf redistribute bgp enable
```

Viewing advertised AS_External LSAs

An ASBR advertises routes (such as the RIP routes from the previous example), as AS_external LSAs (LSA Type 5).

To display the advertised AS_external LSAs, use the following show command:

```
show ip ospf ase
```

Figure 39 shows sample output the `show ip ospf ase` command.

Figure 39 show ip ospf ase command

```
Passport-8610:6# show ip ospf ase

=====
                        Ospf AsExternal Lsas
=====

LSTYPE      LINKSTATEID ADV_ROUTER E_METRIC  ASE_FWD_ADDR  AGE  SEQ_NBR      CSUM
-----
AsExternal  0.0.0.0      1.1.1.1      1 100      0.0.0.0      276  0x8000015c  0x2fdc
AsExternal  15.15.15.0   1.1.1.2      1 2        10.1.1.18    262  0x800000ed  0xaa2
AsExternal  172.2.2.0    1.1.1.3      0 10       0.0.0.0      236  0x800000be  0x769d
```

You can also use the following show command to view all the LSAs, including AS_External LSAs:

```
show ip ospf lsdb
```

Configuration example — Controlling NSSA external routes advertised

In an OSPF NSSA Area, the NSSA N/b-bit (in the OSPF Hello packets Options field) is used to tell the ABR which external routes can be advertised to other areas.


When the NSSA N/p-bit is set true, the ABR exports the external route. This is the default setting for the Passport 8600 switch.

When the NSSA N/p-bit is **not** set true, the ABR drops the external route. You can create a route policy on the Passport 8600 switch to manipulate the N/p-bit value.

For example, Figure 40 shows a RIP network located in NSSA Area 2. If you want to only advertise the 15.15.15.0/24 network to Area 0, the following tasks are required:

- Enable R1 as an OSPF ASBR
- Create NSSA Area 2
- Create a Route Policy to advertise OSPF and direct interfaces to RIP
- Create a Route Policy to only advertise RIP network 15.15.15.0/24 to Area 0 by using the NSSA N/p-bit

Figure 40 Controlling external routes advertised



11036fa

To configure R1 to only advertise the 15.15.15.0/24 network to Area 0, use the commands shown in the following steps:

1 Configure the RIP interface:

The following commands configure port 1/48 as a brouter port and enables RIP on this interface.

```
Passport-8610:5# config ethernet 1/48 ip create
20.1.1.2/30 2091
Passport-8610:5# config ethernet 1/48 ip rip enable
```

2 Enable RIP globally and configure a RIP interface for RIPv2:

The following commands globally enable RIP and configure a RIP interface for RIPv2.

```
Passport-8610:5# config ip rip enable
Passport-8610:5# config ip rip interface 20.1.1.2
send-mode rip2
Passport-8610:5# config ip rip interface 20.1.1.2
receive-mode rip2
```

3 Configure the OSPF interface:

The following commands configure port 2/6 as a brouter port and enable OSPF on this interface.

```
Passport-8610:5# config ethernet 2/6 ip create
10.1.1.18/30 2090
Passport-8610:5# config ethernet 2/6 ip ospf enable
```

4 Assign the CLIP address:

The following commands assign the CLIP address to R1, which is used for the OSPF Router-ID.

```
Passport-8610:5# config ip circuitless-ip-int 1 create
1.1.1.5/255.255.255.255
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

5 Enable OSPF:

The following commands configure R1 as an ASBR, assign the CLIP created in Step 4 as the OSPF Router-ID, create OSPF NSSA area 2, and add the OSPF interfaces to Area 2.

```
Passport-8610:5# config ip ospf as-boundary-router enable
Passport-8610:5# config ip ospf router-id 1.1.1.5
Passport-8610:5# config ip ospf enable
Passport-8610:5# config ip ospf area 0.0.0.2 create
Passport-8610:5# config ip ospf area 0.0.0.2 nssa true
Passport-8610:5# config ip ospf interface 10.1.1.18
area 0.0.0.2
Passport-8610:5# config ip ospf interface 1.1.1.5
area 0.0.0.2
```

6 Configure a route policy to distribute direct interfaces and OSPF to RIP:

The following commands create a Route Policy named "Rip_Dist" that distribute directly connected and OSPF routes into RIP.

```
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
create
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
enable
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
action-permit
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
match-protocol local|ospf
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
set-metric-type type1
Passport-8610:5# config ip route-policy "Rip_Dist" seq 1
set-nssa-pbit enable
```

7 Apply Route Policy to RIP Out-Policy:

```
Passport-8610:5# config ip interface 20.1.1.2 out-policy
"Rip_Dist"
```

8 Configure Prefix Lists:

The following commands add two prefix lists ("15.15.15.0" and "14.14.14.0") that are associated with the network addresses from the RIPv2 network.

```
Passport-8610:5# config ip prefix-list "15.15.15.0"
add-prefix 15.15.15.0/24
Passport-8610:5# config ip prefix-list "14.14.14.0"
add-prefix 14.14.14.0/24
```

9 Configure a route policy to set NSSA p-bit:

The following commands create a Route Policy named "P_bit" that sets the NSSA N/P-bit only for the Prefix List named "15.15.15.0".

```
Passport-8610:5# config ip route-policy "P_bit" seq 1
create
Passport-8610:5# config ip route-policy "P_bit" seq 1
enable
Passport-8610:5# config ip route-policy "P_bit" seq 1
action permit
Passport-8610:5# config ip route-policy "P_bit" seq 1
match-network "15.15.15.0"
```

```
Passport-8610:5# config ip route-policy "P_bit" seq 1
match-protocol ospf
Passport-8610:5# config ip route-policy "P_bit" seq 1
set-nssa-pbit enable
Passport-8610:5# config ip route-policy "P_bit" seq 2
create
Passport-8610:5# config ip route-policy "P_bit" seq 2
enable
Passport-8610:5# config ip route-policy "P_bit" seq 2
action permit
Passport-8610:5# config ip route-policy "P_bit" seq 2
match-network "14.14.14.0"
Passport-8610:5# config ip route-policy "P_bit" seq 2
match-protocol ospf
Passport-8610:5# config ip route-policy "P_bit" seq 2
set-nssa-pbit disable
```

10 Configure OSPF route distribution parameters:

The following commands configure OSPF route distribution to distribute RIP routes as AS-external-LSA Type 1.

```
Passport-8610:5# config ip ospf redistribute rip create
Passport-8610:5# config ip ospf redistribute rip
metric-type type1
Passport-8610:5# config ip ospf redistribute rip
route-policy "P_bit"
Passport-8610:5# config ip ospf redistribute rip enable
```

Displaying configuration files

You can use the following show command to display the configuration commands and parameters used to create the topology shown in [Figure 40 on page 163](#):

```
Passport-8610:5# show config
```



Note: You can copy and paste the command outputs shown here to update your configuration files.

Configuration file for R1

```
# PORT CONFIGURATION - PHASE II
#
ethernet 1/48 ip create 20.1.1.2/255.255.255.252 2091 mac_offset 2
ethernet 1/48 ip rip enable
ethernet 2/6 ip create 10.1.1.18/255.255.255.252 2090 mac_offset 1
ethernet 2/6 ip ospf enable
ethernet 2/6 ip ospf priority 0
#
# IP PREFIX LIST CONFIGURATION
#
ip prefix-list "15.15.15.0" add-prefix 15.15.15.0/24 maskLenFrom 24
maskLenTo 24
ip prefix-list "14.14.14.0" add-prefix 14.14.14.0/24 maskLenFrom 24
maskLenTo 24
#
# IP ROUTE POLICY CONFIGURATION
#
ip route-policy "Rip_Dist" seq 1 create
ip route-policy "Rip_Dist" seq 1 enable
ip route-policy "Rip_Dist" seq 1 action permit
ip route-policy "Rip_Dist" seq 1 match-protocol local|ospf
ip route-policy "Rip_Dist" seq 1 set-metric-type type1
ip route-policy "Rip_Dist" seq 1 set-nssa-pbit enable
ip route-policy "P_bit" seq 1 create
ip route-policy "P_bit" seq 1 enable
ip route-policy "P_bit" seq 1 action permit
ip route-policy "P_bit" seq 1 match-network "15.15.15.0"
ip route-policy "P_bit" seq 1 match-protocol ospf
ip route-policy "P_bit" seq 1 set-metric-type type2
ip route-policy "P_bit" seq 1 set-nssa-pbit enable
ip route-policy "P_bit" seq 2 create
ip route-policy "P_bit" seq 2 enable
ip route-policy "P_bit" seq 2 action permit
ip route-policy "P_bit" seq 2 match-network "14.14.14.0"
ip route-policy "P_bit" seq 2 match-protocol ospf
ip route-policy "P_bit" seq 2 set-metric-type type2
ip route-policy "P_bit" seq 2 set-nssa-pbit disable
#
ip rip enable
ip rip interface 20.1.1.2 send-mode rip2
ip rip interface 20.1.1.2 receive-mode rip2
#
# CIRCUITLESS IP INTERFACE CONFIGURATION
#
ip circuitless-ip-int 1 create 1.1.1.5/255.255.255.255
ip circuitless-ip-int 1 ospf enable
```


```
#
# OSPF CONFIGURATION
#
ip ospf admin-state enable
ip ospf as-boundary-router enable
ip ospf router-id 1.1.1.5
ip ospf enable
ip ospf area 0.0.0.2 create
ip ospf area 0.0.0.2 nssa true
ip ospf interface 10.1.1.18 area 0.0.0.2
ip ospf interface 10.1.1.18 add-message-digest-key 1 md5-key Test
ip ospf interface 1.1.1.5 area 0.0.0.2
ip ospf interface 172.3.3.1 area 0.0.0.2
#
# IP REDISTRIBUTION CONFIGURATION
#
ip ospf redistribute rip create
ip ospf redistribute rip metric-type type1
ip ospf redistribute rip route-policy "P_bit"
ip ospf redistribute rip enable
#
# RIP POLICY CONFIGURATION
#
ip rip interface 20.1.1.2 out-policy "Rip_Dist"
```


Configuration example — Multi-area complex

The multi-area complex configuration example described in this section uses five Passport 8600 switches (R1 to R5) in a multi-area configuration (Figure 41).

Many of the concepts and topology descriptions that are used in this example configuration are described in the previous sections of this chapter. The concepts shown in those examples are combined in this example configuration to show real-world topology, with command descriptions.

Figure 41 Multi-area complex configuration example



11037fa

For this configuration example, Passport 8600 switches R1 through R5 are configured as follows:

- R1 is an OSPF ASBR that is associated with OSPF Area 0 and OSPF Area 3. R1 is configured to distribute a default route for Internet traffic.
- R2 is an OSPF Stub ABR for OSPF Area 2 and ABR to OSPF Area 3.
- R3 is an OSPF ASBR and is configured to distribute OSPF to RIP and RIP to OSPF.
- R4 is an OSPF internal router in Area 3.
- R5 is an internal OSPF Sub router in Area 2.
- All OSPF interfaces are brouter ports, with the exception of R5.

For network 172.3.3.0/24 on R5, a VLAN configuration is used in place of a brouter port. The reason this example uses brouter ports rather than VLANs, is because the spanning tree algorithm is disabled by default when using brouter interfaces.

- All interfaces used for this configuration are Ethernet, therefore the OSPF interfaces are broadcast, with the exception of the Circuitless IP interfaces which are passive.
- The interface priority value on R5 is set to 0, therefore R5 cannot become a Designated Router (DR).
- Configure the OSPF Router Priority so that R1 becomes the DR (priority = 100) and R2 becomes Backup Designated Router (BDR) with a priority value (priority = 50).

The reason for using Stub Areas or NSSA Areas is to reduce the LSDB size by not including external LSAs. The Sub ABR will advertise a default route into the Stub Area for all external routes.

Displaying configuration files

You can use the following show command to display the configuration commands and parameters used to create the topology shown in [Figure 41 on page 169](#):

```
Passport-8610:5# show config
```



Note: You can copy and paste the command outputs shown here to update your configuration files.

Configuration file for R1

```
#
# PORT CONFIGURATION - PHASE II
#
ethernet 1/2 auto-negotiate disable
ethernet 1/2 speed 100
ethernet 1/2 duplex full
ethernet 1/2 ip create 131.168.1.2/255.255.255.252 2065
mac_offset 1
ethernet 2/1 ip create 10.1.1.21/255.255.255.252 2190 mac_offset 6
ethernet 2/1 ip ospf enable
ethernet 2/1 ip ospf priority 100
ethernet 2/7 ip create 10.1.1.1/255.255.255.252 2134 mac_offset 0
ethernet 2/7 ip ospf enable
ethernet 2/7 ip ospf priority 100
#
# IP PREFIX LIST CONFIGURATION
#
ip prefix-list "Default_Prefix" add-prefix 0.0.0.0/0 maskLenFrom 0
maskLenTo 0
#
# IP ROUTE POLICY CONFIGURATION
#
ip route-policy "Default_Route" seq 1 create
ip route-policy "Default_Route" seq 1 enable
ip route-policy "Default_Route" seq 1 action permit
ip route-policy "Default_Route" seq 1 set-injectlist
"Default_Prefix"
ip route-policy "Default_Route" seq 1 set-metric 100
ip route-policy "Default_Route" seq 1 set-metric-type type2
ip route-policy "Default_Route" seq 1 set-nssa-pbit enable
#
```

```
# CIRCUITLESS IP INTERFACE CONFIGURATION
#
ip circuitless-ip-int 1 create 1.1.1.1/255.255.255.255
ip circuitless-ip-int 1 ospf enable
#
# OSPF CONFIGURATION
#
ip ospf admin-state enable
ip ospf as-boundary-router enable
ip ospf router-id 1.1.1.1
ip ospf enable
ip ospf area 0.0.0.3 create
ip ospf interface 10.1.1.1 area 0.0.0.3
#
# BGP CONFIGURATION
#
ip bgp local-as 65500
ip bgp enable
ip bgp neighbor 131.168.1.1 create
ip bgp neighbor 131.168.1.1 remote-as 65503
ip bgp neighbor 131.168.1.1 route-advertisement-interval 30 add
ip bgp neighbor 131.168.1.1 admin-state enable
#
# IP REDISTRIBUTION CONFIGURATION
#
ip ospf redistribute bgp create
ip ospf redistribute bgp metric 1
ip ospf redistribute bgp route-policy "Default_Route"
ip ospf redistribute bgp enable
```

Configuration file for R2

Because R2 is associated with three areas, including one that is a stub area, it is configured as an NSSA ABR.

```
# PORT CONFIGURATION - PHASE II
#

ethernet 2/1 ip create 10.1.1.22/255.255.255.252 2201 mac_offset 6
ethernet 2/1 ip ospf enable
ethernet 2/1 ip ospf priority 50
ethernet 2/6 ip create 10.1.1.17/255.255.255.252 2200 mac_offset 5
ethernet 2/6 ip ospf enable
ethernet 2/6 ip ospf priority 50
ethernet 2/7 ip create 10.1.1.9/255.255.255.252 2198 mac_offset 1
ethernet 2/7 ip ospf enable
ethernet 2/7 ip ospf priority 50

#
# CIRCUITLESS IP INTERFACE CONFIGURATION
#

ip circuitless-ip-int 1 create 1.1.1.2/255.255.255.255
ip circuitless-ip-int 1 ospf enable

#
# OSPF CONFIGURATION
#

ip ospf admin-state enable
ip ospf router-id 1.1.1.2
ip ospf enable
ip ospf area 0.0.0.2 create
ip ospf area 0.0.0.2 stub true
ip ospf area 0.0.0.3 create
ip ospf interface 10.1.1.17 area 0.0.0.2
ip ospf interface 10.1.1.9 area 0.0.0.3
```

Configuration file for R3

```
#
# PORT CONFIGURATION - PHASE II
#
ethernet 1/31 ip create 10.1.1.41/255.255.255.252 2136 mac_offset 8
ethernet 1/31 ip rip enable
ethernet 1/31 ip rip default-supply enable
ethernet 2/1 ip create 10.1.1.25/255.255.255.252 2190 mac_offset 4
ethernet 2/1 ip ospf enable
ethernet 2/7 ip create 10.1.1.2/255.255.255.252 2134 mac_offset 1
ethernet 2/7 ip ospf enable
#
# IP ROUTE POLICY CONFIGURATION
#
ip route-policy "Allow" seq 1 create
ip route-policy "Allow" seq 1 enable
ip route-policy "Allow" seq 1 action permit
ip route-policy "Allow" seq 1 match-protocol local|ospf
ip route-policy "Allow" seq 1 set-metric-type type2
ip route-policy "Allow" seq 1 set-nssa-pbit enable
#
ip rip enable
ip rip interface 10.1.1.41 send-mode rip2
ip rip interface 10.1.1.41 receive-mode rip2
#
# CIRCUITLESS IP INTERFACE CONFIGURATION
#
ip circuitless-ip-int 1 create 1.1.1.3/255.255.255.255
ip circuitless-ip-in 1 ospf enable
#
# OSPF CONFIGURATION
#
ip ospf admin-state enable
ip ospf as-boundary-router enable
ip ospf router-id 1.1.1.3
ip ospf enable
ip ospf area 0.0.0.3 create
ip ospf interface 10.1.1.41 create broadcast
ip ospf interface 10.1.1.25 area 0.0.0.3
ip ospf interface 10.1.1.2 area 0.0.0.3
ip ospf interface 1.1.1.3 area 0.0.0.3
#
# IP REDISTRIBUTION CONFIGURATION
#
ip ospf redistribute rip create
ip ospf redistribute rip metric 10
ip ospf redistribute rip enable
```

```
#
# RIP POLICY CONFIGURATION
#
ip rip interface 10.1.1.41 out-policy "Allow"
```

Configuration file for R4

```
#
# PORT CONFIGURATION - PHASE II
#
ethernet 2/1 ip create 10.1.1.26/255.255.255.252 2190 mac_offset 6
ethernet 2/1 ip ospf enable
ethernet 2/7 ip create 10.1.1.10/255.255.255.252 2134 mac_offset 1
ethernet 2/7 ip ospf enable
#
# CIRCUITLESS IP INTERFACE CONFIGURATION
#
ip circuitless-ip-int 1 create 1.1.1.4/255.255.255.255
ip circuitless-ip-int 1 ospf enable
#
# OSPF CONFIGURATION
#
ip ospf admin-state enable
ip ospf router-id 1.1.1.4
ip ospf enable
ip ospf area 0.0.0.3 create
ip ospf interface 10.1.1.26 area 0.0.0.3
ip ospf interface 10.1.1.10 area 0.0.0.3
ip ospf interface 1.1.1.4 area 0.0.0.3
```

Configuration file for R5

```
# VLAN CONFIGURATION
#
vlan 1 ports remove 1/1-1/48,2/1-2/8 member portmember
vlan 2 create byport 1 color 1
vlan 2 ports remove 1/3-1/48,2/1-2/8 member portmember
vlan 2 ports add 1/1-1/2 member portmember
vlan 2 ip create 172.3.3.1/255.255.255.0 mac_offset 0
vlan 2 ip ospf enable
vlan 2 ip ospf priority 0
#
# PORT CONFIGURATION - PHASE II
#
ethernet 2/6 ip create 10.1.1.18/255.255.255.252 2090 mac_offset 1
ethernet 2/6 ip ospf enable
ethernet 2/6 ip ospf priority 0
#
```

```
# CIRCUITLESS IP INTERFACE CONFIGURATION
#
ip circuitless-ip-int 1 create 1.1.1.5/255.255.255.255
ip circuitless-ip-int 1 ospf enable
#
# OSPF CONFIGURATION
#
ip ospf admin-state enable
ip ospf router-id 1.1.1.5
ip ospf enable
ip ospf area 0.0.0.2 create
ip ospf area 0.0.0.2 stub true
ip ospf interface 10.1.1.18 area 0.0.0.2
ip ospf interface 1.1.1.5 area 0.0.0.2
ip ospf interface 172.3.3.1 area 0.0.0.2
```


VRRP configuration examples

You can use Virtual Router Redundancy Protocol (VRRP) to eliminate single points of failure by providing dual-homed connectivity in routed environments.

VRRP uses an election process to select a *master* router that hosts use as the default gateway. If the master router (the default gateway) fails, the VRRP backup router automatically replaces the master router and becomes the new default gateway. In either case, the default gateway IP address and MAC address does not change, thereby providing transparent operation.

For load balancing applications that use Split-MLT (SMLT), the Passport 8600 switch can be configured in a Master-Master configuration, which allows both switches to respond to ARPs and forward traffic.

For more information about how to configure VRRP operations that use SMLT, refer to [“VRRP configuration example—VRRP operation with SMLT” on page 189](#).

You can configure the Passport 8600 switch’s VRRP Priority setting to select the VRRP master router for a specified VLAN. The VRRP Priority setting is an integer value, in the range 0 and 255, where the highest value is used to elect the VRRP master router. If two or more switches have the same priority value, the switch with the highest numerical IP address value is selected and becomes the VRRP master. The host is oblivious to the entire process.

When a host sends traffic to a different subnet, it sends an ARP request for the MAC address of the default gateway. In this case, the Passport 8600 VRRP master router replies with its *virtual* MAC address. The benefit of using a virtual MAC address is that, if the master router fails, the VRRP backup router uses the same virtual MAC address.

The virtual MAC address does not have to be configured on the Passport 8600.

On Passport 8600 switches, the virtual MAC address is automatically set for:

```
00-00-5E-00-01-<VRID>
```

where:

VRID = an integer value between 1 and 255 that represents the virtual router identification.

The virtual MAC address is assigned when you configure VRRP on a switch port or a VLAN, for example:

```
config vlan 2 ip vrrp 199 address 10.1.20.1
```

where:

199 is the VRID; therefore, the VRRP MAC address becomes:

```
00-00-5E-00-01-199.
```



Note: You should always try to load balance the VRRP master between the Passport 8600 switches.

This section includes the following topics:

- [“VRRP configuration example—Normal operation,”](#) next
- [“VRRP configuration example—VRRP operation with SMLT”](#) on page 189


VRRP configuration example—Normal operation

The following configuration example shows how you can provide VRRP service for two edge host locations (Figure 42). In this example, R1 is the VRRP master for S2 while R2 is the VRRP master for S1. For this example, we will use enable VRRP with OSPF as the routing protocol on R1 and R2.

As shown in Figure 42, the VRRP priority setting is used to select the VRRP master. The higher priority value becomes the VRRP Master. Note that if the vrrp priority settings for both switches have the same values, the higher IP address wins; therefore, it is very important to set the correct vrrp priority value.

VRRP Fast Advertisement is also enabled to allow for fast fail-over detection.

Figure 42 VRRP example



11038fa

The following sections provide step-by-step procedures that show how to configure R1 and R2 for this example.

Configuring R1

This section describes how to configure R1 to create the topology shown in [Figure 42 on page 179](#).

Configure R1 for VLAN 2 access

To configure R1 for VLAN 2 access, complete the following steps:

1 Configure VLAN 2 on R1:

The following command creates VLAN = 2 using spanning tree group = 1. If you want to use another STG, create the new STG group first, then add port 1/48 to the new STG group.

```
Passport-8610:5# config vlan 2 create byport 1
```

2 Configure the access port for VLAN 2 on R1:

The following command adds access port 1/48 to VLAN 2.

```
Passport-8610:5# config vlan 2 ports add 1/48
```

3 Configure an IP address for VLAN 2:

The following command adds IP address 10.1.20.2/24 to VLAN 2

```
Passport-8610:5# config vlan 2 ip create 10.1.20.2/24
```

4 Configure an OSPF interface on R1 VLAN 2:

The following command enables OSPF on R1 VLAN 2 and enables it as a passive interface.

```
Passport-8610:5# config vlan 2 ip ospf interface-type  
passive
```

```
Passport-8610:5# config vlan 2 ip ospf enable
```

5 Configure VRRP on R1 VLAN 2:

The following commands add the VRRP VIP address of 10.1.20.1 to VLAN 2 using VRID = 1. Note that for this example, the VRRP priority is not configured here; it is left at factory default of 100. Instead, the priority setting on R2 will be set to a higher value when R2 is configured.

Note also that fast advertisement is set to enable. This is proprietary to Nortel Networks to support an advertisement interval from 200 to 1000 ms with a default of 200. If you require normal vrrp, set fast advertisement to disable.

```
Passport-8610:5# config vlan 2 ip vrrp 1 address
10.1.20.1
Passport-8610:5# config vlan 2 ip vrrp 1 fast-adv-enable
enable
Passport-8610:5# config vlan 2 ip vrrp 1 enable
```

6 Disable spanning tree on access port 1/48:

The following command disables spanning tree on the port level.

```
Passport-8610:5# config ethernet 1/48 stg 1 stp disable
```

Configure R1 for VLAN 3 access

To configure R1 for VLAN 3 access, complete the following steps:

1 Configure VLAN 3 on R1:

The following command creates VLAN = 3 using spanning tree group = 1. If you want to use another STG, create the new STG group first, then add port 1/47 to the new STG group.

```
Passport-8610:5# config vlan 3 create byport 1
```

2 Configure the access port for VLAN 3 on R1:

The following command adds access port 1/47 to VLAN 3.

```
Passport-8610:5# config vlan 3 ports add 1/47
```

3 Configure an IP address for VLAN 3:

The following command adds IP address 10.1.21.2/24 to VLAN 3

```
Passport-8610:5# config vlan 3 ip create 10.1.21.2/24
```

4 Configure an OSPF interface on R1 VLAN 3:

The following command enables OSPF on R1 VLAN 3 and enables it as a passive interface.

```
Passport-8610:5# config vlan 3 ip ospf interface-type  
passive
```

```
Passport-8610:5# config vlan 3 ip ospf enable
```

5 Configure VRRP on R1 VLAN 3:

Note also that fast advertisement is set to enable. This is proprietary to Nortel Networks to support an advertisement interval from 200 to 1000 ms with a default of 200. If you require normal VRRP, set fast advertisement to disable.

```
Passport-8610:5# config vlan 3 ip vrrp 2 address  
10.1.21.1
```

```
Passport-8610:5# config vlan 3 ip vrrp 2 fast-adv-enable  
enable
```

```
Passport-8610:5# config vlan 3 ip vrrp 2 enable
```

6 Disable spanning tree on access port 1/47:

The following command disables spanning tree on the port level.

```
Passport-8610:5# config ethernet 1/47 stg 1 stp disable
```

Configuring R2

This section describes how to configure R2 to create the topology shown in [Figure 42 on page 179](#).

Configure R2 for VLAN 2 access

To configure R2 for VLAN 2 access, complete the following steps:

1 Configure VLAN 2 on R2:

The following command creates VLAN = 2 using spanning tree group = 1. If you want to use another STG, create the new STG group first, then add port 1/48 to the new STG group.

```
Passport-8610:5# config vlan 2 create byport 1
```

2 Configure the access port for VLAN 2 on R2:

The following command adds access port 1/48 to VLAN 2.

```
Passport-8610:5# config vlan 2 ports add 1/48
```

3 Configure an IP address for VLAN 2:

The following command adds IP address 10.1.20.3/24 to VLAN 2

```
Passport-8610:5# config vlan 2 ip create 10.1.20.3/24
```

4 Configure an OSPF interface on R2 VLAN 2:

The following command enables OSPF on R2 VLAN 2 and enables it as a passive interface.

```
Passport-8610:5# config vlan 2 ip ospf interface-type  
passive  
Passport-8610:5# config vlan 2 ip ospf enable
```

5 Configure VRRP on R2 VLAN 2:

The following commands add VRRP VIP address of 10.1.21.1 to VLAN 2. Note that for this example the VRRP priority value is set to 200, which allows it to be elected as the VRRP master router.

Note also that fast advertisement is set to enable. This is proprietary to Nortel Networks to support an advertisement interval from 200 to 1000 ms with a default of 200. If you require normal vrrp, set fast advertisement to disable.

```
Passport-8610:5# config vlan 2 ip vrrp 1 address
10.1.20.1
Passport-8610:5# config vlan 2 ip vrrp 1 priority 200
Passport-8610:5# config vlan 2 ip vrrp 1 fast-adv-enable
enable
Passport-8610:5# config vlan 2 ip vrrp 1 enable
```

6 Disable spanning tree on access port 1/48:

The following command disables spanning tree on the port level.

```
Passport-8610:5# config ethernet 1/48 stg 1 stp disable
```

Configure R2 for VLAN 3 access

To configure R2 for VLAN 3 access, complete the following steps:

1 Configure VLAN 3 on R2:

The following command creates VLAN = 3 using spanning tree group = 1. If you want to use another STG, create the new STG group first, then add port 1/47 to the new STG group.

```
Passport-8610:5# config vlan 3 create byport 1
```

2 Configure the access port for VLAN 3 on R2:

The following command adds access port 1/47 to VLAN 3.

```
Passport-8610:5# config vlan 3 ports add 1/47
```


3 Configure an IP address for VLAN 3:

The following command adds IP address 10.1.21.3/24 to VLAN 3

```
Passport-8610:5# config vlan 3 ip create 10.1.21.3/24
```

4 Configure an OSPF interface on R2 VLAN 3:

The following command enables OSPF on R2 VLAN 3 and enables it as a passive interface.

```
Passport-8610:5# config vlan 3 ip ospf interface-type  
passive  
Passport-8610:5# config vlan 3 ip ospf enable
```

5 Configure VRRP on R2 VLAN 3:

The following commands add VRRP VIP address of 10.1.20.1 to VLAN 3.

Note that for this example, the VRRP priority is not configured here; it is left at factory default of 100. Instead, the priority setting on R1 will be set to a higher value when R1 is configured.

Note also that fast advertisement is set to enable. This is proprietary to Nortel Networks to support an advertisement interval from 200 to 1000 ms with a default of 200. If you require normal vrrp, set fast advertisement to disable.

```
Passport-8610:5# config vlan 3 ip vrrp 2 address  
10.1.21.1  
Passport-8610:5# config vlan 3 ip vrrp 2 fast-adv-enable  
enable  
Passport-8610:5# config vlan 3 ip vrrp 2 enable
```

6 Disable spanning tree on access port 1/47:

The following command disables spanning tree on the port level.

```
Passport-8610:5# config ethernet 1/47 stg 1 stp disable
```

Viewing the VRRP status

After the Passport switches are configured, you can view the VRRP status for each switch.

To view the VRRP status, use the following show command:

```
show ip vrrp info
```

Figure 43 shows the sample output for R1, using the show ip vrrp info command.

Figure 43 show ip vrrp info command for R1





Figure 44 shows sample output for R2, using the show ip vrrp info command.

Figure 44 show ip vrrp info command for R2



Displaying VLAN configuration files

You can use the following show command to display the configuration commands and parameters used to create the topology shown in [Figure 42 on page 179](#):

```
Passport-8610:5# show config module vlan
```



Note: You can copy and paste the command outputs shown here to update your configuration files.

VLAN configuration file for R1

```
# VLAN CONFIGURATION
#
vlan 2 create byport 1
vlan 2 ports remove 1/1-1/47,2/1-2/8,3/1-3/8 member portmember
vlan 2 ports add 1/48 member portmember
```

```
vlan 2 ip create 10.1.20.2/255.255.255.0 mac_offset 0
vlan 2 ip ospf interface-type passive
vlan 2 ip ospf enable
vlan 2 ip vrrp 1 address 10.1.20.1
vlan 2 ip vrrp 1 enable
vlan 3 create byport 1
vlan 3 ip create 10.1.21.2/255.255.255.0 mac_offset 1
vlan 3 ip ospf interface-type passive
vlan 3 ip ospf enable
vlan 3 ip vrrp 2 address 10.1.21.1
vlan 3 ip vrrp 2 priority 200
vlan 3 ip vrrp 2 enable
#
# PORT CONFIGURATION - PHASE II
#
ethernet 1/47 stg 1 stp disable
ethernet 1/48 stg 1 stp disable
```

VLAN configuration file for R2

```
# VLAN CONFIGURATION
#
vlan 2 create byport 1
vlan 2 ports remove 1/1-1/47,2/1-2/8,3/1-3/8 member portmember
vlan 2 ports add 1/48 member portmember
vlan 2 ip create 10.1.20.3/255.255.255.0 mac_offset 0
vlan 2 ip ospf interface-type passive
vlan 2 ip ospf enable
vlan 2 ip vrrp 1 address 10.1.20.1
vlan 2 ip vrrp 1 priority 200
vlan 2 ip vrrp 1 enable
vlan 3 create byport 1


vlan 3 ports remove 1/1-1/46,1/48,2/1-2/8,3/1-3/8 member portmember
vlan 3 ports add 1/47 member portmember
vlan 3 ip create 10.1.21.3/255.255.255.0 mac_offset 1
vlan 3 ip ospf interface-type passive
vlan 3 ip ospf enable
vlan 3 ip vrrp 2 address 10.1.21.1
vlan 3 ip vrrp 2 enable
#
# PORT CONFIGURATION - PHASE II
#
ethernet 1/47 stg 1 stp disable
ethernet 1/48 stg 1 stp disable
```

VRRP configuration example—VRRP operation with SMLT

This configuration example shows how you can provide high availability for a Layer 2 edge switch feeding into a Layer 3 core. As shown in Figure 45, both R1 and R2 switches are configured with a port-based VLAN (VLAN 2) with SMLT and VRRP set to enable. This topology provides fail-over protection and load-balancing.

The BPS-2000 (BPS1) is configured with one port-based VLAN and one MultiLink Trunk (MLT) group for the aggregate uplink ports. Passport 8600 switches (R1 and R2) are configured with backup-master enabled so that both switches can reply to ARP.

Figure 45 VRRP example with SMLT



11039fa

The following sections provide step-by-step procedures that show how to configure R1 and R2 for this example.

Configuring R1 for VRRP and SMLT

This section describes how to configure R1 to create the topology shown in [Figure 45 on page 189](#).

Configure the IST VLAN configuration for R1

To configure the Inter Switch Trunk (IST) VLAN configuration for R1, complete the following commands:

1 Configure IST VLAN 3999 on R1:

The following command creates the IST VLAN = 3999 under the default STG = 1 group.

```
Passport-8610:5# config vlan 3999 create byport 1
Passport-8610:5#vlan 3999 ip create 2.1.1.1/24
```

2 Configure the IST MLT on R1:

The following command creates the IST MLT and adds the IST ports 2/1 and 2/2 using MLT ID = 1.

```
Passport-8610:5# config mlt 1 create
Passport-8610:5# mlt 1 add ports 2/1-2/2
Passport-8610:5# mlt 1 perform-tagging enable
```

3 Add the IST to VLAN 3999:

The following command adds the newly created IST to VLAN 3999.

```
Passport-8610:5# config vlan 3999 add-mlt 1
```

4 Configure an IST peer for R1:

The following command creates the IST peer and enables the IST link.

```
Passport-8610:5# config mlt 1 ist create ip 2.1.1.2
vlan-id 3999
Passport-8610:5# mlt 1 ist enable
```

Configure the IST Ethernet port configuration for R1

Nortel Networks recommends that you set the cp-limit value for all ports added to the IST group to disable.

To configure the IST Ethernet ports for R1, complete the following step:

- Disable cp-limit for both IST links on R1:

The following commands disable cp-limits for both IST links on R1.

```
Passport-8610:5# config ethernet 2/1 cp-limit disable
Passport-8610:5# config ethernet 2/2 cp-limit disable
```

Configure VRRP and SMLT for access VLAN to BPS1

To configure VRRP and SMLT for access to BPS1, complete the following steps:

- 1 Configure VLAN 2 on R1:

The following command creates VLAN = 2 using spanning tree group = 1. If you want to use another STG, create the new STG group first, then add port 1/48 to the new STG group.

```
Passport-8610:5# config vlan 2 create byport 1
```

- 2 Configure the access port for VLAN 2 on R1:

The following command adds access port 1/48 to VLAN 2.

```
Passport-8610:5# config vlan 2 ports add 1/48
```

- 3 Create SMLT on R1:

The following commands create SMLT (ID = 1) on R1

```
Passport-8610:5# config mlt 2 create
Passport-8610:5# config mlt 2 smlt create smlt-id 1
```

- 4 Add VLAN 2 to the IST and SMLT groups:

The following commands add VLAN 2 to the IST and SMLT groups.

```
Passport-8610:5# config vlan 2 add-mlt 1
Passport-8610:5# config vlan 2 add-mlt 2
```

5 Add port 1/48 to the SMLT:

```
Passport-8610:5# config mlt 2 add ports 1/48
```

6 Create an IP address for VLAN 2:

The following command adds IP address 10.1.20.2/24 to VLAN 2.

```
Passport-8610:5# config vlan 2 ip create 10.1.20.2/24
```

7 Enable a passive OSPF interface for VLAN 2 on R1:

The following commands enable OSPF on R1 and configures it as a passive interface for VLAN 2.

```
Passport-8610:5# config vlan 2 ip ospf interface-type  
passive
```

```
Passport-8610:5# config vlan 2 ip ospf enable
```

8 Configure a VRRP VIP address for VLAN 2 on R1:

The following commands add the VRRP VIP address of 10.1.20.1 to VLAN 2 with backup-master enabled, allowing both R1 and R2 to respond to ARP.

Note, that fast advertisement is enabled. This is proprietary to Nortel to support an advertisement interval from 200 to 1000 ms with default of 200. If normal vrrp is required, disable fast advertisement.

```
Passport-8610:5# config vlan 2 ip vrrp 1 address  
10.1.20.1
```

```
Passport-8610:5# config vlan 2 ip vrrp 1 backup-master  
enable
```

```
Passport-8610:5# config vlan 2 ip vrrp 1 enable
```


Configuring R2 for VRRP and SMLT

This section describes how to configure R2 to create the topology shown in [Figure 45 on page 189](#).

Configure the IST VLAN configuration for R2

To configure the Inter Switch Trunk (IST) VLAN configuration for R2, complete the following commands:

- 1 Configure IST VLAN 3999 on R2:

The following commands create the IST VLAN = 3999 under the default STG = 1 group.

```
Passport-8610:5# config vlan 3999 create byport 1
Passport-8610:5#vlan 3999 ip create 2.1.1.2/24
```

- 2 Configure the IST MLT on R2:

The following commands create the IST MLT and adds the IST ports 2/1 and 2/2 using MLT ID = 1.

```
Passport-8610:5# config mlt 1 create
Passport-8610:5# mlt 1 add ports 2/1-2/2
Passport-8610:5# mlt 1 perform-tagging enable
```

- 3 Add the IST to VLAN 3999:

The following command adds the newly created IST to VLAN 3999.

```
Passport-8610:5# config vlan 3999 add-mlt 1
```

- 4 Configure an IST peer for R2:

The following commands create the IST peer and enables the IST link.

```
Passport-8610:5# config mlt 1 ist create ip 2.1.1.1
vlan-id 3999
Passport-8610:5# mlt 1 ist enable
```

Configure the IST Ethernet port configuration for R2

Nortel Networks recommends that you set the cp-limit value for all ports added to the IST group to disable.

To configure the IST Ethernet ports for R2, complete the following step:

- Disable cp-limit for both IST links on R2:

The following commands disable cp-limits for both IST links on R2.

```
Passport-8610:5# config ethernet 2/1 cp-limit disable
Passport-8610:5# config ethernet 2/2 cp-limit disable
```

Configure VRRP and SMLT for access VLAN to BPS1

To configure VRRP and SMLT for access to BPS1, complete the following steps:

- 1 Configure VLAN 2 on R2:

The following command creates VLAN = 2 using spanning tree group = 1. If you want to use another STG, create the new STG group first, then add port 1/48 to the new STG group.

```
Passport-8610:5# config vlan 2 create byport 1
```

- 2 Configure the access port for VLAN 2 on R2:

The following command adds access port 1/48 to VLAN 2.

```
Passport-8610:5# config vlan 2 ports add 1/48
```

- 3 Create SMLT on R2:

The following commands create SMLT (ID = 1) on R2

```
Passport-8610:5# config mlt 2 create
Passport-8610:5# config mlt 2 smlt create smlt-id 1
```

- 4 Add VLAN 2 to the IST and SMLT groups:

The following commands add VLAN 2 to the IST and SMLT groups.

```
Passport-8610:5# config vlan 2 add-mlt 1
Passport-8610:5# config vlan 2 add-mlt 2
```

5 Add port 1/48 to the SMLT:

```
Passport-8610:5# config mlt 2 add ports 1/48
```

6 create an IP address for VLAN 2:

The following command adds IP address 10.1.20.3/24 to VLAN 2.

```
Passport-8610:5# config vlan 2 ip create 10.1.20.3/24
```

7 Enable a passive OSPF interface for VLAN 2 on R2:

The following commands enable OSPF on R2 and configure it as a passive interface for VLAN 2.

```
Passport-8610:5# config vlan 2 ip ospf interface-type  
passive
```

```
Passport-8610:5# config vlan 2 ip ospf enable
```

8 Configure a VRRP VIP address for VLAN 2 on R2:

The following commands add the VRRP VIP address of 10.1.20.1 to VLAN 2 with backup-master enabled, allowing both R1 and R2 to respond to ARP.

Note, that fast advertisement is enabled. This is proprietary to Nortel to support an advertisement interval from 200 to 1000 ms with default of 200. If normal vrrp is required, disable fast advertisement.

```
Passport-8610:5# config vlan 2 ip vrrp 1 address  
10.1.20.1
```

```
Passport-8610:5# config vlan 2 ip vrrp 1 backup-master  
enable
```

```
Passport-8610:5# config vlan 2 ip vrrp 1 enable
```

Displaying VLAN configuration files

You can use the following show command to display the configuration commands and parameters used to create the topology shown in [Figure 45 on page 189](#):

```
Passport-8610:5# show config
```



Note: You can copy and paste the command outputs shown here to update your configuration files.

Configuration file for R1

```
# MLT CONFIGURATION
#
mlt 1 create
mlt 1 add ports 2/1-2/2
mlt 1 perform-tagging enable
mlt 1 ist create ip 2.1.1.2 vlan-id 3999
mlt 1 ist enable
mlt 2 create
mlt 2 add ports 1/48
mlt 2 smlt create smlt-id 1

# VLAN CONFIGURATION
#
vlan 1 ports remove 1/48,2/1-2/2 member portmember
vlan 1 ip igmp mrdisc mrdisc-enable disable
vlan 2 create byport 1
vlan 2 add-mlt 1
vlan 2 add-mlt 2
vlan 2 ports remove 1/1-1/47,2/3-2/8,3/1-3/8 member portmember
vlan 2 ports add 1/48,2/1-2/2 member portmember
vlan 2 ip create 10.1.20.2/255.255.255.0 mac_offset 1
vlan 2 ip ospf interface-type passive
vlan 2 ip ospf enable
vlan 2 ip vrrp 1 address 10.1.20.1
vlan 2 ip vrrp 1 backup-master enable
vlan 2 ip vrrp 1 fast-adv-enable enable
vlan 2 ip vrrp 1 enable
vlan 3999 create byport 1
vlan 3999 add-mlt 1
vlan 3999 ports remove 1/1-1/48,2/3-2/8,3/1-3/8 member portmember
vlan 3999 ports add 2/1-2/2 member portmember
vlan 3999 ip create 2.1.1.1/255.255.255.0 mac_offset 0
```

```
# PORT CONFIGURATION - PHASE II
#
ethernet 2/1 default-vlan-id 3999
ethernet 2/1 cp-limit disable multicast-limit 15000 broadcast-limit
10000
ethernet 2/1 stg 1 stp disable
ethernet 2/2 default-vlan-id 3999
ethernet 2/2 cp-limit disable multicast-limit 15000 broadcast-limit
10000
ethernet 2/2 stg 1 stp disable
```

VLAN configuration file for R2

```
# MLT CONFIGURATION
#
mlt 1 create
mlt 1 add ports 2/1-2/2
mlt 1 perform-tagging enable
mlt 1 ist create ip 2.1.1.1 vlan-id 3999
mlt 1 ist enable
mlt 2 create
mlt 2 add ports 1/48
mlt 2 smlt create smlt-id 1
# VLAN CONFIGURATION
#
vlan 1 ports remove 1/48,2/1-2/2 member portmember
vlan 1 ip igmp mrdisc mrdisc-enable disable
vlan 2 create byport 1
vlan 2 add-mlt 1
vlan 2 add-mlt 2
vlan 2 ports remove 1/1-1/47,2/3-2/8,3/1-3/8 member
portmember
vlan 2 ports add 1/48,2/1-2/2 member portmember
vlan 2 ip create 10.1.20.3/255.255.255.0 mac_offset 1
vlan 2 ip ospf interface-type passive
vlan 2 ip ospf enable
vlan 2 ip vrrp 1 address 10.1.20.1
vlan 2 ip vrrp 1 backup-master enable
vlan 2 ip vrrp 1 fast-adv-enable enable
vlan 2 ip vrrp 1 enable
vlan 3999 create byport 1
vlan 3999 add-mlt 1
vlan 3999 ports remove 1/1-1/48,2/3-2/8,3/1-3/8 member
portmember
vlan 3999 ports add 2/1-2/2 member portmember
```

```
vlan 3999 ip create 2.1.1.2/255.255.255.0 mac_offset 0
# PORT CONFIGURATION - PHASE II
#
ethernet 2/1 default-vlan-id 3999
ethernet 2/1 cp-limit disable multicast-limit 15000
broadcast-limit 10000
ethernet 2/1 stg 1 stp disable
ethernet 2/2 default-vlan-id 3999
ethernet 2/2 cp-limit disable multicast-limit 15000
broadcast-limit 10000
ethernet 2/2 stg 1 stp disable
```

Chapter 3

Configuring IP routing using Device Manager

This chapter describes how to use Device Manager to perform basic IP routing interface configurations and management tasks. It discusses the basic IP router interface configuration required before any routing protocols, such as ARP, RIP and OSPF, can be configured.

- For conceptual information about interface configuration and router management, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter includes the following topics:

Topic	Page
Router interface types	200
Enabling or disabling per-port routing	205
Globally enabling IP routing features	206
IP router management	211
IP static route table overview	217
Configuring circuitless IP	228
Configuring ICMP router discovery	231

Router interface types

The 8000 series switch supports two types of router interfaces:

- Brouter ports

A brouter port is a single-port VLAN that can route IP packets as well as bridge all nonroutable traffic.

The difference between a brouter port and a standard IP protocol-based VLAN that is configured to do routing is: the routing interface of the brouter port is not subject to the spanning tree state of the port.

A brouter port is actually a one-port VLAN; therefore, each brouter port decreases the number of available VLANs by one and uses one VLAN ID.

- Virtual router interface

Virtual router interfaces correspond to routing on a virtual port that is associated with a VLAN. A virtual router interface allows routing of IP traffic to, and from, a VLAN. Because a given port can belong to multiple VLANs (some of which are configured for routing on the switch and some of which are not), there is no longer a one-to-one correspondence between the physical port and the router interface.

For VLAN routing, the router interface for the VLAN is called a *virtual router interface* because the IP address is assigned to an interface on the routing entity in the switch. This initial interface has a one-to-one correspondence with a VLAN on any given switch.

This section includes the following topics:

- [“Assigning an IP address on a brouter port,”](#) next
- [“Assigning an IP address to a virtual routing port”](#) on page 203

Assigning an IP address on a brouter port

A brouter port is a single-port VLAN that can route IP packets as well as bridge all non-routable traffic. The difference between a brouter port and a standard IP protocol-based VLAN (that is configured for routing), is the brouter port's routing interface is not affected by the port's spanning tree state. Therefore, when you use a brouter port, the spanning tree protocol is eliminated from the backbone network.

When you assign an IP address to a brouter port, note the following:

- You cannot edit the IP address, and you can assign only one IP address to any router interface (brouter or virtual).

Attempting to assign a second IP address returns an invalid IP address error.

- You also cannot assign an IP address to a brouter port that is a member of a routed VLAN. To assign an IP address to the brouter port, you must first remove it from the routed VLAN.
- You can assign a new IP address to a VLAN or a brouter port that already has an IP address, by first removing the existing IP address and then inserting the new IP address.

To configure an IP address on a brouter port:

- 1 From the Device Manager menu bar, select IP Routing > IP.

The IP dialog box opens with the Globals tab displayed ([Figure 51 on page 207](#)).

- 2 Select forwarding in the Forwarding check box (this action enables routing on the device).
- 3 On the device view, open the Port dialog box by completing any *one* of the following actions:
 - Double-click a port.
 - Right-click a port, and then choose Edit from the shortcut menu.
 - Select a port, and then choose Edit > Port from the Device Manager menu bar.

- Select a port, and then click the Edit Selected button from the Device Manager menu bar.




Edit Selected button

The Port dialog box opens with the Interface tab displayed (Figure 59 on page 226).

- 4 Click the IP Address tab.

The IP Address tab opens as in Figure 46 on page 202.


Figure 46 Port dialog box—IP Address tab



- 5 Click Insert.

The Port, Insert IP Address dialog box opens (Figure 47).

Figure 47 Port, Insert IP Address dialog box



- 6 Enter the IP address, Netmask, and VlanID.

- 7 Click Insert.

Table 7 describes the fields in the Port, Insert IP Address dialog box.


Table 7 Port, Insert IP Address dialog box fields

Field	Description
IpAddress	The IP address of the brouter interface on this port. Note that only one IP address can be defined on a given port interface.
NetMask	The subnet mask of the brouter interface on this port.
VlanId	The ID of the VLAN associated with the brouter port. This field is used for tagging ports.

Assigning an IP address to a virtual routing port

To specify an IP address for a virtual routing port:

- 1 From the Device Manager menu bar, select IP Routing > IP.
The IP dialog box opens with the Globals tab displayed (Figure 51).
- 2 Click forwarding to enable routing on the device.
- 3 Click Apply
- 4 Click Close
- 5 From the Device Manager menu bar, choose VLAN > VLANs.
The VLAN dialog box opens with the Basic tab displayed (Figure 48).

Figure 48 VLAN dialog box — Basic tab


- 6 Select a VLAN.

The IP button becomes highlighted.

7 Click IP.

The IP, VLAN dialog box opens with the IP Address tab displayed (Figure 49).


Figure 49 IP, VLAN dialog box—IP Address tab



8 Click Insert.

The IP, VLAN, Insert IP Address dialog box opens (Figure 50).

Figure 50 IP, VLAN, Insert IP Address dialog box



9 Enter the IP address and netmask.



Note: You can assign only one IP address to any router interface (brouter or VLAN). Attempting to assign a second IP address returns an invalid IP address error.

You cannot assign an IP address to a VLAN if a brouter port is a member of the VLAN. To assign an IP address to the VLAN, you must first remove the brouter port member.

10 Click Insert.

The new IP address and netmask appears in the IP, VLAN, Insert IP Address dialog box.

Enabling or disabling per-port routing

You can enable or disable routing capabilities on specified switch ports. The specified port can be part of a routed VLAN, while routing is disabled only on that port. The default setting for this feature is enable.

To enable or disable a port for routing:

- 1** On the device view, open the Port dialog box by completing any *one* of the following actions:
 - Double-click a port.
 - Right-click a port, and then choose Edit from the shortcut menu.
 - Select a port, and then choose Edit > Port from the Device Manager menu bar.
 - Select a port, and then click the Edit Selected button from the Device Manager menu bar.



Edit Selected button

The Port dialog box opens with the Interface tab displayed ([Figure 59 on page 226](#)).

- 2** In the AdminRouting field, click enable to set the port for routing; or click disable to set the port for bridging (and disable routing on this port).
- 3** Click Apply.
- 4** Click Refresh.

The OperRouting field (read only) changes to show the new configuration setting.

Globally enabling IP routing features

This section describes how to enable IP routing features globally, and contains the following topics:

- [“Enabling IP forwarding globally, next](#)
- [“Enabling ECMP globally” on page 209](#)
- [“Enabling alternative routes globally” on page 210](#)

Enabling IP forwarding globally

In Device Manager, the IP address of any physical or virtual router interface can be used for IP-based network management (SNMP, Telnet, and Web).

To enable IP forwarding:

- 1** From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed [Figure 51 on page 207](#).

Figure 51 IP dialog box—Globals tab

The screenshot shows the '192.168.151.163 - IP' dialog box with the 'Globals' tab selected. The dialog has a title bar with a close button. Below the title bar are tabs for 'ARP', 'Multicast ARP', 'Route Pref', 'Router Discovery', and 'Circuitless IP'. The 'Globals' tab is active, showing a 'Forwarding' section with radio buttons for 'forwarding' (selected) and 'not-forwarding'. Below this are input fields for 'DefaultTTL' (255), 'ReasmTimeout' (30 sec), 'ARPLifeTime' (360), and 'ArpThreshold' (500). There are also checkboxes for 'ICMPUnreachableMsgEnable', 'ICMPRedirectMsgEnable', 'AlternativeEnable' (checked), and 'RouteDiscoveryEnable'. A section for 'Ecmp' has a checkbox for 'EcmpEnable' and an 'EcmpMaxPath' field (1). Below these are four 'EcmpPathList' fields (1 through 4) and a checkbox for 'EcmpPathListApply'. At the bottom are 'Apply', 'Refresh', 'Close', and 'Help...' buttons.

192.168.151.163 - IP

ARP Multicast ARP Route Pref Router Discovery Circuitless IP

Globals Addresses Routes Static Routes

Forwarding: ☒ forwarding ☐ not-forwarding

DefaultTTL: 255 1..255

ReasmTimeout: 30 sec

ARPLifeTime: 360 1..32767 min

ArpThreshold: 500 50..1000

☐ ICMPUnreachableMsgEnable

☐ ICMPRedirectMsgEnable

☒ AlternativeEnable

☐ RouteDiscoveryEnable

☐ EcmpEnable

EcmpMaxPath: 1 1..4

Ecmp1PathList: ...

Ecmp2PathList: ...

Ecmp3PathList: ...

Ecmp4PathList: ...

☐ EcmpPathListApply

Apply Refresh Close Help...

- 2 Select forwarding in the Forwarding check box.
- 3 Click Apply.

Table 8 describes the Globals tab fields.

Table 8 IP dialog box—Globals tab fields

Field	Description
Forwarding	Sets the switch for forwarding (routing) or non forwarding. The default value is forwarding.
DefaultTTL	Sets the default Time-To-Live (TTL) value for a routed packet. TTL indicates the maximum number of seconds elapsed before a packet is discarded. Enter an integer between 1 and 255. The default value of 255 is inserted in the TTL field whenever one is not supplied in the datagram header.
ReasmTimeout	Read-only field—The maximum number of seconds that received fragments are held while they are waiting for reassembly at this entity. The default value is 30 seconds.
ARPLifeTime	The lifetime of an ARP entry within the system, global to the switch. The default value is 360 minutes. The range for this value is 1 through 32767 minutes.
ARPThreshold	ARP Threshold limits the number of unresolved ARP entries that can be stored on the switch. The default number of entries is 500 and it can vary between 50 and 1000 which is configured by the user.
ICMPNetUnreachableEnable	If checked, enables the generation of Internet Control Message Protocol (ICMP) net unreachable messages if the destination network is not reachable from this router. These messages assist in determining if the routing switch is reachable over the network. The default is disabled (not checked).
ICMPRedirectMsgEnable	Allows you to enable or disable the switch from sending ICMP destination redirect messages.
AlternativeEnable	Allows you to enable or disable the alternative-route feature globally. For more information about alternative routes, see Chapter 1, "IP routing concepts," on page 31 . Note: If the alternative-route parameter is disabled, all existing alternative routes are removed. When the parameter is enabled all alternative routes are added back.
RouteDiscoveryEnable	If checked, enables ICMP Route Discovery feature. The default is disabled (not checked).

Table 8 IP dialog box—Globals tab fields (continued)

Field	Description
EcmpEnable	Used to globally enable or disable the Equal Cost Multipath (ECMP) feature. The default is disabled. Note: When ECMP is disabled, the EcmpMaxPath is reset to the default value of 1.
EcmpMaxPath	Used to globally configure the maximum number of ECMP paths (1-4). The default value is 1. This feature cannot be configured unless ECMP is enabled globally.
EcmpPathList 1 - 4	Allows you to select a preconfigured ECMP path. To select a pathname: 1. Click the ellipses button, which appears to the right of the field. 2. Select the pathname in the EcmpPath dialog box. 3. Click Ok. 4. The EcmpPath dialog box closes. 5. In the IP dialog box, click Apply. The selected pathname appears in the specified (2 -4) EcmpPathList 1 - 4 field.
EcmpPathListApply	Click this field to apply any changes in the ECMP path list configuration or in the prefix-lists configured to be used as path list.

Enabling ECMP globally

The Equal Cost MultiPath (ECMP) feature allows routers to determine up to four equal cost paths to the same destination prefix. The multiple paths can be used for load sharing of traffic and allows faster convergence to alternate paths. By maximizing load sharing among equal-cost paths, you can use your links between routers more efficiently when sending IP traffic.

For more information about the ECMP feature, see [Chapter 1, “IP routing concepts,” on page 31](#).

To configure the ECMP feature:

- 1 From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).

- 2 In the EcmpEnable check box, click to enable (checked) or to disable (not checked) the ECMP feature. The default value is disabled (not checked).
- 3 Enter your preferred number of equal cost paths. You can configure up to four equal cost paths to the same destination prefix.

The default value is 1 when the EcmpEnable field is disabled. When the EcmpEnable field is enabled, the default is 4. The range for this value is 1 through 4 paths.

[Table 8 on page 208](#) describes the IP dialog box—Globals tab fields.

Enabling alternative routes globally

This section includes the following topics:

- [“Alternative routes overview,”](#) next
- [“Globally enabling alternative routes” on page 211](#)

Alternative routes overview

Software can execute several routes to a given destination network through several protocols. If the alternate route is enabled, it stores all of these routes sorted in order of preference/cost. The best route according to the preference/cost is used for the data forwarding. The remaining routes are referred to as alternate routes.

To avoid traffic interruption, the alternative-route feature can be enabled globally to replace best routes with the next-best route if the best route becomes unavailable. The alternate route concept is applied between routing protocols; for example if an OSPF route becomes unavailable and an alternate RIP route is available it is immediately activated without waiting for an update interval to expire.

For more information about alternative routes, see [Chapter 1, “IP routing concepts,” on page 31](#).

Globally enabling alternative routes

To enable alternative routes:

- 1 From the Device Manager menu bar, choose IP Routing > IP.
The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).
- 2 Select AlternativeEnable. If the alternative-route parameter is disabled, all existing alternative routes are removed. When the parameter is enabled all alternative routes are added back.
- 3 Click Apply.

[Table 8 on page 208](#) describes the Globals tab fields.

IP router management

In Device Manager, most of the dialog boxes related to managing the IP router are found under the IP Routing menu.

This section includes the following topics:

- “[Configuring a router’s IP protocol stack](#),” next
- “[Viewing IP addresses and their associated router interfaces](#)” on page 213
- “[Viewing and managing the system routing table](#)” on page 214

Configuring a router’s IP protocol stack

The IP dialog box contains parameters for configuring the router’s IP protocol stack.

To configure the router’s IP protocol stack:

- ➔ From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).

[Table 8 on page 208](#) describes the Globals tab fields.

Viewing IP addresses and their associated router interfaces

You can view IP addresses and their associated router interfaces in one central location.

To view IP addresses and their associated router interfaces:


- 1 From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).

- 2 Click the Addresses tab.

The Addresses tab opens ([Figure 52](#)).

Figure 52 IP dialog box—Addresses tab



[Table 9](#) describes the Addresses tab fields.

Table 9 Addresses tab fields

Field	Description
Interface	The router interface. - Virtual router interfaces are identified by the name of the VLAN followed by the VLAN designation. - Brouter interfaces are identified by the slot/port number of the brouter port.
IpAddress	The IP address of the router interface.

Table 9 Addresses tab fields (continued)

Field	Description
NetMask	The subnet mask of the router interface.
BcastAddrFormat	The IP broadcast address format used on this interface; that is, whether zero (0) or one (1) is used for the broadcast address. The Passport 8000 switch uses 1.
ReasmMaxSize	The size of the largest IP datagram that this entity can reassemble from incoming IP fragmented datagrams received on this interface (not editable).
VlanId	A value that uniquely identifies the virtual LAN associated with this entry. This value corresponds to the lower 12 bits in the IEEE 802.1Q VLAN tag.
BrouterPort	Used to indicate whether this entry corresponds to a brouter port (as opposed to a routable VLAN). This value cannot be changed after the row is created.
MacOffset	A user-assigned MAC address. This MAC address is used in place of the default MAC address.

Viewing and managing the system routing table

You can view the contents of the system routing table and delete a route whether it is a static or a dynamically learned route from RIP or OSPF. (Exercise care when deleting entries from the route table.)

To view or manage the system routing table:

- 1 From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).

- 2 Click the Routes tab.

The Routes tab opens ([Figure 53](#)).


Figure 53 IP dialog box—Routes tab

Table 10 describes the Routes tab fields.

Table 10 Routes tab fields

Field	Description
Dest	The destination IP network of this route. An entry with a value of 0.0.0.0 is considered a default route. Multiple routes to a single destination can appear in the table, but access to such multiple entries is dependent on the table access mechanisms defined by the network management protocol in use.
Mask	Indicate the network mask to be logically ANDed with the destination address before being compared to the value in the ipRouteDest field.
NextHop	The IP address of the next hop of this route.
AltSequence	The alternative route sequence. The value of 0 denotes the best route.
HopOrMetric	The primary routing metric for this route. The semantics of this metric are specific to different routing protocols.
Interface	The router interface for this route. - Virtual router interfaces are identified by the VLAN number of the VLAN followed by the (VLAN) designation. - Brouter interfaces are identified by the slot/port number of the brouter port.
Proto	The routing mechanism through which this route was learned: - local = directly learned - netmgmt = a static route - RIP - OSPF
Age	The number of seconds since this route was last updated or otherwise determined to be correct.

Table 10 Routes tab fields (continued)

Field	Description
PathType	The type of route: • direct • indirect Note that the values direct and indirect refer to the notion of direct and indirect routing in the IP architecture.
Pref	The Preference value.

IP static route table overview

The Static Route table is separate from the System Routing Table that the router uses to make forwarding decisions. The Static Route Table allows you to change static routes directly. Although the tables are separate, the Static Route Table Manager entries are automatically reflected in the System Routing Table if the next hop address in the static route is reachable, and if the static route is enabled.

The Static Route table is indexed by three attributes:

- Destination Network
- Destination Mask
- Next Hop

The maximum number of entries is 500. You can insert Static routes using the Static Route Table, and you can delete static routes by using either the Static Route Table or the System Routing Table.



Note: Only active static routes with a “best route” preference are displayed in the System Routing Table. A static route is active only if the route is enabled and the next hop address is reachable (for example, if there is a valid ARP entry for the next hop).

You can enter multiple routes (for example, multiple default routes) that have different costs, and the lowest-cost route that is reachable will be used in the routing table. Note that if you enter multiple next hops for the same route with the same cost, the software does not replace the existing route. If you enter the same route with the same cost and a different next hop, the first route is used. However, should that first route become unreachable, the second route (with a different next hop) is activated with no loss of connectivity.

Static routes that are configured for the management port are applied with the natural mask of the network. As traffic that originates from the switch refers to these routes before checking the IP routing table, the switch management traffic may be incorrectly forwarded out the management port, even though a more specific route exists in the routing table.

For more in-depth information about static routes, see [Chapter 1, “IP routing concepts,” on page 31](#).

This section includes the following topics:

- “Creating IP static routes,” next
- “Creating a static default route” on page 220
- “Creating a Black hole static route” on page 221
- “Deleting a static route” on page 222
- “Configuring IP route preferences” on page 223
- “Flushing routing tables” on page 224

Creating IP static routes

Static routes provide a way to create routes to destination IP address prefixes manually.

To create a static IP route:


- 1 From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).

- 2 Click the Static Routes tab.


The IP dialog box, Static Routes tab opens ([Figure 54](#)).

Figure 54 IP dialog box—Static Routes tab



- 3 Click Insert.

The IP, Insert Static Routes dialog box opens ([Figure 55](#)).

Figure 55 IP, Insert Static Routes dialog box

The dialog box is titled "134.177.229.235 - IP, Insert Static Routes". It contains the following fields and controls:

- Dest:** A text input field.
- Mask:** A text input field.
- NextHop:** A text input field.
- Metric:** A text input field containing the value "1", followed by the text "(number)".
- Preference:** A text input field containing the value "5", followed by the range "1..255".
- Enable:** A checked checkbox.
- LocalNextHop:** A checked checkbox.
- Buttons:** "Insert", "Close", and "Help..." at the bottom.

- 4 In the IP, Insert Static Routes dialog box Dest field, type the IP address.
- 5 In the Mask field, type the mask.
- 6 In the NextHop field, type the IP address of the router through which the specified route is accessible.
- 7 In the Metric field, type the HopOrMetric value.
- 8 In the Preference field, select the route preference.
- 9 Check the enable option.
- 10 Check the LocalNextHop option.
This field is used when creating L3 static routes.
- 11 Click Insert.

The new route appears in the IP dialog box, Static Routes tab.

Table 11 describes the fields in the IP dialog box, Static Routes tab.

Table 11 IP dialog box, Static Routes tab fields

Field	Description
Dest	Shows the destination network address.
Mask	Shows the destination mask.
NextHop	Displays the next hop IP address. When creating a black hole static route, set this field to 255.255.255.255.
Metric	Displays the primary routing metric for this route. If this metric is not used, set the value to 1.
IfIndex	The route index of the Next Hop.
Preference	This is the routing preference of the destination IP address.
Enable	Sets whether the configured static route is available on the port. The default is enable. Note: If a static route is disabled, it must be enabled before it can be added to the system routing table.
Status	Status of the route.
LocalNextHop	The IP address of the next hop of this route.

Creating a static default route

The default route is used to specify a route to all networks for which there are no explicit routes in the Forwarding Information Base or the routing table. This route is a route with the prefix length of zero (RFC1812). The routing switches can be configured with the default route statically, or they can learn it through a dynamic routing protocol.

**Note:** To create a default static route, the destination address and subnet mask must be set to 0.0.0.0.

To create a static default route:

- 1 From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).

- 2 Click the Static Routes tab.

The IP dialog box, Static Routes tab opens (see [Figure 54 on page 218](#)).

- 3 In the IP dialog box, Static Routes tab, click Insert.

The IP, Insert Static Routes dialog box opens (see [Figure 55 on page 219](#)).

- 4 In the IP, Insert Static Routes dialog box Dest field, type 0.0.0.0.
- 5 In the IP, Insert Static Routes dialog box Mask field, type 0.0.0.0.
- 6 In the NextHop field, select the router through which the specified route is accessible.
- 7 In the Metric field, type the HopOrMetric value.
- 8 Click Insert.

The default route record is created in the routing table.

Creating a Black hole static route

While aggregating or injecting routes to other routers, a router may not have a route to the aggregated destination, which causes a “black hole.” To avoid routing loops, you can configure a black hole static route to the destination it is advertising. A black hole route is a route with invalid next hop, so that the data packets destined to this network will be dropped by the switch.



Note: To create a black hole static route, the NextHop field must be set to 255.255.255.255.

To create a black hole static route:

- 1 From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).

- 2 Click the Static Routes tab.

The IP dialog box, Static Routes tab opens (see [Figure 54 on page 218](#)).

- 3 In the IP dialog box, Static Routes tab, click Insert.

The IP, Insert Static Routes dialog box opens (see [Figure 55 on page 219](#)).

- 4 In the IP, Insert Static Routes dialog box Dest field, enter the IP address.
- 5 In the IP, Insert Static Routes dialog box Mask field, enter the network mask.
- 6 In the NextHop field, type 255.255.255.255 as the IP address of the router through which the specified route is accessible.
- 7 In the Metric field, type the HopOrMetric value.
- 8 In the Preference field, select the route preference.
- 9 Check the enable option.
- 10 Click Insert.

The black hole static route record is created in the routing table.

Deleting a static route

To delete a static route:


- 1 From the Device Manager menu bar, choose IP routing> IP.

The IP dialog box opens with the Globals tab displayed.

- 2 Click the Static routes tab.

The Static Routes tab opens ([Figure 56](#)).

Figure 56 The Static Routes tab



- 3 Select a static route entry you wish to delete.
- 4 Click Delete.

The static route is removed from the Static Routes tab.

- 5 Click Close.

Configuring IP route preferences

The RoutePref tab displays the protocol, default, and configured IP global route preface information. You can use the RoutePref tab to edit IP route preference entries.



Note: Changing route preferences is a process-oriented operation that can affect system performance and network reachability while performing the procedures. Therefore, Nortel Networks recommends that if you want to change default preferences for routing protocols, you should do so before enabling the protocols.

To edit an IP route preference:

- 1 From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).

- 2 Click the RoutePref tab.

The RoutePref tab opens ([Figure 57](#)).

Figure 57 IP dialog box—RoutePref tab




Table 12 describes the RoutePref tab dialog box fields.

Table 12 RoutePref tab dialog box fields

Field	Description
Protocol	This is the name given to the protocol.
Default	This is the default preference value for the given protocol.
Configured	Allows you to change the default preference value for the given protocol.

Flushing routing tables

For administrative and troubleshooting purposes, it is sometimes necessary to flush the routing tables.

You can use Device Manager to flush the routing tables in two contexts:

Flushing by VLAN

To set flushing by VLAN:


- 1 From the Device Manager menu bar, choose VLAN > VLANs.

The VLAN dialog box opens with the Basic tab displayed [Figure 48 on page 203](#).

- 2 Click the Advanced tab.

The Advanced tab opens ([Figure 58](#)).

Figure 58 VLAN dialog box—Advanced tab



- 3 In the Vlan Operation Action field, select a flush option.

In a VLAN context, all entries associated with the VLAN will be flushed. The ARP entries and IP routes for the VLAN can be flushed.

Flushing by port


To set flushing by port:

- 1 On the device view, open the Port dialog box by completing any *one* of the following actions:
 - Double-click a port.
 - Right-click a port, and then choose Edit from the shortcut menu.
 - Select a port, and then choose Edit > Port from the Device Manager menu bar.
 - Select a port, and then click the Edit Selected button from the Device Manager menu bar.



Edit Selected button

The Port dialog box opens with the Interface tab displayed ([Figure 59 on page 226](#)).

Figure 59 Port dialog box—Interface tab


192.168.151.163 - Port 1/11

IGMP | OSPF | RIP | PIM | PGM | VRRP | Router Discovery | IPX BRouter

Remote Mirroring | Mroute Stream Limit | Fdb Protect | IP Address | ARP | DHCP | DVMRP

Interface | VLAN | STG | MAC Learning | Rate Limiting | Test | SMLT | PCAP | EAPOL | LACP | VLACP

Index: 74

Name:

Descr: 10/100BaseTX Port 1/11 Name

Type: rc100BaseTX

Mtu: 1950

PhysAddress: 00:04:38:7e:84:0a

VendorDescr:

AdminStatus: ☒ up ☐ down ☐ testing

OperStatus: down

LastChange: 08h:45m:45s

LinkTrap: ☒ enabled ☐ disabled

AutoNegotiate: ☒ true ☐ false

AdminDuplex: ☒ half ☐ full

OperDuplex: full

AdminSpeed: ☒ mbps10 ☐ mbps100

OperSpeed: 0

QosLevel: ☐ level0 ☒ level1 ☐ level2 ☐ level3 ☐ level4 ☐ level5 ☐ lev

☐ DiffServEnable

DiffServType: ☐ none ☐ access ☒ core

Apply Refresh Close Help...

2 Select flushAll.

In a port context, all entries associated with the port will be flushed. The ARP entries and IP routes for a port can be flushed.



Note: After you flush a routing table, it is not automatically repopulated. The time delay depends on the routing protocols in use.

Configuring circuitless IP

This section describes how to configure the circuitless IP feature.



Note: You can configure a maximum of 32 Circuitless IP interfaces on each device.

This section includes the following topics:

- “Configuring a circuitless IP interface,” next
- “Enabling OSPF on a circuitless IP interface” on page 229
- “Deleting a circuitless IP interface” on page 231


For conceptual information about the Circuitless IP feature, see [Chapter 1, “IP routing concepts,”](#) on page 31.

Configuring a circuitless IP interface

To configure a circuitless IP interface:

- 1 From the Device Manager menu bar, choose IP routing > IP.
The IP dialog box opens with the Globals tab displayed ([Figure 51](#)).
- 2 Click the Circuitless IP tab.
The Circuitless IP tab opens ([Figure 60](#)).


Figure 60 IP dialog box—Circuitless IP tab



3 Click Insert.

The IP, Insert Circuitless dialog box opens (Figure 61).

Figure 61 IP, Insert Circuitless dialog box



4 Enter an integer value in the Interface field (in the range 1 and 32).

5 Enter the IP address.

6 Enter the network Mask.

7 Click Insert.

The new interface is created and appears in the Circuitless IP tab (see Figure 60 on page 228).

Table 13 describes the Circuitless IP tab fields .

Table 13 IP dialog box, Circuitless IP tab fields

Field	Description
Interface	Displays the number assigned to the interface.The range is 1...32.
IP Address	Displays the IP address of the interface you are specifying as circuitless.
Net Mask	Displays the Net Mask address of the interface you are specifying as circuitless.

Enabling OSPF on a circuitless IP interface

To enable OSPF on an interface:

- 1 Select the interface (CLIP1, CLIP2, etc.) in the Circuitless IP tab dialog box.




Note: You must enable OSPF for Circuitless IP to function.

- 2 Click OSPF.

The OspfCircuitless dialog box opens.

Figure 62 OspfCircuitless dialog box



- 3 Click the Enable check box.

- 4 Click Apply to enable OSPF.



Note: When OSPF is enabled, the Circuitless IP interface is configured to OSPF backbone AreaId (0.0.0.0) until you change the configuration.

- 5 To change the OSPF backbone AreaId:
 - a Choose IP Routing > OSPF from the Device Manager menu bar.
 - b Click the Interfaces tab.
 - c Click in the current AreaID field to make the change to the OSPF backbone area.
- 6 Close the dialog box.

Deleting a circuitless IP interface

To delete a Circuitless IP interface:

- 1 From the Device Manager menu bar, choose IP routing > IP.

The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).
- 2 Click the Circuitless IP tab.

The Circuitless IP tab opens (see [Figure 60 on page 228](#)).
- 3 In the Interface column, select the CLIP number of the interface you want to delete.
- 4 Click Delete.

The new interface is deleted from the list of interfaces.
- 5 Close the dialog box.

Configuring ICMP router discovery

Internet Control Message Protocol (ICMP) router discovery specifies an extension to enable hosts attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers.

This section includes the following topics:

- [“Enabling ICMP router discovery globally,”](#) next
- [“Viewing the ICMP router discovery table”](#) on page 232
- [“Configuring router discovery on a VLAN”](#) on page 234
- [“Configuring router discovery on a port”](#) on page 236

Enabling ICMP router discovery globally

To enable ICMP router discovery globally on the switch:

- 1 From the Device Manager menu bar, choose IP Routing > IP.
The IP dialog box opens with the Globals tab displayed [Figure 51 on page 207](#).
- 2 Click RouteDiscoveryEnable.
- 3 Click Apply.
- 4 Close the dialog box.

Viewing the ICMP router discovery table

To view the ICMP router discovery table:

- 1 From the Device Manager menu bar, choose IP Routing > IP.
The IP dialog box opens with the Globals tab displayed (see [Figure 51 on page 207](#)).
- 2 Click the Router Discovery tab.
The Router Discovery tab opens ([Figure 63](#)).

Figure 63 IP dialog box—Router Discovery tab

Interface	AdvAddress	AdvFlag	AdvLifetime	MaxAdvInterval	MinAdvInterval	PreferenceLevel
5/1	255.255.255.255	true	1800	600	450	0

1 row(s)

Table 14 describes the Router Discovery tab fields.

Table 14 IP dialog box—Router Discovery tab fields

Item	Description
Interface	VLAN ID or the port.
AdvAddress	The IP destination address to be used for broadcast or multicast router advertisements sent from the interface. The accepted values are the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255. The default value is 255.255.255.255.
AdvFlag	A flag indicating whether (True) or not (False) the address is to be advertised on the interface. The default value is TRUE (advertise address).
AdvLifetime	The value (TTL) of router advertisements (in seconds) sent from the interface. The accepted value is no less than the MaxAdvInterval and no greater than 9000 seconds. The default value is 1800 seconds.
MaxAdvInterval	The maximum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The accepted values are no less than 4 seconds and no greater than 1800 seconds. The default value is 600 seconds.


Table 14 IP dialog box—Router Discovery tab fields (continued)

Item	Description
MinAdvInterval	The minimum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The accepted values are no less than 3 seconds and no greater than the MaxAdvInterval. The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet. The accepted values are -2147483648 to 2147483647. The default value is 0.

Configuring router discovery on a VLAN

To configure router discovery on a VLAN:

- 1 From the Device Manager menu bar, choose **VLAN > VLANS**.
The VLAN dialog box opens with the Basic tab displayed.
- 2 Click on the VLAN ID that you want to configure with router discovery.
Several buttons on the bottom of the dialog box become available.
- 3 Click **IP**.
The IP, VLAN dialog box opens with the IP Address tab displayed [Figure 49 on page 204](#).
- 4 Click the Router Discovery tab.
The IP, VLAN—Router Discovery tab opens ([Figure 64](#)).

Figure 64 IP, VLAN—Router Discovery tab


192.32.96.82 - IP, VLAN 10

IP Address | ARP | DHCP | DVMRP | IGMP | OSPF | RIP | PIM | PGM | VRRP | Router Discovery | Direct Broadcast

AdvAddress: 255.255.255.255

☒ AdvFlag

AdvLifetime: 1800 4..9000

MaxAdvInterval: 600 4..1800

MinAdvInterval: 450 3..1800

PreferenceLevel: 0 -2147483648..2147483647

Apply Refresh Close Help...

Table 15 describes the IP, VLAN— Router Discovery tab fields.

Table 15 IP, VLAN—Router Discovery tab fields

Field	Description
AdvAddress	The IP destination address to be used for broadcast or multicast router advertisements sent from the interface. The accepted values are the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255. The default value is 255.255.255.255.
AdvFlag	A flag indicating whether (True) or not (False) the address is to be advertised on the interface. The default value is TRUE (advertise address).
AdvLifetime	The value (TTL) of router advertisements (in seconds) sent from the interface. The accepted value is no less than the MaxAdvInterval and no greater than 9000 seconds. The default value is 1800 seconds.
MaxAdvInterval	The maximum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The accepted values are no less than 4 seconds and no greater than 1800 seconds. The default value is 600 seconds.

Table 15 IP, VLAN—Router Discovery tab fields (continued)

Field	Description
MinAdvInterval	The minimum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The accepted values are no less than 3 seconds and no greater than the MaxAdvInterval. The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet. The accepted values are -2147483648 to 2147483647. The default value is 0.

Configuring router discovery on a port


To configure router discovery on a port:

- 1 On the device view select a port.
- 2 From the menu, select Edit > Port.

The Port dialog box opens with the Interface tab displayed ([Figure 59](#)).

- 3 Click the Router Discovery tab.

The Router Discovery tab opens ([Figure 65](#)).

Figure 65 Port dialog box—Router Discover tab


192.32.96.82 - Port 1/1

Interface VLAN STG MAC Learning Rate Limiting Test IP Address ARP DHCP
DVMRP IGMP OSPF RIP PIM PGM VRRP Router Discovery IPX BRouter

AdvAddress: 255.255.255.255

☒ AdvFlag

AdvLifetime: 1800 4..9000

MaxAdvInterval: 600 4..1800

MinAdvInterval: 450 3..1800

PreferenceLevel: 0 -2147483648..2147483647

Apply Refresh Close Help...

- 4 Edit the parameters, using the parameter descriptions in Table 5. Make sure that the AdvFlag box is selected to indicate that you want the address advertised on the interface.

Table 16 describes the Port dialog box—Router Discovery tab fields.

Table 16 Port dialog box—Router Discovery tab fields

Field	Description
AdvAddress	The IP destination address to be used for broadcast or multicast router advertisements sent from the interface. The accepted values are the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255. The default value is 255.255.255.255.
AdvFlag	A flag indicating whether (True) or not (False) the address is to be advertised on the interface. The default value is TRUE (advertise address).
AdvLifetime	The value (TTL) of router advertisements (in seconds) sent from the interface. The accepted value is no less than the MaxAdvInterval and no greater than 9000 seconds. The default value is 1800 seconds.

Table 16 Port dialog box—Router Discovery tab fields (continued)

Field	Description
MaxAdvInterval	The maximum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The accepted values are no less than 4 seconds and no greater than 1800 seconds. The default value is 600 seconds.
MinAdvInterval	The minimum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The accepted values are no less than 3 seconds and no greater than the MaxAdvInterval. The default value is 450 seconds.
PreferenceLevel	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet. The accepted values are -2147483648 to 2147483647. The default value is 0.

Chapter 4

Configuring IP routing using the CLI

This chapter describes CLI commands that you use to configure layer 3 (routing) functions in your Passport 8000 switch. The chapter includes sections about the following command groups that you use to configure routing characteristics:

- For conceptual information about layer 3 routing functions, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter includes the following topics:

Command	Page
Roadmap of IP commands	240
IP routing commands	245
Show IP commands	270
Enabling or disabling per-port routing	274
Configuring Ethernet IP commands	275
VLAN IP commands	283
Configuring circuitless IP	288

Roadmap of IP commands

The following roadmap lists some of the IP commands and their parameters. Use this list as a quick reference or click on any command or parameter entry for more information.

Command	Parameter
<code>config ip</code>	<code>info</code> <code>alternative-route <enable disable></code> <code>icmp-unreach-msg <enable disable></code> <code>ecmp <enable disable></code> <code>ecmp-<1 2 3 4>-pathlist <prefix-list-name></code> <code>ip-supernet <enable disable></code> <code>icmp-redirect-msg <enable disable></code> <code>default-ttl <seconds></code>
<code>config ip forwarding</code>	<code>info</code> <code>disable</code> <code>enable</code>
<code>config ip route</code>	<code>info</code> <code>delete <ipaddr/mask> next-hop <value></code>
<code>config ip route preference</code>	<code>info</code> <code>protocol <protocol> <value></code>
<code>config ip route-discovery</code>	<code>info</code> <code>disable</code> <code>enable</code>
<code>config ip route-policy <policy name> seq <seq number></code>	<code>info</code> <code>action <permit deny> action <permit deny></code> <code>create</code> <code>delete</code>

Command	Parameter
	disable
	enable
	match-interface <prefix-list> [clear] match-interface <prefix-list> [clear]
	match-metric <metric> [clear]
	match-network <prefix-list> [clear]
	match-next-hop <prefix-list> [clear] match-next-hop <prefix-list> [clear]
	match-protocol <protocol name> [clear] match-protocol <protocol name> [clear]
	match-route-src <prefix-list> [clear] match-route-src <prefix-list> [clear]
	match-route-type <route-type> match-route-type <route-type>
	name <policy name>
	set-injectlist <prefix-list> [clear] set-injectlist <prefix-list> [clear]
	set-mask <ipaddr>
	set-metric <metric-value> [clear]
	set-metric-type <metric-type> [clear]
	set-preference <pref-value> [clear] set-preference <pref-value> [clear]
config ip static-route	info
	create <ipaddr/mask> next-hop <value> cost <value> [preference <value>] [local-next-hop <value>]

Command	Parameter
	<code>delete <ipaddr/mask> next-hop <value></code>
	<code>disable <ipaddr/mask> next-hop <value></code>
	<code>enable <ipaddr/mask> next-hop <value></code>
	<code>local-next-hop <true false> <ipaddr/mask> next-hop <value></code>
	<code>preference <value> <ipaddr/mask> next-hop <value></code>
<code>config ip mroute interface</code>	<code>info</code>
	<code>ttl <ttl></code>
<code>config ip mroute static-source-group</code>	<code>info</code>
	<code>create <SourceAddress/ SubnetMask></code>
	<code>delete <SourceAddress/ SubnetMask></code>
<code>config ethernet <ports> routing <enable disable></code>	
<code>config ethernet <ports> ip</code>	<code>info</code>
	<code>create <ipaddr/mask> <vid> [mac_offset <value>]</code>
	<code>delete <ipaddr/mask></code>
<code>config ethernet <ports> ip directed-broadcast</code>	<code>info</code>
	<code>disable</code>
	<code>enable</code>
<code>config ethernet <slot/port> ip route-discovery</code>	<code>info</code>
	<code>advertisement-address <value></code>
	<code>advertise-flag <true false></code>
	<code>advertisement-lifetime <seconds></code>

Command	Parameter
	max-advertisement-interval <seconds>
	min-advertisement-interval <seconds>
	preference-level <preference-level value>
config vlan <vlan-id> ip route-discovery	info
	advertisement-address <value>
	advertise-flag <true false>
	advertisement-lifetime <seconds>
	max-advertisement-interval <seconds>
	min-advertisement-interval <seconds>
	preference-level <preference-level value>
config vlan <vid> ip	info
	create <ipaddr/mask> [mac_offset <value>]
	delete <ipaddr>
config vlan <vid> ip directed-broadcast	info
	disable
	enable
config ip circuitless-ip-int	info
	area <ipaddr>
	create <ipaddr/mask>
	delete <ipaddr>
	ospf <enable/disable>
show config [verbose]	

Command	Parameter
---------	-----------

<code>show ip route preference info</code>	
<code>show ip route-policy info</code>	
<code>show ip forwarding</code>	
<code>show ip interface</code>	
<code>show ip route-discovery</code>	
<code>show ip route info</code>	
<code>show ip static-route info</code>	
<code>show ports info ip [<ports>]</code>	
<code>show port info route-discovery</code>	
<code>show vlan info ip [<vid>]</code>	
<code>show vlan info route-discovery</code>	
<code>show ip circuitless-ip-int info</code>	

IP routing commands

The general IP routing commands allow you to enable and disable IP forwarding (routing) on the switch, ports, and/or VLANs).

This section includes the following topics:

- “Configuring global parameters,” next
- “Configuring alternative routes” on page 247
- “Configuring IP forwarding” on page 248
- “Configuring IP routes” on page 248
- “Configuring IP route preferences” on page 249
- “Showing IP route preference information” on page 250
- “Configuring route discovery” on page 251
- “Configuring IP route policies” on page 253
- “Configuring IP static routes” on page 262
- “Creating Layer 3 static routes” on page 268
- “Creating a black hole static route” on page 269
- “Configuring an IP mroute interface” on page 269
- “Configuring an IP mroute static-source-group” on page 270

Configuring global parameters

The global `config ip` command includes the following options.

config ip followed by:	
info	Displays current config ip info command output (Figure 66).
alternative-route <enable disable>	Allows you to enable or disable alternative routes. the default value is enabled (see “Configuring alternative routes” on page 247). Note: If the alternative-route parameter is disabled, all existing alternative routes are removed. When the parameter is enabled all alternative routes are added back.

config ip followed by:	
icmp-unreach-msg <enable disable>	When enabled, allows the generation of Internet Control Message Protocol (ICMP) net unreachable messages if the destination network is not reachable from this router. These messages assist in determining if the routing switch is reachable over the network. The default is disabled.
ecmp <enable disable>	Allows you to enable or disable the ECMP feature. The default is disabled. Note: If the ecmp parameter is disabled, all existing ECMP routes are removed. When <i>ecmp</i> is enabled all ECMP routes are added back.
ecmp-<1 2 3 4>-pathlist <prefix-list-name>	Allows you to configure up to four equal cost paths to the same destination prefix. The default value is 1 when the ECMP is disabled. When ECMP is enabled, the default is 4. The range for this value is 1 to 4 paths. Note: This parameter cannot be set unless the ECMP feature is enabled globally.
ip-supernet <enable disable>	Allows you to enable or disable the switch supernet IP route. Note: If the <i>ip-supernet</i> feature is globally enabled, the switch can learn routes with a route mask less than eight bits. Routes with a mask length less than eight bits cannot have ECMP paths, even if the ECMP feature is globally enabled.
icmp-redirect-msg <enable disable>	Allows you to enable or disable the switch from sending ICMP destination redirect messages.
default-ttl <seconds>	Sets the default time to live (ttl) value for a routed packet. It is the maximum number of seconds before a packet is discarded. • <seconds> is a number between 1 and 255. The default value of 255 is inserted in the ttl field whenever one is not supplied in the datagram header.

Configuring alternative routes

You can use the `config ip` commands to enable or disable the alternative route and ECMP features. The maximum number of ECMP paths are set using this command. The Passport 8000 switch can learn multiple routes to a given destination network through several protocols. If you enable an alternative route, the switch stores all routes sorted in order of preference/cost. The best route is used for data forwarding according to the preference/cost. The remaining routes are referred to as alternate routes.

To avoid traffic interruption, you can enable the alternative-route feature globally to replace best routes with the next-best route if the best route becomes unavailable. The alternate route concept is applied between routing protocols. For example, if an OSPF route becomes unavailable and an alternate RIP route is available, it is immediately activated without waiting for an update interval to expire.

For more information about alternative routes, see [Chapter 1, “IP routing concepts,” on page 31](#).

[Figure 66](#) shows sample output for the `config ip info` command.

Figure 66 `config ip info` command output

```
Passport-8606:6# config ip info

Sub-Context: clear config dump monitor show test trace wsm
Current Context:

        alternative-route : enable
                default-ttl : 255 (sec.)
                ecmp : enable
        ecmp-1-pathlist :
        ecmp-2-pathlist :
        ecmp-3-pathlist :
        ecmp-4-pathlist :
ecmp-max-path-number : 4
        icmp-redirect-msg : disable
        icmp-unreach-msg : disable
        ip-supernet : disable
```

Configuring IP forwarding

The **config ip forwarding** command enable or disables IP forwarding (routing) on the entire switch. You can use this command to disable IP forwarding, thus allowing you to manage an Passport 8000 switch over a network without forcing the switch to also perform routing.

The **config ip forwarding** command includes the following options:

config ip forwarding followed by:	
info	Displays current config ip info command output (Figure 66).
disable	Disables IP forwarding (routing) on the entire switch.
enable	Enables IP forwarding (routing) on the entire switch. Default is enable.

Configuring IP routes

The **config ip route** command allows you display route information and to delete an IP route path.

The **config ip route** command includes the following options:

config ip route followed by:	
info	Displays route information.
delete <ipaddr/mask> next-hop <value>	Deletes a route. • <ipaddr/mask> is the IP address and mask for the route's destination. • <next-hop> - the next hop ip address for the route

Configuring IP route preferences

The `config ip route preference` command allows you to configure the route preference by protocol. This allows you to override default route preferences and gives preference to routes learned for a specific protocol.



Note: ECMP must be disabled before route preferences can be configured.



Note: Changing route preferences is a process-oriented operation that can affect system performance and network accessibility while performing the procedures. Therefore, Nortel Networks recommends that you change a prefix list or a routing protocol before enabling the protocols.

The `config ip route preference` command includes the following options:

<code>config ip route preference</code> followed by:	
<code>info</code>	Displays the route preference configured for different protocols (see Figure 67 on page 250).
<code>protocol <protocol> <value></code>	Sets the preference value for the specified protocol. If two protocols have the same configured value the default value is used for tie-breaking. • <code><protocol></code> must be set to one of the following: static, ospf-intra, ospf-inter, rip, ospf-external1, or ospf-external2. • <code><value></code> is from 1 to 255.

[Figure 67](#) shows sample out for the `config ip route preference` command.

Figure 67 config ip route preference command

```
Passport-8606:6# config ip route preference info

Sub-Context: clear config dump monitor show test trace wsm
Current Context:

protocol :
            LOCAL - 0
            STATIC - 5
            OSPF_INTRA - 20
            OSPF_INTER - 25
            EBGp - 45
            RIP - 100
            OSPF_E1 - 120
            OSPF_E2 - 125
            IBGP - 175
```

Showing IP route preference information

The **show ip route preference info** command displays information about IP route preferences.

The command uses the syntax:

```
show ip route preference info
```

[Figure 68 on page 251](#) shows sample output for this command.

Figure 68 show ip route preference command output

```
Passport-8606:6# show ip route preference info
```

```
=====
                                Ip Route Preference
=====
```

PROTOCOL	DEFAULT	CONFIG
LOCAL	0	0
STATIC	5	5
OSPF_INTRA	20	20
OSPF_INTER	25	25
EBGP	45	45
RIP	100	100
OSPF_E1	120	120
OSPF_E2	125	125
IBGP	175	175

Configuring route discovery

The **config route-discovery** command allows you to enable and disable route discovery.

The **config route-discovery** command includes the following options:

config ip route-discovery followed by:	
info	Displays the global status of the router discovery feature.
disable	Disables ICMP router discovery globally on the switch.
enable	Enables ICMP router discovery globally on the switch.

Configuration Example

The following configuration example uses the above command to:

- Disables ICMP router discovery globally
- Enables ICMP router discovery globally

After configuring the parameters, use the info command to show a summary of the results.

```
Passport-8010:6# config ip route-discovery
```

```
Passport-8010:6/config/ip/route-discovery# ?
```

Sub-Context:

Current Context:

disable

enable

info

```
Passport-8010:6/config/ip/route-discovery# enable
```

```
Passport-8010:6/config/ip/route-discovery# info
```

Sub-Context:

Current Context:

enable : true

```
Passport-8010:6/config/ip/route-discovery# disable
```

```
Passport-8010:6/config/ip/route-discovery# info
```

Sub-Context:

Current Context:

enable : false

```
Passport-8010:6/config/ip/route-discovery#
```

Configuring IP route policies

In the Passport 8000 switch, the behavior of IP route policies has been restructured to accommodate new scalability requirements. You can now form a unified database of route policies that can be used by the protocols RIP or OSPF for any type of filtering purpose. A policy is identified by a name or an ID.

Under a given policy you can have several sequence numbers, each of which is equal to one policy in the old convention. If you do not configure a field in a policy, it appears as 0 or “any” when it is displayed using the CLI `info` command. This indicates that the switch ignores the field in the match criteria. The `clear` option can be used to remove existing configurations for any field.



Note: Each policy sequence number contains a set of fields. Only a subset of those fields are used when the policy is applied in a certain context. For example, if a policy has a set-preference field set, it will be used only when the policy is applied for accept purposes. This field will be ignored when the policy is applied for announce/redistribute purpose.

You can apply one policy for one purpose, for example, RIP Announce, on a given RIP interface. In this case, all sequence numbers under the given policy apply to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

The `config ip route-policy <policy name> seq <seq number>` context includes the following commands that you can use to configure a route policy.

config ip route-policy <policy name> seq <seq number> followed by:	
info	Displays current configuration information about this policy sequence number.
action <permit deny>	This field specifies the action to be taken when a policy is selected for a specific route. This can be permit or deny. Permit allows the route, deny ignores the route.
create	Creates a route policy with a policy name and a sequence number. Note: When creating a route policy in the CLI, the ID is internally generated using an automated algorithm. When you create a route policy in Device Manager, you can manually assign the ID number.
delete	Deletes a route policy with a policy name and a sequence number.
disable	Disables a route policy with a policy name and a sequence number.
enable	Enables a route policy with a policy name and a sequence number.
match-as-path <as-list> [clear]	If configured, the switch matches the as-path attribute of the BGP routes against the contents of the specified as-lists. This field is used only for BGP routes and ignored for all other route types. • <as-list> specifies the list IDs of up to 4 as-lists, separated by a comma. • [clear] removes the configured value for match-as-path.
match-community <community-list> [clear]	If configured, the switch matches the community attribute of the BGP routes against the contents of the specified community-lists. This field is used only for BGP routes and ignored for all other route types. • <community-list> specifies the list IDs of up to four defined community-lists, separated by a comma. • [clear] removes the configured value for match-community.

config ip route-policy <policy name> seq <seq number> followed by:	
match-community-exact <enable disable> [clear]	When disabled, match-community results in a match when the community attribute of the BGP routes matches any entry of any community-list specified in match-community. When enabled, match-community results in a match when the community attribute of the BGP routes matches all of the entries of all the community-lists specified in match-community.
match-interface <prefix-list> [clear]	If configured, the switch matches the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other type of route. • <prefix-list> specify the name of up to four defined prefix list separated by a comma. • [clear] removes the configured value for match-interface.
match-metric <metric> [clear]	If configured, the switch matches the metric of the incoming advertisement or existing route against the specified value. If 0, then this field is ignored. • <metric> is 1 to 65535. The default is 0. • [clear] removes the configured value for match-metric.
match-network <prefix-list> [clear]	If configured, the switch matches the destination network against the contents of the specified prefix list(s). • <prefix-list> specify the name of up to four defined prefix list by name separated by a comma. • [clear] removes the configured value for match-network.
match-next-hop <prefix-list> [clear]	If configured, matches the next hop IP address of the route against the contents of the specified prefix list. This field applies only to non-local routes. • <prefix-list> specify the name of up to four defined prefix list by name separated by a comma. • [clear] removes the configured value for match-next-hop.
match-protocol <protocol name> [clear]	If configured, matches the protocol through which the route is learned. This field is used only for RIP announce purposes.

config ip route-policy <policy name> seq <seq number> followed by:	
match-route-src <prefix-list> [clear]	If configured, matches the next hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option ignored for all other route types. • <prefix-list> specify the name of up to four defined prefix list by name separated by a comma. • [clear] removes the configured value for match-route-src.
match-route-type <route-type>	Sets a specific route-type to be matched (applies only to OSPF routes). • <route-type> External-1 and External-2 specifies OSPF routes of the specified type only (any other value is ignored).
match-tag <tag> [clear]	Specifies a list of tag(s) that will be used during the match criteria process. Contains one or more tag values. • tag is a value from 0 to 256. • [clear] removes the configured value(s) for match-tag.
name <policy name>	This command is used to rename a policy once it has been created. This command changes the name field for all sequence numbers under the given policy.
set-as-path <as-list> [clear]	If configured, the switch adds the as number of the as-list to the BGP routes that match this policy. • <as-list> specifies the list id of up to four defined as-lists separated by a comma. • [clear] removes the configured value for set-as-path.
set-as-path-mode <tag prepend> [clear]	prepend is the default configuration. The switch prepends the as number of the as-list specified in set-as-path to the old as-path attribute of the BGP routes that match this policy.
set-automatic-tag <enable disable> [clear]	Sets the tag automatically. This option is used for BGP routes only.

config ip route-policy <policy name> seq <seq number> followed by:	
set-community <community-list> [clear]	If configured, the switch adds the community number of the community-list to the BGP routes that match this policy. • <i><community-list></i> specifies the list ID of up to four defined community -lists separated by a comma. • <i>[clear]</i> removes the configured value for <i>set-community</i>.
set-injectlist <prefix-list> [clear]	If configured, the switch replaces the destination network of the route that matches this policy with contents of the specified prefix list. • <i><prefix-list></i> specify one prefix list by name. • <i>[clear]</i> removes the configured value for <i>set-injectlist</i>.
set-community-mode <additive none> [clear]	Sets the community mode. • <i>additive</i> -- the switch prepends the community number of the community-list specified in <i>set-community</i> to the old community path attribute of the BGP routes that match this policy. • <i>none</i> --the switch removes the community path attribute of the BGP routes that match this policy to the specified value. • <i>[clear]</i> removes the configured value for <i>set-community-mode</i>.
set-local-pref <pref-value> [clear]	A value used during route decision process in the BGP protocol. Applicable to BGP only.
set-mask <ipaddr>	If configured, the switch sets the mask of the route that matches this policy. This applies only to RIP accept policies. <i><ipaddr></i> is a valid contiguous IP mask.
set-metric <metric-value> [clear]	If configured, the switch sets the metric value for the route while announcing a redistributing. The default is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or default-import-metric is used.
set-metric-type <metric-type> [clear]	If configured, sets the metric type for the routes to be announced into the OSPF domain that matches this policy. The default is type 2. This field is applicable only for OSPF announce policies.

config ip route-policy <policy name> seq <seq number> followed by:	
set-nssa-pbit <enable disable>	Sets the not-so-stubby-area (nssa) translation P bit. Applicable to OSPF announce policies only.
set-next-hop <ipaddr> [clear]	Specifies the IP address of the next hop router. Ignored for DVMRP routes.
set-origin <origin> [clear]	If configured, the switch changes the origin path attribute of the BGP routes that match this policy to the specified value.
set-origin-egp-as <origin-egp-as> [clear]	Indicates the remote autonomous sys number. Applicable to BGP only.
set-preference <pref-value> [clear]	Setting the preference greater than zero, specifies the route preference value to be assigned to the routes which matches this policy. This applies to accept policies only. • <pref-value> set from 0 to 255. The default is 0. If the default is configured, the global preference value is used. • [clear] removes the configured value for set-preference.
set-tag <tag> [clear]	Sets the tag of the destination routing protocol. If not specified, forward the tag value in the source routing protocol. A value of zero indicates that this parameter is not set.
set-weight <weight> [clear]	The weight value for the routing table. For BGP, this value will override the weight configured through NetworkTableEntry or FilterListWeight or NeighborWeight. Used for BGP only. A value of zero indicates that this parameter is not set.

Figure 69 displays sample output for this command.

Figure 69 config ip route-policy <policy name> seq <seq number> command

```

Passport-8606:6/config/ip/route-policy/test/seq/5#

Sub-Context:
Current Context:

    action <permit|deny>
    create
    delete
    disable
    enable
    info
    match-as-path <as-list> [clear]
    match-community <community-list> [clear]
    match-community-exact <enable|disable> [clear]
    match-interface <prefix-list> [clear]
    match-metric <metric> [clear]
    match-network <prefix-list> [clear]
    match-next-hop <prefix-list> [clear]
    match-protocol <protocol name> [clear]
    match-route-src <prefix-list> [clear]
    match-route-type <route-type>
    match-tag <tag> [clear]
    name <policy name>
    set-as-path <as-list-id> [clear]
    set-as-path-mode <tag|prepend> [clear]
    set-automatic-tag <enable|disable> [clear]
    set-community <community-list> [clear]
    set-community-mode <unchanged|additive|none> [clear]
    set-injectlist <prefix-list> [clear]
    set-local-pref <pref-value> [clear]
    set-mask <ipaddr>
    set-metric <metric-value> [clear]
    set-metric-type <metric-type> [clear]
    set-nssa-pbit <enable|disable>
    set-next-hop <ipaddr> [clear]
    set-origin <origin> [clear]
    set-origin-egp-as <origin-egp-as> [clear]
    set-preference <pref-value> [clear]
    set-tag <tag> [clear]
    set-weight <weight> [clear]

```

Figure 70 displays sample output for this command.

Figure 70 config ip route-policy <policy name> seq <seq number> info command

```
Passport-8606:6/config/ip/route-policy/policy-1/seq/1001# info
```

```
Sub-Context:
```

```
Current Context:
```

```

        id : 1
        seq : 1001
        name : policy-1
        enable : disable
        mode : permit
        match-protocol : N/A
        match-as-path :
        match-community :
        match-community-exact : disable
        match-interface :
        match-metric : 0
        match-network : prefix-2
        match-next-hop :
        match-route-type : any
        match-route-src :
        match-tag :
        set-as-path :
        set-as-path-mode : prepend
        set-automatic-tag : disable
        set-community :
        set-community-mode : unchanged
        set-local-pref : 0
        set-injectlist :
        set-mask : 0.0.0.0
        set-metric : 0
        set-metric-type : type2
        set-nssa-pbit : enable
        set-metric-type-internal : 0
        set-next-hop : 0.0.0.0
        set-origin : unchanged
        set-origin-egg-as : 0
        set-preference : 0
        set-tag : 0
        set-weight : 0
```

To display route policy information, use the following command:

```
show ip route-policy info
```

Figure 71 displays sample output for this command.

Figure 71 show ip route-policy info command

```
Passport-8606:6# show ip route-policy info
```

```
=====
```

```
                                Route Policy
```

```
=====
```

NAME	SEQ	MODE	EN
policy-1	1001	PRMT	DIS
policy-2	1002	PRMT	DIS
policy-3	1003	PRMT	DIS
policy-4	1004	PRMT	DIS
policy-5	1005	PRMT	DIS
policy-6	1011	PRMT	DIS
policy-8	1008	PRMT	DIS
policy-9	1009	PRMT	DIS
bob	1100	PRMT	EN
junky	10	PRMT	DIS

Figure 72 shows sample output for the show ip route-policy info ? command.

Figure 72 show ip route-policy info ? command

```
Passport-8010:5# show ip route-policy info ?
```

```
show boot flags
```

```
Optional parameters:
```

```
name <value> = <value> {string length 1..64}
```

```
seq <value> = <value> {0..65535}
```

```
all                = long format information of route policy
```

```
Command syntax:
```

```
info [name <value>] [seq <value>] [all]
```

Configuring IP static routes

The **config ip static-route** command allows you to create a new static route, or to modify existing static route parameters.

The **config ip static-route** command includes the following options:

config ip static-route followed by:	
info	Displays characteristics of the created static route (Figure 73).
create <ipaddr/mask> next-hop <value> cost <value> [preference <value>] [local-next-hop <value>]	Adds a static or default route to the switch. <i>ipaddr/mask</i> is the IP address and mask for the route's destination. <i>next-hop <value></i> is the IP address of the next hop router; the next router at which packets must arrive on this route. When creating a black hole static route, set this field to 255.255.255.255 as the IP address of the router through which the specified route is accessible. <i>cost <value></i> is the metric of the route.
delete <ipaddr/mask> next-hop <value>	Deletes a static route. <i>ipaddr/mask</i> is the IP address and mask for the route's destination. <i>next-hop <value></i> is the IP address of the next hop router; the next router at which packets must arrive on this route.
disable <ipaddr/mask> next-hop <value>	Disables a static route. <i>ipaddr/mask</i> is the IP address and mask for the route's destination. <i>next-hop <value></i> is the IP address of the next hop router; the next router at which packets must arrive on this route.
enable <ipaddr/mask> next-hop <value>	Enables a static route. <i>ipaddr/mask</i> is the IP address and mask for the route's destination. <i>next-hop <value></i> is the IP address of the next hop router; the next router at which packets must arrive on this route.

config ip static-route followed by:	
local-next-hop <true/false> <ipaddr/mask> next-hop <value>	Modify static route local-next-hop.
preference <value> <ipaddr/mask> next-hop <value>	Modify static route preference.

Figure 73 shows sample output for this command.

Figure 73 config ip static-route info command output

```

Passport-8606:6# config ip static-route info

create :
        - 1.0.0.0/255.0.0.0
        next-hop - 10.10.40.1
        cost - 1
        preference - 5
        local-next-hop - TRUE
        status - INACTIVE
        enable - TRUE
        - 6.6.6.0/255.255.255.0
        next-hop - 66.77.88.100
        cost - 1
        preference - 5
        local-next-hop - TRUE
        status - INACTIVE
        enable - TRUE

disable : N/A
delete : N/A
enable : N/A

```

The following configuration example uses the above command to:

- Adds a static or default route to the switch
- Deletes a static route
- Disables a static route
- Enables a static route

- Modify static route local-next-hop
- Modify static route preference

After configuring the parameters, use the info command to show a summary of the results.

```
Passport-8010:6/config/ip/static-route#  
Passport-8010:6/config/ip/static-route# ?
```

Sub-Context:

Current Context:

```
    create <ipaddr/mask> next-hop <value> cost <value> [preference  
<value>] [local-next-hop <value>]  
    delete <ipaddr/mask> next-hop <value>  
    disable <ipaddr/mask> next-hop <value>  
    enable <ipaddr/mask> next-hop <value>  
    info  
    local-next-hop <true|false> <ipaddr/mask> next-hop <value>  
    preference <value> <ipaddr/mask> next-hop <value>
```

```
Passport-8010:6/config/ip/static-route# create 0.0.0.0/0 next-hop  
60.1.60.51 cost 10  
Passport-8010:6/config/ip/static-route# info
```

```
create :  
        - 0.0.0.0/0.0.0.0  
    next-hop - 60.1.60.51  
        cost - 10  
    preference - 5  
    local-next-hop - TRUE  
        status - ACTIVE  
        enable - TRUE
```

```
disable : N/A
```


delete : N/A

enable : N/A

```
Passport-8010:6/config/ip/static-route# disable 0.0.0.0/0 next-hop  
60.1.60.51
```

```
Passport-8010:6/config/ip/static-route# info
```

```
create :  
          - 0.0.0.0/0.0.0.0  
next-hop - 60.1.60.51  
cost     - 10  
preference - 5  
local-next-hop - TRUE  
status   - INACTIVE  
enable   - FALSE
```

disable : N/A

delete : N/A

enable : N/A

```
Passport-8010:6/config/ip/static-route# local-next-hop false 0.0.0.0/0  
next-hop 60.1.60.51
```

```
Passport-8010:6/config/ip/static-route# info
```

```
create :  
          - 0.0.0.0/0.0.0.0  
next-hop - 60.1.60.51  
cost     - 10  
preference - 5  
local-next-hop - FALSE
```

```
status - INACTIVE
enable - FALSE
```

```
disable : N/A
```

```
delete : N/A
```

```
enable : N/A
```

```
Passport-8010:6/config/ip/static-route# preference 10 0.0.0.0/0 next-hop
60.1.60.51
```

```
Passport-8010:6/config/ip/static-route# info
```

```
create :
          - 0.0.0.0/0.0.0.0
next-hop - 60.1.60.51
cost     - 10
preference - 10
local-next-hop - TRUE
status   - INACTIVE
enable   - FALSE
```

```
disable : N/A
```

```
delete : N/A
```

```
enable : N/A
```

```
Passport-8010:6/config/ip/static-route# enable 0.0.0.0/0 next-hop
60.1.60.51
```

```
Passport-8010:6/config/ip/static-route# info
```

```
create :  
        - 0.0.0.0/0.0.0.0  
      next-hop - 60.1.60.51  
        cost - 10  
    preference - 10  
  local-next-hop - TRUE  
        status - ACTIVE  
        enable - TRUE
```

```
disable : N/A
```

```
delete : N/A
```

```
enable : N/A
```

```
Passport-8010:6/config/ip/static-route# delete 0.0.0.0/0 next-hop  
60.1.60.51
```

```
Passport-8010:6/config/ip/static-route# info
```

```
create : not created  
disable : N/A  
delete : N/A  
enable : N/A
```

```
Passport-8010:6/config/ip/static-route#
```

```
Passport-8010:6/config/ip/static-route# create default next-hop  
60.1.60.51 cost 10
```

```
Passport-8010:6/config/ip/static-route# info
```

```
create :  
        - 0.0.0.0/0.0.0.0  
      next-hop - 60.1.60.51  
        cost - 10  
    preference - 5  
  local-next-hop - TRUE
```

```
status - ACTIVE  
enable - TRUE
```

```
disable : N/A
```

```
delete : N/A
```

```
enable : N/A
```

Creating Layer 3 static routes

Layer 3 (L3) redundancy supports the creation of static routes to enhance network stability. When you configure a static route in primary SSF cards, the secondary SSF cards have the same setup through synchronization. You can configure a static route with local next hop or without local next hop by using the local-next-hop option.



Note: L3 redundancy supports only ARP and static route. None-local next-hop of static route supports only none-local next-hop configured by static ARP. No other dynamic routing protocols provide none-local next-hop.

To configure an L3 static route on the switch, use the **config ip static-route** command.

Figure 74 shows sample output for creating an L3 static route.

Figure 74 creating an L3 static route

```
8610:5#/config/ip/static-route# Create 172.2.0.0 next-hop  
172.2.3.3 cost 15
```

Creating a black hole static route

While aggregating or injecting routes to other routers, a router may not have a route to the aggregated destination, which causes a “black hole.” To avoid routing loops, you can configure a black hole static-route to the destination it is advertising.

A black hole route is a route with invalid next hop, so that the data packets destined to this network will be dropped by the switch.

When you specify a route preference, be sure that you configure the preference value appropriately so that when the black-hole route is used, it gets elected as the best route. Before adding the black hole route a check is made to ensure that no other static route to that identical destination in an enabled state exists. If such a route exists, then you are not allowed to add the black hole route, and an error message is generated.

However, if there is an enabled black hole route, then you will not be allowed to add another static route to that destination. You must first delete or disable the black hole route before you can add a regular static route to that destination.

Figure 75 shows sample output for creating a black hole static route.

Figure 75 creating a black hole static route

```
Passport-8610# config ip static-route create 10.10.0.0/16
next-hop 255.255.255.255 cost 1
```

Configuring an IP mroute interface

The **config ip mroute interface** command allows you to display current IP multicast route settings and set a default time-to-live for the interface.

The **config ip mroute interface** command includes the following options:

config ip mroute interface followed by:	
info	Displays IP multicast route settings.
ttl <ttl>	Sets the default time-to-live for the multicast route interface.

Configuring an IP mroute static-source-group

The **config ip mroute static-source-group** command allows you to display current IP multicast route settings and create or delete timed prune list entries.

The **config ip mroute static-source-group** command includes the following options:

config ip mroute static-source-group followed by:	
info	Displays IP multicast route settings.
create <SourceAddress/ SubnetMask>	Create a timed prune list entry.
delete <SourceAddress/ SubnetMask>	Delete the timed prune list entry created

Show IP commands

The show IP commands display the general IP characteristics of the switch.

This section includes the following topics:

- [“Showing IP forwarding status,” next](#)
- [“Showing IP interfaces” on page 271](#)
- [“Showing IP route discovery status” on page 272](#)

- [“Showing IP route table information” on page 272](#)
- [“Showing IP static-route information” on page 273](#)

Showing IP forwarding status

To display the status of IP forwarding (routing) on the switch, use the following command:

```
show ip forwarding
```

[Figure 76](#) shows sample output for this command.

Figure 76 show ip forwarding command output

```
Passport-8606:6# show ip forwarding

IP Forwarding is enabled
IP ECMP feature is enabled
Maximum ECMP paths number is 4
ECMP 1 pathlist :
ECMP 2 pathlist :
ECMP 3 pathlist :
ECMP 4 pathlist :
IP Alternative Route feature is enabled
IP ICMP Unreachable Message is disabled
IP Supernetting is disabled
IP Icmp-redirect-msg is disabled
IP Default TTL is 255 seconds
IP ARP life time is 360 minutes
```

Showing IP interfaces

To display the IP interfaces on the switch, use the following command:

```
show ip interface
```

[Figure 77](#) shows sample output for this command.

Figure 77 show ip interface command output

```
Passport-8606:6# show ip interface
```

=====						
Ip Interface						
=====						
INTERFACE	IP ADDRESS	NET MASK	BCASTADDR FORMAT	REASM MAXSIZE	VLAN ID	BROUTER PORT

Port6/1	10.10.54.27	255.255.255.0	ones	1500	0	false
Vlan2	200.1.1.1	255.255.255.0	ones	1500	--	false
Vlan3	111.111.111.111	255.255.255.0	ones	1500	--	false
Vlan4	66.77.88.99	255.255.255.0	ones	1500	--	false
Vlan5	55.66.77.88	255.255.255.0	ones	1500	--	false
Vlan7	5.5.5.5	255.255.255.0	ones	1500	--	false
Vlan11	33.33.33.33	255.255.255.0	ones	1500	--	false
Vlan14	78.67.67.77	255.255.255.0	ones	1500	--	false

Showing IP route discovery status

To show whether or not route discovery is enabled on the device, use the following command:

```
show ip route-discovery
```

Showing IP route table information

The following command displays the existing IP route table for the switch or for a specific net or subnet:

```
show ip route info
```

This command uses the syntax:

```
show ip route info [<ip address>] [-s <value>]
```

where:

- *ip address* is the specific net (1.2. = 1.2.0.0) {a.b.c.d}.
- *-s <value>* is the specific subnet {a.b.c.d/x | a.b.c.d/x.x.x.x | default}.

Figure 78 shows sample output for this command.

Figure 78 show ip route info command output

```
Passport-8606:6# show ip route info

=====
                                Ip Route
=====
      DST            MASK            NEXT COST VLAN  PORT PROT AGE TYPE PRF
-----
      200.1.1.0      255.255.255.0      200.1.1.1    1    2   -/-   LOC   0 DB
0

1 out of 1 Total Num of Dest Networks,1 Total Num of Route Entries displayed.
-----
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route, E=Ecmp
Route, U=Unresolved Route, N=Not in HW
```

Showing IP static-route information

To display the existing IP static routes for the switch or for a specific net or subnet, use the following command:

```
show ip static-route info
```

This command uses the syntax:

```
show ip static-route info [<ip address>] [-s <value>]
```

where:

- *ip address* is the specific net (1.2. = 1.2.0.0) {a.b.c.d}.
- *-s <value>* is the specific subnet {a.b.c.d/x | a.b.c.d/x.x.x.x | default}.

Figure 79 shows sample output for this command.

Figure 79 show ip static-route info command output

```
Passport-8606:6# show ip static-route info

Total number of static routes: 2

=====
                                Ip Static Route
=====
=====
DEST                MASK                NEXT                COST  PREF  LCLNHOP  STATUS  ENABLE
-----
1.0.0.0             255.0.0.0             10.10.40.1          1     5     TRUE     INACTV  TRUE
6.6.6.0             255.255.255.0         66.77.88.100        1     5     TRUE     INACTV  TRUE
Total 2
```

Enabling or disabling per-port routing

You can enable or disable routing capabilities on specified switch ports. The specified port can be part of a routed VLAN, while routing is disabled only on that port. The default setting for this feature is enable.

To enable or disable a port for routing, use the following command:

```
config ethernet <ports> routing <enable|disable>
```

Where:

- *ports* indicates the slot/port number of the port you are configuring.
- *enable | disable* allows you to enable or disable routing for the specified port.

Configuring Ethernet IP commands

This section describes some of the generic port-related IP routing commands. Other port commands are included in sections of this manual that describe commands that are used with a specific protocol or feature (for example, DHCP).



Note: You must enable ip forwarding on the switch to allow the **config ethernet <ports> ip** commands to take effect.

Use the following command to enable IP forwarding:

```
config ip forwarding enable
```

This section includes the following topics:

- [“Configuring Ethernet IP addresses,” next](#)
- [“Creating a brouter port” on page 276](#)
- [“Configuring a directed broadcast on a port” on page 277](#)
- [“Showing routing IP information” on page 278](#)
- [“Configuring route discovery on a port” on page 279](#)
- [“Showing ICMP router discovery information for all interfaces” on page 281](#)
- [“Showing ICMP router discovery information for all VLANs” on page 281](#)
- [“Showing ICMP router discovery information for all ports” on page 282](#)

Configuring Ethernet IP addresses

The `config ethernet <ports> ip` command includes the following options:

<code>config ethernet <ports> ip</code> followed by:	
<code>info</code>	Displays configured IP characteristics on the port (Figure 81).
<code>create <ipaddr/mask> <vid> [mac_offset <value>]</code>	Assigns an IP address to a port. Assigning an IP address to a port creates a brouter port (see Creating a brouter port , next). • <code><ipaddr/mask></code> is the IP address and mask {a.b.c.d}. • <code><vid></code> is the VLAN ID {1..4094}. • <code>mac_offset <value></code> is a user-assigned MAC address. This MAC address is used in place of the default MAC address.
<code>delete <ipaddr/mask></code>	Deletes an IP address from a brouter port.

Creating a brouter port

To create a brouter port, you must first create a routed IP policy-based single-port VLAN. You can then create the brouter port by assigning an IP address to the port and specifying a VLAN ID for that port.

To create the brouter port and display the brouter port information for the associated VLANs, enter the following command sequence:

```
config ethernet <ports> ip create <ipaddr/mask> <vid>

show ports info brouter-port
```

[Figure 80](#) shows sample output for this command.

Figure 80 show ports info brouter-port command output

```
Passport-8610# show ports info brouter-port
```

Port	Vlan Id
====	=====
1/1	2

Figure 81 shows sample output for the **config ethernet ip info** command.

Figure 81 config ethernet ip info command output

```
Passport-8610# config ethernet 9/13 ip info
```

```
Sub-Context:
```

```
Current Context:
```

```
Port 9/13 :
```

```
create :5.5.5.5/255.0.0.0 Vlan5 mac_offset 0
```

```
delete : N/A
```

Configuring a directed broadcast on a port

A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. Directed broadcast suppression protects hosts from possible denial of service (DOS) attacks.

The **config ethernet <ports> ip directed-broadcast** command allows you to enable or disable the directed broadcast suppression configuration settings.



Note: When directed broadcast suppression is enabled (the default setting), the CPU does not receive a copy of the directed broadcast. As a result, the switch does not respond to a subnet broadcast ping sent from a remote subnet.

The `config ethernet <ports> ip directed-broadcast` command includes the following options:

<code>config ethernet <ports> ip directed-broadcast</code> followed by	
info	Displays information about the directed broadcast suppression settings (Figure 82).
disable	Disables directed broadcast suppression on the specified port or ports. By disabling or suppressing directed broadcasts on an interface, you cause all frames sent to the subnet broadcast address for a local router interface to be dropped.
enable	Enables directed broadcast suppression on the specified port or ports. The default setting is enabled.

Showing routing IP information

To display routing (IP) information about the specified port or for all ports, use the following command:

```
show ports info ip [<ports>]
```

Figure 82 shows sample output for this command.

Figure 82 show ports info ip command output

Passport-8606:6# show ports info ip

Port Ip						
PORT NUM	IP_ADDRESS	NET_MASK	BROADCAST	REASM MAXSIZE	ADVERTISE WHEN_DOWN	DIRECT BCAST
1/15	222.222.222.222	255.0.0.0	ones	1500	disable	enable

Configuring route discovery on a port

The `config ethernet <slot/port> ip route-discovery` command includes the following options:

config ethernet <slot/port> ip route-discovery followed by:	
info	Displays current configuration information about the ICMP router discovery parameters.
advertisement-address <value>	The IP destination address to be used for broadcast or multicast router advertisements sent from the interface. The accepted values are the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255. The default value is 255.255.255.255.
advertise-flag <true false>	A flag indicating whether (True) or not (False) the address is to be advertised on the interface. The default value is TRUE (advertise address).
advertisement-lifetime <seconds>	The value (TTL) of router advertisements (in seconds) sent from the interface. The accepted value is no less than the MaxAdvInterval and no greater than 9000 seconds. The default value is 1800 seconds.
max-advertisement-interval <seconds>	The maximum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the router interface. The accepted values are no less than 4 seconds and no greater than 1800 seconds. The default value is 600 seconds.
min-advertisement-interval <seconds>	The minimum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The accepted values are no less than 3 seconds and no greater than the Max AdvInterval. The default value is 450 seconds.
preference-level <preference-level value>	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet. The accepted values are -2147483648 to 2147483647. The default value is 0.

Figure 83 shows a configuration example that uses the commands described above to configure ICMP router discovery on a VLAN and an ethernet port. After configuring the parameters, use the `info` command to show a summary of the results.

Figure 83 Route-discovery configuration examples

```

8610:5/config/vlan/1/ip/route-discovery#
8610:5/config/vlan/1/ip/route-discovery# advertisement-address 255.255.255.255
8610:5/config/vlan/1/ip/route-discovery# advertise-flag true
8610:5/config/vlan/1/ip/route-discovery# advertisement-lifetime 1800
8610:5/config/vlan/1/ip/route-discovery# max-advertisement-interval 600
8610:5/config/vlan/1/ip/route-discovery# min-advertisement-interval 450
8610:5/config/vlan/1/ip/route-discovery# preference-level 0
8610:5/config/vlan/1/ip/route-discovery# info

Sub-Context:
Current Context:
advertisement-address : 255.255.255.255
                    advertise-flag : true
                    advertisement-lifetime : 1800
                    max-advertisement-interval : 600
                    min-advertisement-interval : 450
                    preference-level : 0
8610:5/config/vlan/1/ip/route-discovery# box
8610:5# config ethernet 1/3 ip route-discovery
8610:5/config/ethernet/1/3/ip/route-discovery# advertisement-address
255.255.255.255
8610:5/config/ethernet/1/3/ip/route-discovery# advertise-flag true
8610:5/config/ethernet/1/3/ip/route-discovery# advertisement-lifetime 1800
8610:5/config/ethernet/1/3/ip/route-discovery# max-advertisement-interval 600
8610:5/config/ethernet/1/3/ip/route-discovery# min-advertisement-interval 450
8610:5/config/ethernet/1/3/ip/route-discovery# preference-level 0
8610:5/config/ethernet/1/3/ip/route-discovery# info

Sub-Context:
Current Context:
advertisement-address : 255.255.255.255
                    advertise-flag : true
                    advertisement-lifetime : 1800
                    max-advertisement-interval : 600
                    min-advertisement-interval : 450
                    preference-level : 0
8610:5/config/ethernet/1/3/ip/route-discovery#

```


Showing ICMP router discovery information for all interfaces

To show information about the parameters configured on the interfaces, use the following command:

```
show config [verbose]
```

where:

- `verbose` shows all interface-specific parameters for the interface, including those that do not differ from their default values.

Figure 84 shows sample output for this command.

Figure 84 show config verbose command output

```
8610:5#show config verbose
.
.
.
vlan 1 ip route-discovery advertisement-address 255.255.255.255
vlan 1 ip route-discovery advertise-flag true
vlan 1 ip route-discovery advertisement-lifetime 1800
vlan 1 ip route-discovery max-advertisement-interval 600
vlan 1 ip route-discovery min-advertisement-interval 450
vlan 1 ip route-discovery preference-level 0
.
.
.
8610:5#
```

Showing ICMP router discovery information for all VLANs

To show ICMP router discovery information for all VLANs, use the following command:

```
show vlan info route-discovery
```

To show all router discovery parameters for a specific VLAN, use the following command:


```
show vlan info route-discovery <vlan-id>
```

where:

vlan-id is the unique number that identifies the VLAN (1 to 4094).

Figure 85 shows sample output for this command.

Figure 85 show vlan info route-discovery command output



```
8610:5# show vlan info route-discovery 1
```

Vlan Ip Icmp Route Discovery						
VLAN_ID	ADV_ADDRESS	ADV_FLAG	LIFETIME	MAX_INT	MIN_INT	PREF_LEVEL
1	255.255.255.255	true	1800	600	450	0

8610:5#

Showing ICMP router discovery information for all ports

To show ICMP router discovery information for all ports, use the following command:

```
show port info route-discovery
```

To show router discovery information for one or more specific ports, use the following command:

```
show port info route-discovery <slot/port>
```

where:

slot/port specifies the port for which you are entering the command. To enter more than one port, use the form *slot/port, slot/port* [...].

Figure 86 shows sample output for this command.

Figure 86 show port info route-discovery command output

```
8610:5# show port info route-discovery 1/1

=====
Port Ip Icmp Route Discovery
=====
PORT_NUM ADV_ADDRESS      ADV_FLAG LIFETIME  MAX_INT  MIN_INT  PREF_LEVEL
-----
1/1      255.255.255.255  true    1800      600      450      0

8610:5#
```

VLAN IP commands

The VLAN IP commands are the general routing commands for the VLAN. Other VLAN commands are included in the sections of this manual that describe commands that are used with a specific protocol or feature (for example, DHCP).

This section includes the following topics:

- [“Configuring a VLAN”](#), next
- [“Configuring a directed-broadcast on a VLAN”](#) on page 285
- [“Configuring route discovery on a VLAN”](#) on page 285
- [“Showing VLAN information”](#) on page 287

Configuring a VLAN

The general `config vlan ip <vid>` command requires that you enter a VLAN ID (VID) along with the command. The range is 1 to 4094.

The `config vlan ip <vid>` command includes the following options:

<code>config vlan <vid> ip</code> followed by:	
<code>info</code>	Displays VLAN routing characteristics (Figure 87).
<code>create <ipaddr/mask></code> <code>[mac_offset <value>]</code>	Assigns an IP address and subnet mask to the VLAN. • <code><ipaddr/mask></code> is the IP address and mask {a.b.c.d}. • <code>mac_offset <value></code> is a user-assigned MAC address. This MAC address is in place of the default MAC address.
<code>delete <ipaddr></code>	Deletes the specified VLAN IP address.

Figure 87 shows sample output for this command.

Figure 87 config vlan <vid> ip info command output

```
Passport-8606:6# config vlan 5 ip info

Sub-Context: clear config dump monitor show test trace wsm
Current Context:

                                create : 55.66.77.88/255.255.255.0
mac_offset 2                    delete : N/A
```

Configuring a directed-broadcast on a VLAN

A directed broadcast is a frame sent to the subnet broadcast address on a remote IP subnet. Directed broadcast suppression protects hosts from possible denial of service (DOS) attacks.

The `config vlan <vid> ip directed-broadcast` command allows you to enable or disable the directed broadcast suppression configuration settings.



Note: When directed broadcast suppression is enabled (the default setting), the CPU does not receive a copy of the directed broadcast. As a result, the switch does not respond to a subnet broadcast ping sent from a remote subnet.

The `config vlan <vid> ip directed-broadcast` command includes the following options:

<code>config vlan <vid> ip directed-broadcast</code> followed by	
info	Displays information about the directed broadcast suppression settings.
disable	Disables directed broadcast suppression on the specified VLAN. By disabling or suppressing directed broadcasts on an interface, you cause all frames sent to the subnet broadcast address for a local router interface to be dropped.
enable	Enables directed broadcast suppression on the specified VLAN. • The default setting is enabled.

Configuring route discovery on a VLAN

The `config vlan <vid-id> ip route-discovery` command enables and disables ip route-discovery features on a VLAN. It also displays ip route-discovery status on the VLAN.

The **config vlan <vid-id> ip route-discovery** command includes the following options:

config vlan <vlan-id> ip route-discovery followed by:	
info	Displays current configuration information about the VLAN ICMP router discovery parameters.
advertisement-address <value>	The IP destination address to be used for broadcast or multicast router advertisements sent from the interface. The accepted values are the all-systems multicast address, 224.0.0.1, or the limited-broadcast address, 255.255.255.255. The default value is 255.255.255.255.
advertise-flag <true false>	A flag indicating whether (True) or not (False) the address is to be advertised on the interface. The default value is TRUE (advertise address).
advertisement-lifetime <seconds>	The value (TTL) of router advertisements (in seconds) sent from the interface. The accepted value is no less than the MaxAdvInterval and no greater than 9000 seconds. The default value is 1800 seconds.
max-advertisement-interval <seconds>	The maximum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the router interface. The accepted values are no less than 4 seconds and no greater than 1800 seconds. The default value is 600 seconds.
min-advertisement-interval <seconds>	The minimum time (in seconds) allowed between sending unsolicited broadcast or multicast router advertisements from the interface. The accepted values are no less than 3 seconds and no greater than the Max AdvInterval. The default value is 450 seconds.
preference-level <preference-level value>	Specifies the preference value (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet. The accepted values are -2147483648 to 2147483647. The default value is 0.

Showing VLAN information

To display the routing (IP) configuration for all VLANs on the switch or for a specified VLAN, use the following command:

```
show vlan info ip [<vid>]
```

[Figure 88](#) shows sample output for this command.

Figure 88 show vlan info ip command output

```
Passport-8606:6# show vlan info ip
```

=====						
Vlan Ip						
=====						
VLAN ID	IP ADDRESS	NET MASK	BCASTADDR FORMAT	REASM MAXSIZE	ADVERTISE WHEN_DOWN	DIRECTED BROADCAST

2	200.1.1.1	255.255.255.0	ones	1500	disable	enable
3	111.111.111.111	255.255.255.0	ones	1500	disable	enable
4	66.77.88.99	255.255.255.0	ones	1500	disable	enable
5	55.66.77.88	255.255.255.0	ones	1500	disable	enable
7	5.5.5.5	255.255.255.0	ones	1500	disable	enable
11	33.33.33.33	255.255.255.0	ones	1500	disable	enable
14	78.67.67.77	255.255.255.0	ones	1500	disable	enable

Configuring circuitless IP

This section describes how to configure the circuitless IP feature.



Note: You can configure a maximum of 32 Circuitless IP interfaces on each device.

This section includes the following topics:

- [“Configuring circuitless IP on an interface,”](#) next
- [“Showing circuitless IP information”](#) on page 290

For conceptual information about the Circuitless IP feature, see [Chapter 1, “IP routing concepts,”](#) on page 31.

Configuring circuitless IP on an interface

To configure circuitless IP, use the following command:

```
config ip circuitless-ip-int <id>
```

where:

<id> is an integer value in the range 1 to 32 that indicates the identification number for the specific circuitless ip interface.

This command includes the following options:

config ip circuitless-ip-int followed by:	
info	Displays the configured parameters for the circuitless IP interface (Figure 89)
area <ipaddr>	Designates an area for the circuitless IP interface • <ipaddr> is the IP address of the OSPF area that is associated with the circuitless IP interface (CLIP).
create <ipaddr/mask>	Creates a circuitless IP interface • <ipaddr/mask> is the IP address and Net Mask of the circuitless-IP interface.

config ip circuitless-ip-int

followed by:

<code>delete <ipaddr></code>	Deletes the specified circuitless IP interface <code><ipaddr></code> is the IP address of the Circuitless IP interface to be deleted.
<code>ospf <enable/disable></code>	Configures OSPF in passive mode for the circuitless IP interface. <code><enable/disable></code> enables or disables the option.

Configuration example

The following configuration example uses the above commands to configure circuitless IP, assign an interface number to the CLIP interface, and enable OSPF support:

```
Passport-8010:5 config ip circuitless-ip-int 1 create
11.126.205.1/255.0.0.0
Passport-8010:5 config ip circuitless-ip-int 1 area
134.177.1.0
Passport-8010:5 config ip circuitless-ip-int 1 ospf enable
```

To display information about the CLIP setup, use the following command:

```
Passport-8010:5 config ip circuitless-ip-int 1 info
```

Figure 89 shows sample output for this command.

Figure 89 config ip circuitless-ip-int info command output

```
Passport-8010:5 config ip circuitless-ip-int 1 info
Sub-Context:
Current Context:
    Clip 1 :
    area : 134.177.1.0
    create : 11.126.205.1/255.0.0.0
    delete : N/A
    ospf : enabled
```


Showing circuitless IP information

To display information about the current Circuitless IP configuration, use the following command:

```
show ip circuitless-ip-int info
```

Figure 90 shows sample output for this command.

Figure 90 show ip circuitless-ip-int info command output



Chapter 5

Configuring ARP using Device Manager

Network stations using the IP protocol need both a physical address and an IP address to transmit a packet. In situations where the station knows only the network host's IP address, the Address Resolution Protocol (ARP) enables the network station to determine a network host's physical address by binding a 32-bit IP address to a 48-bit MAC address. A network station can use ARP across a single network only, and the network hardware must support physical broadcasts. If a network station wants to send a packet to a host but knows only the host's IP address, the network station uses ARP to determine the host's physical address.

- For conceptual information about ARP management, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter includes the following topics:

Topic	Page
Enabling or disabling ARP on the routing interface	291
Enabling or disabling ARP on the router port	292
Viewing and managing ARP	293
Creating static ARP entries	294
Configuring Proxy ARP	296

Enabling or disabling ARP on the routing interface

After the IP address is assigned, ARP can be configured. By default, ARP Response is enabled and Proxy ARP is disabled.

To configure ARP on a routing interface:

- 1 From the Device Manager menu bar, choose Edit > Port.
The Port dialog box opens with the Interface tab displayed (Figure 59).
- 2 Click the ARP tab.
The Port dialog box—ARP tab opens (Figure 91).

Figure 91 Port dialog box—ARP tab




Table 17 describes the Port dialog box—ARP tab fields.

Table 17 Port dialog box—ARP tab fields

Field	Description
DoProxy	Sets the switch to respond to an ARP request from a locally attached host or end station for a remote destination. The default value is disable.
DoResp	Sets the switch to send ARP responses for this IP interface address. The default value is enable.

Enabling or disabling ARP on the brouter port

To enable or disable ARP on a port:

- 1 Select a port.

- 2 From the Device Manager menu bar, choose Edit > Port.

The Port dialog box opens with the Interface tab displayed ([Figure 59 on page 226](#)).

- 3 Click the ARP tab.

The Port dialog box—ARP tab opens ([Figure 91](#)).

- 4 In the DoProxy field, click enable to enable Proxy ARP function (see “[Configuring Proxy ARP](#)” on [page 296](#) for an explanation of the option).

The default is disabled.

- 5 In the DoResp field, click disable or enable to select whether or not to respond to an ARP. The default is enabled.

- 6 Click Apply.



Note: Use the ARP dialog box when setting the ARP response behavior on a brouter port. To configure the ARP response for a routing VLAN, use VLAN > VLANs > Basic > IP > ARP. The ARP dialog box is not applicable unless the port or VLAN is routed, that is, it is assigned an IP address.

Viewing and managing ARP

You can view and manage known MAC address to IP address associations. In addition, you can create or delete individual ARP entries.

To view and manage known MAC address to IP address associations, or to create or delete individual ARP entries:

- 1 From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed ([Figure 51 on page 207](#)).

- 2 Click the ARP tab.

The IP dialog box—ARP tab opens ([Figure 91](#)).

Figure 92 IP dialog box—ARP tab




Table 18 describes the IP dialog box—ARP tab fields.

Table 18 IP dialog box—ARP tab fields

Field	Description
Interface	The router interface for this ARP entry: • Brouter interfaces are identified by the slot/port number of the brouter port. • For virtual router interfaces, the brouter slot/port and the name of the VLAN followed by the (VLAN) designation are specified.
MacAddress	The media-dependent physical address (that is, the Ethernet address).
IpAddress	The IP address corresponding to the media-dependent physical address.
Type	Type of ARP entry: • local—a locally configured ARP entry • static—a statically configured ARP entry • dynamic—a learned ARP entry

Creating static ARP entries

To create static ARP entries:

- 1 From the Device Manager menu bar, choose IP Routing > IP.

The IP dialog box opens with the Globals tab displayed (Figure 51 on page 207).


- 2 Click the ARP tab.

The IP dialog box—ARP tab opens (see Figure 92 on page 294).

- 3 In the IP dialog box—ARP tab, click Insert.

The IP, Insert ARP dialog box opens (Figure 93).

Figure 93 IP, Insert ARP dialog box



- 4 In the Interface field, click the ellipses button to select the router interface from the ipNetToMediaIfIndex dialog box (Figure 93).

- 5 In the ipNetToMediaIfIndex dialog box, click OK.

This action specifies the interface connected to the station for which a static ARP entry is being defined.

- 6 In the Port in VLAN field, use the pull-down menu to associate the router port with a VLAN, from the VLAN dialog box (Figure 93).

- 7 Click OK.

This action specifies the VLAN interface connected to the station for which a static ARP entry is being defined.

- 8 In the IPAddress box, type the IP address.
- 9 In the MacAddress box, type the MAC address.
- 10 Click Insert.

Configuring Proxy ARP

Proxy ARP allows an Passport 8000 switch to respond to an ARP request from a locally attached host or end station for a remote destination. It does so by sending an ARP response back to the local host with its own MAC address of the router interface for the subnet on which the ARP request was received. The reply is generated only if the switch has an active route to the destination network.

To configure proxy ARP:

- 1 From the Device Manager menu bar, choose VLAN > VLANs.


The VLAN dialog box opens, with the Basic tab displayed ([Figure 48 on page 203](#)).

- 2 Choose a VLAN.
- 3 Click the IP button.

The IP, VLAN dialog box opens with the IP Address tab displayed ([Figure 49 on page 204](#)).

- 4 Select the ARP tab.

The ARP tab opens ([Figure 94](#)).

Figure 94 IP, VLAN dialog box—ARP tab

5 Click the DoProxy enable button.

6 Click Apply.

Proxy ARP is enabled for the VLAN.

ARP Threshold

ARP Threshold limits the number of unresolved ARP entries that can be stored on the switch. The default number of entries is 500 and it can vary between 50 and 1000, which is configured by the user.

1 From the Device Manager, choose IP Routing > IP

The Global tab opens ([Figure 95](#)).

Figure 95 IP, Global tab

The screenshot shows a window titled "134.177.229.236 - IP". The "Globals" tab is selected. The settings are as follows:

- Forwarding: ☒ forwarding ☐ not-forwarding
- DefaultTTL: 1..255
- ReasmTimeout: 30 sec
- ARPLifeTime: 1..32767 min
- ArpThreshold: 50..1000
- ☐ ICMPUnreachableMsgEnable
- ☐ ICMPRedirectMsgEnable
- ☐ AlternativeEnable
- ☐ RouteDiscoveryEnable
- ☐ EcmpEnable
- EcmpMaxPath: 1..4
- Ecmp1PathList: ...
- Ecmp2PathList: ...
- Ecmp3PathList: ...
- Ecmp4PathList: ...
- ☐ EcmpPathListApply

At the bottom, there are four buttons: Apply, Refresh, Close, and Help...

- 2 Enter an ARP Threshold entry number.
- 3 Click Apply.

Chapter 6

Configuring ARP using the CLI

Network stations that use IP protocol require both a physical address and an IP address to transmit packets. In situations where the station knows only the network host IP address, the Address Resolution Protocol (ARP) enables the network station to determine a network host physical address by binding a 32-bit IP address to a 48-bit MAC address.

A network station can use ARP across a single network only, and the network hardware must support physical broadcasts. If a network station wants to send a packet to a host but knows only the host IP address, the network station uses ARP to determine the host physical address.

- For conceptual information about ARP management, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter includes the following topics:

Topic	Page
Roadmap of IP commands	300
Configuring ARP on a port	301
Configuring ARP on a VLAN	303
Configuring IP ARP	306

Roadmap of IP commands

The following roadmap lists the IP ARP commands and their parameters. Use this list as a quick reference or click on any command or parameter entry for more information.

Command	Parameter
<code>config ethernet <ports> ip arp-response</code>	info disable enable
<code>config ethernet <ports> ip proxy</code>	info disable enable
<code>config vlan <vid> ip arp-response</code>	info disable enable
<code>config vlan <vid> ip proxy</code>	info disable enable
<code>config ip arp</code>	info add ports <value> ip <value> mac <value> [vlan <value>] aging <minutes> delete <ipaddr> multicast-mac-flooding <enable disable>
<code>show ports info arp [<ports>]</code>	
<code>show vlan info arp [<ports>]</code>	
<code>show ip arp info [<ip address>] [-s <value>]</code>	

Configuring ARP on a port

You can configure your switch to enable or disable ARP responses on a specified port. You can also enable proxy ARP on a port, which allows a router to answer a local ARP request for a remote destination.

This section includes the following topics:

- “Configuring an ARP proxy on a port,” next.
- “Showing ARP port information” on page 302.

The **config ethernet <ports> ip arp-response** command allows you to configure IP ARP on specific ports and includes the following options:

config ethernet <ports> ip arp-response followed by:	
info	Displays ARP response status on the port (Figure 96).
disable	Disables ARP responses on the port.
enable	Enables ARP responses on the port.

Figure 96 shows sample output for the **config ethernet <ports> ip arp-response info** command.

Figure 96 config ethernet <ports> ip arp-response info command output

```
Passport-8610# config ethernet 9/2 ip arp-response info
Sub-Context:
Current Context:

Port 9/2 :
          arp-response : enable
```

Configuring an ARP proxy on a port

The `config ethernet <ports> ip proxy` command includes the following options:

<code>config ethernet <ports> ip proxy</code> followed by:	
<code>info</code>	Displays ARP proxy status on the port.
<code>disable</code>	Disables proxy ARP on the port.
<code>enable</code>	Enables proxy ARP on the port, allowing a router to answer a local ARP request for a remote destination.

Figure 97 shows sample output for the `config ethernet <ports> ip arp-response info` command.

Figure 97 `config ethernet <ports> ip arp-response info` command output

```
Passport-8610# config ethernet 9/2 ip arp-response info

Sub-Context:
Current Context:

Port 9/2 :
           arp-response : enable
```

Showing ARP port information

To display ARP information about the specified port or for all ports, use the following command:

```
show ports info arp [<ports>]
```

[Figure 98](#) shows sample output for the `show ports info arp` command.

Figure 98 show ports info arp command (partial output)

```
Passport-8610# show ports info arp
```

```
=====
                                Port Arp
=====
PORT_NUM  DOPROXY    DORESP
-----
9/1       false     true
9/2       false     true
9/3       false     true
9/4       false     true
9/5       false     true
9/6       false     true
9/7       false     true
9/8       false     true
9/9       false     true
9/10      false     true
```

Configuring ARP on a VLAN

You can configure your switch to enable or disable ARP VLAN responses on a specified port. You can also enable proxy ARP on the VLAN, which allows a router to answer a local ARP request for a remote destination.

This section includes the following topics:

- [“Configuring an ARP proxy on a VLAN,”](#) next.
- [“Showing ARP VLAN information”](#) on page 305.

The **config vlan <vid> ip arp-response** command allows you to configure IP ARP on a VLAN and includes the following options:

config vlan <vid> ip arp-response followed by:	
info	Displays ARP response status on the VLAN (Figure 99).
disable	Disables ARP responses on the VLAN.
enable	Enables ARP responses on the VLAN.

Figure 99 shows sample output for the **config vlan <vid> ip arp-response info** command.

Figure 99 config vlan <vid> ip arp-response info command output

Passport-8606:6# config vlan 1 ip arp-response info	
Sub-Context: clear config dump monitor show test trace wsm	
Current Context:	
	resp : enable

Configuring an ARP proxy on a VLAN

The **config vlan <vid> ip proxy** command enables and disables proxy ARP on the VLAN. It also displays ARP proxy status on the VLAN.

The **config vlan <vid> ip proxy** command includes the following options:

config vlan <vid> ip proxy followed by:	
info	Displays ARP proxy status on the VLAN (Figure 100).

config vlan <vid> ip proxy followed by:	
disable	Disables proxy ARP on the VLAN.
enable	Enables proxy ARP on the VLAN, allowing a router to answer a local ARP request for a remote destination.

Figure 100 shows sample output for the **config vlan <aid> ip proxy info** command.

Figure 100 config vlan <aid> ip proxy info command output

```
Passport-8610# config vlan 9/2 ip proxy info

Sub-Context:
Current Context:

                proxy : enable
```

Showing ARP VLAN information

To display ARP information about the specified port or for all ports, use the following command:

```
show vlan info arp [<ports>]
```

Figure 101 shows sample output for the **show vlan info arp** command.

Figure 101 show vlan info arp command

```
Passport-8610# show vlan info arp

=====
Vlan Arp
=====
VLAN ID    DOPROXY    DORESP
-----
1           false      true
2           false      true
```

Configuring IP ARP

The ARP commands enable you to add and delete static entries in the ARP table and to display the ARP table. The ARP table maps MAC addresses to IP addresses. If you add an ARP entry for a VLAN, the VLAN is associated with the MAC address you specify. When you display the ARP table, all entries (static and dynamic) are displayed. Before you can add an ARP entry to a port or port-based VLAN, you must first assign an IP address to the port or VLAN and enable routing.

The only way to change a static ARP to another static ARP is to delete the old static ARP entry and create a new one with new information. When you create a static ARP entry using an IP address that belongs to another static ARP entry and then execute the `show config module ip` CLI command, the output displays your new entry.

This section includes the following topics:

- [“Configuring ARP static entries,”](#) next.
- [“ARP Threshold”](#) on page 313.

Configuring ARP static entries

The **config ip arp** command allows you to display ARP characteristics and modify the ARP parameters on the switch.

The **config ip arp** command include the following options:

config ip arp followed by:	
info	Displays ARP characteristics (Figure 102).
add ports <value> ip <value> mac <value> [vlan <value>]	Adds a static entry to the ARP table. - ports <value> are the port numbers, shown as slot/port. - ip <value> is the IP address {a.b.c.d}. - mac <value> is the 48-bit hardware MAC address in the format {0x00:0x00:0x00:0x00:0x00:0x00}. - vlan <value> is the name or number of a VLAN.
aging <minutes>	Sets the length of time in seconds an entry remains in the ARP table before timeout. <minutes> is a number between 1 and 32767.
delete <ipaddr>	Removes an entry from the ARP table. <ipaddr> is the IP address in dotted-decimal notation {a.b.c.d}.
multicast-mac-flooding <enable / disable>	Allows you to choose whether ARP entries for multicast MAC addresses are associated with the VLAN or the port interface on which it was learned. This is useful if multiple end stations or servers are sharing a multicast MAC address as is the case with certain Microsoft network load balancing applications, wherein the traffic is flooded to the VLAN to ensure that every end station using this virtual multicast MAC address is receiving a copy of the stream. Default is disable. This option is not dynamic, in that if the setting of this feature is changed it will not dynamically reprogram all previously learned ARP entries from multicast MAC addresses.

Figure 102 shows sample output for the `config ip arp info` command.

Figure 102 `config ip arp info` command (partial output)

```
Passport-8610# config ip arp info

Sub-Context: clear config dump monitor show test trace wsm
Current Context:

    multicast-mac-flooding : disable
        aging : 360 (min)
        delete : N/A
        add :

                                ports - N/A
                                ip - 200.1.1.1
                                mac - 00:80:2d:c1:ce:05
                                vlan - 2

                                ports - N/A
                                ip - 200.1.1.15
                                mac - 00:00:5e:00:01:01
                                vlan - 2

                                ports - N/A
                                ip - 200.1.1.255
                                mac - ff:ff:ff:ff:ff:ff
                                vlan - 2
```

Configuration Example

The following configuration example uses the above command to:

- Add static entry to an ARP table
- Sets the length of time in seconds an entry remains in the ARP table before timeout
- Removes an entry from the ARP table
- Allows you to choose whether ARP entries for multicast MAC addresses are associated with the VLAN or the port interface on which it was learned.

After configuring the parameters, use the `info` command to show a summary of the results.

Passport-8010:6/config/ip/arp# ?

Sub-Context: static-mcastmac

Current Context:

```
add ports <value> ip <value> mac <value> [vlan <value>]
aging <minutes>
arpreqthreshold <integer>
delete <ipaddr>
info
multicast-mac-flooding <enable|disable>
```

Passport-8010:6/config/ip/arp# add ports 1/8 ip 58.1.58.51 mac
00:80:2d:39:02:01 vlan 58

Passport-8010:6/config/ip/arp# info

Sub-Context: static-mcastmac

Current Context:

```
multicast-mac-flooding : disable
      aging : 360 (min)
      arpreqthreshold : 500
      delete : N/A
      add :
            ports - 1/8
            ip - 58.1.58.51
            mac - 00:80:2d:39:02:01
            vlan - 58
```

Passport-8010:6/config/ip/arp# aging 60

Passport-8010:6/config/ip/arp# info

Sub-Context: static-mcastmac

Current Context:

```
multicast-mac-flooding : disable
                        aging : 60 (min)
                        arpreqthreshold : 500
                        delete : N/A
```

```
Passport-8010:6/config/ip/arp# delete 58.1.58.51
```

```
Passport-8010:6/config/ip/arp# info
```

Sub-Context: static-mcastmac

Current Context:

```
multicast-mac-flooding : disable
                        aging : 60 (min)
                        arpreqthreshold : 500
                        delete : N/A
                        add :
```

```
Passport-8010:6/config/ip/arp# multicast-mac-flooding en
```

```
Passport-8010:6/config/ip/arp# info
```

Sub-Context: static-mcastmac

Current Context:

```
multicast-mac-flooding : enable
                        aging : 60 (min)
                        arpreqthreshold : 500
                        delete : N/A
```

Configuring ARP Threshold

The **config ip arp arpreqthreshold** command allows you to set a limit on the number of unresolved ARP entries that can be stored on the switch and reach a threshold.

config ip arp arpreqthreshold followed by:	
info	Displays ARP characteristics (Figure 102).
add ports <value> ip <value> mac <value> [vlan <value>]	Adds a static entry to the ARP table. - ports <value> are the port numbers, shown as slot/port. - ip <value> is the IP address {a.b.c.d}. - mac <value> is the 48-bit hardware MAC address in the format {0x00:0x00:0x00:0x00:0x00:0x00}. - vlan <value> is the name or number of a VLAN.
aging <minutes>	Sets the length of time in seconds an entry remains in the ARP table before timeout. <minutes> is a number between 1 and 32767.
delete <ipaddr>	Removes an entry from the ARP table. <ipaddr> is the IP address in dotted-decimal notation {a.b.c.d}.

config ip arp arpreqthreshold

followed by:

<code>multicast-mac-flooding</code> <code><enable disable></code>	Allows you to choose whether ARP entries for multicast MAC addresses are associated with the VLAN or the port interface on which it was learned. This is useful if multiple end stations or servers are sharing a multicast MAC address as is the case with certain Microsoft network load balancing applications, wherein the traffic is flooded to the VLAN to ensure that every end station using this virtual multicast MAC address is receiving a copy of the stream. Default is disable. This option is not dynamic, in that if the setting of this feature is changed it will not dynamically reprogram all previously learned ARP entries from multicast MAC addresses.
<code>arpreqthreshold</code> <code><integer></code>	Describes the threshold parameters. • <code><integer></code> is the maximum number of unresolved ARP entries that can be stored on a switch.

Figure 103 ARP Threshold

```

Passport-8606:5/config/ip/arp# arpreqthreshold
Not enough required parameters entered
Max number of Outstanding Unresolved ARP REQ
Required parameters:
<integer>          = Max number of Unresolved ARP Entry Req
{50..1000}
Command syntax:
arpreqthreshold <integer>
Passport-8606:5/config/ip/arp#

Passport-8606:5/config/ip/arp# arpreqthreshold 501
Passport-8606:5/config/ip/arp# info

Sub-Context: static-mcastmac
Current Context:

    multicast-mac-flooding :  disable
                          aging : 360 (min)
                          arpreqthreshold : 501
                          delete : N/A
                          add :

                                ports - 2/5
                                ip - 10.1.1.1

```

Showing ARP information

To display the ARP table, use the following command:

```
show ip arp info [<ip address>] [-s <value>]
```

where:

<ip address> is the specific net IP address for the table .

-s <value> is the specific subnet in the format
(a.b.c.d/x|a.b.c.d/x.x.x.x|default).

Figure 104 shows sample output for the `show ip arp info` command with no IP address or subnet specified. In the TTL column, the output is measured in seconds.

Figure 104 show ip arp info command output

Passport-8610# show ip arp info

Ip Arp					
IP_ADDRESS	MAC_ADDRESS	VLAN	PORT	TYPE	TTL
161.69.150.10	00:e0:16:ff:01:3a	2	1/1	DYNAMIC	2133
161.69.150.1	00:80:2d:23:02:00	2	-	LOCAL	2160
161.69.150.255	ff:ff:ff:ff:ff:ff	2	1/1	LOCAL	2160
161.69.100.255	ff:ff:ff:ff:ff:ff	2	1/1	LOCAL	2160
Total 4					

Chapter 7

Configuring RIP using Device Manager

In a routed environment, routers communicate with one another to keep track of available routes. Routers can learn about available routes dynamically using the Routing Information Protocol (RIP). The Passport 8000 Series Switch software implements standard RIP for exchanging TCP/IP route information with other routers.

- For conceptual information about RIP management, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter describes how you use Device Manager to configure and manage the RIP on a brouter port or routed VLAN in an Passport 8000 switch.

This chapter includes the following topics:

Topic	Page
Configuration prerequisites	316
Enabling RIP globally	316
Enabling and configuring RIP on a brouter port	318
Enabling and configuring RIP on a VLAN	321
Viewing RIP protocol statistics	323
Configuring RIP interface parameters	324
Configuring Advanced featured on a RIP interface	326

Configuration prerequisites

To configure RIP on an interface, you must already have performed the following steps:

- 1 Created a router interface (either a brouter port or a virtual routing interface)
- 2 Assigned an IP address to the interface.

If an IP address has not been assigned, refer to [“Assigning an IP address on a brouter port” on page 201](#), for information about assigning IP addresses.



Note: A RIP configuration will not take affect unless RIP is configured globally and on the interface (see [Enabling RIP globally](#)).

- 3 Enabled RIP globally on the interface.


Enabling RIP globally

In the Passport 8000 switch, the RIP global parameters are used by all router interfaces using RIP. Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters.

To enable RIP globally:

- 1 From the Device Manager main menu, choose IP Routing > RIP.

The RIP dialog box opens with the Globals tab displayed ([Figure 105](#)).

Figure 105 RIP dialog box—Globals tab

- 2 In the Operation check box, select enable box or click to clear the RIP Globals option.



Note: You can configure RIP on the interfaces with RIP globally disabled, thus having the flexibility to configure all interfaces before turning on RIP for the switch.

Table 19 describes the Globals tab fields.

Table 19 Globals tab fields

Field	Description
Operation	Enables or disables the operation of RIP on all interfaces. The default is disabled.
UpdateTime	The time interval between RIP updates on all interfaces. It is a global parameter for the box; that is, it applies to all interfaces and cannot be set individually for each interface. The default is 30 seconds.
Timeoutinterval	The time out interval between RIP update and all interfaces.

Table 19 Globals tab fields (continued)

Field	Description
RouteChanges	The number of route changes made to the IP Route Database by RIP; does not include the refresh of a route's age.
Queries	The number of responses sent to RIP queries from other systems.
HoldDownTime	Sets the length of time that RIP will continue to advertise a network after determining it is unreachable. The range is 0 to 360 seconds. The default is 120 seconds.
DefImportMetric	Sets the value of the default import metric to import a route into a RIP domain. For announcing OSPF internal routes into a RIP domain, if the policy does not specify a metric value, the default import metric should be used. For OSPF external routes, the external cost is used.

Enabling and configuring RIP on a router port

To access the RIP enable and configuration parameters for a router port:

- 1 From the Device Manager menu bar, choose Edit > Port.

The Port dialog box opens with the Interface tab displayed ([Figure 59 on page 226](#)).

- 2 Click the RIP tab.

The RIP tab opens ([Figure 106](#)).

Figure 106 Port dialog box—RIP tab

192.168.151.163 - Port 1/13

VLACP | Remote Mirroring | Mroute Stream Limit | Fdb Protect | IP Address | ARP | DHCP
Interface | VLAN | STG | MAC Learning | Rate Limiting | Test | SMLT | PCAP | EAPOL | LACP
DVMP | IGMP | OSPF | **RIP** | PIM | PGM | VRRP | Router Discovery | IPX BRouter

☐ Enable

Supply: ☐ disable ☒ enable

Listen: ☐ disable ☒ enable

Poison: ☒ disable ☐ enable

☐ DefaultSupply

☐ DefaultListen

☐ TriggeredUpdateEnable

☐ AutoAggregateEnable

☐ AdvertiseWhenDown

InPolicy: ...

OutPolicy: ...

Cost: 1..15

HolddownTime: 0..360

TimeoutInterval: 15..259200

Apply Refresh Close Help...

Table 20 describes the Port dialog box—RIP tab fields.

Table 20 Port dialog box—RIP tab fields

Field	Description
Enable	If selected, enables RIP on the port.
Supply	Specifies that the routing switch will advertise RIP routes through the interface. The default is enable.

Table 20 Port dialog box—RIP tab fields (continued)

Field	Description
Listen	Specifies that the routing switch will learn RIP routes through this interface. The default is enable.
Poison	If disabled, split horizon is invoked, meaning that IP routes learned from an immediate neighbor are not advertised back to the neighbor from which the routes were learned. If enabled, the RIP update sent to a neighbor from which a route is learned is “poisoned” with a metric of 16. In this manner, the route entry is not passed along to the neighbor, because historically 16 is “infinity” in terms of hops on a network. The default is disable.
DefaultSupply	Set the value to true if a default route must be advertised out this interface. The default is false. Note: The default route will be advertised only if it exists in the routing table.
DefaultListen	Set value to true if default route should be learned on this interface when advertised by another router connected to the interface. The default is false.
TriggeredUpdateEnable	Allows you to enable or disable triggered RIP updates. The default is false (disabled).
AutoAggregateEnable	Allows you to enable or disable RIP automatic aggregation. RIP2 automatically aggregates routes to their natural mask. Auto aggregation can be enabled only in RIP2 mode or RIP1 compatibility mode. The default is false.
AdvertiseWhenDown	If true, the network on this interface will be advertised as up, even if the port is down. The default is false. Note: When you configure a port without any link and enable AdvertiseWhenDown, it will not advertise the route until the port is active. Then the route will be advertised even when the link is down. To disable advertising based on link states, disable AdvertiseWhenDown.
InPolicy	Right click in the InPolicy name field and select the policy name to be applied from the PolicyName dialog box (Figure 106). This policy will determine whether to learn a route on this interface. It also specifies the parameters of the route when it is added to the routing table.
OutPolicy	Right click in the OutPolicy name field and select the policy name to be applied from the PolicyName dialog box (Figure 106). This policy will determine whether to advertise a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. Enter a value between 1 to 15. The default is 1.

Table 20 Port dialog box—RIP tab fields (continued)

Field	Description
HolddownTime	Indicates the rip holddown timer for this interface. Enter a value between 0 to 360.
TimeoutInterval	Indicates the rip timeout interval for this interface. Enter a value between 15 to 259200.

Enabling and configuring RIP on a VLAN

In the Passport 8000 switch, the RIP global parameters are used by all router interfaces using RIP. Both brouter ports and VLAN virtual routing interfaces use the same RIP global parameters. Before you configure RIP on a VLAN you must first set the RIP global parameters.


To access the RIP enable and configuration parameters for a virtual router:

- 1 From the Device Manager menu bar, choose VLAN > VLANs.

The VLAN dialog box opens with the Basic tab displayed ([Figure 49 on page 204](#)).

- 2 Select a VLAN.
- 3 Click IP.

The IP, VLAN dialog box opens with the IP Address tab displayed ([Figure 107](#)).

Figure 107 IP, VLAN dialog box—IP Address tab

- 4 Click the IP, VLAN dialog box—RIP tab.

The IP, VLAN dialog box—RIP tab opens ([Figure 108](#)).

Figure 108 IP, VLAN dialog box—RIP tab

192.168.151.163 - IP, VLAN 1

IP Address	ARP	DHCP	DVMRP	IGMP	OSPF
RIP	PIM	PGM	VRRP	Router Discovery	Direct Broadcast
					RSMLT

☐ Enable

Supply: ☐ disable ☒ enable

Listen: ☐ disable ☒ enable

Poison: ☒ disable ☐ enable

☐ DefaultSupply

☐ DefaultListen

☐ TriggeredUpdateEnable

☐ AutoAggregateEnable

☐ AdvertiseWhenDown

InPolicy: ...

OutPolicy: ...

Cost: 1 1..15

HolddownTime: 120 0..360

TimeoutInterval: 180 15..259200

Apply Refresh Close Help...



Note: The screen captures in this section are for a virtual router interface for a VLAN. The screens for configuring a brouter port have the same parameters, and the parameters function the same.

Viewing RIP protocol statistics

To view the RIP protocol statistics:

- 1 From the Device Manager main menu, choose IP Routing > RIP.

The RIP dialog box opens with the Globals tab displayed ([Figure 105 on page 317](#)).

- 2 Click the RIP dialog box—Status tab.

The RIP dialog box—Status tab opens ([Figure 109](#)).

Figure 109 RIP dialog box—Status tab

Address	RcvBadPackets	RcvBadRoutes	SentUpdates
128.125.201.2	0	0	0
128.125.202.1	0	0	0
128.125.203.1	0	0	0
128.125.204.1	0	0	0

[Table 21](#) describes the RIP dialog box—Status tab fields.

Table 21 RIP dialog box—Status tab fields

Field	Description
Address	The IP address of the router interface.
RcvBadPackets	The number of RIP response packets received by the RIP process that were subsequently discarded for any reason (Examples: a version 0 packet or an unknown command type).

Table 21 RIP dialog box—Status tab fields (continued)

Field	Description
RcvBadRoutes	The number of routes, in valid RIP packets, that were ignored for any reason (Examples: unknown address family or invalid metric).
SentUpdates	The number of triggered RIP updates actually sent on this interface. This field explicitly does <i>not</i> include full updates sent containing new information.

Configuring RIP interface parameters

You can specify the RIP version to use on interfaces configured to send (supply) or receive (listen to) RIP updates.

 **Note:** The AuthType and AuthKey parameters are not supported.

To configure the RIP version:

- 1 From the Device Manager main menu, choose IP Routing > RIP.
The RIP dialog box opens with the Globals tab displayed ([Figure 105](#)).
- 2 Click the RIP dialog box—Interface tab.
The RIP dialog box—Interface tab opens.

The RIP dialog box—Interface tab opens ([Figure 110](#)).


Figure 110 RIP dialog box—Interface tab

Table 22 describes the RIP dialog box—Interface tab fields.

Table 22 RIP dialog box—Interface tab fields

Field	Description
Address	The IP address of the router interface.
Domain	The value inserted into the Routing Domain field of all RIP packets sent on this interface.
AuthType	The type of authentication used on this interface.
AuthKey	The value to be used as the Authentication Key whenever the corresponding instance of rip2IfConfAuthType has a value other than noAuthentication.
Send	What the router sends on this interface (selected from a pull-down menu): - DoNotSend—no RIP updates sent on this interface - ripVersion1—RIP updates compliant with RFC 1058 - rip1Compatible—broadcast RIP-2 updates using RFC 1058 route subsumption rules - ripVersion2—multicasting RIP-2 updates The default is rip1compatible.
Receive	Indicates which versions of RIP updates are to be accepted: - rip1 - rip2 - rip1OrRip2 The default is rip1OrRip2. Note that rip2 and rip1OrRip2 imply reception of multicast packets.

- 3 In the Send field, use the pull-down menu to select which RIP version the router sends.
- 4 In the Receive field, use the pull-down menu to select which RIP version the router listens for.
- 5 Click Apply.

Configuring Advanced featured on a RIP interface


You can edit the RIP version parameters using the Interface Advanced tab

To edit RIP version parameters:

- 1 From the Device Manager main menu, choose IP Routing > RIP.
The RIP dialog box opens with the Globals tab displayed.(Figure 105)
- 2 Click the RIP dialog box—Interface Advance tab.
The RIP dialog box—Interface Advance tab opens.

The RIP dialog box—Interface Advance tab opens (Figure 111).

Figure 111 RIP dialog box—Interface Advance tab



[Table 23](#) describes the RIP dialog box—Interface Advance tab fields.

Table 23 RIP dialog box—Interface Advance tab fields

Field	Description
Address	Display the address of the entry in the IP RIP Interface Table.
IfIndex	The index value of the RIP interface.
Enable	Displays if the RIP interface is enabled or disabled.
Supply	Enables (true) or disables (false) the switch to send out RIP updates on this interface.
Listen	Configures whether (true) or not (false) the switch will learn routes on this interface.
Poison	Sets whether (true) or not (false) RIP routes on the interface learned from a neighbor are advertised back to the neighbor. If disabled, split horizon is invoked and IP routes learned from an immediate neighbor are not advertised back to the neighbor. If enabled, the RIP updates sent to a neighbor from which a route is learned are "poisoned" with a metric of 16. Therefore, the receiver neighbor will ignore this route because the metric 16 indicates infinite hops in the network.
DefaultSupply	Enables (true) or disables (false) an advertisement of a default route on this interface. This command takes effect only if a default route exists in the routing table.
DefaultListen	Enables (true) or disables (false) the switch to accept the default route learned through RIP on this interface. The default is disabled.
TriggeredUpdate	Enables (true) or disables (false) the switch to send out RIP updates on this interface.
AutoAggregate	Enables (true) or disables (false) automatic route aggregation on this interface. When enabled, the switch automatically aggregates routes to their natural mask when they are advertised on an interface*. The default is disabled. *In previous software releases, this configuration changed the mask for all routes. Now, this configuration aggregates only the routes with a mask length longer than natural mask.
InPolicy	Right click in the InPolicy name field and select the policy name to be applied from the PolicyName dialog box (Figure 110 and Figure 111). This policy will determine whether to learn a route on this interface. It also specifies the parameters of the route when it is added to the routing table.

Table 23 RIP dialog box—Interface Advance tab fields (continued)

Field	Description
OutPolicy	Right click in the OutPolicy name field and select the policy name to be applied from the PolicyName dialog box (Figure 110 and Figure 111). This policy will determine whether to advertise a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.
Cost	Indicates the RIP cost for this interface. Enter a value between 1 to 15. The default is 1.
HoldDownTime	Indicate the rip holddown timer for this interface. Enter a value between 0 to 360.
TimeoutInterval	Indicates the rip timeout intervals for this interface. Enter a value between 15 to 259200.

Chapter 8

Configuring RIP using the CLI

This chapter describes the Run-Time CLI commands that are used to configure the Routing Information Protocol (RIP) in the Passport 8000 Series Switch. You can configure RIP on a port or on a VLAN, but you must first globally enable RIP.

- For conceptual information about RIP management, see [Chapter 1, “IP routing concepts,”](#) on page 31.
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,”](#) on page 93.

This chapter includes the following topics:

Topic	Page
Roadmap of IP commands	330
Configuring RIP global parameters	332
Configuring RIP parameters on an interface	336
Showing RIP global configuration information	339
Showing information on a RIP interface	340
Configuring RIP on a port	341
Showing RIP information on a port	344
Setting RIP parameters for a VLAN	345
Showing RIP information for VLANs	349

Roadmap of IP commands

The following roadmap lists some of the IP commands and their parameters. Use this list as a quick reference or click on any command or parameter entry for more information.

Command	Parameter
<code>config ip rip</code>	<code>info</code> <code>disable</code> <code>enable</code> <code>default-import-metric <metric></code> <code>holddown <seconds></code> <code>updatetime <seconds></code>
<code>config ip rip interface <ipaddr></code>	<code>info</code> <code>auto-aggr <enable disable></code> <code>cost <cost></code> <code>default-listen <enable disable></code> <code>default-supply <enable disable></code> <code>disable</code> <code>enable</code> <code>in-policy <policy name></code> <code>listen <enable disable></code> <code>out-policy <policy name></code> <code>poison <enable disable></code> <code>receive-mode <mode></code> <code>send-mode <mode></code> <code>supply <enable disable></code> <code>trigger <enable disable></code>
<code>config ethernet <ports> ip rip</code>	<code>info</code>

Command	Parameter
	advertise-when-down <enable disable>
	auto-aggr <enable disable>
	default-listen <enable disable>
	cost <cost>
	default-supply <enable disable>
	disable
	enable
	listen <enable disable>
	in-policy <policy name>
	manualtrigger
	out-policy <policy name>
	poison <enable disable>
	supply <enable disable>
	trigger <enable disable>
config vlan <vid> ip rip	info
	advertise-when-down <enable disable>
	auto-aggr <enable disable>
	cost <cost>
	default-listen <enable disable>
	default-supply <enable disable>
	disable
	enable
	in-policy <policy name>
	listen <enable disable>
	manualtrigger
	out-policy <policy name>
	poison <enable disable>
	supply <enable disable>

Command	Parameter
	<code>trigger <enable disable></code>
<code>show ip rip info</code>	
<code>show ports info rip [<ports>]</code>	
<code>show vlan info rip [<vid>]</code>	

Configuring RIP global parameters

To enable or disable RIP globally on the switch and configure holddown and update timers, use the following command:

```
config ip rip
```

This command includes the following parameter options:

config ip rip followed by:	
info	Displays the global RIP configuration settings (Figure 112).
disable	Globally disables RIP on the switch.
enable	Globally enables RIP on the switch.
default-import-metric <metric>	Sets the value of default import metric to import a route into RIP domain. For announcing OSPF internal routes into RIP domain, if the policy does not specify a metric value, this value is used. For OSPF external routes, the external cost is used <metric> sets the RIP Default Import Metric {0..15} with a default value of 8.
holddown <seconds>	Sets the RIP holddown timer value, the length of time (in seconds) that RIP will continue to advertise a network after determining that it is unreachable. <seconds> is an integer between 0 and 360, with a default of 120.
updatetime <seconds>	Sets RIP update timer, the time interval between RIP updates. <seconds> is an integer between 0 and 360, with a default of 30 seconds.

Figure 112 shows sample output for this command.

Figure 112 config ip rip info command output

```
Passport-8606:6# config ip rip info

Sub-Context: clear config dump monitor show test trace wsm
Current Context:

        default-import-metric : 8
                enable : false
                holddown : 120
                updatetime : 30
```

Configuration Example

The following configuration example uses the above command to:

- Globally disables RIP on the switch
- Globally enables RIP on the switch
- Sets the value of default import metric to import a route into RIP domain
- Sets the RIP holddown timer value
- Sets RIP update timer, the time interval between RIP updates

After configuring the parameters, use the info command to show a summary of the results.

```
Passport-8010:6/config/ip/rip# ?
```

```
Sub-Context: interface
```

```
Current Context:
```

```
default-import-metric <metric>
disable
enable
holddown <seconds>
info
```

```
timeout <seconds>
updatetime <seconds>
```

```
Passport-8010:6/config/ip/rip# enable
Passport-8010:6/config/ip/rip# info
```

```
Sub-Context: interface
Current Context:
```

```
default-import-metric : 8
        enable : true
        holddown : 120
        timeout : 180
        updatetime : 30
```

```
Passport-8010:6/config/ip/rip# default-import-metric 10
Passport-8010:6/config/ip/rip# info
```

```
Sub-Context: interface
Current Context:
```

```
default-import-metric : 10
        enable : true
        holddown : 120
        timeout : 180
        updatetime : 30
```

```
Passport-8010:6/config/ip/rip# holddown 150
```

```
Holddown timer value will take effect on all the rip interfaces
Passport-8010:6/config/ip/rip# info
```

```
Sub-Context: interface
Current Context:
```

```
default-import-metric : 10
    enable : true
    holddown : 150
    timeout : 180
    updatetime : 30
```

```
Passport-8010:6/config/ip/rip# updatetime 50
```

```
Passport-8010:6/config/ip/rip# info
```

```
Sub-Context: interface
```

```
Current Context:
```

```
default-import-metric : 10
    enable : true
    holddown : 150
    timeout : 180
    updatetime : 50
```

```
Passport-8010:6/config/ip/rip# disable
```

```
Passport-8010:6/config/ip/rip# info
```

```
Sub-Context: interface
```

```
Current Context:
```

```
default-import-metric : 10
    enable : false
    holddown : 150
    timeout : 180
    updatetime : 50
```

```
Passport-8010:6/config/ip/rip#
```

Configuring RIP parameters on an interface

To configure RIP parameters on a specified interface, use the

```
config ip rip interface <ipaddr>
```

This command uses the interface IP address to specify the interface for which you are entering the command.

There are no CLI `delete` or `remove` commands to remove an IP In or Out policy after it has been added. To remove the In or Out policy using the CLI, enter one of the following commands:

```
config ip rip interface <ipaddr> in-policy <policy name>
```

```
config ip rip interface <ipaddr> out-policy <policy name>
```

The interface-based RIP command includes the following options:

config ip rip interface <ipaddr> followed by:	
<code>info</code>	Displays RIP configurations on this interface.
<code>auto-aggr <enable disable></code>	Enables or disables automatic route aggregation on this interface. When enabled, the switch automatically aggregates routes to their natural mask when they are advertised on an interface*. The default is disabled. This configuration does not change the mask for all routes, instead it aggregates only the routes with a mask length longer than natural mask.
<code>cost <cost></code>	Indicates the RIP cost for this interface. • <code><cost></code> is 1 to 15. The default is 1.
<code>default-listen <enable disable></code>	Enables or disables the switch to accept the default route learned through RIP on this interface. The default is disabled.
<code>default-supply <enable disable></code>	Enables or disables an advertisement of a default route on this interface. This command takes effect only if a default route exists in the routing table.

config ip rip interface <ipaddr> followed by:	
disable	Disables RIP on the interface. The default is disabled.
enable	Enables RIP on the interface.
in-policy <policy name>	The policy name for inbound filtering on this RIP interface. This policy will determine whether to learn a route on this interface. It also specifies the parameters of the route when it is added to the routing table.
listen <enable/disable>	Configures whether or not the switch will learn routes on this interface.
out-policy <policy name>	The policy name for outbound filtering on this RIP interface. This policy will determine whether to advertise a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.
poison <enable disable>	Sets whether (enabled) or not (disabled) RIP routes on the interface learned from a neighbor are advertised back to the neighbor. If disabled, split horizon is invoked and IP routes learned from an immediate neighbor are not advertised back to the neighbor. If enabled, the RIP updates sent to a neighbor from which a route is learned are "poisoned" with a metric of 16. Therefore, the receiver neighbor will ignore this route because the metric 16 indicates infinite hops in the network.
receive-mode <mode>	Indicates which version of RIP updates are to be accepted on this interface. • <mode> is rip1, rip2, or rip1orrip2.
send-mode <mode>	Indicates which version of RIP updates the router sends on this interface. ripVersion1 implies sending RIP updates compliant with RFC 1058. rip1Compatible implies broadcasting RIP-2 updates using RFC 1058 route sub assumption rules. The default is rip1Compatible. • <mode> is notsend, rip1, rip2, or rip1comp.

config ip rip interface <ipaddr> followed by:	
supply <enable disable>	Enables or disables the switch to send out RIP updates on this interface.
trigger <enable disable>	Enables or disables automatic triggered updates for RIP on this interface.

Figure 113 shows sample output for this command.

Figure 113 config ip rip interface command

```
Passport-8606:6# config ip rip interface ?

Sub-Context:
Current Context:

    auto-aggr <enable|disable>
    cost <cost>
    default-listen <enable|disable>
    default-supply <enable|disable>
    disable
    domain <value>
    enable
    info
    in-policy <policy name>
    listen <enable|disable>
    out-policy <policy name>
    poison <enable|disable>
    receive-mode <mode>
    send-mode <mode>
    supply <enable|disable>
    trigger <enable|disable>
```

Showing RIP global configuration information

To display information about the RIP global configuration information on the switch, use the following command:

```
show ip rip info
```

[Figure 114](#) shows sample output for this command.

Figure 114 show ip rip *info* command output

```
Passport-8610:5# show ip rip info
```

```
=====
                        Rip Global
=====
Default Import Metric : 8
      Domain : 0
HoldDown Time : 100
      Queries : 0
      Rip : Enabled
Route Changes : 0
      Update Time : 60
```

Showing information on a RIP interface

To display information about a specific RIP interface or all RIP interfaces on the switch, use the following command:


```
show ip rip interface [<ipaddr>]
```

Where:

ipaddr is the interface IP address.

Figure 115 shows sample output for this command.

Figure 115 show ip rip interface command output



Configuring RIP on a port

To display RIP on a specified port, use the following command:

```
config ethernet <ports> ip rip
```

Where:

ports is the slot/port number you want to configure.

You must enable RIP globally on the switch for this command to take effect.

This command includes the following options:

config ethernet <ports> ip rip followed by:	
info	Displays RIP characteristics on the port (Figure 116).
advertise-when-down <enable disable>	If enabled, the network on this interface is advertised as up, even if the port is down. The default is disabled. Note: When you configure a port without any link and enable advertise-when-down, it will not advertise your route until the port is active. Then the route will be advertised even when the link is down. To disable advertising based on link status, this parameter should be disabled.
auto-aggr <enable disable>	Enables or disables automatic route aggregation on the port. When enabled, the router switch automatically aggregates routes to their natural mask when they are advertised on an interface in a different class network. The default is disable.
default-listen <enable disable>	Enables or disables RIP listen to accept the default route via RIP.
cost <cost>	Sets the RIP cost at this port.
default-supply <enable disable>	Enables or disables an advertisement of a default route only if one exists in the routing table.

config ethernet <ports> ip rip followed by:	
disable	Disables RIP on the port. This setting is the default.
enable	Enables RIP on the port.
listen <enable disable>	Configures whether or not the switch will listen for a default route without listening for all routes.
in-policy <policy name>	Sets the port RIP in-policy. • <i>policy name</i> is a string length {0..64}.
manualtrigger	Allows you to manually issue a RIP update.
out-policy <policy name>	Sets the port RIP out-policy. • <i>policy name</i> is a string length {0..64}.
poison <enable disable>	Indicates whether or not RIP routes on the port learned from a neighbor are advertised back to the neighbor. If disabled, split horizon is invoked and IP routes learned from an immediate neighbor are not advertised back to the neighbor. If enabled, the RIP updates sent to a neighbor from which a route is learned are "poisoned" with a metric of 16. Therefore, the receiver neighbor will ignore this route because the metric 16 indicates infinite hops in the network.
supply <enable disable>	Enables or disables the switch to supply RIP routes with including the default routes.
trigger <enable disable>	Enables or disables automatic triggered updates for RIP.

Figure 116 shows sample output for this command.

Figure 116 config ethernet ip rip info command output

```
Passport-8610# config ethernet 1/2 ip rip info

Sub-Context: clear config dump monitor show test trace
Current Context:

Port 1/2 :
        rip : disable
    advertise-when-down : disable
    auto-aggregation : disable
        cost : 1
    default-listen : disable
    default-supply : disable
        in-policy : N/A
        out-policy : N/A
    triggered-update : disable
        listen : enable
    manualtrigger : N/A
        poison : disable
        supply : enable
```

Table 24 indicates the relationship between switch action and the RIP supply and listen settings.

Table 24 RIP supply and listen settings and switch action

RIP supply settings		RIP listen settings		Switch action
Supply	Default supply	Listen	Default listen	
Disabled	Disabled			Sends no RIP updates.
Enabled	Disabled			Sends RIP updates except the default.
Disabled	Enabled			Sends only the default (default route must exist in routing table).
Enabled	Enabled			Sends RIP updates including the default route (if it exists).

Table 24 RIP supply and listen settings and switch action (continued)

RIP supply settings		RIP listen settings		Switch action
Supply	Default supply	Listen	Default listen	
		Disabled	Disabled	Does not listen for RIP updates.
		Enabled	Disabled	Listens for all RIP updates except the default.
		Disabled	Enabled	Listens only for the default.
		Enabled	Enabled	Listens for RIP updates including the default route (if it exists).

Showing RIP information on a port

To display information about RIP parameters on a specified port, use the following command:

```
show ports info rip [<ports>]
```

Where:

ports is the slot/port number you want to configure.

[Figure 117](#) show sample output for this command.

Figure 117 show ports info rip command (partial output)

Port Rip								
PORT	DEFAULT	DEFAULT	TRIGGERED	AUTOAGG				
NUM	ENABLE	SUPPLY	LISTEN	UPDATE	ENABLE	SUPPLY	LISTEN	POISON
1/1	false	false	false	false	false	true	true	false
1/2	false	false	false	false	false	true	true	false
1/3	false	false	false	false	false	true	true	false
1/4	false	false	false	false	false	true	true	false
1/5	false	false	false	false	false	true	true	false
1/6	false	false	false	false	false	true	true	false
1/7	false	false	false	false	false	true	true	false
1/8	false	false	false	false	false	true	true	false

Setting RIP parameters for a VLAN

To set RIP parameters for a VLAN, use the following command:

```
config vlan <vid> ip rip
```

where:

vid is a unique integer value in the range 1 and 4094 that represents the VLAN ID.

This command includes the following options:

config vlan <vid> ip rip followed by:	
info	Displays RIP characteristics on the VLAN (Figure 118).
advertise-when-down <enable disable>	If enabled, the network on this interface will be advertised as up, even if the port is down. The default is disabled. Note: When you configure a port without any link and enable advertise-when-down, it will not advertise your route until the port is active. Then the route is advertised even when the link is down. To disable advertising based on link status, this parameter should be disabled.
auto-aggr <enable disable>	Enables or disables automatic route aggregation on the VLAN. When enabled, the router switch automatically aggregates routes to their natural mask when they are advertised on an interface in a different class network. The default is disable.
cost <cost>	Sets the vlan RIP interface cost.
default-listen <enable disable>	Allows the user to enable or disable setting RIP listen to accept the default route via RIP.
default-supply <enable disable>	Allows the user to send a default route only if one exists in the routing table.
disable	Disables RIP on the VLAN. This is the default setting.
enable	Enables RIP on the VLAN.
in-policy <policy name>	Sets the VLAN RIP in-policy. • <i>policy name</i> is a string length {0..64}.
listen <enable disable>	Configures whether or not the switch will listen for RIP routes.
manualtrigger	Allows you to manually issue RIP updates.
out-policy <policy name>	Sets the VLAN RIP out -policy. • <i>policy name</i> is a string length {0..64}.

config vlan <vid> ip rip followed by:	
poison <enable/disable>	Sets whether or not RIP routes on the VLAN learned from a neighbor are advertised back to the neighbor. If disabled, split horizon is invoked and IP routes learned from an immediate neighbor are not advertised back to the neighbor. If enabled, the RIP updates sent to a neighbor from which a route is learned are “poisoned” with a metric of 16. Therefore, the receiver neighbor ignores this route because the metric 16 indicates infinite hops in the network.
supply <enable disable>	Enables or disables the switch to supply RIP updates.
trigger <enable /disable>	Enables or disables automatic triggered updates for RIP.

Refer to [Table 24 on page 343](#) for actions that result from RIP supply and listen settings.

[Figure 118 on page 348](#) shows sample output for this command.

Figure 118 config vlan ip rip info command output

```
Passport-8610# config vlan 1 ip rip info

Sub-Context: clear config dump monitor show test trace
Current Context:

                rip : disable
advertise-when-down : disable
auto-aggregation : disable
                cost : 1
    default-listen : disable
    default-supply : disable
            in-policy : N/A
            out-policy : N/A
    triggered-update : disable
                listen : enable
    manualtrigger : N/A
                poison : disable
                supply : enable
```

Showing RIP information for VLANs

To display RIP parameters for all VLANs or a the specified VLAN, use the following command:

```
show vlan info rip [<vid>]
```

where:

vid is a unique integer value in the range 1 and 4094 that represents the VLAN ID.

Figure 119 shows sample output for this command.

Figure 119 show vlan info rip command output

Passport-8606:6# show vlan info rip

=====								
Vlan Rip								
=====								
VLAN	DEFAULT	DEFAULT	TRIGGERED	AUTOAGG				
ID	ENABLE	SUPPLY	LISTEN	UPDATE	ENABLE	SUPPLY	LISTEN	POISON

1	false	false	false	false	false	true	true	false
2	false	false	false	false	false	true	true	false
3	false	false	false	false	false	true	true	false
4	false	false	false	false	false	true	true	false
5	false	false	false	false	false	true	true	false
6	false	false	false	false	false	true	true	false
7	false	false	false	false	false	true	true	false
8	false	false	false	false	false	true	true	false
9	false	false	false	false	false	true	true	false

Chapter 9

Configuring OSPF using Device Manager

The Open Shortest Path First (OSPF) protocol is a link-state protocol. The state of a link, or interface on a router, is a description of that interface and its relationship to its neighboring routers. The link-state description includes, for example, the IP address of the interface, the mask, the type of network it is connected to, the routers connected to that network and so on. The collection of all these link-states form the link-state database. OSPF uses this link-state database to build and calculate the shortest path to all known destinations.



Note: You can configure OSPF parameters only on an interface that has an IP address assigned to it.

- For conceptual information about OSPF management, see [Chapter 1, “IP routing concepts,”](#) on page 31.
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,”](#) on page 93.

This chapter includes the following topics:


Topic	Page
Viewing general OSPF routing information	352
Enabling or disabling OSPF on a router	355
Manually initiating a SPF run	356
Configuring OSPF interfaces	357
Managing an OSPF brouter port interface	366
Managing an OSPF VLAN interface	372
Managing OSPF areas information	375
Creating a virtual link	378

Topic	Page
Specifying ASBRs	384
Configuring metric speed	385
Viewing stub area metrics	387
Viewing advertisements in the Link State Database	388
Viewing characteristics in the Ext. Link State database	390
Inserting OSPF area aggregate ranges	391
Configuring an OSPF redistribute policy	394

Viewing general OSPF routing information

To view general OSPF information:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.
The OSPF dialog box opens with the [General tab displayed \(Figure 120\)](#).

Figure 120 OSPF dialog box—General tab

The image shows a screenshot of the OSPF dialog box in a network configuration tool. The title bar reads "192.168.151.163 - OSPF". The "General" tab is selected, showing various configuration options for OSPF. The "RouterId" is set to "56.126.132.0". The "AdminStat" is set to "disabled". The "VersionNumber" is "version2". The "AreaBdrRtrStatus" is "false". The "ASBdrRtrStatus" checkbox is unchecked. The "ExternLsaCount" is "0". The "ExternLsaCksumSum" is "0". The "OriginateNewLsas" is "00". The "RxNewLsas" is "00". The "10MbpsPortDefaultMetric" is "100". The "100MbpsPortDefaultMetric" is "10". The "1000MbpsPortDefaultMetric" is "1". The "10000MbpsPortDefaultMetric" is "1". The "TrapEnable" checkbox is unchecked. The "AutoVirtLinkEnable" checkbox is unchecked. The "SpfHoldDownTime" is "10" with a range of "3..60". The "OspfAction" is set to "none". The "LastSpfRun" is "none". At the bottom, there are buttons for "Apply", "Refresh", "Close", and "Help...".

192.168.151.163 - OSPF

Hosts | Link State Database | Ext. Link State Database | Area Aggregate | Redistribute

General | Areas | Stub Area Metrics | Interfaces | If Metrics | Neighbors | Virtual If | Virtual Neighbors

RouterId: 56.126.132.0

AdminStat: ☐ enabled ☒ disabled

VersionNumber: version2

AreaBdrRtrStatus: false

☐ ASBdrRtrStatus

ExternLsaCount: 0

ExternLsaCksumSum: 0

OriginateNewLsas: 00

RxNewLsas: 00

10MbpsPortDefaultMetric: 100 (number)

100MbpsPortDefaultMetric: 10 (number)

1000MbpsPortDefaultMetric: 1 (number)

10000MbpsPortDefaultMetric: 1 (number)

☐ TrapEnable

☐ AutoVirtLinkEnable

SpfHoldDownTime: 10 3..60

OspfAction: ☒ none ☐ runSpf

LastSpfRun: none

Apply Refresh Close Help...

Table 25 describes the General tab fields.

Table 25 General tab fields

Field	Description
RouterID	The Router ID, which in OSPF has the same format as an IP address but identifies the router independent of other routers in the OSPF domain.
AdminStat	The administrative status of OSPF in the router. The value “enabled” denotes that the OSPF process is active on at least one interface; “disabled” disables it on all interfaces. The default is disabled.
VersionNumber	Current version number of OSPF.
AreaBdrRtrStatus	A flag to note if this router is an area border router (ABR). Note: AreaBdrRtrStatus value must be <i>true</i> to create a virtual router interface.
ASBdrRtrStatus	When the ASBdrRtrStatus option is selected, the router is configured as an autonomous system boundary router (ASBR).
ExternLsaCount	The number of external (LS type 5) link state advertisements in the link state database.
ExternLsaChecksumSum	The 32-bit unsigned sum of the LS checksums of the external link state advertisements contained in the link state database. This sum is used to determine if there has been a change in a router's link state database and to compare the link state databases of two routers.
OriginateNewLsas	The number of new link state advertisements that have been originated. This number is incremented each time the router originates a new LSA.
RxNewLsas	The number of link state advertisements received that are determined to be new instantiations. This number does not include newer instantiations of self-originated link state advertisements.
10MbpsPortDefaultMetric	Indicates the default cost to be applied to the 10 Mb/s interface (port). The default is 100.
100MbpsPortDefaultMetric	Indicates the default cost to be applied to the 100 Mb/s interface (port). The default is 10.
1000MbpsPortDefaultMetric	Indicates the default cost to be applied to the 1000 Mb/s interface (port). The default is 1.
1000MbpsPortDefaultMetric	Indicates the default cost to be applied to the 10000 Mb/s interface (port). The default is 1.
TrapEnable	Indicates whether or not to enable traps relating to the OSPF. The default is false.

Table 25 General tab fields (continued)

Field	Description
AutoVirtLinkEnable	Enables or disables automatic creation of virtual links. The default is false.
SpfHoldDownTime	Allows you to change the OSPF hold down timer value (3 to 60 seconds). The default is 10 seconds.
OspfAction	Allows you to initiate a new SPF run to update the routing table. The default is none.
LastSpfRun	Used to indicate the time (SysUpTime) since the last SPF calculated by OSPF.

Enabling or disabling OSPF on a router

When configuring an interface for OSPF protocol, you must first enable OSPF globally on the router and then assign an IP address.

For instructions on assigning an IP address, see one of the following topics:

Topic	Page
Assigning an IP address to a brouter port interface	367
Assigning an IP address to a VLAN interface	372

To enable or disable OSPF globally on a router:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the [General tab](#) displayed ([Figure 120 on page 353](#)).



Note: Notice that the name or IP address of the device is always displayed in the upper left corner of the title bar.

- 2 In the AdminStat option box, select enabled to activate OSPF, or disabled to deactivate OSPF.

3 Click Apply.

The OSPF protocol is enabled (or disabled) on this router.

[Table 25 on page 354](#) describes the General tab fields.

Manually initiating a SPF run

From the OSPF > General tab, you can manually initiate a SPF, or Dijkstra, run to immediately update the OSPF link-state database. This is useful, for example:

- When you need to immediately restore a deleted OSPF-learned route.
- As a debug mechanism when the routing table's entries and the link-state database are out of sync.

To force an SPF calculation:


1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the [General tab displayed \(Figure 120 on page 353\)](#).

2 In the OSPF Action field, click runSpf.**3** Click Apply.

Device Manager prompts you to confirm if you want to initiate the SPF run ([Figure 121](#)).

Figure 121 Force SPF run dialog box

**4** Click Yes to confirm the forced SPF run.

The router performs the SPF run and the OSPF link state database is updated.



Note: After initiating an SPF run, wait 10 seconds before initiating another SPF run.

[Table 25 on page 354](#) describes the General tab fields.

Configuring OSPF interfaces

An OSPF interface, or link, is configured on an IP interface. In the Passport 8600, an IP interface can be either a single link (brouter port) or a logical interface configured on a VLAN (multiple ports). The state information associated with the interface is obtained from the underlying lower level protocols and the routing protocol itself.

Before you can configure OSPF protocol on a router interface, you must first [enable OSPF globally](#) on the router, and [assign an IP address](#) to the interface.

For more information, see [“Enabling or disabling OSPF on a router” on page 355](#).

When you enable an OSPF interface, you designate it as one of the following types:

- broadcast (active)
- non-broadcast multiaccess (NBMA)
- passive



Note: When an OSPF interface is enabled, you cannot change its interface type. You must first disable the interface. You can then change its type and re-enable it. If it is an NBMA interface, you must also first delete its manually-configured neighbors.

This section includes the following topics:

- [“Viewing OSPF interface information,”](#) next
- [“Creating an OSPF interface” on page 360](#)


- “Changing an OSPF interface type” on page 362
- “Configuring OSPF NBMA interfaces” on page 363

Viewing OSPF interface information

To view OSPF interface information:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.
The OSPF dialog box opens with the [General tab displayed \(Figure 120\)](#).
- 2 Click the Interfaces tab.
The OSPF dialog box—Interfaces tab opens ([Figure 122](#)).

Figure 122 OSPF dialog box—Interfaces tab



[Table 26](#) describes the OSPF dialog box—Interfaces tab fields.

Table 26 OSPF dialog box—Interfaces tab fields

Field	Description
IpAddress	IP address of the current OSPF interface
AddressLessIf	Designates whether an interface has an IP address. Interfaces with an IP address = 0 Interfaces without IP address = ifIndex

Table 26 OSPF dialog box—Interfaces tab fields

Field	Description
AreaId	Dotted decimal value to designate the OSPF area name. For VLANs keeping the default area setting on the interface causes the LSDB to be inconsistent. Note: The area name is not related to an IP address. You can use any value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdminStat	Current administrative state of the OSPF interface (enabled or disabled).
State	Current DR state of the OSPF interface (DR, BDR, OtherDR)
Rtrpriority	OSPF priority for the interface during the election process for the designated router. The interface with the highest priority number is the designated router. The interface with the second-highest priority becomes the backup designated router. If the priority is 0, the interface cannot become the designated router or the backup. The priority is used only during election of the designated router and backup designated router. The range is 0 to 255. The default is 1.
DesignatedRouter	IP address of the router elected by the Hello Protocol to send link state advertisements on behalf of the NBMA network.
BackupDesignatedRouter	IP address of the router elected by the Hello Protocol to send link state advertisements on behalf of the NBMA network if the designated router fails.
Type	Type of OSPF interface (broadcast, nbma, or passive)
AuthType	Type of authentication required for the interface. • none = No authentication required. • simple password = All OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey field. • MD5 authentication = All OSPF updates received by the interface must contain the md5 key.
AuthKey	Key (up to 8 characters) required when simple password authentication is specified in the interface AuthType field.
HelloInterval	Length of time, in seconds, between hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds. Note: When you change the Hello interval values, you must save the configuration file and reboot the switch for the values to be restored and checked for consistency.

Table 26 OSPF dialog box—Interfaces tab fields

Field	Description
TransitDelay	Length of time, in seconds between 1 and 3600, required to transmit an LSA update packet over the interface.
RetransInterval	Length of time, in seconds between 1 and 3600, required between LSA retransmissions.
RtrDeadInterval	Interval used by adjacent routers to determine if the router has been removed from the network. This interval must be identical on all routers on the subnet and a minimum of four times the Hello Interval. To avoid interpretability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.
PollInterval	Length of time, in seconds, between hello packets sent to an inactive OSPF router.
Events	Number of state changes or error events that have occurred through all interfaces.
MtuIgnore	Enable or disable Mtuignore flag for ignoring the mtu checking in ospf bdb.

Creating an OSPF interface

To create an OSPF interface:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.
The OSPF dialog box opens with the [General tab displayed \(Figure 120 on page 353\)](#).
- 2 Click the Interfaces tab.
The [Interfaces tab](#) opens ([Figure 122 on page 358](#)).
- 3 Click Insert.
The OSPF, Insert Interfaces dialog box opens ([Figure 123](#)).

Figure 123 OSPF Insert Interfaces dialog box

192.168.151.163 - OSPF, Insert Interfaces

IP Address: 192.168.151.163

AddressLessIf: 0 (number)

AreaId: 0.0.0.0

AdminStat: ☒ enabled ☐ disabled

RtrPriority: 1 0..255

Type: ☐ broadcast ☐ nbma ☐ passive

AuthType: ☒ none ☐ simplePassword ☐ md5

AuthKey:

HelloInterval: 10 1..65535

TransitDelay: 1 0..3600

RetransInterval: 5 0..3600

RtrDeadInterval: 40 0..2147483647

PollInterval: 120 0..2147483647 sec

MtuIgnore: ☐ enable ☒ disable

Insert Close Help...

- 4 Select the IP address for the interface from the IP Address pull-down list.
- 5 In the Type field, click the type of OSPF interface you want to create (broadcast, NMBA, or passive).
- 6 To designate a router priority, in the RtrPriority field, highlight the current value and type in a new value.
- 7 To change their values, highlight the current HelloInterval, RtrDeadInterval, or PollInterval and type in new values for the network.
- 8 To enable authentication, in the AuthType field, click either simplePassword or MD5.
- 9 If you chose simplePassword, in the AuthKey field, type in a password of up to eight characters.
- 10 Click Insert.

The OSPF Insert Interfaces dialog box closes.

- 11 On the Interfaces tab, click Apply.

The interface is created.



Note: When an OSPF interface is enabled, you cannot change its interface type. You must first disable the interface, change its type, and then re-enable it. If it is an NBMA interface, you must also first delete its manually-configured neighbors.

Changing an OSPF interface type

When an OSPF interface is enabled, you cannot change its interface type. You must first disable the interface, change its type, and then re-enable it. If it is an NBMA interface, you must also first delete its manually-configured neighbors.


To change an OSPF interface type:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the [General tab displayed \(Figure 120 on page 353\)](#). If the interface is currently an NBMA interface with manually-configured neighbors, go to step 2. If not, go to step 3.

- 2 To first delete the manually-configured neighbors on an NBMA interface, click the Neighbors tab ([Figure 124](#)).

Figure 124 Neighbors tab—NBMA manually-configured neighbors



- 3 Select the neighbors with a value of permanent in the ospfNbmaNbrPermanence column.
- 4 Click Delete.

The manually-configured neighbors are deleted.

- 5 Click the Interfaces tab.

The Interfaces tab opens (see [Figure 122 on page 358](#)).

- 6 To disable the interface, click in the AdminStat field, and choose disabled from the pull-down list.

- 7 Click Apply.

The interface is disabled.

- 8 To change the interface type, click in the Type field, and choose the new interface type (broadcast, passive, or nbma) from the pull-down list.

- 9 Click Apply.

The interface type is changed.

- 10 To enable the interface, click in the AdminStat field, and choose enabled from the pull-down list.

- 11 Click Apply.

The interface is enabled as the new type.

Configuring OSPF NBMA interfaces

In contrast to a broadcast network, where some OSPF protocol packets are multicast (sent to AllSPFRouters and AllDRouters), NBMA packets are replicated and sent to each neighboring router as unicast. NBMA networks drop all OSPF packets with destination addresses AllSPFRouters and AllDRouters. Because the NBMA network does not broadcast, you must manually configure a list of neighbors and their priorities for all routers in the network that are eligible to become the DR (those with a positive, non-zero router priority).

Before you begin this configuration, identify the following:

- Specific interfaces to be included in the NBMA network
- IP address for each interface
- Router priority for each interface
- HelloInterval for the network
- RtrDeadInterval for the network
- PollInterval for the network

After you gather the above information, you can configure the interfaces, and add neighbors for each interface that is eligible to become the DR (those with a positive, non-zero router priority).

This section includes the following topics:

- “Adding NBMA neighbors,” next
- “Viewing OSPF neighbor information” on page 365


Adding NBMA neighbors

An NBMA interface that has a positive, non-zero router priority is eligible to become the DR for the NBMA network and is configured with the identification of all attached routers, their IP addresses, and their router priorities.

To add neighbors for an OSPF NBMA interface:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.
The OSPF dialog box opens with the [General tab](#) displayed ([Figure 120 on page 353](#)).
- 2 Click the Interfaces tab.
The [Interfaces tab](#) opens (see [Figure 122 on page 358](#)).
- 3 Select an NBMA interface with a positive, non-zero router priority.
- 4 Click the Neighbors tab.
The [Neighbors tab](#) opens ([Figure 125](#)).

Figure 125 OSPF dialog box—Neighbors tab



5 Click Insert.

The OSPF, Insert Neighbors dialog box opens ([Figure 126](#)).

Figure 126 OSPF, Insert Neighbors dialog box



6 Enter the IP address and priority for the first neighbor.

7 Click Insert.

The neighbor is added to the Neighbors tab.

8 Repeat step 6 for all neighbors.

9 Click Apply.

The neighbors are configured for this NBMA interface.

10 To configure neighbors for other NBMA interfaces eligible to become DR (those with a positive, non-zero router priority), repeat Steps 1-8.

The neighbors are configured for the NBMA network.

Viewing OSPF neighbor information

Two routers that have interfaces to a common network are called neighbors and appear on each other's Neighbors tab. Neighbor relationships are maintained by, and usually dynamically discovered by, OSPF's Hello protocol.

The exception is that, in an NBMA network, permanent neighbors are [manually configured](#) on each router eligible to become the DR.

To view the OSPF neighbors:

1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the [General tab](#) displayed (see [Figure 120 on page 353](#)).

2 Click the Neighbors tab.

The **Neighbors tab** opens (see [Figure 125 on page 364](#)).

[Table 27](#) describes the Neighbors tab fields.

Table 27 Neighbors tab fields

Field	Descriptions
IpAddr	IP address.
AddressLessIndex	On an interface having an IP address, zero. On addressless interfaces, the corresponding value of ifIndex in the Internet standard MIB. On row creation, this value is derived from the instance.
Router	The router ID of the neighboring router, which in OSPF has the same format as an IP address but identifies the router independent of its IP address.
Options	A bit mask corresponding to the neighbor's options field.
Priority	Assignment of preferential treatment to place the transmitted packets in queues and possible selection of the priority field in the data link header when the packet is forwarded.
State	The OSPF Interface state.
Events	The number of state changes or error events that have occurred between the OSPF router and the neighbor router.
Retransmission Queue Length	Retransmission Queue Length.
ospfNbmaNbrPermanence	Indicates whether the neighbor is a manually-configured NBMA neighbor.
HelloSuppressed	This variable indicates whether Hellos are being suppressed to a neighbor.

Managing an OSPF brouter port interface

From the Edit > Port dialog box you can assign an IP address to an OSPF port, and make specific OSPF interface configurations. When configuring an interface for OSPF protocol, you must first enable OSPF globally on the router and then assign an IP address.

For instructions on globally enabling OSPF, see [“Enabling or disabling OSPF on a router” on page 355](#).

This section includes the following topics:

- [“Assigning an IP address to a brouter port interface,”](#) next
- [“Configuring OSPF on a brouter port interface”](#) on page 368


Assigning an IP address to a brouter port interface

To assign an IP address to an interface:

- 1 On the Device View, select a port.
- 2 From the Device Manager menu bar, choose Edit > Port.
The Port dialog box opens with the Interface tab displayed.
- 3 Click the IP Address tab.

The IP Address tab opens ([Figure 127](#)).


Figure 127 Port dialog box — IP Address tab



- 4 Click Insert.

The Port, Insert IP Address dialog box opens ([Figure 128](#)).

Figure 128 Port, Insert IP Address dialog box



- 5 In the IpAddress field, type the interface IP Address.

- 6 To automatically enter the default netmask, press the Tab key.
- 7 In the VlanId text box, select the Vlan ID.
- 8 Click Insert.

The IP Address is assigned to the selected port.


Configuring OSPF on a router port interface

Before you configure OSPF on a port, make sure to [enable OSPF globally](#) on the router and [assign an IP address](#) to the interface.

To configure OSPF on a port interface:

- 1 On the Device View, select a port.
- 2 From the Device Manager menu bar, choose Edit > Port.
The Port dialog box opens with the Interface tab displayed.
- 3 Click the OSPF tab.

The OSPF tab opens ([Figure 129](#)).

Figure 129 Port dialog box — OSPF tab


192.168.151.163 - Port 1/13

VLACP	Remote Mirroring	Mroute Stream Limit	Fdb Protect	IP Address	ARP	DHCP			
Interface	VLAN	STG	MAC Learning	Rate Limiting	Test	SMLT	PCAP	EAPOL	LACP
DVMRP	IGMP	OSPF	RIP	PIM	PGM	VRRP	Router Discovery	IPX BRouter	

☐ Enable

HelloInterval: 10 sec

RtrDeadInterval: 40 sec

DesigRtrPriority: 1 (minimum=0)

Metric: 0 (use port speed=0)

AuthType: ☒ none ☐ simplePassword ☐ md5

AuthKey:

AreaId: 0.0.0.0

☐ AdvertiseWhenDown

IfType: ☒ broadcast ☐ nbma ☐ passive

PollInterval: 120 0..2147483647

Graph... Apply Refresh Close Help...



Note: Use the Edit > Port > OSPF tab to configure OSPF on a router port. To configure OSPF on a VLAN, use VLAN > VLANs > Basic > IP > OSPF. OSPF must be globally enabled before the configuration takes effect.

The OSPF tab is not applicable unless the port or VLAN is routed, that is, it is assigned an IP address.

- 4 In the IfType field, click the interface type you want to create (broadcast, NMBA, or passive).
- 5 Click Enable.

- 6 To designate a router priority, in the DesigRtrPriority field, highlight the current value and type in a new value.
- 7 To change their values, highlight the current HelloInterval, RtrDeadInterval, or PollInterval and type in new values for the network.
- 8 To enable authentication, in the AuthType field, click either simplePassword or MD5.
- 9 If you chose simplePassword, in the AuthKey field, type in a password of up to eight characters.
- 10 Click Apply.

OSPF is configured for the port.



Note: When an OSPF interface is enabled, you cannot change its interface type. You must first disable the interface. You can then change its type and re-enable it. If it is an NMBA interface, you must also first delete its manually-configured neighbors.

Table 28 describes the fields on the OSPF tab.

Table 28 OSPF tab fields

Field	Description
Enable	Enable or disable OSPF routing on the specified interface. The default is false.
HelloInterval	Length of time, in seconds, between hello packets. This value must be the same for all routers attached to a common network. The default is 10 seconds. Note: When you change the Hello interval values, you must save the configuration file and reboot the switch for the values to be restored and checked for consistency.
RtrDeadInterval	Interval used by adjacent routers to determine if the router has been removed from the network. This interval must be identical on all routers on the subnet and a minimum of four times the Hello Interval. To avoid inter operability issues, the RtrDeadInterval value for the OSPF interface needs to match with the RtrDeadInterval value for the OSPF virtual interface. The default is 40 seconds.

Table 28 OSPF tab fields (continued)

Field	Description
DesigRtrPriority	The priority of this interface. Used in multiaccess networks. This field is used in the designated router election algorithm. The value 0 indicates the router is not eligible to become the designated router on this particular network. In the event of a tie in this value, routers will use their router id as a tie breaker. The default is 1.
Metric	The metric for this type of service (TOS) on this interface. The value of the TOS metric is $(10^9 / \text{interface speed})$. The default is 1. FFFF= There is no route for this TOS. POS/IPCP links = defaults to 0. 0 = The interface speed is used as the metric value when the state of the interface is up.
AuthType	Type of authentication required for the interface. • none = No authentication required. • simple password = All OSPF updates received by the interface must contain the authentication key specified in the interface AuthKey field. • MD5 authentication = All OSPF updates received by the interface must contain the md5 key.
AuthKey	Key (up to 8 characters) required when simple password authentication is specified in the interface AuthType field.
AreaId	Dotted decimal value to designate the OSPF area name. Note: The area name is not related to an IP address. You can use any value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
AdvertiseWhenDown	If true, the network on this interface will be advertised as up, even if the port is down. The default is false. Note: When you configure a port without any link and enable AdvertiseWhenDown, it will not advertise the route until the port is active. Then the route will be advertised even when the link is down. To disable advertising based on link states, disable AdvertiseWhenDown.
IfType	Type of OSPF interface (broadcast, NBMA, or passive). Note: Before changing an OSPF interface type, you must first disable the interface. If it is an NBMA interface, you must also delete all configured neighbors.
PollInterval	Length of time, in seconds, between hello packets sent to an inactive OSPF router. Neighbors must have the same PollInterval value.

Managing an OSPF VLAN interface

From the VLAN dialog box you can assign an IP address to an OSPF port, and make specific OSPF interface configurations. When configuring an interface for OSPF protocol, you must first [enable OSPF globally](#) on the router and then [assign an IP address](#).

For instructions on globally enabling OSPF, see “[Enabling or disabling OSPF on a router](#)” on page 355.

This section includes the following topics:

- “[Assigning an IP address to a VLAN interface](#),” next
- “[Configuring OSPF on a VLAN interface](#)” on page 373

For OSPF configuration examples, refer to [Chapter 2, “IP routing configuration examples,”](#) on page 93.


Assigning an IP address to a VLAN interface

To assign an IP address to an VLAN interface:

- 1 On the Device View, select a port.
- 2 From the Device Manager menu bar, choose VLAN > VLANs.

The VLAN dialog box opens with the Basic tab displayed ([Figure 130](#)).

Figure 130 VLAN dialog box—Basic tab




- 3 Select a VLAN.

4 Click IP.

The IP VLAN dialog box opens with the IP Address tab displayed (Figure 131).


Figure 131 IP, VLAN dialog box—IP Address tab



5 Click Insert.

The VLAN, Insert IP Address dialog box opens (Figure 132).

Figure 132 IP, VLAN dialog box—Insert IP Address dialog box



6 In the IpAddress field, type the interface IP Address.

7 To automatically enter the default netmask, press the Tab key.

8 Click Insert.

The IP Address is assigned to the selected VLAN interface.

Configuring OSPF on a VLAN interface

To enable and configure OSPF on a VLAN interface:

1 From the Device Manager menu bar, choose VLAN > VLANs.

The VLAN dialog box opens with the Basic tab displayed (see Figure 130 on page 372).

2 Select a VLAN.

3 Click IP.

The IP, VLAN dialog box opens with the IP Address tab displayed (see [Figure 131 on page 373](#)).

4 Click the OSPF tab.

The IP, VLAN, OSPF tab opens ([Figure 133](#)).

Figure 133 IP, VLAN dialog box—OSPF tab



Note: Use the VLAN > VLANs > Basic > IP > OSPF tab to configure OSPF on a VLAN. To configure OSPF on a port, use the Edit > Port > OSPF tab. OSPF must be globally enabled before the configuration takes effect.

The OSPF tab is not applicable unless the port or VLAN is routed, that is, it is assigned an IP address.

5 In the IfType field, click the interface type you want to create (broadcast, NMBA, or passive).

- 6 To enable OSPF on the VLAN interface, click Enable.
- 7 To designate a router priority, in the DesigRtrPriority field, highlight the current value and type in a new value.
- 8 To change their values, highlight the current HelloInterval, RtrDeadInterval, or PollInterval and type in new values for the network.
- 9 To enable authentication, in the AuthType field, click either simplePassword or MD5.
- 10 If you chose simplePassword, in the AuthKey field, type in a password of up to eight characters.
- 11 Click Apply.

OSPF is configured for the VLAN.



Note: When an OSPF interface is enabled, you cannot change its interface type. You must first disable the interface. You can then change its type and re-enable it. If it is an NMBA interface, you must also first delete its manually-configured neighbors.

Managing OSPF areas information

OSPF allows collections of contiguous networks and hosts to be grouped together. Such a group, together with the routers having interfaces to any of the included networks, is called an area. Each area runs a separate copy of the basic link-state routing algorithm. This means that each area has its own link-state database.

This section includes the following topics:

- [“Viewing OSPF areas information,”](#) next
- [“Creating a stub area or NSSAs”](#) on page 377

Viewing OSPF areas information

To view information about OSPF areas:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the [General tab displayed](#) (see [Figure 120 on page 353](#)).

- 2 Click the Areas tab

The [Areas tab](#) opens ([Figure 134](#)).

Notice that the backbone ID is always displayed as 0.0.0.0.

Figure 134 OSPF dialog box—Areas tab

Virtual Neighbors		Hosts	Link State Database		Ext. Link State Database		Area Aggregate	Redistribute
General		Areas	Stub Area Metrics		Interfaces	If Metrics	Neighbors	Virtual If
Araid	ImportAsExtern	SpfRuns	AreaBdrRtrCount	AsBdrRtrCount	AreaLsaCount	AreaLsaChecksumSum	AreaSummary	ActiveIfCount
0.0.0.0	importExternal	00	0	0	0	0	sendAreaSummary	0
12.12.12.12	importExternal	00	0	0	0	0	sendAreaSummary	0
22.22.22.22	importExternal	00	0	0	0	0	sendAreaSummary	0

[Table 29](#) describes the Areas tab fields.

Table 29 Areas tab fields

Field	Description
Araid	A 32-bit integer uniquely identifying an area. Area ID 0.0.0.0 is used for the OSPF backbone. For VLANs keeping the default area setting on the interface causes the LSDB to be inconsistent.
ImportAsExtern	The area's support for importing AS external link state advertisements. Could be importExternal (default), importNotExternal, or importNssa (not so stubby area).
SpfRuns	Used to indicate the number of SPF calculations performed by OSPF.
AreaBdrRtrCount	The total number of area border routers reachable within this area. The value, initially zero, is calculated in each SPF Pass.

Table 29 Areas tab fields (continued)

Field	Description
AsBdrRtrCount	The total number of autonomous system border routers reachable within this area. The value, initially zero, is calculated in each SPF Pass.
AreaLsaCount	The total number of link state advertisements in this area's link state database, excluding AS External LSAs.
AreaLsaCksumSum	The 32-bit unsigned sum of the link state advertisements. This sum excludes external (LS type 5) link state advertisements. The sum is used to determine if there has been a change in a router's link state database and to compare the link state database of two routers.
AreaSummary	The area's support for Summary advertisements in a stub area.
ActiveifCount	The number of active interfaces in this area.

Creating a stub area or NSSAs

A stub area does not receive advertisements for external routes, which reduces the size of the link state database. A stub area has only one area border router. Any packets destined outside the area are simply routed to that area border exit point, examined by the area border router, and forwarded to a destination.

A not so stubby area (NSSA) also prevents the flooding of AS-External Link State advertisements into the area by replacing them with a default route. The added feature of NSSAs is the ability to import small stub (non-OSPF) routing domains into OSPF.

To create a stub area or NSSA:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.
The OSPF dialog box opens with the General tab displayed (see [Figure 120 on page 353](#)).
- 2 Click the Areas tab.
The Areas tab opens ([Figure 135](#)).

Figure 135 OSPF dialog box—Areas tab

10.10.54.27 - OSPF								
Virtual Neighbors		Hosts	Link State Database		Ext. Link State Database		Area Aggregate	Redistribute
General		Areas	Stub Area Metrics		Interfaces	If Metrics	Neighbors	Virtual If
AreaId	ImportAsExtern	SpfRuns	AreaBdrRtrCount	AsBdrRtrCount	AreaLsaCount	AreaLsaCksumSum	AreaSummary	ActiveIfCount
0.0.0.0	ImportExt...	00	0	0	0	0	sendAreaSummary	0
12.12.12.12	ImportExternal	00	0	0	0	0	sendAreaSummary	0
22.22.22.22	ImportNoExterna	00	0	0	0	0	sendAreaSummary	0
	ImportNssa							

3 row(s)

Apply Refresh Insert... Delete [Icons] Close Help...

- 3 Under the ospfImportASExtern field, select the area you want to change to a stub area or NSSA; use the pull-down menu to select importExternal, ImportNoExternal, or importNssa.
- 4 Under the ospfImportASExtern field, select the area you want to change to a stub area or NSSA; use the pull-down menu to select importExternal, ImportNoExternal, or importNssa.
- 5 Click Apply.

Creating a virtual link

When using OSPF, Passport 8000 switches, which are ABRs, need to be connected directly to the backbone. If they are not directly connected, they need to have a virtual link. In an Passport 8000 switch, you can specify that virtual links be automatically created, or you can manually configure a virtual link.

When automatic virtual linking is enabled, it acts as insurance. A virtual link will be created for vital traffic paths in your OSPF configuration if something goes amiss, such as when an interface cable providing connection to the backbone (either directly or indirectly) becomes disconnected from the switch. Specifying automatic virtual linking ensures that a link will be created via another switch. When you specify automatic virtual linking, it is always ready to create a virtual link. If automatic virtual linking uses more resources than you want to expend, creating a manual virtual link may be the better solution. This approach lets you conserve resources while having specific control of where virtual links are placed in your OSPF configuration.

OSPF behavior has been modified according to OSPF standards so that OSPF routes cannot be learned through an area border router (ABR) unless it is connected to the backbone or through a virtual link.

This section includes the following topics:

- [“Managing an automatic virtual link,”](#) next
- [“Configuring a manual virtual link”](#) on page 379
- [“Viewing virtual links on neighboring devices”](#) on page 381
- [“Managing router hosts”](#) on page 382

Managing an automatic virtual link

To specify that virtual links be automatically created:

- 1 Choose IP Routing > OSPF.

The OSPF dialog box opens with the [General tab](#) displayed (see [Figure 120 on page 353](#)).

- 2 Select true in the AutoVirtLinkEnable field.

By default, this feature is set to false, and virtual links are not automatically created.

- 3 Click Apply.

Configuring a manual virtual link

To manually configure a virtual link:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.


The OSPF dialog box opens with the [General tab](#) displayed (see [Figure 120 on page 353](#)).

- 2 Click the Virtual If tab.

The Virtual IF tab opens.

The [Virtual If tab](#) opens ([Figure 136](#)).


Figure 136 OSPF dialog box—Virtual If tab



3 Click Insert.

The OSPF, Insert Virtual If dialog box opens (Figure 137).

Figure 137 OSPF, Insert Virtual If dialog box



4 Specify the area ID of the transit area.

The transit area is the common area between two ABRs.

5 Specify the neighbor ID.

The neighbor ID is the IP router ID of the ABR that the other ABR needs to go through to get to the backbone.

6 Click Insert.

7 To verify that the virtual link is active, refresh the Virtual If tab (see Figure 136 on page 380) and check the state column.

If the state displays “point to point,” the virtual link is active. If the state column displays “down,” the virtual link is configured incorrectly.

Table 31 describes the Virtual If tab fields.

Table 30 OSPF dialog box—Virtual If tab fields

Field	Description
Areald	The Transit Area Id that the virtual link traverses.
Neighbor	The router ID of the virtual neighbor.
TransitDelay	The estimated number of seconds it takes to transmit a link- state update packet over this interface.
Retrans Interval	The number of seconds between link-state, advertisement, and retransmissions, for adjacencies belonging to this interface. This value is also used when retransmitting database description and link-state request packets. This value should be well over the expected round- trip time.
HelloInterval	The length of time, in seconds, between the Hello packets that the router sends on the interface. This value must be the same for the virtual neighbor.
RtrDeadInterval	The number of seconds that a router's Hello packets have not been seen before it's neighbors declare the router down. This should be some multiple of the Hello interval. This value must be the same for the virtual neighbor.
State	The OSPF virtual interface states.
Events	The number of state changes or error events on this Virtual Link.
AuthType	The authentication type specified for a virtual interface. Additional authentication types may be assigned locally.
AuthKey	If Authentication Type is a simple password, the device will left adjust and zero fill to 8 octets. Note that unauthenticated interfaces need no authentication key and simple password authentication can not use a key of more than 8 octets.

Viewing virtual links on neighboring devices

You can check the Virtual Neighbor tab to view the area and virtual link configuration for the neighboring device.

To view the virtual neighbor:


- 1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the General tab displayed (see [Figure 120 on page 353](#)).

- 2 Click the Virtual Neighbor tab.

The Virtual Neighbor tab opens ([Figure 138](#)).

Figure 138 OSPF dialog box—Virtual Neighbor tab



[Table 31](#) describes the Virtual Neighbor tab fields.

Table 31 OSPF dialog box—Virtual Neighbor tab fields

Field	Description
Area	The subnetwork in which the virtual neighbor resides.
RtrId	A 32-bit integer (represented as a type IpAddress) uniquely identifying the neighboring router in the autonomous system.
IpAddr	The IP address of the virtual neighboring router.
Options	A bit mask corresponding to the neighbor's options field.
State	The OSPF Interface state.
Events	The number of state changes or error events that have occurred between the OSPF router and the neighbor router.
LsRetransQLen	The number of elapsed seconds between advertising retransmissions of the same packet to a neighbor.
HelloSuppressed	This field indicates whether or not Hellos are being suppressed to the neighbor.

Managing router hosts

You can specify which hosts are directly attached to the router, and the metrics and types of service that should be advertised for them.

To manage router hosts:


- 1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the General tab displayed ([Figure 120 on page 353](#)).

- 2 Click the Hosts tab.

The Hosts tab opens ([Figure 139](#)).

Figure 139 OSPF dialog box—Hosts tab



[Table 32](#) describes the Hosts tab fields.

Table 32 Host tab fields

Field	Description
IpAddress	The IP address of the host used to represent a point of attachment in a TCP/IP internetwork.
TOS	The type of service of the route being configured.
Metric	The metric advertised to other areas. The value indicates the distance from the OSPF router to any network in the range.
AreaID	Area where the host is found. By default, the area that is submitting the OSPF interface is in 0.0.0.0.

Specifying ASBRs

ASBRs advertise non-OSPF routes into OSPF domains so that they can be passed along throughout the OSPF routing domain. A router can function as an ASBR if one or more of its interfaces is connected to a non-OSPF network (for example, RIP, BGP, or EGP).

To conserve resources, you may want to limit the number of ASBRs in your network or to specifically control which routers perform as ASBRs to control traffic flow.

To specify whether or not a router should be an ASBR:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.
The OSPF dialog box opens with the General tab displayed (see [Figure 120 on page 353](#)).
- 2 From the ASBdrRtrStatus field, select true to designate the router as an ASBR or false to remove ASBR status from the router.
- 3 Click Apply.

Configuring metric speed

You can configure the metric speed globally or for specific ports and interfaces on your network. In addition, you can control redistribution options between non-OSPF interfaces and OSPF interfaces.

This section includes the following topics:

- [“Configuring global default metric speed,”](#) next
- [“Managing metrics with the peer layer interface”](#) on page 385

Configuring global default metric speed

To change the default metric speed on specific port types:

- 1 Choose IP Routing > OSPF > General.

The OSPF dialog box opens with the General tab displayed as shown in [Figure 120 on page 353](#).

- 2 Change the metric value in one or all of the following fields:

- 10MbpsPortDefaultMetric (default = 100)
- 100MbpsPortDefaultMetric (default = 10)
- 1000MbpsPortDefaultMetric (default = 1)
- 10000MbpsPortDefaultMetric (default = 1)

- 3 Click Apply.

The default port metric speed will be changed on all port types for which you have specified a new metric speed.


Managing metrics with the peer layer interface

The If Metrics tab indicates the metrics associated with the peer layer interface. For finer control over port-specific metric speed, you can specify the metric speed when you enable OSPF on a port or when you edit a port.

To specify the metric speed on a specific port instead of a port type:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.
The OSPF dialog box opens with the General tab displayed (see [Figure 120 on page 353](#)).
- 2 Click the If Metrics tab.
The IF Metrics tab opens ([Figure 140](#)).

Figure 140 OSPF dialog box—If Metrics tab



- 3 Specify a new metric speed in the value field of the If Metrics tab or the metric field of the port OSPF tab.
- 4 Click Apply.

 **Note:** When you enable a port for OSPF routing, the default metric in the port tab is “0.” A value of “0” (zero) means that the port will use the default metrics for port types that are specified on the OSPF general tab.

[Table 33](#) describes the If Metrics tab fields.

Table 33 If Metrics tab fields

Field	Description
IpAddress	The Internet Protocol address of the device used to represent a point of attachment in a TCP/IP internetwork.
AddressLessIf	For the purpose of easing the instancing of addressed and addressless interfaces. This variable takes the value 0 on interfaces with IP addresses and the corresponding value of ifIndex for interfaces having no IP address.
TOS	Type of service is a mapping to the IP type of service flags as defined in the IP forwarding table MIB.
Value	The value advertised to other areas indicating the distance from the OSPF router to any network in the range.
Status	Active or not active. Not configurable.

Viewing stub area metrics

The Stub Area Metrics tab contains the set of metrics that will be advertised by a default area border router into a stub area.


To view the set of stub area metrics:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the General tab displayed (see [Figure 120 on page 353](#)).

- 2 Click the Stub Area Metrics tab.

The Stub Area Metrics tab opens ([Figure 141](#)).

Figure 141 OSPF dialog box—Stub Area Metrics tab

- 3 Specify a new metric speed in the value field of the Interface Metric field or the metric field of the port OSPF tab.
- 4 Click Apply.

 **Note:** When you enable a port for OSPF routing, the default metric in the port tab is “0.” A value of “0” (zero) means that the port will use the default metrics for port types that are specified on the OSPF general tab.

Table 34 describes the Stub Area Metrics tab fields.

Table 34 Stub Area Metrics tab fields

Field	Description
AreaID	The 32-bit identifier for the stub area. On creation, it can be derived from the instance.
TOS	The type of service associated with the metric. On creation, it can be derived from the instance.
Metric	The metric value applied at the indicated type of service. By default, it equals the lowest metric value at the type of service among the interfaces to other areas.
Status	Active or not active. Not configurable.

- 5 Specify a new metric speed in the value field of the If Metrics tab or the metric field of the port OSPF tab.
- 6 Click Apply.

 **Note:** When you enable a port for OSPF routing, the default metric in the port tab is “0.” A value of “0” (zero) means that the port will use the default metrics for port types that are specified on the OSPF General tab.

Viewing advertisements in the Link State Database

To view the advertisements of the areas throughout the link state database:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the General tab displayed (Figure 120 on page 353).

- 2 Click the Link State Database tab.

The Link State Database tab opens (Figure 142).

Figure 142 OSPF dialog box—Link State Database tab




Table 35 describes the Link State Database tab fields.

Table 35 Link State Database tab fields

Field	Description
Areald	A 32-bit integer uniquely identifying an area. Area ID 0.0.0.0 is used for the OSPF backbone.
Type	The OSPF interface type. By way of a default, this field may be intuited from the corresponding value of ifType. Broadcast LANs, such as Ethernet and IEEE 802.5, take the value broadcast; X.25 and similar technologies take the value nbma; and links that are definitively point-to-point take the value pointToPoint.
Lsid	The Link State ID is an LS type-specific field containing either a router ID or an IP address. It identifies the piece of the routing domain that is being described by the advertisement.
RouterId	A 32-bit integer uniquely identifying the router in the autonomous system.
Sequence	The sequence number is a signed 32-bit integer that identifies old and duplicate link state advertisements.

Table 35 Link State Database tab fields (continued)


Field	Description
Age	The age in seconds of the link state advertisement.
Checksum	This field is the checksum of the complete contents of the advertisement, excepting the age field. The age field is excepted so that an advertisement's age can be incremented without updating the checksum. The checksum used is the same that is used for ISO connectionless datagrams. It is commonly referred to as the Fletcher checksum.

Viewing characteristics in the Ext. Link State database

To view the characteristics of the external link state database:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.
The OSPF dialog box opens with the General tab displayed (see [Figure 120 on page 353](#)).
- 2 Click the Ext. Link State DB tab.
The Ext. Link State DB tab opens ([Figure 143](#)).

Figure 143 OSPF dialog box—Ext. Link State DB tab



[Table 36](#) describes the Ext. Link State DB tab fields.

Table 36 Ext. Link State DB tab fields

Field	Description
Type	The OSPF interface type. By way of a default, this field may be intuited from the corresponding value of ifType. Broadcast LANs, such as Ethernet and IEEE 802.5, take the value broadcast; X.25 and similar technologies take the value nbma; and links that are definitively point-to-point take the value pointToPoint.
Lsid	The Link State ID is an LS type-specific field containing either a router ID or an IP address. It identifies the piece of the routing domain that is being described by the advertisement.
RouterId	A 32-bit integer uniquely identifying the router in the autonomous system.
Sequence	The sequence number is a signed 32-bit integer that identifies old and duplicate link state advertisements.
Age	The age in seconds of the link state advertisement.
Checksum	This field is the checksum of the complete contents of the advertisement, excepting the age field. The age field is excepted so that an advertisement's age can be incremented without updating the checksum. The checksum used is the same that is used for ISO connectionless datagrams. It is commonly referred to as the Fletcher checksum.
Advertisement	Hex representation of the entire link state advertisement, including the header.

Inserting OSPF area aggregate ranges


To insert OSPF area aggregate ranges:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.

The OSPF dialog box opens with the General tab displayed (see [Figure 120 on page 353](#)).


- 2 Click the Area Aggregate tab.

The Area Aggregate tab opens ([Figure 144](#)).

Figure 144 OSPF dialog box—Area Aggregate tab

3 Click Insert.

The OSPF, Insert Area Aggregate dialog box opens ([Figure 145](#)).

Figure 145 OSPF, Insert Area Aggregate dialog box

4 Type the Area ID.

5 Select the type of link state database.

- summaryLink—to generate an aggregated summary
- nssaExternalLink—to generate an (NSSA) link summary

6 Select the effect you want:

- advertiseMatching—to advertise the aggregate summary LSA with the same LSID
- doNotAdvertiseMatching—to suppress all networks that fall within the entire range
- advertiseDoNotAggregate—to advertise individual networks

7 In the AdvertiseMetric field, enter a cost value (in the range 0 and 65535) to advertise for the OSPF area range.

8 Click Insert.

[Table 37](#) describes the Area Aggregate tab fields.

Table 37 Area Aggregate tab fields

Field	Description
AreaID	The area in which the address would be found.
LsdbType	One of the following: - summaryLink—aggregated summary link - nssaExternalLink—not so stubby area link
IP Address	The IP Address of the Net or Subnet indicated by the range.
Mask	Network mask for the area range.
Effect	One of the following: - advertiseMatching—advertise the aggregate summary LSA with same LSID. - doNotAdvertiseMatching—suppress all networks that fall within the entire range. - advertiseDoNotAggregate—advertise individual networks.
AdvertiseMetric	Changes the advertised metric cost value of the OSPF area range. Enter an integer value in the range 0 and 65535, which represents the metric cost value for the OSPF area range.

Configuring an OSPF redistribute policy

You can configure a redistribute entry for OSPF to announce routes of a certain source type, for example, static, RIP, or direct. If a route policy field is not configured for a redistribute entry, then the default action is taken on the basis of metric, metric-type, and subnet configured. This is called basic redistribution. Otherwise, you use the route policy specified to perform detailed redistribution. If no redistribution entry is configured, no external LSA is generated for non-OSPF routes.

You can also configure OSPF redistribute policies in the OSPF Redistribute tab of the Policy dialog box. See [“Configuring an OSPF redistribute policy” on page 394](#).



Note: Changing OSPF Redistribute contexts is a process-oriented operation that can affect system performance and network reachability while performing the procedures. Therefore, Nortel Networks recommends that if you want to change default preferences for an OSPF Redistribute context, you should do so before enabling the protocols.

To set up or edit an OSPF redistribute policy:

- 1 From the Device Manager menu bar, choose IP Routing > OSPF.
The OSPF dialog box opens with the General tab displayed ([Figure 120 on page 353](#)).
- 2 Click the Redistribute tab.
The Redistribute tab opens ([Figure 146](#)).

Figure 146 OSPF dialog box—Redistribute tab



3 Click Insert.

The OSPF, Insert OSPF Redistribute dialog box opens (Figure 147).

4 Click Insert.

Figure 147 OSPF, Insert OSPF Redistribute dialog box




Table 38 describes the OSPF, Insert OSPF Redistribute dialog box fields.

Table 38 OSPF, Insert OSPF Redistribute dialog box fields

Field	Description
RouteSource	Select the route source protocol for the redistribution entry.
Enable	Enables (or disables) an OSPF redistribute entry for a specified source type. You can also enable or disable this feature in the OSPF Redistribute tab of the Policy dialog box by clicking in the field and selecting enable or disable from the pulldown menu.
Metric	Set the OSPF route redistribution metric for basic redistribution. The value can be a range between 0 to 65535. If configured as 0, the original cost of the route is used.
MetricType	Set the OSPF route redistribution metric type. The default is Type 2. You can also select your entry in the OSPF Redistribution tab of the Policy dialog box by clicking in the field and selecting any, type1, or type2 from the pulldown menu.

Table 38 OSPF, Insert OSPF Redistribute dialog box fields (continued)

Field	Description
Subnets	Allows or suppresses external subnet routes while being redistributed into an OSPF domain. You can also select your entry in the OSPF Distribution tab of the Policy dialog box by clicking in the field and selecting allow or deny from the pulldown menu.
RoutePolicy	Sets the route policy by name to be used for the detailed redistribution of external routes from a specified source into an OSPF domain. Click the ellipse button and choose from the list in the Route Policy dialog box (Figure 147). To deselect an entry, use the ALT key.

Chapter 10

Configuring OSPF using the CLI

This chapter describes the Run-Time CLI commands that are used to configure the Open Shortest Path First (OSPF) protocol in the Passport 8000 Series Switch. Routers use the Open Shortest Path First (OSPF) protocol to exchange network topology information among themselves, providing each router with a map of the network.

Before you can configure OSPF parameters on an interface, you must first configure IP on that interface.



Note: OSPF behavior has been modified according to OSPF standards so that OSPF routes cannot be learned through an area border router (ABR) unless it is connected to the backbone or through a virtual link.

- For conceptual information about OSPF management, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter includes the following topics:

Topic	Page
Roadmap of IP commands	398
Configuring OSPF global parameters	403
Configuring OSPF host route parameters	404
Configuring an OSPF interface	407
Configuring OSPF areas	410
Configuring OSPF area ranges	412
Configuring OSPF area virtual interface	412

Topic	Page
Configuring OSPF neighbors	414
Show OSPF commands	415
Configuring port-based OSPF parameters	424
Showing OSPF port statistics	427
Configuring OSPF parameters for a VLAN	430
Showing OSPF parameters configured for VLANs	433

Roadmap of IP commands

The following roadmap lists some of the IP commands and their parameters. Use this list as a quick reference or click on any command or parameter entry for more information.

Command	Parameter
<code>config ip ospf</code>	<code>info</code> <code>admin-state <enable disable></code> <code>as-boundary-router <enable disable></code> <code>auto-vlink <enable disable></code> <code>default-metric [ethernet <value>] [fast-ethernet <value>] [gig-ethernet <value>]</code> <code>disable</code> <code>enable</code> <code>holddown <seconds></code> <code>router-id <ipaddr></code> <code>spf-run</code> <code>trap <enable disable></code>
<code>config ip ospf host-route <ipaddr></code>	<code>info</code> <code>create [metric <value>]</code> <code>delete</code>

Command	Parameter
<code>config ip ospf interface <ipaddr></code>	info add-message-digest-key <md5-key-id> md5-key <value> admin-status <enable disable> area <ipaddr> interface_type <if-type> authentication-key <authentication-key> authentication-type <auth-type> dead-interval <seconds> change-primary-md5-key <md5-key-id> create <if-type> delete delete-message-digest- key <md5-key-id> hello-interval <seconds> metric <metric> poll-interval <seconds> priority <priority> transit-delay <seconds>
<code>config ip ospf area <ipaddr></code>	info create delete import-summaries <true false> nssa <true false> stub <true false> stub-metric <stub-metric>
<code>config ip ospf area <ipaddr> range <ipaddr/mask></code>	info create advertise-mode <value> lsa-type <value>

Command	Parameter
	delete
	advertise-mode <mode>
	advertise-metric <cost>
config ip ospf area <ipaddr>	info
virtual-interface <nbr>	
	add-message-digest-key <md5-key-id> md5-key <value>
	authentication-key <authentication-key>
	authentication-type <auth-type>
	create
	dead-interval <seconds>
	delete
	delete-message-digest- key <md5-key-id>
	hello-interval <seconds>
	retransmit-interval <seconds>
	transit-delay <seconds>
config ip ospf area <ipaddr>	info
virtual-interface <nbr>	
	create <priority>
	priority<priority>
config ip ospf neighbor	info
	advertise-when-down <enable disable>
	area <ipaddr>
	authentication-key <string>
	authentication-type <auth-type>
	disable
	interface_type <if-type>
	dead-interval <seconds>
	enable

Command	Parameter
config ethernet <ports> ip ospf	hello-interval <seconds>
	metric <cost>
	priority <integer>
	info
	advertise-when-down <enable disable>
	area <ipaddr>
	authentication-key <string>
	authentication-type <auth-type>
	disable
	interface_type <if-type>
	dead-interval <seconds>
	enable
	hello-interval <seconds>
	metric <cost>
	priority <integer>
config vlan <vid> ip ospf	info
	advertise-when-down <enable disable>
	area <ipaddr>
	authentication-key <string>
	authentication-type <auth-type>
	disable
	dead-interval <seconds>
	poll-interval <seconds>
	enable
	hello-interval <seconds>
	metric <cost>
	priority <integer>

Command	Parameter
---------	-----------

<code>show ip ospf area</code>	
<code>show ip ospf ase [metric-type <value>]</code>	
<code>show ip ospf default-metric</code>	
<code>show ip ospf host-route</code>	
<code>show ip ospf ifstats [mismatch] [detail]</code>	
<code>show ip ospf info</code>	
<code>show ip ospf interface</code>	
<code>show ports error ospf [<ports>]</code>	
<code>show ip ospf int-timers</code>	
<code>show ip ospf lsdb [area <value>] [lsatype <value>] [lsid <value>] [adv_rtr <value>] [detail]</code>	
<code>show ip ospf neighbors</code>	
<code>show ip ospf stats</code>	
<code>show ip ospf stats</code>	
<code>show ports info ospf [<ports>]</code>	
<code>show ports stats ospf main [<ports>]</code>	
<code>show ports stats interface extended [<ports>]</code>	
<code>show vlan info ospf [<vid>]</code>	

Configuring OSPF global parameters

To configure global OSPF parameters for the Passport 8000 switch as follows, use the following commands:

config ip ospf followed by:	
info	Displays the current OSPF configuration on the switch (Figure 148).
admin-state <enable disable>	Globally enables or disables the OSPF administrative status. The default is disable.
as-boundary-router <enable disable>	Enables or disables the OSPF Autonomous System boundary router.
auto-vlink <enable/disable>	Enables or disables automatic creation of OSPF virtual links when required. The default is disable.
default-metric [ethernet <value>] [fast-ethernet <value>] [gig-ethernet <value>]	Sets the OSPF default metrics. The range is 1 to 65535. - ethernet <value> is for 10 Mb/s Ethernet (default is 100). - fast-ethernet <value> is for 100 Mb/s (fast) Ethernet (default is 10). - gig-ethernet <value> is for the Gigabit (gig) Ethernet (default is 1).
disable	Globally disables OSPF on the switch.
enable	Globally enables OSPF on the switch.
holddown <seconds>	Sets the OSPF holddown timer value in seconds. <seconds> is the range of seconds from 3 to 60; default is 10.
router-id <ipaddr>	Sets the OSPF router ID IP address. <ipaddr> is the IP address in dotted decimal format.
spf-run	This option is used to indicate the number of SPF calculations performed by OSPF.
trap <enable disable>	Enables or disables issuing traps relating to OSPF.

Figure 148 shows sample output for the **config ip ospf info** command.

Figure 148 config ip ospf info command output

```
Passport-8610# config ip ospf info

Sub-Context: clear config dump monitor show test trace
Current Context:

        admin-state : disabled
        router-id : 220.116.252.0
        version : 2
        area border : false
as-boundary-router : false
    ext lsa count : 0
    ext lsa chksum : 0
    orig new lsa : 0
    rx new lsa : 0
    default-metric :
        ethernet - 100
        fast-ethernet - 10
        gig-ethernet - 1
    auto-vlink : disable
    holddown : 10
    trap : disable
```

Configuring OSPF host route parameters

To configure OSPF host route parameters for your 8000 series switch, use the following command:

```
config ip ospf host-route <ipaddr>
```

where <ipaddr> is the address of the host router.

This command includes the following options:

config ip ospf host-route <ipaddr> followed by:	
info	Displays the current OSPF host-route configuration on the switch.
create [metric <value>]	Creates an OSPF host route for the IP address and Sets the metric (cost) for the host route. <i>metric</i> is between 1 and 65535.
delete	Deletes an OSPF host route for the IP address.

Configuration Example

The following configuration example uses the above command to:

- Creates an OSPF host route for the IP address
- Deletes an OSPF host route for the IP address

After configuring the parameters, use the info command to show a summary of the results.

Passport-8010:6#config ip ospf host-route

object <ipaddr> not entered

<ipaddr> = ip address {a.b.c.d}

Passport-8010:6#config ip ospf host-route 10.1.10.10

Passport-8010:6/config/ip/ospf/host-route/10.1.10.10# ?

Sub-Context:

Current Context:

create [metric <value>]

delete

info

Passport-8010:6/config/ip/ospf/host-route/10.1.10.10# create metric 10

Passport-8010:6/config/ip/ospf/host-route/10.1.10.10# info

Sub-Context:

Current Context:

```
create :  
delete : N/A  
metric : 10
```

Passport-8010:6/config/ip/ospf/host-route/10.1.10.10# delete

Passport-8010:6/config/ip/ospf/host-route/10.1.10.10# info

Sub-Context:

Current Context:

```
create : not created  
delete : not created  
metric : not created
```

Passport-8010:6/config/ip/ospf/host-route/10.1.10.10#

Configuring an OSPF interface

To configure an OSPF interface, use the following command:

```
config ip ospf interface <ipaddr>
```

The *ipaddr* is represented by an IP address {a.b.c.d}.

This command includes the following options:

config ip ospf interface <ipaddr> followed by:	
info	Displays OSPF characteristics for the interface (Figure 149).
add-message-digest-key <md5-key-id> md5-key <value>	Adds an md5 key to the interface. At most, two md5 keys can be configured to an interface. Multiple md5 key configurations are used for md5 transitions without bringing down an interface.
admin-status <enable disable>	Sets the state (enabled or disabled) of the OSPF interface.
area <ipaddr>	Sets the OSPF interface area. • <ipaddr> is a dotted-decimal notation to specify the area name. Note: The area name is not related to an IP address. You can use any value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).
interface_type <if-type>	Specifies the type of OSPF interface. <if-type> is the ospf interface type {broadcast nbma passive}
authentication-key <authentication-key>	Sets the authentication key for the OSPF interface. • <authentication-key> is a string that specifies the key in up to eight characters.

config ip ospf interface <ipaddr> followed by:	
authentication-type <auth-type>	Sets the OSPF authentication type for the interface. • <auth-type> is none, simple password, or MD5 authentication. If simple, all OSPF updates received by the interface must contain the authentication key specified by the interface authentication-key command. If MD5, they must contain the md5 key.
dead-interval <seconds>	Sets the OSPF dead interval for the interface. • <seconds> is the number of seconds the switch's OSPF neighbors should wait before assuming that this OSPF router is down. The range is from 1 to 2147483647. This value must be at least four times the hello interval value. The default is 40.
change-primary-md5-key <md5-key-id>	Changes the primary key used for encrypting outgoing packets. • <md5-key-id> is ID for the message-digest-key {1..255}
create <if-type>	Creates an OSPF interface. • <if-type> is the ospf interface type {broadcast nbma passive}
delete	Deletes an OSPF interface.
delete-message-digest-key <md5-key-id>	Deletes the specified md5 key ID from the configured md5 keys.
hello-interval <seconds>	Sets the OSPF hello interval for the interface. • <seconds> is the number of seconds between hello packets sent on this interface. The range is 1 to 65535. The default is 10. Note: When you change the hello interval values, you must save the configuration file and reboot the switch for the values to be restored and checked for consistency.
metric <metric>	Sets the OSPF metric for the interface. The switch advertises the metric in router link advertisements. • <metric> is the range 0 to 65535.
poll-interval <seconds>	Sets the polling interval for the OSPF interface in seconds. • <seconds> is between 1 and 2147483647.

config ip ospf interface <ipaddr> followed by:	
priority <priority>	Sets the OSPF priority for the interface during the election process for the designated router. The interface with the highest priority number is the designated router. The interface with the second-highest priority becomes the backup designated router. If the priority is 0, the interface cannot become either the designated router or a backup. The priority is used only during election of the designated router and backup designated router. The range is 0 to 255. The default is 1.
retransmit-interval <seconds>	Sets the retransmit interval for the OSPF interface, the number of seconds between link-state advertisement retransmissions. • <second> is an integer between 1 and 3600.
transit-delay <seconds>	Sets the transit delay time for the OSPF interface, the estimated time in seconds it takes to transmit a link-state update packet over the interface. • <seconds> is an integer between 1 and 3600.

Figure 149 shows sample output for the **config ip ospf interface info** command.

Figure 149 config ip ospf interface info command output

```
Passport-8610/config/ip/ospf/interface/130.1.1.1# info

Sub-Context:
Current Context:

    add-message-digest-key :
        admin-status : enabled
        area : 1.1.1.1
        authentication-key : password
        authentication-type : simple
    delete-message-digest-key : N/A
        hello-interval : 10
        dead-interval : 40
            metric : 200
        poll-interval : 120
            priority : 1
        retransmit-interval : 5
        transit-delay : 1
```

Configuring OSPF areas

To control the OSPF area parameters, use the following command:

```
config ip ospf area
```

where *ipaddr* is the address of an OSPF area. Use dotted-decimal notation to specify the area name.

You can use any value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).

The command includes the following options:

config ip ospf area <ipaddr> followed by:	
info	Displays OSPF area characteristics (Figure 150).
create	Creates an OSPF area.
delete	Deletes an OSPF area.
import-summaries <true false>	Sets the area's support for importing summary advertisements into a stub area. This field should be used only if the area stub is set to true.
nssa <true false>	Sets a not so stubby area (true or false). An NSSA prevents flooding of normal route advertisements into the area by replacing them with a default route.
stub <true false>	Sets the import external option for this area to be stub or not {true false}. A stub area has only one exit point (router interface) out of the area.
stub-metric <stub-metric>	Stub default metric for this stub area, which is the cost from 0 to 16777215. This is the metric value applied at the indicated type of service.

Figure 150 shows sample output for the **config ip ospf area info** command.

Figure 150 config ip ospf area info command output

```

Passport-8610# config ip ospf area 1.0.0.0 info

Sub-Context:
Current Context:

                create :
                delete : N/A
import-summaries : true
                nssa : false
                stub : false
                stub-metric : 1

```

Configuring OSPF area ranges

To control the OSPF area range parameters, use the following command:

```
config ip ospf area <ipaddr> range <ipaddr/mask>
```

where *ipaddr* is the identification of an OSPF area and *<ipaddr/mask>* is the IP address and subnet mask of the range.

This command includes the following options:

config ip ospf area <ipaddr> range <ipaddr/mask> followed by:	
info	Displays information about the OSPF area range settings.
create advertise-mode <value> lsa-type <value>	Creates an OSPF area range with the specified IP address and advertising mode.
delete	Deletes an OSPF area range.
advertise-mode <mode>	Changes the advertise-mode of the range. <i>mode</i> is the mode value {summarize suppress no-summarize}
advertise-metric <cost>	Changes the advertised metric cost value of the OSPF area range. <i>cost</i> is an integer value in the range 0 and 65535, which represents the metric cost value for the OSPF area range.

Configuring OSPF area virtual interface

To configure an OSPF area virtual interface, use the following command:

```
config ip ospf area virtual-interface
```

All of the commands have the following two required parameters:

- <ipaddr>* is the identification of an OSPF area in dotted-decimal notation. You can use any value for the OSPF area name (for example, 1.1.1.1 or 200.200.200.200).

- `virtual-interface <nbr>` is the OSPF router ID of the neighbor.

This command includes the following options:

config ip ospf area <ipaddr> virtual-interface <nbr> followed by:	
<code>info</code>	Displays current OSPF area virtual interface information.
<code>add-message-digest-key <md5-key-id> md5-key <value></code>	Adds an md5 key to the interface. At most, two md5 keys can be configured to an interface. Multiple md5 key configurations are used for md5 transitions without bringing down an interface.
<code>authentication-key <authentication-key></code>	Sets the authentication key. • <code>authentication-key</code> is a string that specifies the key in up to eight characters.
<code>authentication-type <auth-type></code>	Sets the OSPF authentication type for the OSPF area. • <code>auth-type</code> is none, simple password, or MD5 authentication. If simple, all OSPF updates received by the interface must contain the authentication key specified by the area authentication-key command. If MD5, they must contain the md5 key.
<code>create</code>	Creates a virtual interface area identifier.
<code>dead-interval <seconds></code>	Sets the dead interval for the virtual interface, the number of seconds that a router's hello packets have not been seen before its neighbors declare the router down. • <code><seconds></code> is an integer between 1 and 214783647. This value must be at least four times the hello interval value. The default is 60.
<code>delete</code>	Deletes the virtual interface.
<code>delete-message-digest-key <md5-key-id></code>	Deletes the specified md5 key ID from the configured md5 keys.
<code>hello-interval <seconds></code>	Sets the hello interval on the virtual interface for the length of time (in seconds) between the hello packets that the router sends on the interface. • <code><seconds></code> is a value between 1 and 65535. The default is 10.

config ip ospf area <ipaddr> virtual-interface <nbr> followed by:	
retransmit-interval <seconds>	Sets the retransmit interval for the virtual interface, the number of seconds between link-state advertisement retransmissions. • <seconds> is an integer between 1 and 3600.
transit-delay <seconds>	Sets the transmit delay for the virtual interface, the estimated number of seconds it takes to transmit a link-state update over the interface. • <seconds> is an integer between 1 and 3600.



Note: Both sides of the OSPF connection must use the same authentication type and key.

Configuring OSPF neighbors

To create, delete, and obtain information about an OSPF neighbor and to set priorities for an OSPF neighbor, use the following command:

`show ip ospf neighbor`

This command includes the following options:

config ip ospf neighbor followed by:	
info	Displays information about the OSPF neighbor settings.
create <priority>	Creates an OSPF neighbor and assigns a priority level. • <priority> is a value between 0 and 255
delete	Deletes an OSPF neighbor.
priority<priority>	Changes the priority level of the neighbor. • <priority> is a value between 0 and 255.

Show OSPF commands

This section describes how to display OSPF configuration information.

This section includes the following topics:

- [“Showing OSPF areas,” next](#)
- [“Showing OSPF ASE link state advertisements” on page 416](#)
- [“Showing OSPF default metric information” on page 417](#)
- [“Showing OSPF host route configuration” on page 417](#)
- [“Showing OSPF interface statistics” on page 417](#)
- [“Showing OSPF information” on page 418](#)
- [“Showing OSPF interface information” on page 419](#)
- [“Showing OSPF interface timer settings” on page 420](#)
- [“Showing the OSPF link state database table” on page 420](#)
- [“Showing OSPF neighbors” on page 423](#)
- [“Showing OSPF range statistics” on page 423](#)

Showing OSPF areas

To display information about OSPF area parameters, use the following command:

```
show ip ospf area
```

[Figure 151](#) shows sample output for this command.

Figure 151 show ip ospf area command output

```
Passport-8610/show/ip/ospf# area
```

Ospf Area					
AREA_ID	STUB_AREA	NSSA	IMPORT_SUM	ACTIVE_IFCNT	
0.0.0.0	false	false	true	0	
1.1.1.1	false	false	true	1	
STUB_COST	SPF_RUNS	BDR_RTR_CNT	ASBDR_RTR_CNT	LSA_CNT	LSACK_SUM
0	8	0	0	0	0
1	8	0	1	3	128484

Showing OSPF ASE link state advertisements

To display the OSPF Autonomous System External (ASE) link state advertisements, use the following command:

```
show ip ospf ase [metric-type <value>]
```

Information is displayed for all metric types or for the type specified.

[Figure 152](#) shows sample output for this command.

Figure 152 show ip ospf ase command output

```
Passport-8610/show/ip/ospf# ase
```

Ospf AsExternal Lsas					
LSTYPE	LINKSTATEID	ADV_ROUTER	E_METRIC	ASE_FWD_ADDR	AGE
SEQ_NBR	CSUM				
AsExternal	199.100.1.0	45.57.236.0	0 10	0.0.0.0	608
0x80000001	0x4fb8				

Showing OSPF default metric information

To display OSPF default metric information for each type of port, use the following command:

```
show ip ospf default-metric
```

Figure 153 shows sample output for this command.

Figure 153 show ip ospf default-metric command output

```
Passport-8606:6# show ip ospf default-metric

=====
                                     Ospf Default Metric
=====
10MbpsPortDefaultMetric: 100
100MbpsPortDefaultMetric: 10
1000MbpsPortDefaultMetric: 1
10000MbpsPortDefaultMetric: 1
```

Showing OSPF host route configuration

To display the OSPF host route configuration, including host IP address, type of service, and the metric used, use the following command:

```
show ip ospf host-route
```

Showing OSPF interface statistics

To display OSPF interface statistics where the parameter mismatch is the number of times the area ID is not matched.

```
show ip ospf ifstats [mismatch] [detail]
```

Figure 154 show sample output for the this command.

Figure 154 show ip ospf ifstats command output

Passport-8610/show/ip/ospf# ifstats

Ospf Interface Statistics										
INTERFACE	---HELLOS---RX	---HELLOS---TX	---DBS---RX	---DBS---TX	-LS REQ--RX	-LS REQ--TX	--LS UDP---RX	--LS UDP---TX	--LS ACK---RX	--LS ACK---Tx
130.1.1.1	86	85	8	3	1	1	13	1	2	1

Showing OSPF information

To display the current OSPF settings for the switch, use the following command:

show ip ospf info

Figure 155 shows sample output for the command.

Figure 155 show ip ospf info command output

Passport-8610/show/ip/ospf# info

Ospf General	
RouterId:	45.57.0.0
AdminStat:	enabled
VersionNumber:	2
AreaBdrRtrStatus:	false
ASBdrRtrStatus:	false
ExternLsaCount:	1
ExternLsaChecksumSum:	20408 (0x4fb8)
TOSSupport:	0
OriginateNewLsas:	3
RxNewLsas:	12
TrapEnable:	false
AutoVirtLinkEnable:	false
SpfHoldDownTime:	10

Showing OSPF interface information

To display information about the OSPF interface, use the following command:

```
show ip ospf interface
```

Figure 156 shows sample output for the **show ip ospf interface** command.

Figure 156 show ip ospf interface command output

Passport-8606:6# show ip ospf interface

Ospf Interface							
INTERFACE	AREAID	ADM	IFST	MET	PRIO	DR/BDR	TYPE AUTHTYPE
66.77.88.99 message-digest	0.0.0.0	dis	Down	10	1	0.0.0.0	brdc
						0.0.0.0	
200.1.1.1	0.0.0.0	en	Down	10	1	0.0.0.0	brdc none
						0.0.0.0	
111.111.111.111 message-digest	22.22.22.22	en	Down	10	1	0.0.0.0	brdc
						0.0.0.0	
Ospf Virtual Interface							
AREAID	NBRIPADDR	STATE		AUTHTYPE			
22.22.22.22	6.7.8.9			none			

Showing OSPF interface timer settings

To display OSPF interface timer settings, use the following command:

```
show ip ospf int-timers
```

Figure 157 show sample output for this command.

Figure 157 show ip ospf int-timers command output

Passport-8606:6# show ip ospf int-timers

Ospf Interface Timer						
INTERFACE	AREAID	TRANSIT DELAY	RETRANS INTERVAL	HELLO INTERVAL	DEAD INTERVAL	POLL INTERVAL
66.77.88.99	0.0.0.0	1	5	10	40	120
200.1.1.1	0.0.0.0	1	5	10	40	120
111.111.111.111	22.22.22.22	1	5	10	40	120

Showing the OSPF link state database table

To display the OSPF link state database (lsdb) table, use the following command:

```
show ip ospf lsdb [area <value>] [lsatype <value>] [lsid  
<value>] [adv_rtr <value>] [detail]
```

You can optionally specify an area string, link state advertisement type (0 to 5), link state ID, or advertising router. Adding the `detail` option to the command provides more details.

Figure 158 shows sample output without any variables for the `show ip ospf lsdb` command.

Figure 158 show ip ospf lsdb command output

```

Passport-8610/show/ip/ospf# lsdb

=====
                        Ospf Lsdb
=====

  Router Lsas in Area 1.1.1.1

LSTYPE      LINKSTATEID      ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----
Router      45.57.0.0                45.57.0.0        1028 0x80000003 0x8be3
Router      45.57.236.0              45.57.236.0      586  0x8000000a 0xa402

  Network Lsas in Area 1.1.1.1

LSTYPE      LINKSTATEID      ADV_ROUTER      AGE  SEQ_NBR      CSUM
-----
Network     130.1.1.2          45.57.236.0      1034 0x80000001 0xc5ff
  
```

Figure 159 shows partial output of the `show ip ospf lsdb` command with the `detail` option.

Figure 159 show ip ospf lsdb detail command output

```

Passport-8610/show/ip/ospf# lsdb detail

=====
                        Ospf Lsdb
=====

Router Link LSA :
Area: 1.1.1.1 (0x1010101)
Age : 1123
Opt: true (External Routing Capability)
Type: 1
LsId: 45.57.0.0 (0x2d390000)
Rtr: 45.57.0.0
Seq : -2147483645 (0x80000003)
Csum: 35811 (0x8be3)
Len : 36
ABR: false
ASBR: false
Vlnk: false (endpoint of active Vlink)
#Lnks: 1
[1]
Id : 130.1.1.2 (0x82010102)
Data: 130.1.1.1 (0x82010101)
Type: (conn-to-transmit-net) (Id=DR-Addr, Data=Rtr-Addr)
Met : 10

Router Link LSA :
Area: 1.1.1.1 (0x1010101)
Age : 697
.
.
.

Network Link LSA :
Area: 1.1.1.1 (0x1010101)
Age : 1156
Opt: true (External Routing Capability)

```


Showing OSPF neighbors

To display OSPF neighbors configuration information, use the following command:

```
show ip ospf neighbors
```

Figure 160 show sample command output for the **show ip ospf neighbors** command.

Figure 160 show ospf neighbors command output



Showing OSPF range statistics

To display the OSPF range statistics, including area ID, range network address, range subnet mask, range flag, and LSDB type, use the following command:

```
show ip ospf stats
```

Figure 161 shows sample output for this command.

Figure 161 show ip ospf stats command output

```
Passport-8606:6# show ip ospf stats

=====
                                     Ospf Statistics
=====

      NumBufAlloc: 0
      NumBufFree: 0
NumBufAllocFail: 0
      NumBufFreeFail: 0
      NumTxPkt: 0
      NumRxPkt: 0
      NumTxDropPkt: 0
      NumRxDropPkt: 0
      NumRxBadPkt: 0
      NumSpfRun: 0
      LastSpfRun: 0x0
      LsdbTblSize: 0
NumAllocBdDDP: 0
NumFreeBdDDP: 0
      NumBadLsReq: 0
NumSeqMismatch: 0
```

Configuring port-based OSPF parameters

To configure port-based OSPF parameters for specified ports, use the following command:

```
config ethernet <ports> ip ospf
```

where:

ports is the port you are configuring.

This command includes the following options:

<code>config ethernet <ports> ip ospf</code> followed by:	
<code>info</code>	Displays OSPF characteristics on the port (Figure 162).
<code>advertise-when-down</code> <code><enable disable></code>	If enabled, the network on this interface is advertised as up, even if the port is down. The default is disabled. Note: When you configure a port without any link and enable <code>advertise-when-down</code>, the route is not advertised until the port is active. Then the route is advertised even when the link is down. To disable advertising based on link status, this parameter should be disabled.
<code>area <ipaddr></code>	Sets the OSPF identification number for the area, typically formatted as an IP address.
<code>authentication-key <string></code>	Sets the authentication key for the port (OSPF interface). <code>string</code> specifies the key as a simple password with eight characters.
<code>authentication-type</code> <code><auth-type></code>	Sets the OSPF authentication type for the port: none, simple password, or MD5 authentication. If simple, all OSPF updates received by the interface must contain the authentication key specified by the area <code>authentication-key</code> command. If MD5, they must contain the md5 key.
<code>disable</code>	Disables OSPF on the port.
<code>interface_type <if-type></code>	Specifies the type of OSPF interface <code><if-type></code> is the ospf interface type {broadcast nbma passive}
<code>dead-interval <seconds></code>	Sets the router OSPF dead interval—the number of seconds the switch's OSPF neighbors should wait before assuming that the OSPF router is down. <code><seconds></code> is a value from 1 to 2147836437; the default is 40. The value must be at least four times the hello interval.
<code>enable</code>	Enables OSPF on the port.

config ethernet <ports> ip ospf followed by:	
hello-interval <seconds>	Sets the OSPF hello interval, which is the number of seconds between hello packets sent on this interface. • <seconds> is a value from 1 to 65535. The default is 10.
metric <cost>	Sets the OSPF metric associated with this interface and advertised in router link advertisements. • <cost> is in the range from 0 to 65535; the default is 0.
priority <integer>	Sets the OSPF priority for the port during the election process for the designated router. The port with the highest priority number is the best candidate for the designated router. If you set the priority to 0, the port cannot become either the designated router or a backup designated router. • <integer> is between 0 and 255. The default is 1.



Note: Both sides of the OSPF connection must use the same authentication type and key.

[Figure 162](#) shows sample output for the `config ethernet ip ospf info` command.

Figure 162 config ethernet ip ospf info command output

```
Passport-8610# config ethernet 1/2 ip ospf info#  
  
Sub-Context:  
Current Context:  
  
Port 1/2 :  
    advertise-when-down : disable  
        ospf : disable  
        if-type : broadcast  
    hello-interval : 10  
    dead-interval : 40  
    poll-interval : 120  
        priority : 1  
        metric : 0  
    authentication-type : none  
    authentication-key :  
        area : 0.0.0.0
```

Showing OSPF port statistics

This section describes commands that display OSPF parameters and statistics for a port or all ports.

This section includes the following topics:

- [“Showing OSPF errors on a port,” next](#)
- [“Showing OSPF configuration settings on a port” on page 428](#)
- [“Showing basic OSPF information on a port” on page 429](#)
- [“Showing extended OSPF information” on page 430](#)


Showing OSPF errors on a port

To display extended information about OSPF errors for the specified port or for all ports, use the following command:

```
show ports error ospf [<ports>]
```

[Figure 163](#) shows sample output for the **show ports error ospf** command.

Figure 163 show ports error ospf command output



```
Passport-8610# show ports error ospf
```

Port Ospf Error									
PORT	VERSION	AREA	AUTHTYPE	AUTH	NET_MASK	HELLOINT	DEADINT	OPTION	
NUM	MISMATCH	MISMATCH	MISMATCH	FAILURES	MISMATCH	MISMATCH	MISMATCH	MISMATCH	
9/1	0	0	0	0	0	0	0	0	
9/13	0	0	0	0	0	0	0	0	

Showing OSPF configuration settings on a port

To display information about the OSPF parameters of the specified port or all ports, use the following command:

```
show ports info ospf [<ports>]
```

[Figure 164](#) show sample output for the **show ports info ospf** command.

Figure 164 show ports info ospf command (partial output)

Passport-8610# show ports info ospf

Port Ospf								
PORT	HELLO	RTRDEAD	OSPF					
NUM	ENABLE	INTVAL	INTVAL	PRIORITY	METRIC	AUTHTYPE	AUTHKEY	AREA_ID
9/1	false	10	40	1	0	none		0.0.0.0
9/2	true	10	40	1	0	none		1.0.0.0
9/3	false	10	40	1	0	none		0.0.0.0
9/4	false	10	40	1	0	none		0.0.0.0
9/5	false	10	40	1	0	none		0.0.0.0
9/6	false	10	40	1	0	none		0.0.0.0
9/7	false	10	40	1	0	none		0.0.0.0
9/8	false	10	40	1	0	none		0.0.0.0
9/9	false	10	40	1	0	none		0.0.0.0

Showing basic OSPF information on a port

To display basic OSPF information about the specified port or for all ports, use the following command:

```
show ports stats ospf main [<ports>]
```

Figure 165 shows sample output for this command.

Figure 165 show ports stats ospf main command output

Passport-8610# show ports stats ospf main

Port Stats Ospf						
PORT_NUM	RX_HELLO	TX_HELLO	RXDB_DESCR	TXDB_DESCR	RXLS_UPDATE	TXLS_UPDATE
9/2	0	0	0	0	0	0
9/3	0	0	0	0	0	0


Showing extended OSPF information

To display extended OSPF information about the specified port or for all ports, use the following command:

```
show ports stats interface extended [<ports>]
```

Figure 166 shows sample output for this command.

Figure 166 show ports stats interface extended command output



Configuring OSPF parameters for a VLAN

To configure OSPF parameters for a specified VLAN, use the following command:

```
config vlan <vid> ip ospf
```

where:

vid is a unique integer value in the range 1 and 4094 that identifies the VLAN you are configuring.

The command include the following options:

config vlan <vid> ip ospf followed by:	
info	Displays OSPF characteristics on the VLAN (Figure 167).
advertise-when-down <enable disable>	If enabled, the network on this interface is advertised as up, even if no ports in the VLAN are active. The default is disabled. Note: When you create a VLAN with no active ports and enable advertise-when-down, the route is not advertised until a port is active. Then the route is advertised even when the link is down. To disable advertising based on link status, disable this parameter.
area <ipaddr>	Sets the OSPF interface area ID for the VLAN.
authentication-key <string>	Sets the authentication key for the VLAN. • <string> is key of a string with up to eight characters.
authentication-type <auth-type>	Sets the OSPF authentication type for the VLAN. • <auth-type> is none, simple password, or MD5 authentication. If simple, all OSPF updates received by the VLAN must contain the authentication key specified by the area authentication-key command. If MD5, they must contain the md5 key.
disable	Disables OSPF on the VLAN.
dead-interval <seconds>	Sets the OSPF dead interval for the VLAN, the number of seconds the switch's OSPF neighbors should wait before assuming that this OSPF router is down. • <seconds> is the range from 1 to 2147483647. This value must be at least four times the hello interval value. The default is 40.
poll-interval <seconds>	Sets the OSPF poll interval for the VLAN, the number of seconds the switch's OSPF neighbors should wait before sending the next poll. <seconds> is the range from 1 to 2147483647.
enable	Enables OSPF on the VLAN.
hello-interval <seconds>	Sets the OSPF hello interval for a VLAN, the number of seconds between hello packets sent on the VLAN. • <seconds> is the range from 1 to 65535. The default is 10.

config vlan <vid> ip ospf followed by:	
metric <cost>	Sets the OSPF metric for the VLAN. The switch advertises the metric in router link advertisements. • <i><seconds></i> is the range from 0 to 65535. The default is 0.
priority <integer>	Sets the OSPF priority for the VLAN during the election process for the designated router. The VLAN with the highest priority number is the best candidate for the designated router. If the priority is 0, the VLAN cannot become either the designated router or a backup. The priority is used only during election of the designated router and backup designated router. • <i><integer></i> is the range from 0 to 255. The default is 1.



Note: Both sides of the OSPF connection must use the same authentication type and key.

Figure 167 shows the output of the **config vlan ip ospf info** command.

Figure 167 config vlan ip ospf info command output

```
Passport-8610# config vlan 2 ip ospf info

Sub-Context: clear config dump monitor show test trace
Current Context:

    advertise-when-down : disable
                ospf : enable
    hello-interval : 10
    dead-interval : 40
        priority : 1
        metric : 10
authentication-type : none
authentication-key :
                area : 0.0.0.0
```


Showing OSPF parameters configured for VLANs

To display OSPF parameters configured for all VLANs or the specified VLAN, use the following command:

```
show vlan info ospf [<vid>]
```

Figure 168 shows sample output for this command.

Figure 168 show vlan info ospf command output

Passport-8610# show vlan info ospf

=====								
Vlan Ospf								
=====								
VLAN	HELLO	RTRDEAD	DESIGRTR					
ID	ENABLE	INTERVAL	INTERVAL	PRIORITY	METRIC	AUTHTYPE	AUTHKEY	AREAID

1	false	10	40	1	0	none		0.0.0.0
2	true	10	40	1	10	none		0.0.0.0
3	false	10	40	1	0	none		0.0.0.0
4	false	10	40	1	0	none		0.0.0.0

Chapter 11

Configuring VRRP using Device Manager

End stations are often configured with a static default gateway IP address. Loss of the default gateway router can have catastrophic results. Virtual Router Redundancy Protocol (VRRP), RFC 2338, is designed to eliminate this single point of failure in a routed environment by introducing the concept of a virtual IP address (transparent to users) shared between two or more routers connecting the common subnet to the enterprise network. With the virtual IP address as the default gateway on end hosts, VRRP provides a dynamic default gateway redundancy in the event of a failure.

The current implementation of VRRP allows you to have one active master switch per IP subnet. All other VRRP interfaces in a network are in backup mode.

On a Passport 8000 Series switch, you cannot directly check or set the virtual IP address on the standby CPU module. In order to check or set the virtual IP address on the standby CPU, you must configure the virtual IP address on the master CPU, save it to the config.cfg file, and then copy that file to the standby CPU module

If you have VRRP and IP routing protocols (for example, OSPF) configured on the same IP physical interface, selecting the interface address as the VRRP virtual IP address (logical IP address) is not supported. Use a separate dedicated IP address for VRRP.

The timer delays the preemption of the master over the backup, when the master becomes available. This timer is called the Hold Down Timer, and it has a default value of 0 second. Nortel Networks recommends that you set all your routers to the identical number of seconds for the Hold Down Timer.

In addition, you can manually force the preemption of the master over the backup before the delay timer expires.

This chapter describes configuring and managing VRRP in Device Manager. Use the Hold Down Timer to modify the behavior of the VRRP failover mechanism by allowing the router enough time to detect and update the OSPF or RIP routes.

- For conceptual information about VRRP, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter includes the following topics:

Topic	Page
Configuration prerequisites	436
VRRP and Split-MLT	437
Configuring VRRP for the interface	437
Configuring VRRP secondary features	440
Configuring VRRP on a port	442
Configuring VRRP on a VLAN (or brouter port)	444
Configuring Fast Advertisement Interval on a Port	447
Configuring Fast Advertisement Interval on a VLAN	447

Configuration prerequisites

Before your VRRP configurations can take effect, you must perform the following step:

- Assign an IP address to the interface

VRRP and Split-MLT

The current implementation of VRRP allows you to have one active master switch per IP subnet. All other VRRP interfaces in a network are in backup mode.

A deficiency occurs when Split-MLT is used. Users who, access switches which are aggregated into two Split-MLT switches, send their traffic load shared (based on the MLT traffic distribution algorithm) on all uplinks towards the Split-MLT aggregation switches.

VRRP however only has one active routing interface enabled. All other interfaces are in backup mode and therefore in standby mode. In this case, all traffic is forwarded over the Inter Switch Trunk (IST) link towards the master VRRP switch. Potentially there will be not enough bandwidth on the IST to carry all the aggregated riser traffic.

An enhancement in VRRP overcomes this deficiency. The enhancement makes sure that the IST trunk is not used in such case.

Configuring VRRP for the interface

You can manage and configure VRRP parameters for the routing interface.

To configure the VRRP Interface:


- 1 From the Device Manager menu bar, choose IP Routing > VRRP.

The VRRP dialog box opens with the Globals tab displayed [Figure 169 on page 438](#).

- 2 Click the Interface tab.

The Interface tab opens ([Figure 169](#)).

Figure 169 VRRP dialog box—Interface tab



- 3 Select a VLAN row and make the appropriate changes.
- 4 Click Apply.

Table 39 describes the fields in the Interface tab.

Table 39 Interface tab fields

Field	Description
Vrid	A number that uniquely identifies a virtual router on a given VRRP router. The virtual router acts as the default router for one or more assigned addresses (1 to 255).
Interface	Interface of the VRRP router.
IpAddr	The assigned IP addresses that a virtual router is responsible for backing up.
VirtualMacAddr	MAC address of the virtual router interface.
State	The state of the virtual router interface: • initialize—waiting for a startup event • backup—monitoring availability and state of the master router • master—functioning as the forwarding router for the virtual router IP address(es)
Control	Whether VRRP is enabled or disabled for the port (or VLAN).
Priority	Priority value to be used by this VRRP router. Set a value from 1 to 255, where 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.

Table 39 Interface tab fields (continued)

Field	Description
MasterIpAddr	The IP address of the physical interface of the master virtual router that has the responsibility of forwarding packets sent to the virtual IP address(es) associated with the virtual router.
FastAdvertisementEnable	Enables or disables the Fast Advertisement Interval. When disabled the regular advertisement interval is used. Default is disable.
AdvertisementInterval	The time interval (in seconds) between sending advertisement messages. Set from 1 to 255 seconds with a default of 1 second. Only the master router sends advertisements.
VirtualRouterUpTime	The time interval (in hundredths of a second) since the virtual router was initialized.
FastAdvertisementInterval	Sets the Fast Advertising Interval, the time interval between sending VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. The values must be entered in multiples of 200 milliseconds.
CriticalIpAddrEnable	Sets the IP interface on the local router to enable or disable the backup.
CriticalIPAddr	An IP interface on the local router configured so that a change in its state would cause a role switch in the virtual router (for example, from master to backup) in case the interface stops responding.

Configuring VRRP secondary features

You can manage and configure VRRP parameters for the routing secondary features.

To configure the VRRP secondary features:

- 1 From the Device Manager menu bar, choose IP Routing > VRRP.

The VRRP dialog box opens with the Globals tab displayed (Figure 170).

- 2 Click the Secondary Features tab.

The Secondary Features tab opens (Figure 170).

Figure 170 VRRP dialog box—Secondary Feature tab



- 3 Click the HoldDownTimer text box, and enter the desired number of seconds for the timer.

The HoldDownState field displays active when the Hold Down Timer is counting down and preemption will occur; the text box displays dormant when preemption is not pending. When the Hold Down Timer is active, the HoldDownTimeRemaining field displays the seconds remaining before preemption.

Use the OperAction field to manually override the delay timer and to force preemption. When you click the heading, an arrow appears. Click the text box, and a list opens. Choose preemption to preempt the timer, or choose none to allow the timer to keep working.

Use the BackUpMaster field to enable or disable the backup master feature.

- 4 Click Apply.

Table 40 describes the fields in the Secondary Feature tab.

Table 40 Secondary Feature tab fields

Field	Description
VrId	A number that uniquely identifies a virtual router on a given VRRP router. The virtual router acts as the default router for one or more assigned addresses (1 to 255).
HoldDownTimer	The time interval (in seconds) a router is delayed for the following conditions: - The VRRP holddown timer is executed when the switch transitions from Init to backup to master. This occurs only on a switch bootup. - The VRRP holddown timer is NOT executed under the following condition: In a non-bootup condition the Backup switch will become master after the Master Downtime Interval. (3 * hello interval), if the master VR goes down. - The VRRP holddown timer also applies to the VRRP BackupMaster feature.
HoldDownState	Status is active when the Hold Down Timer is counting down and preemption will occur; the text box displays dormant when preemption is not pending.
HoldDownTimeRemaining	The seconds remaining before preemption.
OperAction	Use the action list to manually override the delay timer and force preemption: - preemption—preempt the timer - none—allow the timer to keep working
BackUpMaster	Indicates if the VRRP backup master is enabled or disabled. This option is not recommended for non Split-MLT ports.
BackUpMastrState	Displays the BackupMaster operational state. When VRRP is enabled on a switch in a master state the BackUpMaster state is DOWN. When VRRP is enabled on a switch that is in a backup state the BackUpMaster state is UP. States: - up: in BackupMaster state - down: original state


Configuring VRRP on a port

You can configure VRRP on a port or brouter port (or on a VLAN) only if the port or brouter port (or VLAN) is assigned an IP address.

To configure VRRP parameters on a port:

- 1 From the Device Manager menu bar, choose Edit > Port.
The Port dialog box opens with the Interface tab displayed (Figure 59).
- 2 Click the VRRP tab.
The VRRP tab opens (Figure 171).

Figure 171 Port dialog box—VRRP tab



- 3 Click Insert.
The Port, Insert VRRP dialog box opens (Figure 172).

Figure 172 Port, Insert VRRP dialog box

10.10.40.27 - Port 1/1, Insert VRRP

Vrid: 1..255

IpAddr:

Control: ☒ enabled ☐ disabled

Priority: 1..254

FastAdvertisementEnable: ☐ enabled ☒ disabled

AdvertisementInterval: 1..255 (sec)

FastAdvertisementInterval: 200..1000 (milliseconds)

☐ CriticallyIpAddrEnable

CriticallyIpAddr:

HoldDownTimer: 0..21600 (sec)

OperAction: ☒ none ☐ preemptHoldDownTimer

BackUpMaster: ☐ enabled ☒ disabled

Table 41 describes the Port, Insert VRRP dialog box fields.

Table 41 Port, Insert VRRP dialog box fields

Field	Description
Vrid	A number that uniquely identifies a virtual router on a given VRRP router. The virtual router acts as the default router for one or more assigned addresses (1 to 255).
IpAddr	IP address of the virtual router interface.
Control	Whether VRRP is enabled or disabled for the port or VLAN.
Priority	Priority value to be used by this VRRP router. Set a value from 1 to 255, where 255 is reserved for the router that owns the IP addresses associated with the virtual router. The default is 100.
FastAdvertisementEnable	Enables or disables the Fast Advertisement Interval. When disabled the regular advertisement interval is used. Default is disable.
AdvertisementInterval	The time interval (in seconds) between sending advertisement messages. Set from 1 to 255 seconds with a default of 1 second. Only the master router sends advertisements.

Table 41 Port, Insert VRRP dialog box fields (continued)

Field	Description
FastAdvertisementInterval	Sets the Fast Advertising Interval, the time interval between sending VRRP advertisement messages. The interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. The values must be entered in multiples of 200 milliseconds.
CriticalIpAddrEnable	Sets the IP interface on the local router to enable or disable the backup.
CriticalIpAddr	Indicates if a user-defined critical IP address should be enabled. There is no effect if a user-defined IP address does not exist. No—use the default IP address (0.0.0.0)
HoldDownTimer	The time interval (in seconds) a router is delayed for the following conditions: - The VRRP holddown timer is executed when the switch transitions from Init to backup to master. This occurs only on a switch bootup. - The VRRP holddown timer is NOT executed under the following condition: In a non-bootup condition the Backup switch will become master after the Master Downtime Interval. (3 * hello interval), if the master VR goes down. - The VRRP holddown timer also applies to the VRRP BackupMaster feature.
OperAction	Use the action list to manually override the delay timer and force preemption: - preemptHoldDownTimer—preempt the timer - none—allow the timer to keep working
BackUpMaster	Enables or disables the VRRP backup master feature. This option is only supported on Split-MLT ports.

Configuring VRRP on a VLAN (or brouter port)


Before you configure VRRP on a VLAN you must first set VRRP globally. You can configure VRRP on a VLAN or brouter port only if the port or VLAN is assigned an IP address.

To configure VRRP parameters on a VLAN or brouter port:

- 1 From the Device Manager menu bar, choose VLAN > VLANs.

The VLAN dialog box opens with the Basic tab displayed (Figure 173).

Figure 173 VLAN dialog box—Basic tab




Id	Name	Color Identifier	Type	StgId	PortMembers	ActiveMembers	StaticMembers	NotAllowToJoin	ProtocolId	SubnetAddr	SubnetMask
4	VLAN-4	blue	byPort	1	3/5	3/5			none	N/A	N/A
5	VLAN-5	yellow	byPort	1	3/5	3/5			none	N/A	N/A
3	VLAN-3	green	byPort	1	1/33-1/48, 2/1-2/2	1/33-1/48, 2/1-2/2			none	N/A	N/A
1	Default	white	byPort	1	1/1-1/32, 3/1-3/5, 4/1-4/8	1/1-1/32, 3/1-3/5, 4/1-4/8			none	N/A	N/A
2	VLAN-2	red	byPort	1					none	N/A	N/A

5 row(s)

- 2 Select a VLAN
- 3 Click IP.
The IP, VLAN dialog box opens with the IP Address tab displayed.
- 4 Select the VRRP tab.
The VRRP tab opens (Figure 174).

Figure 174 IP, VLAN dialog box—VRRP tab



- 5 Click Insert.
The IP, VLAN, Insert VRRP dialog box opens (Figure 175).

Figure 175 IP, VLAN, Insert VRRP dialog box




Table 41 on page 443 describes the IP, VLAN, Insert VRRP dialog box fields.

Configuring Fast Advertisement Interval on a Port

To configure the Fast Advertisement Interval:

- 1 Select a port.
- 2 From the Device Manager menu bar, choose Edit > Port > VRRP.
The Port dialog box opens with the VRRP tab displayed ([Figure 171 on page 442](#)).
- 3 Click Insert.
The Port, Insert VRRP dialog box opens ([Figure 172 on page 443](#)).
- 4 Click, Fast Advertisement Enable.
Set to enable.
- 5 Enter a Fast Advertisement Interval value.
You must set this value using multiples of 200 milliseconds.
- 6 Click Insert.
The new entry appears in the VRRP tab of the Port dialog box.

[Table 41 on page 443](#) describes the VRRP Insert fields.

Configuring Fast Advertisement Interval on a VLAN

To configure the Fast Advertisement Interval:

- 1 Select a port.
- 2 From the Device Manager menu bar, choose VLAN > VLANs.
The VLAN dialog box opens with the Basic tab displayed ([Figure 173 on page 445](#)).
- 3 Select a VLAN, click IP > VRRP.

The IP, VLAN dialog box opens with the VRRP tab displayed ([Figure 174 on page 446](#)).

- 4** Click Insert.

The IP, VLAN, Insert VRRP dialog box opens ([Figure 175 on page 446](#)).

- 5** Click, Fast Advertisement Enable.

Set to enable.

- 6** Enter a Fast Advertisement Interval value.

You must set the value using multiples of 200 milliseconds.

- 7** Click Insert.

The new entry appears in the VRRP tab of the IP, VLAN dialog box.

Refer to [Table 41 on page 443](#) for a description of the VRRP Insert fields.

Chapter 12

Configuring IP VRRP using the CLI

This chapter describes the VRRP commands that allow you to configure VRRP on a port or VLAN.

- For conceptual information about VRRP, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter includes the following topics:

Topic	Page
Roadmap of IP commands	450
Configuring VRRP on a port	450
Showing VRRP port information	453
Configuring VRRP on a VLAN	454
Showing vlan info vrrp extended command	457
Showing VRRP interface information	458
Dependencies and rules	459

Roadmap of IP commands

The following roadmap lists some of the IP commands and their parameters. Use this list as a quick reference or click on any command or parameter entry for more information.

Command	Parameter
<code>show ip vrrp info [vrid <vrid>] [ip <ipaddr>]</code>	
<code>show ip vrrp info [<vrid>] [<ipaddr>]</code>	
<code>show vlan info vrrp extended [<vid>]</code>	

Configuring VRRP on a port

Use the following command to configure VRRP on a port:

```
config ethernet <ports> ip vrrp <vrid>
```

Where:

- *ports* specifies the ports for which you are entering the command in the form port list {slot/port[-slot/port][, ...]}.
- *vrid* is a unique integer value that represents the virtual router ID in the range 1 and 255. The virtual router acts as the default router for one or more assigned addresses.

The commands use the following options:

config ethernet <ports> ip vrrp <vrid> followed by:	
info	Displays the current port VRRP configuration (Figure 176).
action <action choice>	Use the action choice to manually override the Hold Down Timer and force preemption. <i>action choice</i> can be set to preemption to preempt the timer or set to none to allow the timer to keep working.
address <ipaddr>	Sets the IP address of the physical interface of the VRRP that has the responsibility of forwarding packets sent to the virtual IP address(es) associated with the virtual router. <i>ipaddr</i> is the IP address of the master VRRP.
adver-int <seconds>	Sets the advertising interval, the time interval between sending VRRP advertisement messages. <i>seconds</i> the interval can be between 1 and 255 seconds, and it must be the same on all participating routers. The default is 1.
backup-master <enable disable>	Enables or disables the VRRP backup master. This option is supported only on SMLT ports.
critical-ip <ipaddr>	Sets the critical IP address for VRRP. <i>ipaddr</i> is the IP address on the local router, which is configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup in case the interface went down).
critical-ip-enable <enable disable>	Enables or disables the critical IP address option.
delete	Deletes VRRP from the port.
disable	Disables VRRP on the port.
enable	Enables VRRP on the port.
fast-adv-enable <enable / disable>	Enables or disables the Fast Advertisement Interval. Default is disable. <i>enable</i>, means use the Fast Advertisement Interval. <i>disable</i>, means use the regular Advertisement interval.

config ethernet <ports> ip vrrp <vrid> followed by:	
<code>fast-adv-int</code> <code><milliseconds></code>	Sets the Fast Advertising Interval, the time interval between sending VRRP advertisement messages. • <i>milliseconds</i> the interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. • values must be entered in multiples of 200 milliseconds.
<code>holddown-timer <seconds></code>	Modifies the behavior of the VRRP failover mechanism by allowing the router enough time to detect the OSPF or RIP routes. • The time interval (in seconds) a router is delayed when changing to master state.
<code>priority <prio></code>	Sets the port VRRP priority. • <i>prio</i> is the value (between 1 and 254) used by the VRRP router. The default is 100. The value 255 is assigned to the router that owns the IP address associated with the virtual router.

Figure 176 shows sample output for the **config ethernet ip vrrp info** command.

Figure 176 config ethernet ports ip vrrp info command output

```

Passport-8606:6# config ethernet 1/1 ip vrrp 1 fast-adv-enable
enable

Passport-8606:6# config ethernet 1/1 ip vrrp 1 fast-adv-int 200

Passport-8606:6# config ethernet 1/1 ip vrrp 1 info

Sub-Context: clear config dump monitor show test trace wsm
Current Context:

Port 1/1 :
            action : none
            address : 11.11.11.11
            adver-int : 1
            backup-master : disable
            critical-ip : 0.0.0.0
            critical-ip-enable : disable
            delete : N/A
            fast-adv-int : 200
            fast-adv-enable : enable
            vrrp : enable
            holddown-timer : 0
            priority : 255

```

Showing VRRP port information


The **show ip vrrp info** command displays basic VRRP configuration information about the specified port or for all ports.

The command uses the syntax:

```
show ip vrrp info [vrid <vrid>] [ip <ipaddr>]
```

[Figure 177](#) shows sample output for the **show ip vrrp info** command.

Figure 177 show ip vrrp info command output



Configuring VRRP on a VLAN

Use the following command to configure VRRP on a VLAN:

```
config vlan <vid> ip vrrp <vrid>
```

Where:

- *vid* is the VLAN ID (1 to 4094).
- *vrid* is the virtual router ID (1 to 255), a number that uniquely identifies a virtual router on a given VRRP router. The virtual router acts as the default router for one or more assigned addresses.

The VLAN VRRP commands include the following options:

config vlan <vid> ip vrrp <vrid> followed by:	
info	Displays the current VLAN VRRP settings (Figure 178).
action <action choice>	Sets the manual override of the delay timer for the virtual router interface.
address <ipaddr>	Sets the IP address of the virtual router interface.
adver-int <seconds>	Sets the advertising interval (in seconds), the time interval between sending advertisement messages. • <seconds> is the range 1 to 255, and the default is 1.
backup-master <enable disable>	Enables or disables the VRRP backup master for a VLAN. This option is only supported on Split-MLT ports.
critical-ip <ipaddr>	Sets the critical IP address for VRRP. • <ipaddr> is the IP address of the interface on the local router configured so that a change in its state causes a role switch in the virtual router (for example, from master to backup in case the interface went down).
critical-ip-enable <enable disable>	Enables or disables the critical IP address option.
delete	Deletes the VRRP from the VLAN.
disable	Disables the VRRP on the VLAN.
enable	Enables VRRP on the VLAN.
fast-adv-enable <enable disable>	Enables or disables the Fast Advertisement Interval. Default is disable. • <i>enable</i>, means use the Fast Advertisement Interval. • <i>disable</i>, means use the regular Advertisement interval.
fast-adv-int <milliseconds>	Sets the Fast Advertising Interval, the time interval between sending VRRP advertisement messages. • <i>milliseconds</i> the interval can be between 200 and 1000 milliseconds, and it must be the same on all participating routers. The default is 200. • values must be entered in multiples of 200 milliseconds.

config vlan <vid> ip vrrp <vrid> followed by:	
holddown-timer <seconds>	Sets the time interval (in seconds) that a router is delayed when changing to master state.
priority <prio>	Sets the port VRRP priority value to be used by this VRRP router. • <prio> is between 1 and 254. The default is 100. The value 255 is assigned to the router that owns the IP address associated with the virtual router.

Figure 178 shows sample output for the **config vlan ip vrrp info** command.

Figure 178 config vlan ip vrrp info command output

```
Passport-8606:6# config vlan 2 ip vrrp 1 fast-adv-enable enable

Passport-8606:6# config vlan 2 ip vrrp 1 fast-adv-int 400

Passport-8606:6# config vlan 2 ip vrrp 1 info

Sub-Context: clear config dump monitor show test trace wsm
Current Context:

        action : none
        address : 200.1.1.15
        adver-int : 1
        backup-master : disable
        critical-ip : 0.0.0.0
        critical-ip-enable : disable
        fast-adv-int : 400
        fast-adv-enable : enable
        delete : N/A
        vrrp enable : enable
        holddown-timer : 0
        priority : 100
```


Showing vlan info vrrp extended command

The **show vlan info vrrp extended** command displays the extended VRRP configuration for all VLANs on the switch or for the specified VLAN.

The command uses the syntax:

```
show vlan info vrrp extended [<vid>]
```

[Figure 179](#) shows sample output for the **show vlan info vrrp extended** command.

Figure 179 show vlan info vrrp extended command output

```
Passport-8606:6# show vlan info vrrp extended
```

```
=====
                                Vlan Vrrp Extended
=====
VID  STATE          CONTROL  PRIORITY  MASTER  ADVERTISE CRITICAL
      IPADDR          INTERVAL IPADDR
-----
2    initialize enable   100      0.0.0.0    1      0.0.0.0

VID  HOLDDOWN_TIME  ACTION    CRITICAL IP  BACKUP  BACKUP  FAST ADV  FAST ADV
      STATE          ENABLE      MASTER  MASTER  INTERVAL  ENABLE
-----
2    0              none      disable  disable  down    400      enable
```

Legend:

State =The current state of the virtual router. Values are: initialize - waiting for a start up event,
master - forwarding IP addresses associated with this virtual router, or backup - monitoring the state or availability of the master router.

Showing VRRP interface information


The `show ip vrrp info` command displays VRRP information on the interface. If a virtual router ID or an IP address is entered, the information is displayed only for that VRID or for that interface; if not, all VRRP interfaces are listed.

This command uses the syntax:

```
show ip vrrp info [<vrid>] [<ipaddr>]
```

Figure 180 shows sample output for the `show ip verb info` command.

Figure 180 show ip vrrp info command output



Dependencies and rules

When the Fast Advertisement Interval option is used to configure a master and backup switch, the Fast Advertisement Interval option must be enabled on both switches for VRRP to work correctly. If one is configured with the regular advertisement interval and the other with the Fast Advertisement Interval it will cause an unstable state and advertisements will be dropped.

Chapter 13

Configuring IP policies using Device Manager

Prior to Passport 8000 Series software release 3.2, you configured separate policy databases for RIP accept, RIP announce, OSPF accept, and OSPF announce filtering purposes. Now, you can form a unified database of route policies that can be used by the protocols (RIP or OSPF or BGP) for any type of filtering task.

A policy is identified by a name or an ID. Under a given policy you can have several sequence numbers, each of which is equal to one policy in the old convention. If a field in a policy is not configured, it will appear as 0 or any when it is displayed in Device Manager, as this implies that the field is to be ignored in the match criteria. The clear option can be used to remove existing configurations for any field.

Each policy sequence number contains a set of fields. Only a subset of those fields are used when the policy is applied in a certain context. For example, if a policy has a set-preference field set, it will be used only when the policy is applied for accept purposes. This field will be ignored when the policy is applied for announce/redistribute purpose.

You can apply one policy for one purpose, for example, RIP Announce, on a given RIP interface. In that case, all sequence numbers under the given policy will be applicable for that filter. A sequence number also acts as an implicit preference, a lower sequence number is preferred.

- For conceptual information about IP Policies, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter includes the following topics:

Topic	Page
Route Policy configuration prerequisites	462
Configuring the prefix list	462
Creating and editing the As-Path-List	465
Creating and editing a Community List	467
Creating and editing a route policy	469
Applying routing policies	476
Configuring an OSPF accept policy	477
Configuring an OSPF redistribute policy	479
Configuring inbound/outbound filtering policies on a RIP interface	483
Deleting inbound/outbound filtering policies on a RIP interface	484
Configuring inbound/outbound filtering policies on a DVMRP interface	485

Route Policy configuration prerequisites

Before you can configure a route policy to a protocol, you must do the following:

- Define the prefix list.
- Define a route policy.
- Apply the route policy to the routing protocol.

Configuring the prefix list

The prefix lists are lists of routes that can be applied to one or more route policies. They contain a set of contiguous or non-contiguous routes. Prefix lists are referenced by name from within the routing policies.


You can create one or more IP prefix lists and apply that list to any IP route policy. Prior to the inception of the prefix list, two databases, the address-list and the net0lst, were used by all protocols for different types of policies. The prefix list combines these two databases. A prefix list with a 32 bit mask is equivalent to an address. A prefix list with a mask less than 32 bits can be used as a network. If you configure the MaskLenFrom field to be less than MaskLenUpto field, it can also be used as a range.

To set up or edit a route policy prefix list:

- 1 From the Device Manager menu bar, choose IP Routing > Policy.

The Policy dialog box opens with the Prefix List tab displayed (Figure 181).

Figure 181 Policy dialog box—Prefix List tab



- 2 Click Insert.

The Policy, Insert Prefix List dialog box opens (Figure 182).

- 3 Click Insert.

Figure 182 Policy, Insert Prefix List dialog box




Table 42 describes the Policy, Insert Prefix List dialog box fields.

Table 42 Policy, Insert Prefix List dialog box fields

Field	Description
ID	The list identifier.
Prefix	The IP address.
PrefixMaskLen	This is the specified length of the prefix mask. Note: You must enter the full 32-bit mask in order to exact a full match of a specific IP address (for example, such as when creating a policy to match on next-hop).
Name	The name command is used to name a specified prefix list during the creation process or to rename the specified prefix list. The name length can be from 1 to 64 characters.
MaskLenFrom	The lower bound of the mask length. The default is the mask length. Note: Lower bound and higher bound mask lengths together can define a range of networks.
MaskLenUpto	The higher bound mask length. The default is the mask length. Note: Lower bound and higher bound mask lengths together can define a range of networks.

Creating and editing the As-Path-List

The As-Path-List list is used with route policies and contains one or multiple as-path entries. Each as-path entry contains one or multiple AS numbers with the mode deny or permit. You can use the As-Path-List list to filter a route based on its as-path attribute.

To create or edit the as-path-list:


- 1 From the Device Manager menu bar, choose IP Routing > Policy.

The Policy dialog box opens with the Prefix List tab displayed (Figure 183).

- 2 Click the As Path List tab.

The As Path List tab opens (Figure 183).

Figure 183 Policy dialog box—As Path List tab



3 In the As Path List tab, click Insert.

The Policy, Insert As Path List dialog box opens (Figure 184).

Figure 184 Policy, Insert As Path List dialog box




Table 43 describes the Policy, Insert As Path List dialog box fields.

Table 43 Policy, Insert As Path List dialog box fields

Field	Description
Id	This is the ID of an entry in the As Path list table.
MemberId	The identifier given to the entry of Ip As Path Access List table.
Mode	This field specifies the action to be taken when a policy is selected for a specific route. Select permit (allow the route) or deny (ignore the route).
AsRegularExpression	This field specifies the expression that is to be used for path.

4 Enter the appropriate information for you configuration.

5 Click Insert.

Creating and editing a Community List

The Community-list is used with route policies and contains one or multiple Community-list entries. Each Community-list entry contains one or multiple community numbers with the mode deny or permit. You can use the As-Path-List list to filter a route based on its as-path attribute.

To create or edit the community list:


- 1 From the Device Manager menu bar, choose IP Routing > Policy.

The Policy dialog box opens with the Prefix List tab displayed (Figure 185).

- 2 Click the Community List tab.

The Community List tab opens (Figure 185).

Figure 185 Policy dialog box—Community List tab



3 In the Community List tab, click Insert.

The Policy, Insert Community List dialog box opens (Figure 186).

Figure 186 Policy, Insert Community List dialog box

The image shows a Windows-style dialog box titled "10.10.54.27 - Policy, Insert Community List". It has a standard Windows icon in the top-left corner and a close button (X) in the top-right corner. The dialog contains several input fields: "Id:" with a text box containing "1..1024", "MemberId:" with a text box containing "0..65535", "Mode:" with two radio buttons labeled "permit" and "deny", and "Community:" with a large empty text box. At the bottom of the dialog, there are three buttons: "Insert", "Close", and "Help...".

Table 43 describes the Policy, Insert Community List dialog box fields.

Table 44 Policy, Insert Community List dialog box fields

Field	Description
Id	This is the ID of an entry in the Community list table.
MemberId	The identifier given to the entry of the Community List table.
Mode	This field specifies the action to be taken when a policy is selected for a specific route. Select permit (allow the route) or deny (ignore the route).
Community	The IP Community Access List Community string. Can be 0 to 256 characters.

4 Enter information.

5 Click Insert.

Creating and editing a route policy

When you create a route-policy using Device Manager, you have the option of selecting the ID number. When you create a route-policy using the CLI, the route-policy ID is automatically generated.

You can configure route policies to be used for In, Out, and Redistribute purposes by all protocols.



Note: Changing route preferences is a process-oriented operation that can affect system performance and network reachability while performing the procedures. Therefore, Nortel Networks recommends that if you want to change a prefix list or a routing protocol, you should configure all route policies and prefix lists before enabling the protocols.

Table 45 displays accept and announce policies for RIP, OSPF, and BGP protocols. It displays which matching criteria are applicable for a certain routing policy.

Table 45 Protocol Route Policy table

Criteria	RIP					OSPF					BGP					
	Announce				Accept	Redistribute				Accept	Redistribute				Accept	Announce
	OSPF	Direct	RIP	BGP	RIP	Direct	Static	RIP	BGP	OSPF	OSPF	Static	RIP	Direct	BGP	BGP
Match Protocol	Yes	Yes	Yes	Yes												
Match Network	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Match IpRoute Source	Yes ¹		Yes ²					Yes ²					Yes ²		Yes	
Match NextHop	Yes	Yes	Yes	Yes	Yes		Yes	Yes	Yes		Yes	Yes	Yes		Yes	Yes
Match Interface			Yes					Yes					Yes			
Match Route Type	Yes									Yes ³	Yes					Yes
Match Metric	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MatchAs Path															Yes	Yes

Table 45 Protocol Route Policy table (continued)

Criteria	RIP					OSPF					BGP					
	Announce				Accept	Redistribute				Accept	Redistribute				Accept	Announce
	OSPF	Direct	RIP	BGP	RIP	Direct	Static	RIP	BGP	OSPF	OSPF	Static	RIP	Direct	BGP	BGP
Match Community															Yes	Yes
Match Community Exact															Yes	Yes
MatchTag				Yes					Yes							
NssaPbit																
SetRoute Preference					Yes					Yes						
SetMetric TypeInternal																
SetMetric	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
SetMetric Type						Yes	Yes	Yes	Yes							
SetNextHop									Yes						Yes	Yes
SetInject NetList	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes						
SetMask					Yes											
SetAsPath															Yes	Yes
SetAsPath Mode															Yes	Yes
Set Automatic Tag																
Set Community Number															Yes	Yes
Set Community Mode															Yes	Yes
SetOrigin																Yes
SetLocal Pref															Yes	Yes
SetOrigin EggAs																
SetTag																
SetWeight															Yes	

1 advertise router

2 RIP gateway

3 externaltype1 and externaltype2 are the only options.

To create or edit a route policy:


- 1 From the Device Manager menu bar, choose IP Routing > Policy.

The Policy dialog box opens with the Prefix List tab displayed [Figure 181 on page 463](#).

- 2 Click the Route Policy tab.

The Route Policy tab opens ([Figure 187](#)).

Figure 187 Policy dialog box—Route Policy tab



- 3 Click Insert.

The Policy, Insert Route Policy dialog box opens [Figure 188 on page 472](#).

Figure 188 Policy, Insert Route Policy dialog box

Table 46 describes the Policy, Insert Route Policy dialog box fields.

Table 46 Policy, Insert Route Policy dialog box fields

Field	Description
Id	This is the ID of an entry in the Prefix list table.
SequenceNumber	A second index used to identify a specific policy within a route policy group.
Name	This command is used during the creation process, or to rename a policy once it has been created. This command changes the name field for all sequence numbers under the given policy.
Enable	This field indicates whether this policy sequence number is enabled or disabled. If disabled the policy sequence number is ignored.
Mode	This field specifies the action to be taken when a policy is selected for a specific route. Select permit (allow the route) or deny (ignore the route).
MatchProtocol	Select the appropriate protocol. If configured, matches the protocol through which the route is learned. This field is used only for RIP announce purposes.

Table 46 Policy, Insert Route Policy dialog box fields (continued)

Field	Description
MatchNetwork	If configured, the switch matches the destination network against the contents of the specified prefix list. Click the ellipse button and choose from the list in the MatchNetwork dialog box (see Figure 184 on page 466). You can select up to four entries. To deselect an entry, use the ALT key.
MatchIpRouteSource	If configured, matches the next hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option ignored for all other route types. Click the ellipse button and choose from the list in the Match Route Source dialog box (see Figure 184 on page 466). You can select up to four entries. To deselect an entry, use the ALT key. Note: This field can also be changed in the Route Policy tab of the Policy dialog box as shown in (see Figure 183 on page 465).
MatchNextHop	If configured, matches the next hop IP address of the route against the contents of the specified prefix list. This field applies only to non-local routes. Click the ellipse button and choose from the list in the Match Next Hop dialog box (see Figure 184 on page 466). You can select up to four entries. To deselect an entry, use the ALT key.
MatchInterface	If configured, the switch matches the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other type of route. Click the ellipse button and choose from the list in the Match Interface dialog box (see Figure 184 on page 466). You can select up to four entries. To deselect an entry, use the ALT key.
MatchRouteType	Sets a specific route-type to be matched (applies only to OSPF routes). Externatype1, and Externatype2 specify the OSPF routes of the specified type only. OSPF internal refers to intra and inter area routes.
MatchMetric	If configured, the switch matches the metric of the incoming advertisement or existing route against the specified value (1to 65535). If 0, then this field is ignored. The default is 0.
MatchAsPath	Applicable to BGP protocol only. Match the BGP autonomous system path. This will override the BGP neighbor filter list information.

Table 46 Policy, Insert Route Policy dialog box fields (continued)

Field	Description
MatchCommunity	Applicable to BGP protocol only. This is used to filter incoming and outgoing updates based on a community list.
MatchCommunityExtract	Applicable to BGP protocol only. If enabled, it indicates the match has to exact (i.e., all of the communities specified in the path have to match). Default is disable.
MatchTag	Applicable to BGP protocol only. Specifies a list of tag(s), that will be used during the match criteria process. It contains one or more tag values.
NssaPbit	Set or reset the P bit in specified type 7 LSA. By default the P bit is always set in case the user set it to a disable state for a particular route policy than all type 7. LSAs associated with that route policy will have the P bit cleared with this intact NSSA ABR will not perform translation of these LSAs to type 5. Default is enable.
SetRoutePreference	Setting the preference greater than zero, specifies the route preference value to be assigned to the routes which matches this policy. This applies to Accept policies only. You can set a value from 0 to 255. The default value is 0. If the default is configured, the global preference value is used.
SetMetricTypeInternal	This indicates to set the MED value for routes advertised to BGP numbers to the IGP metric value. Default is 0.
SetMetric	If configured, the switch sets the metric value for the route while announcing or redistributing. The default-import-metric is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or the default value is used.
SetMetricType	Applicable to OSPF protocol only.If configured, sets the metric type for the routes to be announced into the OSPF routing protocol that matches this policy. The default is type 2. This field is applicable only for OSPF announce policies.
SetNextHop	Applicable to BGP protocol only. This is the IP address of the next hop router. It is Ignored for DVMRP routes. Default is 0.0.0.0
SetInjectNetList	If configured, the switch replaces the destination network of the route that matches this policy with the contents of the specified prefix list. Click the ellipse button and choose from the list in the Set Inject NetList dialog box Figure 184 on page 466 .
SetMask	Applicable to RIP protocol only.If configured, the switch sets the mask of the route that matches this policy. This applies only to RIP accept policies.

Table 46 Policy, Insert Route Policy dialog box fields (continued)

Field	Description
SetAsPath	Applicable to BGP protocol only. The AS path value to be used whether the SetAsPathMode field is Tag or Prepend.
SetAsPathMode	Applicable to BGP protocol only. It can be either tag or Prepend tag. It is applicable only while redistributing routes to BGP. It converts the tag of a route into AS path.
SetAutomaticTag	Applicable to BGP protocol only. Default is disable.
SetCommunityNumber	Applicable to BGP protocol only. this value can be a number (1..42949672000) or no-export or no-advertise. Applicable to BGP advertisements.community number.
SetCommunityMode	Applicable to BGP protocol only. This value can be either append, none, or unchanged. Unchanged - keep the community attribute in the route path as it is. None - remove the community in the route path additive. Append - adds the community-number specified in SetCommunityNumber to the community list attribute. Default is unchanged.
SetOrigin	Applicable to BGP protocol only. Set to igp, egp, incomplete, or unchanged. If not set, the system uses the route origin from the Ip routing table (protocol). Default is unchanged.
SetLocalPref	Applicable to BGP protocol only. This value will be used during the route decision process in the BGP protocol. Default is 0.
SetOriginEgpAs	Applicable to BGP protocol only. Indicates the remote autonomous systems number. Default is 0.
SetTag	Applicable to BGP protocol only. This field is to be used for setting the tag of the destination routing protocol. If it is not specified, forward the tag value in the source routing protocol. A value of 0 indicates it is not set. Default is 0.
SetWeight	Applicable to BGP protocol only. This field should be used with match as-path condition. It is the weight value for the routing table. For BGP this value will override the weight configured through the NetworkTableEntry, FilterListWeight, or NeighborWeight. Default is 0.

- 4 Enter the appropriate information for your configuration.
- 5 Click Insert.

Applying routing policies




Note: Changing route policies or prefix lists that affect OSPF accept or redistribute is a process-oriented operation that can affect system performance and network reachability while performing the procedures. Therefore, Nortel Networks recommends that if you want to change a prefix list or a routing protocol, you should configure all route policies and prefix lists before enabling the protocols.

To apply a routing policy:

- 1 From the Device Manager menu bar, choose IP Routing > Policy.
The Policy dialog box opens with the Prefix List tab displayed [Figure 181 on page 463](#).
- 2 Click the Applying Policy tab.
The Applying Policy tab opens ([Figure 189](#)).

Figure 189 Policy dialog box—Applying Policy tab



- 3 Select the type of filter to apply
- 4 Click Apply.

[Table 47](#) describes the Policy, Applying Policy dialog box fields.

Table 47 Policy, Applying Policy dialog box fields


Field	Description
RoutePolicyApply	When selected, allows the configuration changes in the route policy to take effect. This keeps the switch from attempting to apply the changes one-by-one after each configuration change.
RedistributeApply	When selected, allows the configuration changes in the policy to take effect for an OSPF Redistribute context. This keeps the switch from attempting to apply the changes one-by-one after each configuration change.
OspfInFilterApply	When selected, allows the configuration change in a route policy or a prefix list to take effect in an OSPF accept context. This keeps the switch from attempting to apply the change one-by-one after each configuration change.

Configuring an OSPF accept policy

To set up or edit an OSPF accept policy:

- 1 From the Device Manager menu bar, choose IP Routing > Policy.
The Policy dialog box opens with the Prefix List tab displayed [Figure 181 on page 463](#).
- 2 Click the OSPF Accept tab.
The OSPF Accept tab opens [Figure 190 on page 478](#).

Figure 190 Policy dialog box—OSPF Accept tab



3 Click Insert.

The Policy, Insert OSPF Accepts dialog box opens (Figure 191).

4 Click Insert.

Figure 191 Policy, Insert OSPF Accept dialog box



[Table 48](#) describes the Policy, Insert OSPF Accept dialog box fields.

Table 48 Policy, Insert OSPF Accepts dialog box fields

Field	Description
AdvertisingRtr	This field is the routing id of the advertising router.
Enable	Select to enable or disable the advertising router. You can also enable or disable this feature in the OSPF Accept tab of the Policy dialog box by clicking in the field and selecting enable or disable from the pulldown menu.
• MetricType	Select the OSPF external type. This parameter describes which types of OSPF ASE routes match this entry • any means match either ASE type 1 or 2 • type1 means match any external type 1 • type2 means match any external type 2 You can also select your entry in the OSPF Accept tab of the Policy dialog box by clicking in the field and selecting any, type1, or type2 from the pulldown menu.
PolicyName	This field is the name of the OSPF in filter policy. Click the ellipse button and choose from the list in the Policy Name dialog box (Figure 191). To deselect an entry, use the ALT key.

Configuring an OSPF redistribute policy

You can configure a redistribute entry for OSPF to announce routes into OSPF of a certain source type, for example, static, RIP, or direct. If a route policy field is not configured for a redistribute entry, then the default action is taken on the basis of metric, metric-type, and subnet configured. This is called basic redistribution. Otherwise, you use the route policy specified to perform detailed redistribution. If no redistribution entry is configured, no external LSA is generated for non-OSPF routes.

To set up or edit an OSPF redistribute policy:


- 1 From the Device Manager menu bar, choose IP Routing > Policy.

The Policy dialog box opens with the Prefix List tab displayed [Figure 181 on page 463](#).

- 2 Click the OSPF Redistribute tab.


The OSPF Redistribute tab opens [\(Figure 192\)](#).

Figure 192 Policy dialog box—OSPF Redistribute tab



- 3 In the OSPF Redistribute tab, click Insert.

The Policy, Insert OSPF Redistribute dialog box opens [Figure 193 on page 481](#).

Figure 193 Policy, Insert OSPF Redistribute dialog box

- 4 Enter the appropriate information in the Policy, Insert OSPF Redistribute dialog box.

Refer to [Table 49](#), for a description of the screen fields.

- 5 After you enter the appropriate data, click Insert.

Your newly entered configuration information appears in the OSPF Redistribute tab.

[Table 49](#) describes the Policy, Insert OSPF Redistribute dialog box fields.

Table 49 Policy, Insert OSPF Redistribute dialog box fields

Field	Description
RouteSource	Select the route source protocol for the redistribution entry.
Enable	Enables (or disables) an OSPF redistribute entry for a specified source type. You can also enable or disable this feature in the OSPF Redistribute tab of the Policy dialog box by clicking in the field and selecting enable or disable from the pulldown menu.
Metric	Set the OSPF route redistribution metric for basic redistribution. The value can be a range between 0 to 65535. If configured as 0, the original cost of the route is used.

Table 49 Policy, Insert OSPF Redistribute dialog box fields (continued)

Field	Description
MetricType	Sets the OSPF route redistribution metric type. The default is Type 2. You can also select your entry in the OSPF Redistribution tab of the Policy dialog box by clicking in the field and selecting any, type1, or type2 from the pulldown menu.
Subnets	Sets the OSPF route redistribution subnet value (the default value is <i>allow</i>): <i>allow</i> sets the switch to redistribute external subnet routes into an OSPF domain. <i>suppress</i> sets the switch to redistribute external subnet routes into an OSPF domain, with shortened mask lengths. In the advertisement, the external subnet routes mask lengths are shortened to their natural masks. Note: When set to suppress, the switch automatically converts external subnet routes to their natural mask for advertisement on an OSPF interface. You can also select your entry in the OSPF Redistribution tab of the Policy dialog box by clicking in the field and selecting allow or suppress from the pulldown menu.
RoutePolicy	Sets the route policy by name to be used for the detailed redistribution of external routes from a specified source into an OSPF domain. Click the ellipse button and choose from the list in the Route Policy dialog box (see Figure 193 on page 481). To deselect an entry, use the ALT key.

Configuring inbound/outbound filtering policies on a RIP interface

You can configure inbound filtering on a RIP interface. This configured policy determines whether to learn a route on a specified interface. It also specifies the parameters of the route when it is added to the routing table. Conversely, you can configure outbound filtering on a RIP interface. This configured policy determines whether to advertise a route from the routing table on a specified interface. This policy also specifies the parameters of the advertisement.

To configure inbound/outbound filtering on a RIP interface:


- 1 From the Device Manager menu bar, choose IP Routing > Policy.

The Policy dialog box opens with the Prefix List tab displayed [Figure 181 on page 463](#).

- 2 Click the RIP In/Out Policy tab.

The RIP In/Out Policy tab opens ([Figure 194](#)).

Figure 194 Policy dialog box—RIP In/Out Policy tab



- 3 In the desired row, double-click on the InPolicy or OutPolicy column.
The InPolicy or OutPolicy list box opens, displaying preconfigured policies.
- 4 Select a (preconfigured) In/Out Policy and click OK.

[Table 50](#) describes the Policy, RIP In/Out Policy dialog box fields.

Table 50 Policy, RIP In/Out Policy dialog box fields

Field	Description
Address	This field is the IP address of the RIP interface.
Interface	This field is the internal index of the RIP interface.
InPolicy	Right click in the InPolicy name field and select the policy name to be applied from the PolicyName dialog box (see Figure 194 on page 483). The policy name is used for inbound filtering on this RIP interface. This policy will determine whether to learn a route on this interface. It also specifies the parameters of the route when it is added to the routing table.
OutPolicy	Right click in the OutPolicy name field and select the policy name to be applied from the PolicyName dialog box (see Figure 194 on page 483). The policy name is used for outbound filtering on this RIP interface. This policy will determine whether to advertise a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.

Deleting inbound/outbound filtering policies on a RIP interface

To delete a RIP In/Out Policy using Device Manager:

- 1 From the Device Manager menu bar, select IP Routing > Policy.
The Policy dialog box opens with the Prefix List tab displayed ([Figure 194 on page 483](#)).
- 2 Click RIP In/Out Policy.

- 3 In the desired row, double-click on the InPolicy or OutPolicy column for the policy you want to delete.

The InPolicy or OutPolicy dialog box is displayed ([Figure 194 on page 483](#)).

- 4 Press CTRL + Left mouse click on the desired policy to delete.

- 5 Click OK.

The policy is deleted and you are returned to the Rip In/Out Policy tab.

- 6 Click Apply.

Configuring inbound/outbound filtering policies on a DVMRP interface

You can configure inbound filtering on a DVMRP interface. This configured policy determines whether to learn a route on a specified interface. It also specifies the parameters of the route when it is added to the routing table. Conversely, you can configure outbound filtering on a DVMRP interface. This configured policy determines whether to advertise a route from the routing table on a specified interface. This policy also specifies the parameters of the advertisement.

To configure inbound/outbound filtering on a DVMRP interface:


- 1 From the Device Manager menu bar, choose IP Routing > Policy.

The Policy dialog box opens with the Prefix List tab displayed (see [Figure 181 on page 463](#)).

2 In the Policy dialog box, click the DVMRP In/Out Policy tab.

The DVMRP In/Out Policy tab opens (Figure 195).

Figure 195 Policy dialog box—DVMRP In/Out Policy tab



3 In the desired row, double-click on the InPolicy or OutPolicy column.

The InPolicy or OutPolicy list box opens.

4 Select the desired In/Out Policy and click OK.

Table 51 describes the Policy, DVMRP In/Out Policy dialog box fields.

Table 51 Policy, DVMRP In/Out Policy dialog box fields

Field	Description
IfIndex	This field is the internal index of the DVMRP interface.
LocalAddr	This field is the IP address of the DVMRP interface.
Enable	The administrative status of DVMRP in the router. The value 'enabled' denotes that the DVMRP is enabled on the interface; 'disabled' disables it on the interface.

Table 51 Policy, DVMRP In/Out Policy dialog box fields (continued)

Field	Description
InPolicy	Right click in the InPolicy name field and select the policy name to be applied from the PolicyName dialog box (Figure 195 on page 486). The policy name is used for inbound filtering on this DVMRP interface. This policy will determine whether to learn a route on this interface. It also specifies the parameters of the route when it is added to the routing table.
OutPolicy	Right click in the OutPolicy name field and select the policy name to be applied from the PolicyName dialog box (Figure 195 on page 486). The policy name is used for outbound filtering on this DVMRP interface. This policy will determine whether to advertise a route from the routing table on this interface. This policy also specifies the parameters of the advertisement.

Note that in enabling a multimedia filter from a port window, the port becomes a DIFFSERV access port.

Chapter 14

Configuring IP Policies using the CLI

This chapter describes the Run-Time CLI commands that are used to configure IP policies on your 8000 Series switch. In previous releases, you could configure separate policy databases for RIP accept, RIP announce, OSPF accept, and OSPF announce filtering purposes. Now, you can form a unified database of route policies that RIP or OSPF can use for any type of filtering task.

A policy is identified by a name or an ID. Under a given policy you can have several sequence numbers, each of which is equal to one policy in the old convention. Each policy sequence number contains a set of fields. Only a subset of those fields are used when the policy is applied in a certain context. For example, if a policy has a set-preference field set, it will be used only when the policy is applied for accept purposes. This field will be ignored when the policy is applied for announce/redistribute purpose.

You can apply one policy for one purpose, for example, RIP Announce, on a given RIP interface. In that case, all sequence numbers under the given policy can be applied to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

- For conceptual information about IP Policies, see [Chapter 1, “IP routing concepts,” on page 31](#).
- For configuration examples, including the required CLI commands, see [Chapter 2, “IP routing configuration examples,” on page 93](#).

This chapter includes the following topics:

Command	Page
Roadmap of IP commands	490
IP policy commands	493
Showing IP policies	507

Roadmap of IP commands

The following roadmap lists some of the IP commands and their parameters. Use this list as a quick reference or click on any command or parameter entry for more information.

Command	Parameter
<code>config ip prefix-list <prefix-list name></code>	<code>info</code> <code>add-prefix <ipaddr/mask> [maskLenFrom <value>] [maskLenTo <value>]</code> <code>delete</code> <code>name <name></code> <code>remove-prefix <ipaddr/mask></code>
<code>config ip route-policy <policy name> seq <seq number></code>	<code>info</code> <code>action <permit deny></code> <code>create</code> <code>delete</code> <code>disable</code> <code>enable</code> <code>match-as-path <as-list></code> <code>match-community <community-list></code>

Command**Parameter**

```
match-community-exact  
<enable|disable>  
match-interface <prefix-list>  
match-metric <metric>  
match-network <prefix-list>  
match-next-hop <prefix-list>  
match-protocol <protocol name>  
match-route-src <prefix-list>  
match-route-type <route-type>  
match-tag <tag>  
name <policy name>  
set-as-path <as-list-id>  
set-as-path-mode <tag|prepend>  
set-automatic-tag  
<enable|disable>  
set-community <community-list>  
set-community-mode  
<unchanged|additive| none>  
set injectlist <prefix-list>  
set-local-pref <pref-value>  
set-mask <ipaddr>  
set-metric <metric-value>  
set-metric-type <metric-type>  
set-nssa-pbit <enable|disable>  
set-next-hop <ipaddr>  
set-origin <origin>  
set-origin-egp-as  
<origin-egp-as>  
set-preference <pref-value>  
set-tag <tag>
```

Command	Parameter
	set-weight <weight>
config ip ospf accept adv-rtr <ipaddr>	info apply create delete disable enable metric-type <type1 type2 any> route-policy <policy name>
config ip ospf accept	apply
config ip ospf redistribute <source-type>	info apply create disable delete enable metric <metric-value> metric-type <type1 type2> route-policy <policy name> subnets <allow supress>
config ip ospf redistribute <source-type>	info

Command	Parameter
	apply
	create
	disable
	delete
	enable
	metric <metric-value>
	metric-type <type1 type2>
	route-policy <policy name>
	subnets <allow supress>
config ip ospf redistribute	apply
show ip prefix-list	
show ip route-policy info	
show ip ospf accept info	
show ip ospf redistribute info	

IP policy commands

The section describes ip policy commands and includes the following topics:

- [“Configuring prefix-lists,” next](#)
- [“Configuring route policies” on page 496](#)
- [“Configuring a policy for accepting external routes from a router” on page 501](#)
- [“Applying OSPF accept policy changes” on page 503](#)
- [“Configuring OSPF redistribute policies” on page 504](#)
- [“Applying configuration changes to OSPF redistribute policies” on page 506](#)

Configuring prefix-lists

The prefix list is a list of networks used by route policies to define an action. You can create one or more IP prefix lists and apply that list to any IP route policy.

Before the creation of prefix lists, some protocols used two databases for different types of policies: the address-list database, and the net0lst database.

A prefix list combines these two databases:

- A prefix list with a 32-bit mask is equivalent to an address.
- A prefix list with a mask less than 32 bits can be used as a network.

When you configure the masklengthFrom field to be less than the Mask LengthTo field, it can also be used as a range.

For more information about prefix lists, see [Chapter 1, “IP routing concepts,” on page 31](#).

To configure a prefix list, use the following command

```
config ip prefix-list <prefix-list name>
```

This command includes the following options:

config ip prefix-list <prefix-list name> followed by:	
info	Displays all of the prefixes in a given list (see Figure 197 on page 495).
add-prefix <ipaddr/mask> [maskLenFrom <value>] [maskLenTo <value>]	Adds a prefix entry to the prefix list. • <ipaddr/mask> is the IP address and mask. • maskLenFrom <value> is the lower bound of mask length. The default is the mask length. • maskLenTo <value> is the higher bound mask length. The default is the mask length. Note: Lower bound and higher bound mask lengths together can define a range of networks.

config ip prefix-list <prefix-list name> followed by:	
delete	Deletes the prefix list.
name <name>	The name command is used to rename the specified prefix list. The name length can be from 1 to 64 characters.
remove-prefix <ipaddr/mask>	Removes a prefix entry from the prefix list. <i>ipaddr/mask</i> is the IP address and mask.

Figure 196 shows sample output for the **config ip prefix-list** command.

Figure 196 config ip prefix-list command

```
Passport-8010:5# config ip prefix-list ?

Sub-Context:
Current Context:

    add-prefix <ipaddr/mask> [maskLenFrom <value>]
[maskLenTo <value>]
    delete
    info
    name <name>
    remove-prefix <ipaddr/mask>
```

Figure 197 shows sample output for the **config ip prefix-list <name> info** command.

Figure 197 config ip prefix-list <name> info command

```
Passport-8010:5# config ip prefix-list testMore info

add-prefix:
           34.1.1.0/24           (24 , 24 )
    delete: N/A
           name: testMore
    remove-prefix: N/A
```

Configuring route policies

to configure a route policy, use the following command:

```
config ip route-policy <policy name> seq <seq number>
```

This command includes the following options.

config ip route-policy <policy name> seq <seq number> followed by:	
info	Displays current configuration information about this policy sequence number (see Figure 199 on page 501).
action <permit/deny>	This field specifies the action to be taken when a policy is selected for a specific route. This can be permit or deny. Permit allows the route, deny ignores the route.
create	Creates a route policy with a policy name and a sequence number. Note: When creating a route policy in the CLI, the ID is internally generated using an automated algorithm. When you create a route policy in Device Manager, you can manually assign the ID number.
delete	Deletes a route policy with a policy name and a sequence number.
disable	Disables a route policy with a policy name and a sequence number.
enable	Enables a route policy with a policy name and a sequence number.
match-interface <prefix-list>	If configured, the switch matches the IP address of the interface by which the RIP route was learned against the contents of the specified prefix list. This field is used only for RIP routes and ignored for all other type of route. • <prefix-list> specify the name of up to four defined prefix list separated by a comma.
match-as-path <as-list>	Applicable to BGP protocol only. Match the BGP autonomous system path. This will override the BGP neighbor filter list information. • <as-list>

config ip route-policy <policy name> seq <seq number> followed by:	
match-community <community-list>	Applicable to BGP protocol only. This is used to filter incoming and outgoing updates based on a community list. • <community-list>
match-community-exact <enable/disable>	Applicable to BGP protocol only. If enabled, it indicates the match has to exact (i.e., all of the communities specified in the path have to match). • <enable/disable> Default is disable.
match-metric <metric>	If configured, the switch matches the metric of the incoming advertisement or existing route against the specified value. If 0, then this field is ignored. • <metric> is 1 to 65535. The default is 0.
match-network <prefix-list>	If configured, the switch matches the destination network against the contents of the specified prefix list(s). • <prefix-list> specify the name of up to four defined prefix list by name separated by a comma.
match-next-hop <prefix-list>	If configured, matches the next hop IP address of the route against the contents of the specified prefix list. This field applies only to non-local routes. • <prefix-list> specify the name of up to four defined prefix list by name separated by a comma.
match-protocol <protocol name>	If configured, matches the protocol through which the route is learned. This field is used only for RIP announce purposes.
match-route-src <prefix-list>	If configured, matches the next hop IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option ignored for all other route types. • <prefix-list> specify the name of up to four defined prefix list by name separated by a comma.
match-route-type <route-type>	Sets a specific route-type to be matched (applies only to OSPF routes). • <route-type> External-1 and External-2 specifies OSPF routes of the specified type only (any other value is ignored).

config ip route-policy <policy name> seq <seq number> followed by:	
match-tag <tag>	Applicable to BGP protocol only. Specifies a list of tag(s), that will be used during the match criteria process. It contains one or more tag values. • match-tag <tag>
name <policy name>	This command is used to rename a policy once it has been created. This command changes the name field for all sequence numbers under the given policy.
set-as-path <as-list-id>	Applicable to BGP protocol only. The AS path value to be used whether the SetAsPathMode field is Tag or Prepend.
set-as-path-mode <tag/prepend>	Applicable to BGP protocol only. It can be set to either tag or prepend. This will convert the tag of a route into an AS path. Default is prepend.
set-automatic-tag <enable/disable>	Applicable to BGP protocol only. Default is disable.
set-community <community-list>	Applicable to BGP protocol only. This value can be a number from 1 to 42949672000, no-export or no-advertise.
set-community-mode <unchanged/additive/none>	Applicable to BGP protocol only. This value can be either append, none, or unchanged. Unchanged - keep the community attribute in the route path as it is. None - remove the community in the route path. Append- adds the community-number specified in SetCommunityNumber to the community list attribute. Default is unchanged.
set injectlist <prefix-list>	If configured, the switch replaces the destination network of the route that matches this policy with contents of the specified prefix list. • <prefix-list> specify one prefix list by name.
set-local-pref <pref-value>	Applicable to BGP protocol only. This value will be used during the route decision process in the BGP protocol. Default is 0.
set-mask <ipaddr>	If configured, the switch sets the mask of the route that matches this policy. This applies only to RIP accept policies. <ipaddr> is a valid contiguous IP mask.

config ip route-policy <policy name> seq <seq number> followed by:	
set-metric <metric-value>	If configured, the switch sets the metric value for the route while announcing a redistributing. The default is 0. If the default is configured, the original cost of the route is advertised into OSPF; for RIP, the original cost of the route or default-import-metric is used.
set-metric-type <metric-type>	If configured, sets the metric type for the routes to be announced into the OSPF domain that matches this policy. The default is type 2. This field is applicable only for OSPF announce policies.
set-nssa-pbit <enable/disable>	Applicable to BGP protocol only. Enable or disable the P bit in specified type 7 LSA. By default P bit is always enabled. If user sets it to the disable state for a particular route policy, then all type 7 LSAs associated with that route policy will have the P bit cleared. With this intact NSSA ABR will not perform translation of these LSAs to type 5. Default is enable.
set-next-hop <ipaddr>	Applicable to BGP protocol only. Set the IP address of the next hop router. Ignored this for DVMRP routes. Default is 0.0.0.0.
set-origin <origin>	Applicable to BGP protocol only. Set to igp, egp, incomplete, or unchanged. If not set the route origin from the IP routing table (protocol) is used. Default is unchanged.
set-origin-egp-as <origin-egp-as>	Applicable to BGP only. Sets the remote autonomous system number. Default is 0.
set-preference <pref-value>	Setting the preference greater than zero, specifies the route preference value to be assigned to the routes which matches this policy. This applies to accept policies only. • <pref-value> set from 0 to 255. The default is 0. If the default is configured, the global preference value is used.

config ip route-policy <policy name> seq <seq number> followed by:	
set-tag <tag>	Applicable to BGP only. This value is used for setting the tag of the destination routing protocol. If not specified, forward the tag value in the source routing protocol. A value of 0 indicates it is not set. Default is 0.
set-weight <weight>	Applicable to BGP only. Should be used with match as-path condition. This is the weight value for the routing table. For BGP this value will override the weight configured through NetworkTableEntry, FilterListWeight, or NeighborWeight. A value of 0 indicates it is not set. Default is 0.

Figure 198 displays sample output for the **config ip route-policy <policy name> seq <seq number>** command.

Figure 198 config ip route-policy <policy name> seq <seq number> command

```
Passport-8010:5# config ip route-policy test seq 5

Sub-Context:
Current Context:

    action <permit|deny>
    create
    delete
    disable
    enable
    info
    match-interface <prefix-list> [clear]
    match-metric <metric> [clear]
    match-network <prefix-list> [clear]
    match-next-hop <prefix-list> [clear]
    match-protocol <protocol name> [clear]
    match-route-src <prefix-list> [clear]
    match-route-type <route-type>
    name <policy name>
    set-injectlist <prefix-list> [clear]
    set-mask <ipaddr>
    set-metric <metric-value> [clear]
    set-metric-type <metric-type> [clear]
    set-preference <pref-value> [clear]
```

Figure 199 displays sample output for the `config ip route-policy <policy name> seq <seq number> info` command.

Figure 199 `config ip route-policy <policy name> seq <seq number> info` command

```
Passport-8010:5# config ip route-policy test seq 5 info
```

```
Sub-Context:
```

```
Current Context:
```

```

                                id : 4
                                seq : 5
                                name : test
                                enable : disable
                                mode : permit
                                match-protocol : N/A
                                match-interface : N/A
                                match-metric : 0
                                match-network : N/A
                                match-next-hop : N/A
                                match-route-type : any
                                match-route-src : N/A
                                set-injectlist : N/A
                                set-mask : 0.0.0.0
                                set-metric : 0
                                set-metric-type : type2
                                set-preference : 0

```

```
Passport-8010:5#
```

Configuring a policy for accepting external routes from a router

To configure a policy for accepting external routes from a specified advertising router, use the following command:

```
config ip ospf accept adv-rtr <ipaddr>
```

where:

ipaddr is the advertising router ID. If *ipaddr* is equal to 0.0.0.0 it implies all advertising routers. If you do not have an accept entry for a specific advertising router then the default entry is used. When no applicable entry is found, all routes are accepted.

This command includes the following options:

config ip ospf accept adv-rtr <ipaddr> followed by:	
info	Displays OSPF accept configuration information for a specified advertising router.
apply	Applies the OSPF accept policy changes.
create	Creates an OSPF accept entry for a specified advertising router.
delete	Deletes an OSPF accept entry for a specified advertising router.
disable	Disables an OSPF accept entry for a specified advertising router.
enable	Enables an OSPF accept entry for a specified advertising router.
metric-type <type1/type2/any>	Used to indicate the OSPF external type. This parameter describes which types of OSPF external routes match this entry. • <any> means match all external routes. • <type1> means match external type 1 only. • <type2> means match external type 2 only.
route-policy <policy name>	Specifies the name of the route policy to be used for filtering external routes advertised by the specified advertising router before accepting into the routing table.

Figure 200 shows sample output for the `config ip ospf accept adv-rtr <ipaddr>` command.

Figure 200 config ip ospf accept adv-rtr command

```
Passport-8010:5# config ip ospf
accept adv-rtr <ipaddr> apply
accept adv-rtr <ipaddr> create
accept adv-rtr <ipaddr> delete
accept adv-rtr <ipaddr> disable
accept adv-rtr <ipaddr> enable
accept adv-rtr <ipaddr> info
accept adv-rtr <ipaddr> metric-type <type1|type2|any>
accept adv-rtr <ipaddr> route-policy <policy name>
accept apply
```

Applying OSPF accept policy changes

To allow the configuration changes in the policy to take effect for an OSPF Accept context (and to prevent the switch from attempting to apply the changes one-by-one after each configuration change), use the following command:

```
config ip ospf accept
```



Note: Changing OSPF Accept contexts is a process-oriented operation that can affect system performance and network accessibility while performing the procedures. If you want to change default preferences for an OSPF Accept or a prefix-list configuration (as opposed to the default preference), Nortel Networks recommends that you do so before enabling the protocols.

This command includes the following options:

config ip ospf accept followed by:	
apply	Issue this command after modifying any policy configuration that will affect an OSPF accept policy.

Configuring OSPF redistribute policies

Redistribute entries allow OSPF to announce routes of a certain source type, for example, static, RIP, or direct. If you do not configure a route policy field for a redistribute entry, then the default action is taken based on metric, metric-type, and subnet configured. This is called basic redistribution. Otherwise, you use the route policy specified to perform detailed redistribution. If you do not configure a redistribution entry, no external LSA is generated for non-OSPF routes.

To configure a redistribute entry, use the following command:

```
config ip ospf redistribute <source-type>
```

This command includes the following options:

config ip ospf redistribute <source-type> followed by:	
info	Displays OSPF redistribute information for a specified source type.
apply	Applies the OSPF redistribute to the routes from the specified source type to generate or refresh AS external LSAs. For example, RIP static or direct.
create	Creates an OSPF redistribute entry for a specified source type.
disable	Disables an OSPF redistribute entry for a specified source type.
delete	Deletes an OSPF redistribute entry for a specified source type.
enable	Enables an OSPF redistribute entry for a specified source type.

config ip ospf redistribute <source-type> followed by:	
metric <metric-value>	Sets the OSPF route redistribution metric for basic redistribution. • <metric-value> range is 0 to 65535. If configured as 0, the original cost of the route is used.
metric-type <type1/type2>	Sets the OSPF route redistribution metric type for basic redistribution. The default is Type 2.
route-policy <policy name>	Sets the route policy by name to be used for the detailed redistribution of external routes from a specified source into an OSPF domain. <policy name> string length is 0 to 64 characters. A string of length 0 can be used to remove current configuration. If no policy is configured, basic redistribution is performed.
subnets <allow/supress>	Sets the OSPF route redistribution subnet value (the default value is <i>allow</i>): • <i>allow</i> sets the switch to allow external subnet routes to be redistributed into an OSPF domain. • <i>suppress</i> sets the switch to redistribute external subnet routes into an OSPF domain, with shortened mask lengths. In the advertisement, the external subnet routes mask lengths are shortened to their natural masks. Note: When set to <i>allow</i>, the switch automatically converts external subnet routes to their natural mask for advertisement on an OSPF interface. The <i>allow</i> value does not change the mask for all routes; instead it changes the mask for only the subnet routes with mask lengths that are longer than their natural mask.

Figure 201 shows sample output for the **config ip ospf redistribute direct syntax** command.

Figure 201 config ip ospf redistribute direct syntax command

```
Passport:5# config ip ospf redistribute direct
apply
create
disable
delete
enable
info
metric <metric-value>
metric-type <type1|type2>
route-policy <policy name>
subnets <allow|supress>
```

Applying configuration changes to OSPF redistribute policies

To allow the configuration changes in the policy to take effect for OSPF Redistribute context (and to prevent the switch from attempting to apply the changes one-by-one after each configuration change), use the following command:

```
config ip ospf redistribute
```



Note: Changing OSPF Redistribute contexts is a process-oriented operation that can affect system performance and network accessibility while performing the procedures. Therefore, Nortel Networks recommends that if you want to change default preferences for an OSPF Redistribute or a prefix-list configuration (as opposed to the default preference), you should do so before enabling the protocols.

This command includes the following option:

<code>config ip ospf redistribute</code> followed by:	
<code>apply</code>	Issue this command after modifying any policy configuration that will affect an OSPF redistribution.

Showing IP policies

This section describes how to display IP policy characteristics on the Passport 8000 Series and includes the following topics:

- [“Showing prefix lists used by route policies,” next](#)
- [“Showing information about route policies” on page 508](#)
- [“Showing information about OSPF accept policies” on page 509](#)
- [“Showing information about OSPF route redistribute policies” on page 510](#)

Showing prefix lists used by route policies

To display the prefix list of networks used by route policies to define an action, use the following command:

```
show ip prefix-list
```

[Figure 202](#) shows sample output for `show ip prefix-list` command.

Figure 202 show ip prefix-list command

```
Passport-8010:5# show ip prefix-list

=====
                                PrefixList
=====

      PREFIX                MASKLEN FROM TO
-----

List 23   testPref:
      2.3.4.5                8        8   32
      2.3.6.0                8        8    8

List 34   testMore:
      34.1.1.0               24       24   24

-----

Name Appendix for Lists Converted from Old Config:
@A=conv addr list, @N=conv net list, @NR=conv net list modified as range
Passport-8010:5#
```

Showing information about route policies

To display information about the route policies configured on the switch, use the following command:

```
show ip route-policy info
```

Figure 203 displays sample output for the **show ip route-policy info** command.

Figure 203 show ip route-policy info command

```

Passport-8010:5# show ip route-policy info
=====
                        Route Policy
=====
NAME                               SEQ    MODE  EN
-----
ripAnn                             23     PRMT  EN
int-126.100.100.100                1     PRMT  DIS
ripAccept                           23     PRMT  DIS
test                               5     PRMT  DIS

Passport-8010:5#

```

Showing information about OSPF accept policies

To display information about the all configured OSPF entries, use the following command:

```
show ip ospf accept info
```

Figure 204 shows sample output for the **show ip ospf accept info** command.

Figure 204 show ip ospf accept info command output

```
Passport-8010:5 config ip ospf accept adv-rtr 0.0.0.0# show ip os accept info

=====
                                Ospf Accept
=====

ADV_RTR          MET_TYPE  ENABLE  POLICY
-----
0.0.0.0          any       TRUE    xxx
2.2.2.2          any       FALSE   osacc
```

Showing information about OSPF route redistribute policies

To display information about the OSPF redistribution configuration for each route source that is static, direct, and RIP, use the following command:

```
show ip ospf redistribute info
```

[Figure 205](#) shows sample output for the **show ip ospf redistribute info** command.

Figure 205 show ip ospf redistribute command output

```
Passport-8606:6# show ip ospf redistribute info

=====
                                Ospf Redistribute List
=====

SRC  COMM LV  LPRF MET  MTyp  NHOP          ORGN SRCLVL SUBNT    ENABLE
RPolicy
-----
STAT 0    0   0    0   type2 0.0.0.0      0    0    allow    FALSE
```

You can apply one policy for one purpose, for example, RIP Announce, on a given RIP interface. In that case, all sequence numbers under the given policy are applied to that filter. A sequence number also acts as an implicit preference; a lower sequence number is preferred.

Chapter 15

Configuring RSMLT using Device Manager and the CLI

This chapter describes how to configure and display RSMLT information on a VLAN interface using Device Manager and the CLI. For conceptual information about RSMLT, see [Chapter 1, “IP routing concepts,” on page 31](#).

This chapter includes the following topics:

Topic	Page
Configuring RSMLT on a VLAN using Device Manager	511
Configuring RSMLT on a VLAN using the CLI	516

Configuring RSMLT on a VLAN using Device Manager


RSMLT can be configured per IP VLAN interface. The IP routing protocol should be enabled on those layer 3 interfaces. VLANs with those layer 3 interfaces should also participate in SMLT.

To configure RSMLT on a VLAN:


- 1 From the Device Manager menu bar, choose VLAN > VLANs.

The VLAN dialog box opens with the Basic tab displayed ([Figure 206](#)).

- 2 Select a VLAN.
- 3 Click the IP button.

Figure 206 VLAN dialog box—Basic tab

- 4 Click RSMLT tab.
- The RSMLT tab opens (Figure 207).

Figure 207 IP, VLAN2 dialog box—RSMLT tab

- 5 Select Enable.
- 6 In the HoldDownTimer field, enter a hold down timer value.
- 7 In the HoldUpTimer field, enter a hold up timer value.
- 8 Click Apply.


Viewing and editing RSMLT local information

To view and edit RSMLT local VLAN switch information:

- 1 From the Device Manager menu bar, choose IP Routing > RSMLT.

The RSMLT dialog box opens with the Local tab displayed ([Figure 208](#)).

Figure 208 RSMLT dialog box—Local tab



- 2 Enter the appropriate fields.
- 3 Click Apply.

[Table 52](#) describes the RSMLT Local tab fields.


Table 52 RSMLT dialog box—Local tab fields

Field	Description
IfIndex	This is the IP route Smlt operation index.
VlanId	The VLAN ID of the chosen VLAN.
IpAddr	The IP address of the VLAN when RSMLT is enabled.
MacAddr	The MAC address of the selected VLAN.
Enable	This field displays the RSMLT operating status as enabled or disabled.
OperStatus	This field displays the RSMLT operating status as either up or down.
HoldDownTimer	The HoldDownTimer defines for how long the RSMLT switch does not participate in L3 forwarding. It is recommended to configure this value somewhat longer than the anticipated routing protocol convergence. This field displays the hold down timer value, the range of value is from 0 to 3600 seconds.
HoldUpTimer	This field displays the hold up timer value. The HoldUpTimer defines for how long the RSMLT switch maintains forwarding for its peer. The value is a range from 0 to 3600 seconds or 9999. 9999 which means infinity.
SmltId	The id range for the SMLT. A valid range is 1 to 32.
SltdId	The id range for the SLT. A valid range is 1 to 512.

Viewing and editing RSMLT peer information

To view and edit RSMLT peer switch information:

- 1 From the Device Manager menu bar, choose IP Routing > RSMLT.
The RSMLT dialog box opens with the Local tab displayed ([Figure 208 on page 513](#)).
- 2 Click Peer.
The Peer tab opens ([Figure 209](#)).

Figure 209 RSMLT dialog box—Peer tab

- 3 Enter the appropriate fields.
- 4 Click Apply.

[Table 53](#) describes the RSMLT Peer tab fields.

Table 53 RSMLT dialog box—Peer tab fields

Field	Description
IfIndex	This is the IP route Smlt operation index.
VlanId	The VLAN Id of the chosen VLAN.
IpAddr	The IP address of the VLAN when RSMLT is enabled.
MacAddr	The MAC address of the selected VLAN.
Enable	This field displays the RSMLT operating status as enabled or disabled.
OperStatus	This field displays the RSMLT operating status as either up or down.
HoldDownTimer	The HoldDownTimer defines for how long the RSMLT switch does not participate in L3 forwarding. It is recommended to configure this value somewhat longer than the anticipated routing protocol convergence. This field displays the hold down timer value, the range of value is from 0 to 3600 seconds.
HoldUpTimer	This field displays the hold up timer value. The HoldUpTimer defines for how long the RSMLT switch maintains forwarding for its peer. This field displays the hold up timer value. The value is a range from 0 to 3600 seconds or 9999. 9999 which means infinity.
HoldDownTimerRemaining	This field displays the time remaining of the HoldDownTimer.
HoldUpTimerRemaining	This field displays the time remaining of the HoldUpTimer.
SmltId	The id range for the SMLT. A valid range is 1 to 32.
Sltd	The id range for the SLT. A valid range is 1 to 512.

Configuring RSMLT on a VLAN using the CLI

RSMLT can be configured per routed IP VLAN. The IP routing protocol should be enabled on those layer 3 interfaces. VLANs with those layer 3 interfaces should also participate in SMLT.

To create a RSMLT on a VLAN, use the following command:

```
config vlan <vid> ip rsmlt
```

The command includes the following parameters:

config vlan <vid> ip rsmlt followed by:	
< vid >	is the VLAN id.
info	Displays the RSMLT local and peer information.
disable	Disables RSMLT on the VLAN.
enable	Enables RSMLT on the VLAN.
holddown-timer <seconds>	The HoldDownTimer defines for how long the RSMLT switch does not participate in L3 forwarding. It is recommended to configure this value somewhat longer than the anticipated routing protocol convergence. <seconds> the timer value in seconds. The range of the value is from 0 to 3600 seconds.
holdup-timer <seconds>	This field displays the hold up timer value. The HoldUpTimer defines for how long the RSMLT switch maintains forwarding for its peer. <seconds> the timer value in seconds. This field contains the hold up timer value. The value is a range from 0 to 3600 seconds or 9999. 9999 which means infinity.

Figure 210 shows sample output for the **config vlan ip rsmlt info** command.

Figure 210 config vlan <vid> ip rsmlt info command output

```
8610:5# conf vlan 112 ip rsmlt disable
8610:5# conf vlan 112 ip rsmlt info

Sub-Context: clear config dump monitor show test trace wsm asfm
sam
Current Context:

                admin-status : disable
                holddown-timer : 60
                holdup-timer : 180
8610:5#
8610:5# conf vlan 112 ip rsmlt enable
8610:5# conf vlan 112 ip rsmlt info

Sub-Context: clear config dump monitor show test trace wsm asfm
sam
Current Context:

                admin-status : enable
                holddown-timer : 60
                holdup-timer : 180
8610:5#
8610:5# conf vlan 112 ip rsmlt holddown-timer 70
8610:5# conf vlan 112 ip rsmlt info

Sub-Context: clear config dump monitor show test trace wsm asfm
sam
Current Context:

                admin-status : enable
                holddown-timer : 70
                holdup-timer : 180
8610:5# conf vlan 112 ip rsmlt holdup-timer 200
8610:5# conf vlan 112 ip rsmlt info

Sub-Context: clear config dump monitor show test trace wsm asfm
sam
Current Context:

                admin-status : enable
                holddown-timer : 70
                holdup-timer : 200
```

Showing IP RSMLT information

The **show ip rsmlt info** command displays RSMLT information on the interface. If a VLAN ID or an IP address is entered, the information is displayed only for that VID or for that interface; if not, all RSMLT interfaces are listed.

This command uses the syntax:

```
show ip rsmlt info [<local/peer>]
```

[Figure 211](#) shows sample output for the **show ip rsmlt info** command.

Figure 211 show ip rsmlt info local/peer command output

```
TOKYO>:5# show ip rsmlt info local

=====
                                Ip Rsmlt Local Info
=====
VID      MAC                      ADMIN  OPER  HDTMR  HUTMR
-----
41       00:04:38:8c:72:04  Enable  Up    60     180
112      00:04:38:8c:72:03  Enable  Up    60     180

VID      SMLT ID                      SLT ID
-----
41       3
112      1

TOKYO>:5# show ip rsmlt info peer

=====
                                Ip Rsmlt Peer Info
=====
VID      MAC                      ADMIN  OPER  HDTMR  HUTMR
-----
41       00:e0:7b:c9:c6:00  Enable  Up    60     180
112      00:e0:7b:c9:c6:03  Enable  Up    60     180

VID      HDT REMAIN  HUT REMAIN  SMLT ID                      SLT ID
-----
41       60          180          3
112      60          180          1
```

Index

Numbers

10000MbpsPortDefaultMetric field
 OSPF General tab 354

1000MbpsPortDefaultMetric field
 OSPF General tab 354

100MbpsPortDefaultMetric field
 OSPF General tab 354

10MbpsPortDefaultMetric field
 OSPF General tab 354

A

ActiveCount field
 OSPF Areas tab 377

Addr field
 Addresses tab 213

Address field 484
 Configuration tab 325

Address Resolution Protocol
 See ARP.

Addresses tab
 accessing 213
 fields 213

AddressLessIf field 387
 OSPF Interfaces tab 358

AddressLessIndex field
 OSPF Neighbors tab 366

AdminStat field
 OSPF General tab 354
 OSPF Interfaces tab 359

Advertise Metric field 393

Advertisement field 391

AdvertisementInterval field
 Interface tab 439
 Port, Insert VRRP dialog box 443

AdvertiseWhenDown field 320
 OSPF tab, VLAN 371

advertise-when-down option
 OSPF 425, 431

advertising interval, VRRP 451

AdvertisingRtr field 479

Age field
 Ext. Link State DB tab 391
 Link State Database tab 390
 Routes tab 215

alternate route 42

alternative routes 245

alternative routes, enabling 210

AltSequence field
 Routes tab 215

Area field 382

area, not so stubby 411

area, OSPF, stub 411

AreaBdrRtrCount field
 OSPF Areas tab 376

AreaBdrRtrStatus field
 OSPF General tab 354

AreaID field 381
 Area Aggregate tab 393
 Host tab 384

AreaId field
 Link State Database tab 389
 OSPF Interfaces tab 359
 OSPF tab, VLAN 371

- Stub Area Metrics tab 388
- AreaLSACKsumSum field
 - OSPF Areas tab 377
- AreaSummary field
 - OSPF Areas tab 377
- ARP
 - address-resolution cache 53
 - disabling on brouter ports 292
 - enabling on brouter ports 292
 - IP address 53
 - MAC address 53
 - managing 293
 - proxy ARP 55
 - request 53
 - static entries 53
 - table 53
 - viewing 293
- ARP commands
 - configure 306
 - IP 307
- ARP tab
 - accessing 293
 - fields 292, 294
- ARP table, adding static entry 307, 311
- ARP table, displaying 313
- ASBdrRtCount field
 - OSPF Areas tab 377
- ASBdrRtrStatus checkbox
 - OSPF General tab 354
- ASBR, specifying 384
- authentication key
 - OSPF 431
- authentication key, OSPF 407, 413, 425
- authentication type
 - OSPF 425, 431
- authentication type, OSPF 408, 413
- AuthKey field 325
 - OSPF Interfaces tab 359
 - OSPF tab, VLAN 371
- AuthType field 325

- OSPF Interfaces tab 359
- OSPF tab, VLAN 371
- AutoAggregate field 327
- AutoAggregateEnable field 320
- automatic route aggregation 327, 336, 341, 346
- automatic virtual link 379
- Autonomous System boundary router, OSPF 403
- AutoVirtLinkEnable checkbox
 - OSPF General tab 355
- AutoVirtLinkEnable field 379

B

- BackupDesignatedRouter
 - OSPF Interfaces tab 359
- BackUpMaster field 441
 - Port, Insert VRRP dialog box 444
- BackUpMastrState field 441
- BcastAddr field 214
- black hole routes 221
- black hole static routes 40
- broadcast interface, OSPF
 - VLAN option 371
- brouter port 321
 - bridging traffic 38
 - description 38
 - IP routing 38
 - nonroutable traffic 36
 - spanning tree state 38
- brouter port, creating 276
- brouter ports 200

C

- Checksum field
 - Ext. Link State DB tab 391
 - Link State Database fields 390
- config ethernet ip rip commands 336, 341
- config ethernet ip vrrp commands 450, 454
- config ethernet ip arp-response info command 302

- config ethernet ip command 276
- config ethernet ip directed-broadcast command 278
- config ethernet ip info command 277
- config ethernet ip ospf command 425
- config ethernet ip ospf info command 427
- config ethernet ip proxy command 279, 302
- config ethernet ip rip command 341
- config ethernet ip vrrp info command 452
- config ip arp command 307, 311
- config ip arp info command 308
- config ip command 245
- config ip commands 245, 254, 496
- config ip forwarding command 248
- config ip info command 247
- config ip mroute interface command 269
- config ip mroute static-source-group command 270
- config ip ospf accept adv-rtr command 501
- config ip ospf accept apply command 503
- config ip ospf area commands 410
- config ip ospf area info command 411
- config ip ospf area range commands 412
- config ip ospf area virtual-interface commands 412
- config ip ospf command 403
- config ip ospf host-route commands 404
- config ip ospf info command 404
- config ip ospf interface command 407
- config ip ospf interface info command 409, 410
- config ip ospf neighbor command 414
- config ip ospf redistribute apply command 506
- config ip ospf redistribute command 504
- config ip prefix-list command 494
- config ip prefix-list info command 495
- config ip rip command 332
- config ip rip commands 332
- config ip rip info command 333, 338
- config ip rip interface command 336
- config ip route command 248
- config ip route preference command 249
- config ip route-policy seq command 254, 496
- config ip static-route command 262
- config ntp command 288
- config route-discovery command 251
- config vlan ip rip commands 345
- config vlan ip vrrp commands 454, 516
- config vlan ip commands 284
- config vlan ip directed-broadcast command 285
- config vlan ip ospf info command 432
- config vlan ip proxy command 286, 304
- config vlan ip proxy info command 305
- config vlan ip rip command 345
- config vlan ip rip info command 348
- config vlan ip vrrp info command 456, 517
- configuration
 - Vrrp on a VLAN 436, 444
- Configuring ethernet ip commands 275
- Control field
 - Interface tab 438
 - Port, Insert VRRP dialog box 443
- conventions, text 28
- Cost field 320, 328
- CriticalIPAddr field
 - Interface tab 439
- CriticalIpAddr field
 - Port, Insert VRRP dialog box 444
- CriticalIpAddrEnable field
 - Interface tab 439
 - Port, Insert VRRP dialog box 444
- customer support 30

D

- dead interval
 - OSPF 408, 413, 425, 431
- default metric information, OSPF 417
- DefaultListen field 320, 327
- DefaultSupply field 320, 327
- DefImportMetric 318
- deleting, L2/L2 static routes 222
- denial of service 277, 285
- designated router
 - OSPF 69
- DesignatedRouter field
 - OSPF Interfaces tab 359
- DesigRtrPriority field
 - OSPF tab, VLAN 371
- Dest field
 - IP, Insert Static Routes dialog box 220
 - Routes tab 215
 - Static Routes tab 220
- directed broadcast, suppressing
 - port 277, 285
- disabling forwarding 248
- Domain field 325
- DoProxy field 292
- DoResp field 292

E

- ECMP
 - description 42
 - enabling globally 209
- Effect field 393
- Enable field
 - IP, Insert Static Routes dialog box 220
 - OSPF tab, VLAN 370
 - Static Routes tab 220
- Enable RIP field 319
- Equal Cost MultiPath. See ECMP

- Ethernet port commands
 - OSPF show 427
 - show IP info 278
- Events field 381
 - OSPF Interfaces tab 360
 - OSPF Neighbors tab 366
 - Virtual Neighbor tab 382
- ExternalSACKsumSum field
 - OSPF General tab 354
- ExternLsaCount field
 - OSPF General tab 354

F

- FastAdvertisementEnable field 439
 - Port, Insert VRRP dialog box 443
- FastAdvertisementInterval field 439
 - Port, Insert VRRP dialog box 444
- filtering
 - inbound/outbound traffic on a RIP interface 483, 485
- Flows tab
 - fields 224
- flushing routing tables 224

G

- global config ip commands 245
- global parameters, RIP 316, 321
- Globals tab
 - accessing 206
 - fields 208, 211, 212
- Globals tab — RIP 321

H

- hello interval, OSPF
 - interface 408
 - port 426
 - virtual interface 413
 - VLAN 431
- Hello Protocol 69

Hello Suppressed
 OSPF Neighbors tab 366

HelloInterval field
 OSPF Interfaces tab 359
 OSPF tab, VLAN 370
 Virtual Neighbor tab 381

HelloSuppressed field 382

HoldDown Time field 318

holddown timer

 OSPF 403
 RIP 332

HoldDownState field 441

HoldDownTimer field
 Port, Insert VRRP dialog box 444
 Secondary Feature tab 441

HoldDownTimeRemaining field 441

HopOrMetric field 215

I

ICMP router discovery
 configuring using the Device Manager
 globally 232
 on a port 236
 on a VLAN 234
 displaying information for
 all interfaces 281
 ports 282
 VLANs 281
 viewing the router discovery table 232

ID field 464

Id field
 OSPF Areas tab 376

IfIndex field
 IP, Insert Static Routes dialog box 220
 Static Routes tab 220

IfType field
 OSPF tab, VLAN 371

InPolicy field 320, 327, 484, 487

Insert ARP dialog box 295

Interface field
 ARP tab 294
 Interface tab 438
 Routes tab 215

interface information, VRRP, displaying 458, 519

interface statistics, OSPF, displaying 417

Interface tab
 accessing 437
 fields 438
interface, OSPF
 type (broadcast, passive, NBMA)
 VLAN option 371

Interior Gateway Protocol (IGP) 64

IP address
 assigning to a VLAN 203

ip address
 configuring on a brouter port 201

IP commands
 configure 245, 254, 496
 show 270

IP configuration dialog box, accessing 211

IP enhancements and policies
 alternate route 42
 description 41
 ECMP 42
 IP prefix list 50
 IP route policy 50

IP forwarding commands 271

IP forwarding, disabling 248

IP forwarding, enabling 206

IP Globals tab
 fields 210, 229

IP policy commands
 show 507

IP router, managing 211

IP routes 214

IP routing
 address classes 32
 Address Resolution Protocol (ARP) 53

- address, in dotted-decimal notation 32
- alternate routes 42
- brouter ports 38
- CIDR (classless interdomain routing)
 - address 36
- connectivity protocols 53
- IP enhancements and policies 41, 42
- multicast addresses 32
- OSPF benefits 65
- OSPF description 64
- OSPF v2 61
- RIP 62
- RIP v1 61
- RIP v2 61
- static routes 39
- UDP broadcast 56
- unicast addresses 32
- virtual address 58
- VRRP 58

IP static routes, creating 218

IP, Insert Flows dialog box

- fields 224

IP, VLAN, Insert VRRP dialog box

- accessing 444

IpAddr field

- Interface tab 438
- OSPF Neighbors tab 366
- Port, Insert VRRP dialog box 443
- Virtual Neighbor tab 382

IpAddress field

- ARP tab 294
- Host tab 384
- If Metrics tab 387
- OSPF Interfaces tab 358
- Port, Insert IP Address dialog box 203

isolated routing port 201, 203

L

LastSpfRun field

- OSPF General tab 355

link state database 420

Listen field 320, 327

LocalAddr field 486

LocalNextHop field 220

LSACount field

- OSPF Areas tab 377

LsdbType field 393

LSID field

- Ext. Link State DB tab 391
- Link State Database tab 389

LSRetransQLen field 382

M

MacAddress field

- ARP tab 294

manual virtual link 81

Mask field

- Area Aggregate tab 393
- IP, Insert Static Routes dialog box 220
- Routes tab 215
- Static Routes tab 220

MaskLenFrom field 464

MaskLenUpto field 464

MasterIpAddr field 439

MatchAsPath field 473

MatchCommunity field 474

MatchCommunityExtract field 474

MatchInterface field 473

MatchMetric field 473

MatchProtocol field 472

MatchRouteType field 473

MatchTag field 474

md5 key 407, 413

metric

- OSPF 426, 432

Metric field 395, 481

- Host tab 384
- IP, Insert Static Routes tab 220
- Metrics tab 387

- OSPF tab, VLAN 371
- Static Routes tab 220
- Stub Area Metrics tab 388
- metric speed, OSPF 385
- Metric Type field 388
- MetricType field 395, 479, 482

N

- Name field 464
- NBMA interface, OSPF
 - VLAN option 371
- Neighbor field
 - Virtual Neighbor tab 381
- neighbors, OSPF 423
- NetBIOS
 - name service 56
- NetMask field
 - Addresses tab 214
 - Port, Insert IP Address dialog box 203
- NextHop field
 - IP, Insert Static Route dialog box 220
 - Routes tab 215
 - Static Routes tab 220
- non-broadcast multiaccess. *See* NBMA
- NSSA (not so stubby area) 411
- NssaPbit field 474

O

- Open Shortest Path First (OSPF) Protocol. *See* OSPF
- Open Shortest Path First. *See* OSPF
- Open Shortest Path First. *See* OSPF
- OperAction field 444
 - Secondary Feature tab 441
- Operation field 317
- Options field
 - OSPF Neighbors tab 366
 - Virtual Neighbor tab 382
- OriginateNewLSas field
 - OSPF General tab 354
- OSPF 351
 - adjacent routers 70
 - area border router (ABR) 71
 - areas 67
 - AS boundary router (ASBR) 71, 81
 - AS external link advertisement 79
 - AS external routes 80
 - ASBR summary link advertisement 79
 - autonomous system external (ASE) routes 80
 - backbone area 67
 - backup designated router (BDR) 69, 71
 - benefits 65
 - best path 82
 - database description (DD) packets 78
 - definition 351
 - description 64
 - designated router (DR) 69, 71
 - hello packets 78
 - Hello Protocol 69
 - implementation
 - ASBR 384
 - creating a virtual link 378
 - stub area 377
 - importing small stub routing domains 68
 - interfaces description 72
 - internal router (IR) 71
 - IP 77
 - link state acknowledgements 78
 - link state advertisements 79
 - link state request packets 78
 - link state update packets 78
 - link-state database 65
 - link-state information 70
 - manual aggregation 68
 - metric speed 82
- NBMA
 - forming adjacencies 75
 - hello packets 75
 - neighbors list and priorities 74
 - PollInterval 74
- NBMA adjacencies 70
- neighbor adjacencies 70

- neighbors 69
- neighbors on NBMA networks 69
- network links advertisement 79
- network summary link advertisement 79
- not so stubby area 68
- packets 78
- passive interface description 77
- route advertisement 77
- router links advertisement 79
- router types 71
- routing algorithm 66
- specifying ASBRs 81
- stub area 68
- transit areas 68
- variable-length mask 77
- virtual link 70, 80
- virtual links 80
- OSPF Accept policy 394, 477, 479
- OSPF advertise-when-down option 425
- OSPF area
 - configuring 411
 - parameters, displaying 415
 - stub 411
- OSPF Area dialog box 376
- OSPF authentication type 408
- OSPF Autonomous System External (ASE) link
 - state advertisements, displaying 416
- OSPF commands
 - configure 403–414
- OSPF dead interval 408
- OSPF default metrics 403
- OSPF error information, displaying 428
- OSPF hello interval
 - interface 408
 - port 426
 - virtual interface 413
 - VLAN 431
- OSPF host route 405
- OSPF host route configuration, displaying 417
- OSPF interface timer information, displaying 420

- OSPF Neighbors dialog box 365
- OspfAction field
 - OSPF General tab 355
- ospfImportASExtern field
 - OSPF Areas tab 376
- OspfInFilterApply field 477
- ospfNbmaNbrPermanence field
 - OSPF Neighbors tab 366
- OutPolicy field 320, 328, 484, 487

P

- passive interface, OSPF
 - VLAN option 371
- Path Type field
 - Routes tab 216
- Poison field 320, 327
- policies
 - OSPF Accept 394, 477, 479
- PolicyName field 479
- poll interval
 - OSPF 431
- polling interval, OSPF 408
- PollInterval field
 - OSPF Interfaces tab 360
 - OSPF tab, VLAN 371
- port commands
 - show IP info 278
- port OSPF commands
 - show 427
- port RIP commands
 - configure 336, 341
- Port, Insert IP Address dialog box
 - accessing 201
 - fields 202
- Port, Insert VRRP dialog box
 - accessing 442
 - fields 443
- Preference field
 - IP, Insert Static Routes dialog box 220

- Static Routes tab 220
- Prefix field 464
- prefix list, configuring 462
- PrefixMaskLen field 464
- priority
 - OSPF 359, 409, 426, 432
 - VRRP 452
- Priority field
 - Interface tab 438
 - OSPF Neighbors tab 366
 - Port, Insert VRRP dialog box 443
- product support 30
- Proto field 215
- proxy ARP, configuring 296
- proxy ARP, enabling 302, 305
- publications
 - hard copy 29

Q

- Queries field 318

R

- RARP (Reverse Address Resolution Protocol)
 - ARP 57
 - Ethernet type 57
 - protocol-based VLAN 57
 - request 57
 - server 57
 - VLAN 57
- RcvBadPackets field 323
- RcvBadRoutes field 324
- ReasmMaxSize field
 - Addresses tab 214
- Receive field 325
- RedistributeApply field 477
- Retrans Interval field
 - Virtual Neighbor tab 381
- RetransInterval field

- OSPF Interfaces tab 360
- retransmit interval, OSPF 409, 414
- Reverse Address Resolution Protocol. See RARP
- RIP
 - advertise-when-down option 341, 346
 - dialog box 318
 - enabling globally 332
 - filtering inbound/outbound traffic 483, 485
 - global parameters 316
 - poison option 337, 342, 347
 - RcvBadPackets counter 323
 - RcvBadRoutes counter 324
 - supply and listen settings 343
 - triggered updates counter 324
- RIP (Routing Information Protocol)
 - advertisements 62
 - distance vector protocol 62
 - hop count 62
 - metric 62
 - multicast advertisements 63
 - routing table 62
 - TCP/IP route information 62
 - UDP 62
 - variable length subnet masks (VLSM) 63
 - version 1 63
 - version 2 63
- RIP commands
 - IP 332
 - port 336, 341
 - show 339
 - VLAN 345, 349
- RIP global parameters 316, 321
- RIP In/Out Policy tab
 - accessing 483, 485
- RIP interface status 323
- RIP status
 - RcvBadPackets field 323
 - RcvBadRoutes 324
 - SentUpdates 324
- RIP update 342
- RouteChanges field 318

RoutePolicy field 396, 482

router interfaces 200

router tables, flushing 224

RouterID field

Ext. Link State DB tab 391

Link State Database tab 389

OSPF General tab 354

routes

alternative 210

black hole 221

Routes tab

accessing 214

fields 215

RouteSource field 395, 481

Routing Information Protocol. See RIP

Routing Information Protocol. See RIP

Routing Information Protocol. See RIP

routing port

isolated 203

routing port, isolated 201

routing tables

flushing 224

IP 214

routing tables flushing 56

RtrDeadInterval field

OSPF Interfaces tab 360

OSPF VLAN tab 370

Virtual Neighbor tab 381

RtrId field

OSPF Neighbors tab 366

Virtual Neighbor tab 382

Rtrpriority field

OSPF Interfaces tab 359

RxNewLSas field

OSPF General tab 354

S

Secondary Features tab

accessing 440

Send field 325

SentUpdates 324

Sequence field

Ext. Link State DB tab 391

Link State Database tab 389

SetAsPath field 475

SetAsPathMode field 475

SetAutomaticTag field 475

SetCommunityMode field 475

SetCommunityNumber field 475

SetInjectNetList field 474

SetLocalPref field 475

SetMask field 474

SetMetric field 474

SetMetricType field 474

SetMetricTypeInternal field 474

SetNextHop field 474

SetOrigin field 475

SetOriginEgpAs field 475

SetRoutePreference field 474

SetTag field 475

SetWeight field 475

show ip arp info command 313

show IP commands 270

show ip forwarding command 271

show ip interface command 271

show ip ospf area command 415

show ip ospf ase command 416

show ip ospf default-metric command 417

show ip ospf host-route command 417

show ip ospf ifstats command 417

show ip ospf info command 418

show ip ospf interface command 419

show ip ospf int-timers command 420

show ip ospf lsdb command 420

show ip ospf neighbors command 423

- show ip ospf range command 423
 - show ip ospf stats command 423
 - show ip prefix-list command 507
 - show ip rip info command 339
 - show ip rip interface command 340, 509, 510
 - show ip route info command 272
 - show ip route preference info command 250
 - show ip route-discovery command 272
 - show ip route-policy info command 261, 508
 - show ip static-route info command 273
 - show ip vrrp info command 458, 519
 - show port vrrp commands 453
 - show ports error ospf command 428
 - show ports info arp command 302
 - show ports info brouter-port command 276
 - show ports info ip command 278
 - show ports info ospf command 428
 - show ports info rip command 344
 - show ports stats ospf extended command 430
 - show ports stats ospf main command 429
 - show vlan info arp command 305
 - show vlan info ospf command 433
 - show vlan info rip command 349
 - show vlan info vrrp extended command 457
 - SpfHoldDownTime checkbox
 - OSPF General tab 355
 - SpfRuns field
 - OSPF Areas tab 376
 - State field 381, 438
 - OSPF Interfaces tab 359
 - OSPF Neighbors tab 366
 - Virtual Neighbor tab 382
 - static ARP entries 294
 - static default routes, definition 220
 - static entry in ARP table 307, 311
 - static routes 39
 - black hole static routes 40
 - deleting 222
 - overview 217
 - Static Routes tab
 - accessing 218
 - static routes, creating 218
 - static routes, IP, displaying 273
 - Status field
 - If Metrics tab 387
 - IP, Insert Static Routes dialog box 220
 - Static Routes tab 220
 - STP
 - spanning tree convergence 37
 - stub area, creating 377
 - subnet
 - mask 33
 - variable-length 34
 - Subnets field 396, 482
 - supernet
 - address/mask pair 36
 - classless interdomain routing (CIDR)
 - address 36
 - contiguous network addresses 35
 - supply and listen settings, RIP 343
 - Supply field 319, 327
 - support, Nortel Networks 30
- ## T
- technical publications 29
 - technical support 30
 - text conventions 28
 - time to live, setting 208, 246
 - TOS field
 - Host tab 384
 - IF Metrics tab 387
 - Stub Area Metrics tab 388
 - TransitDelay field
 - OSPF Interfaces tab 360
 - Virtual Neighbor tab 381

- transmit delay, OSPF 414
- TrapEnable checkbox
 - OSPF General tab 354
- TriggeredUpdate field 327
- TriggeredUpdateEnable field 320
- Type field
 - ARP tab 294
 - Ext. Link State DB tab 391
 - Link State Database tab 389
 - OSPF Interfaces tab 359

U

- UDP
 - broadcast forwarding 56
 - IP limited broadcast 56
 - MAC-level broadcast 56
 - RIP 62
 - specified protocol 57
 - TTL value 57
- Update Time field 317
- update timer, RIP 332
- User Datagram Protocol. See UDP

V

- VersionNumber field
 - OSPF General tab 354
- VirtIf 379
- virtual interface, OSPF area 412
- virtual link, creating 378
- virtual router interfaces 200
- Virtual Router Redundancy Protocol. See VRRP
- virtual routing interfaces
 - configuring RIP global parameter 321
 - configuring RIP global parameters 316
- VirtualMacAddr field 438
- VirtualRouter UpTime field 439
- VLAN
 - RARP 57

- VLAN IP commands 283, 287
- VLAN OSPF commands
 - show 433
- VLAN RIP commands
 - configure 345
 - show 349
- VlanID 203
- VLANs
 - VRRP 436, 444
- VrId field 438, 441
- Vrid field 443
- VRRP 436, 444
 - advertisement timer 60
 - advertisements 59
 - ARP request 60
 - backup state 60
 - configuring 437, 440, 442, 444
 - controlling state 60
 - default gateway 58
 - dynamic default gateway 58
 - forwarding router 60
 - load sharing 58
 - MAC address 60
 - overview 435
 - primary router 58
 - primary router backup 59
 - router 58, 59
 - virtual primary router 59
 - virtual router ID 59
 - virtual router IP address 59
 - virtual router MAC address 60
- VRRP commands
 - show 458, 519
- VRRP configuration
 - extended, displaying 457
- VRRP failover mechanism 452
- VRRP tab
 - accessing 442, 444
 - fields 443
- VRRPs, Hold Down Timer feature 435