NORTEL

Nortel Ethernet Routing Switch 5000 Series

Configuration — IP Routing Protocols

Release: 6.1
Document Revision: 05.01

www.nortel.com

NN47200-503

Nortel Ethernet Routing Switch 5000 Series
Release: 6.1
Publication: NN47200-503
Document release date: 25 May 2009

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

While the information in this document is believed to be accurate and reliable, except as otherwise expressly agreed to in writing NORTEL PROVIDES THIS DOCUMENT "AS IS" WITHOUT WARRANTY OR CONDITION OF ANY KIND, EITHER EXPRESS OR IMPLIED. The information and/or products described in this document are subject to change without notice.

Nortel, Nortel Networks, the Nortel logo, and the Globemark are trademarks of Nortel Networks.

All other trademarks are the property of their respective owners.

Contents

New in this release 15
Features 15
IGMP Send Query 15

Introduction 17
NNCLI command modes 17

IP Routing Fundamentals 21
IP addressing overview 21
Subnet addressing 23
IP routing 24
IP routing using VLANs 24
Local routes 25
Local and non-local static routes 27
Default routes 27
Management VLAN 27
Multinetting 29
Brouter port 31
Routing Information Protocol (RIP) 32
RIP operation 32
RIP metrics 32
Routing updates 33
Split horizon 34
Triggered updates 34
RIP Send and Receive Modes 34
Supported RIP capabilities on the 5000 switch 36
Open Shortest Path First (OSPF) protocol 37
Overview 37
Autonomous system and areas 38
ASBR and external route advertisements 40
OSPF neighbors 41
Designated routers 41
OSPF operation 42
OSPF route advertisements 42
Router types 42

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

LSA types 43
Area types 44
OSPF virtual link 48
OSPF host route 50
OSPF interfaces 50
OSPF packets 51
OSPF metrics 52
Automatic router ID change in a stack 53
OSPF security mechanisms 53
Equal Cost MultiPath (ECMP) 54
Route policies 54
Route policies in a stack 56
DHCP relay 56
Forwarding DHCP packets 57
Multiple DHCP servers 58
Differences between DHCP and BootP 59
UDP broadcast forwarding 59
Directed broadcasts 60
ARP 61
Static ARP 61
Proxy ARP 62
IP blocking for stacks 63
Virtual Router Redundancy Protocol (VRRP) 64
VRRP operation 65
VRRP topology example 66
Critical IP address 68
VRRP and SMLT 68
VRRP fast advertisement interval 70

IP multicast fundamentals 71
Overview of IP multicast 71
Multicast groups 74
Multicast distribution trees 74
Multicast addresses 76
IP multicast address ranges 76
IP to Ethernet multicast MAC mapping 77
Internet Group Management Protocol 78
IGMPv1 operation 79
IGMPv2 operation 80
IGMPv3 operation 83
Multicast flow over Multi-Link Trunking 84
IGMP requests for comment 84
IGMP snooping 85
IGMPv3 snooping 87

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IGMP proxy 88
Report forwarding 89
Static mrouter port and nonquerier 90
Unknown multicast packet filtering 91
IGMP snooping configuration rules 92
IGMP and stacking 93
Default IGMP values 93
IGMP snooping interworking with Windows clients 93
IGMP Send Query 95
Protocol Independent Multicast-Sparse Mode 95
PIM-SM concepts and terminology 96
PIM-SM shared trees and shortest-path trees 100
Source-to-RP SPT 103
Register suppression timeout 104
Receivers leaving a group 104
PIM assert 104
PIM passive interfaces 104
PIM-SM capabilities and limitations 105
Enabling or disabling routing with IGMP enabled 107
Nonsupported IGMP features 107
Default PIM-SM values 107

IP routing configuration using NNCLI 111
IP routing configuration procedures 111
Configuring global IP routing status 112
Displaying global IP routing status 112
Configuring an IP address for a VLAN 113
Configuring IP routing status on a VLAN 114
Configuring a secondary IP address for a VLAN 114
Displaying the IP address configuration and routing status for a VLAN 115
Displaying IP routes 116
Performing a traceroute 117
Entering Router Configuration mode 118

Brouter port configuration using NNCLI 121
Configuring a brouter port 121
Displaying the brouter port configuration 122

Static route configuration using NNCLI 125
Configuring a static route 125
Displaying static routes 126
Configuring a management route 127
Displaying the management routes 128

OSPF configuration using NNCLI 129
Configuring the OSPF hardware mode 131

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring the router ID 131
Configuring global OSPF status 131
Configuring global OSPF parameters 132
Displaying global OSPF parameters 133
Configuring OSPF area parameters 133
Displaying OSPF area configuration 135
Displaying OSPF area range information 135
Enabling OSPF routing on an interface 136
Assigning an interface to an OSPF area 136
Configuring the OSPF properties for an interface 137
Displaying OSPF interface timers 138
Displaying OSPF interface configurations 139
Displaying OSPF neighbors 139
Specifying a router as an ASBR 140
Configuring the OSPF authentication type for an interface 140
Defining simple authentication keys for OSPF interfaces 141
Defining MD5 keys for OSPF interfaces 141
Displaying OSPF MD5 keys 142
Applying an MD5 key to an OSPF interface 142
Displaying OSPF interface authentication configuration 143
Configuring a virtual link 144
Creating a virtual interface message digest key 145
Configuring automatic virtual links 146
Displaying OSPF virtual links 149
Displaying OSPF virtual neighbors 149
Configuring an OSPF host route 149
Displaying OSPF host routes 150
Displaying the OSPF link state database 150
Displaying the external link state database 151
Initiating an SPF run to update the OSPF LSDB 151
Displaying OSPF default port metrics 152
Displaying OSPF statistics 152
Displaying OSPF interface statistics 152
Clearing OSPF statistics counters 153

OSPF configuration examples using NNCLI 155

Basic OSPF configuration examples 155
- Basic OSPF configuration 155
- Basic ASBR configuration 157
- Configuring ECMP for OSPF 158

Advanced OSPF configuration examples 158
- Configuring an IP OSPF interface 158
- OSPF security configuration example using Message Digest 5 160
- Configuring OSPF network types 161

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring Area Border Routers (ABR) 162
Configuring Autonomous System Border Routers (ASBR) 164
Stub area configuration example 167
NSSA configuration example 170
Controlling NSSA external route advertisements 172
Configuring a multi-area complex 174
Diagnosing neighbor state problems 211

RIP configuration using NNCLI 215
RIP configuration procedures 215
Configuring the global RIP status 216
Configuring the RIP global timeout, holddown timer, and update timer 216
Configuring the default RIP metric value 218
Displaying global RIP information 218
Configuring the RIP status on an interface 219
Configuring RIP parameters for an interface 220
Displaying RIP interface configuration 222
Manually triggering a RIP update 224

RIP configuration examples using NNCLI 225
RIP configuration tasks 225
Configuring RIP 227
Configuring RIP version 2 230
Using RIP accept policies 232
Using RIP announce policies 234

ECMP configuration using NNCLI 237
Configuring the number of ECMP paths allotted for RIP 237
Configuring the number of ECMP paths allotted for OSPF 238
Configuring the number of ECMP paths allotted for static routes 239
Displaying ECMP path information 239
ECMP configuration examples 240
Displaying the IP routing table 242
Displaying global ECMP configuration 242

Route policies configuration using NNCLI 243
Route policies configuration procedures 243
Configuring prefix lists 244
Configuring route maps 244
Displaying route maps 247
Applying a RIP accept (in) policy 247
Applying a RIP announce (out) policy 248
Configuring an OSPF accept policy 249
Applying the OSPF accept policy 250
Displaying the OSPF accept policy 250
Configuring an OSPF redistribution policy 251

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Applying the OSPF redistribution policy 252
Displaying the OSPF redistribution policy 252

DHCP relay configuration using NNCLI 253
DHCP relay configuration procedures 253
Configuring global DHCP relay status 254
Displaying the global DHCP relay status 254
Specifying a local DHCP relay agent and remote DHCP server 255
Displaying the DHCP relay configuration 256
Configuring DHCP relay status and parameters on a VLAN 256
Displaying the DHCP relay configuration for a VLAN 257
Displaying DHCP relay counters 258
Clearing DHCP relay counters for a VLAN 259

UDP broadcast forwarding configuration using NNCLI 261
UDP broadcast forwarding configuration procedures 261
Configuring UDP protocol table entries 262
Displaying the UDP protocol table 262
Configuring a UDP forwarding list 263
Applying a UDP forwarding list to a VLAN 264
Displaying the UDP broadcast forwarding configuration 265
Clearing UDP broadcast counters on an interface 266

Directed broadcasts configuration using NNCLI 269
Configuring directed broadcasts 269
Displaying the directed broadcast configuration 270

Static ARP and Proxy ARP configuration using NNCLI 271
Static ARP configuration 271
Configuring a static ARP entry 271
Displaying the ARP table 272
Displaying ARP entries 272
Configuring a global timeout for ARP entries 273
Clearing the ARP cache 274
Proxy ARP configuration 274
Configuring proxy ARP status 275
Displaying proxy ARP status on a VLAN 275

IP blocking configuration using NNCLI 277
Configuring IP blocking for a stack 277
Displaying IP blocking status 277

VRRP configuration using NNCLI 279
VRRP configuration procedures 279
Configuring global VRRP status 280
Assigning an IP address to a virtual router ID 281
Configuring the router priority for a virtual router ID 281

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring the status of the virtual router 282
Configuring a backup master 282
Configuring the critical IP address 283
Configuring the VRRP critical IP status 283
Configuring the VRRP holddown timer 284
Configuring VRRP holddown action 284
Configuring the VRRP advertisement interval 285
Configuring the fast advertisement interval 285
Configuring fast advertisement status 286
Configuring ICMP echo replies 287
Enabling VRRP traps 287
Displaying VRRP configuration and statistics 288

VRRP configuration examples using NNCLI 291
Configuring normal VRRP operation 291
Configuration command listing 296
Configuring VRRP with SMLT 297
Configuration command listing 300
Configuring VRRP with SLT 301

IGMP snooping configuration using NNCLI 305
IGMP snooping configuration procedures 305
Job aid: Roadmap of IGMP NNCLI commands 306
Configuring IGMP snooping on a VLAN 307
Configuring IGMP send query on a VLAN 308
Prerequisites 308
Procedure steps 308
Configuring IGMP proxy on a VLAN 308
Configuring the IGMP version on a VLAN 309
Configuring static mrouter ports on a VLAN 310
Displaying IGMP snoop, proxy, and mrouter configuration 311
Configuring IGMP parameters on a VLAN 312
Configuring the router alert option on a VLAN 313
Displaying IGMP interface information 314
Displaying IGMP group membership information 316
Configuring unknown multicast packet filtering 317
Displaying the status of unknown multicast packet filtering 318
Specifying a multicast MAC address to be allowed to flood all VLANs 318
Displaying the multicast MAC addresses for which flooding is allowed 319
Displaying IGMP cache information 319
Flushing the router table 320
Configuring IGMP selective channel block 321
Creating an IGMP profile 321
Deleting an IGMP profile 321
Applying the IGMP filter profile on interface 322

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Removing a profile from an interface 322
Displaying an IGMP profile 322

PIM-SM configuration using NNCLI 323
PIM-SM configuration procedures 324
Job aid: Roadmap of PIM-SM configuration commands 325
Enabling and disabling PIM-SM globally 327
Configuring global PIM-SM properties 327
Displaying global PIM-SM properties 328
Enabling PIM-SM on a VLAN 329
Configuring the PIM-SM interface type on a VLAN 330
Displaying PIM-SM neighbors 331
Configuring PIM-SM properties on a VLAN 331
Displaying the PIM-SM configuration for a VLAN 332
Specifying the router as a candidate BSR on a VLAN 333
Specifying a local IP interface as a candidate RP 335
Displaying the candidate RP configuration 336
Displaying the PIM-SM RP set 337
Displaying the active RP per group 337
Enabling and disabling static RP 338
Configuring a static RP 339
Displaying the static RP configuration 340
Specifying a virtual neighbor on an interface 340
Displaying the virtual neighbor configuration 341
Displaying PIM multicast routes 341
Displaying the PIM mode 342
Displaying multicast route information 342

PIM-SM configuration example using NNCLI 345

IP routing configuration using Device Manager 357
IP routing configuration procedures 357
Configuring global IP routing status and ARP lifetime 358
Configuring an IP address and enabling routing for a VLAN 358
Displaying configured IP Addresses 359

Static route configuration using Device Manager 361
Configuring static routes 361
Displaying IP Routes 362
Filtering route information 363
Displaying TCP information for the switch 364
Displaying TCP Connections 365
Displaying TCP Listeners 365
Displaying UDP endpoints 366

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Brouter port configuration using Device Manager 369
Configuring a brouter port 369

OSPF configuration using Device Manager 371
Configuring Global OSPF properties 372
Configuring an OSPF area 374
Configuring an area aggregate range 375
Configuring OSPF stub area metrics 376
Configuring OSPF interfaces 377
Configuring OSPF interface metrics 379
Defining MD5 keys for OSPF interfaces 379
Displaying OSPF neighbor information 380
Configuring an OSPF virtual link 381
Configuring an automatic virtual link 383
Defining MD5 keys for OSPF virtual links 383
Displaying virtual neighbor information 384
Configuring OSPF host routes 385
Displaying link state database information 385
Displaying external link state database information 386
Displaying OSPF statistics 387
Displaying VLAN OSPF statistics 388

RIP configuration using Device Manager 391
RIP configuration procedures 391
Configuring Global RIP properties 392
Configuring a RIP interface 393
Configuring advanced RIP interface properties 394
Displaying RIP Statistics 395
Configuring RIP parameters for a VLAN 396

ECMP configuration using Device Manager 397
Configuring ECMP 397

Route policies configuration using Device Manager 399
Route policies configuration procedures 399
Creating a prefix list 400
Creating a route policy 401
Configuring RIP in and out policies 402
Configuring an OSPF Accept Policy 403
Configuring OSPF redistribution parameters 404
Applying an OSPF accept or redistribution policy 405

DHCP relay configuration using Device Manager 407
DHCP relay configuration procedures 407
Configuring DHCP Relay 408
Configuring DHCP parameters on a VLAN 408

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying and graphing DHCP counters on a VLAN 409

UDP broadcast forwarding configuration using Device Manager 411
UDP broadcast forwarding configuration procedures 411
Configuring UDP protocol table entries 412
Configuring UDP forwarding entries 412
Configuring a UDP forwarding list 413
Applying a UDP forwarding list to a VLAN 414

Static ARP and Proxy ARP configuration using Device Manager 417
Configuring static ARP entries 417
Configuring Proxy ARP 418

VRRP configuration using Device Manager 421
VRRP configuration procedures 421
Configuring global VRRP status and properties 422
Assigning an IP address to a virtual router ID 422
Configuring VRRP interface properties 423
Graphing VRRP interface information 425
Viewing general VRRP statistics 426

IGMP snooping configuration using Device Manager 427
IGMP snooping configuration procedures 427
Configuring IGMP snoop, proxy, and IGMP parameters on a VLAN 428
Configuring IGMP snoop, proxy, and static mrouter ports on a VLAN 429
Configuring IGMP router alert, snoop, and proxy on a VLAN 431
Flushing the IGMP router tables and configuring IGMP router alert 432
Configuring unknown multicast filtering 435
Specifying a multicast MAC address to be allowed to flood all VLANs 435
Displaying IGMP cache information 436
Displaying IGMP group information 437
Displaying multicast route information 438
Displaying multicast next hop information 438
Displaying multicast interface information 439
Configuring IGMP selective channel block using Device Manager 440
Variable definitions 440

PIM-SM configuration using Device Manager 443
PIM-SM configuration procedures 444
Configuring global PIM-SM status and properties 445
Configuring PIM-SM status and properties for a VLAN 447
Configuring PIM-SM VLAN properties from the IP Routing menu 448
Specifying the router as a candidate BSR on a VLAN interface 450
Setting the C-BSR priority using the VLAN menu 450

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Setting the C-BSR priority using the IP Routing menu 451
Displaying the current BSR 451
Specifying a local IP interface as a candidate RP 452
Displaying the active RP 453
Enabling static RP 454
Configuring a static RP 455
Specifying a virtual neighbor on an interface 456
Displaying PIM-SM neighbor parameters 456
Displaying the PIM-SM RP set 457

IGMP snooping configuration using Web-based management 459
Configuring IGMP snooping 459
Displaying multicast membership 460

DHCP Relay configuring using Web-based management 463
DHCP Relay configuration navigation 463
Configuring DHCP relay status 463
Creating a forwarding path 463
Editing a forwarding path 464
Configuring DHCP relay VLAN settings 464

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

New in this release

The following sections detail what's new in *Nortel Ethernet Routing Switch 5000 Series Configuration — IP Routing Protocols* (NN47200-503).

Features

See the following sections for information about feature changes:

- "IGMP Send Query" (page 15)

IGMP Send Query

The IGMP Send Query functionality allows a switch or stack to be configured as an active query router without the need for dedicating a standalone switch in each network to the task. For more information see:

- "IGMP Send Query" (page 95)
- "Configuring IGMP send query on a VLAN" (page 308)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Introduction

This document provides procedures and conceptual information to configure IP routing features on the Nortel Ethernet Routing Switch 5000 Series , including RIP, OSPF, VRRP, static routes, Proxy ARP, DHCP Relay, and UDP forwarding. It also provides procedures and conceptual information to manage multicast traffic using PIM-SM and IGMP snooping.

NNCLI command modes

NNCLI provides the following command modes:

- User EXEC
- Privileged EXEC
- Global Configuration
- Interface Configuration
- Router Configuration

Mode access is determined by access permission levels and password protection.

If no password is set, you can enter NNCLI in User EXEC mode and use the **enable** command to move to the next level (Privileged EXEC mode). However, if you have read-only access, you cannot progress beyond User EXEC mode, the default mode. If you have read-write access you can progress from the default mode through all of the available modes.

With sufficient permission, you can use the rules in the following table to move between the command modes.

Command mode and sample prompt	Entrance commands	Exit commands
User EXEC 5650TD>	No entrance command, default mode	**exit** or **logout**

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Command mode and sample prompt	Entrance commands	Exit commands
Privileged EXEC 5650TD#	**enable**	**exit** or **logout**
Global Configuration 5650TD(config)#	From Privileged EXEC mode, enter: **configure**	To return to Privileged EXEC mode, enter: **end** or **exit** To exit NNCLI completely, enter: **logout**
Interface Configuration 5650TD(config-if)#	From Global Configuration mode, to configure a port, enter: **interface fastethernet <port number>** To configure a VLAN, enter: **interface vlan <vlan number>**	To return to Global Configuration mode, enter: **exit** To return to Privileged EXEC mode, enter: **end** To exit NNCLI completely, enter: **logout**
Router Configuration 5650TD(config-router)#	From Global Configuration mode, to configure OSPF, enter: **router ospf** To configure RIP, enter: **router rip** To configure VRRP, enter:	To return to Global Configuration mode, enter: **exit** To return to Privileged EXEC mode, enter: **end**

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Command mode and sample prompt	Entrance commands	Exit commands
	`router vrrp`	To exit NNCLI completely, enter: `logout`

For more information, see *Nortel Ethernet Routing Switch 5000 Series Fundamentals* (NN47200-104).

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IP Routing Fundamentals

This chapter provides an introduction to IP routing and the IP routing protocols used in the Nortel Ethernet Routing Switch 5000 Series.

Navigation

- “IP addressing overview” (page 21)
- “IP routing” (page 24)
- “Routing Information Protocol (RIP)” (page 32)
- “Open Shortest Path First (OSPF) protocol” (page 37)
- “Equal Cost MultiPath (ECMP)” (page 54)
- “Route policies” (page 54)
- “DHCP relay” (page 56)
- “UDP broadcast forwarding” (page 59)
- “Directed broadcasts” (page 60)
- “ARP” (page 61)
- “Static ARP” (page 61)
- “Proxy ARP” (page 62)
- “IP blocking for stacks” (page 63)
- “Virtual Router Redundancy Protocol (VRRP)” (page 64)

IP addressing overview

An IP version 4 (IPv4) address consists of 32 bits expressed in a dotted-decimal format (XXX.XXX.XXX.XXX). The IPv4 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. The following table lists the breakdown of the IP address space by address range and mask.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Table 1
IP address classifications

Class	Address Range	Mask	Number of Networks	Nodes per Network
A	1.0.0.0 - 127.0.0.0	255.0.0.0	127	16 777 214
B	128.0.0.0 - 191.255.0.0	255.255.0.0	16 384	65 534
C	192.0.0.0 - 223.255.255.0	255.255.255.0	2 097 152	255
D	224.0.0.0 - 239.255.255.254			
E	240.0.0.0 - 240.255.255.255			

Note 1: Although technically part of Class A addressing, network 127 is reserved for loopback.

Note 2: Class D addresses are primarily reserved for multicast operations although the addresses 224.0.0.5 and 224.0.0.6 are used by OSPF and 224.0.0.9 is used by RIP.

Note 3: Class E addresses are reserved for research purposes.

To express an IP address in dotted-decimal notation, each octet of the IP address is converted to a decimal number and separated by decimal points. For example, the 32-bit IP address 10000000 00100000 00001010 10100111 is expressed in dotted-decimal notation as 128.32.10.167.

Each IP address class, when expressed in binary notation, has a different boundary point between the network and host portions of the address, as shown in the following figure. 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.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 1
Network and host boundaries in IP address classes

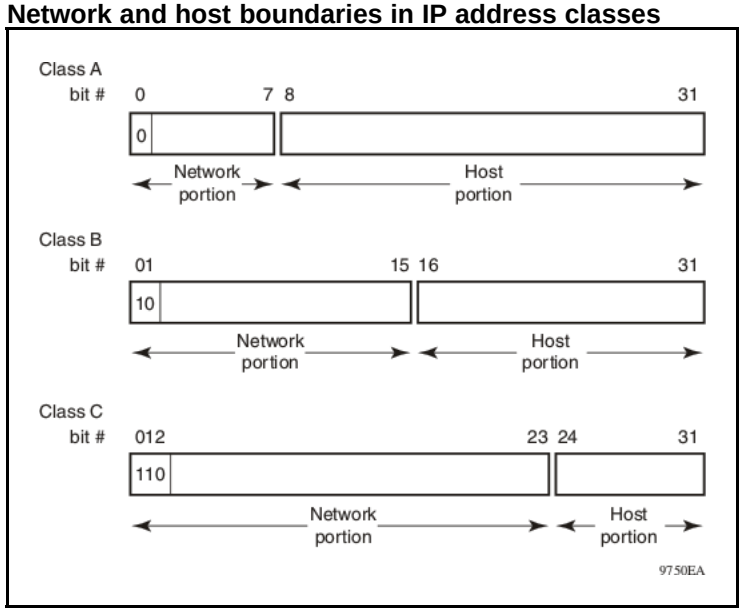

Subnet addressing

Subnetworks (or subnets) are an extension of the IP addressing scheme. Subnets allow 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).

A subnet address is created 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.

The following table illustrates how subnet masks used with Class B and Class C addresses can create differing numbers of subnets and hosts. This example shows the use of the zero subnet, which is permitted on a Nortel Ethernet Routing Switch 5000 Series.

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

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Table 2
Subnet masks for Class B and Class C IP addresses (cont'd.)

Number of bits	Subnet Mask	Number of Subnets (Recommended)	Number of Hosts per Subnet
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) is the ability to divide an intranet into pieces that match network requirements. Routing is based on the longest subnet mask or network that matches.

IP routing

To configure IP routing on the Nortel Ethernet Routing Switch 5000 Series, you must create virtual router interfaces by assigning an IP address to a virtual local area network (VLAN). The following sections provide more details about IP routing functionality.

For a more detailed description about VLANs and their use, see *Nortel Ethernet Routing Switch 5000 Series Configuration — VLANs, Spanning Tree, and Link Aggregation* (), (NN47200-502).

IP routing using VLANs

The Nortel Ethernet Routing Switch 5000 Series supports wire-speed IP routing between VLANs. To create a virtual router interface for a specified VLAN, you must associate an IP address with the VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The virtual router interface is not associated with any specific port. The VLAN IP address can be reached through any of the ports in the VLAN. The assigned IP address also serves as the gateway through which packets are routed out of that VLAN. Routed traffic can be forwarded to another VLAN within the switch or stack.

There is not a one-to-one correspondence between the physical port and the routable interface, because a given port can belong to multiple routable VLANs.

When the Nortel Ethernet Routing Switch 5000 Series is routing IP traffic between different VLANs, the switch is considered to be running in Layer 3 mode; otherwise, the switch runs in Layer 2 mode. When you assign an IP address to a Layer 2 VLAN, the VLAN becomes a routable Layer 3 VLAN. You can assign a unique IP address to each VLAN.

You can configure the global status of IP routing to be enabled or disabled on the Nortel Ethernet Routing Switch 5000 Series. By default, IP routing is disabled.

All IP routing parameters can be configured on the Nortel Ethernet Routing Switch 5000 Series before routing is actually enabled on the switch.

In this release, the Nortel Ethernet Routing Switch 5000 Series supports local routes, static routes, and dynamic routes. With local routing, the switch automatically creates routes to each of the local Layer 3 VLAN interfaces. With static routing, you must manually enter the routes to the destination IP addresses. With dynamic routing, routes are identified using a routing protocol such as RIP or OSPF.

Local routes

With routing globally enabled, if you assign an IP address to a VLAN, IP routing is enabled for that VLAN. In addition, for each IP address assigned to a VLAN interface, the Ethernet Routing Switch adds a directly connected or local route to its routing table based on the IP address/mask assigned.

Local routing example

The following figure shows how the Ethernet Routing Switch can route between Layer 3 VLANs. In this example, the Ethernet Routing Switch has two VLANs configured. IP Routing is enabled globally on the switch and on the VLANs, each of which has an assigned IP address.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 2
Local routes example

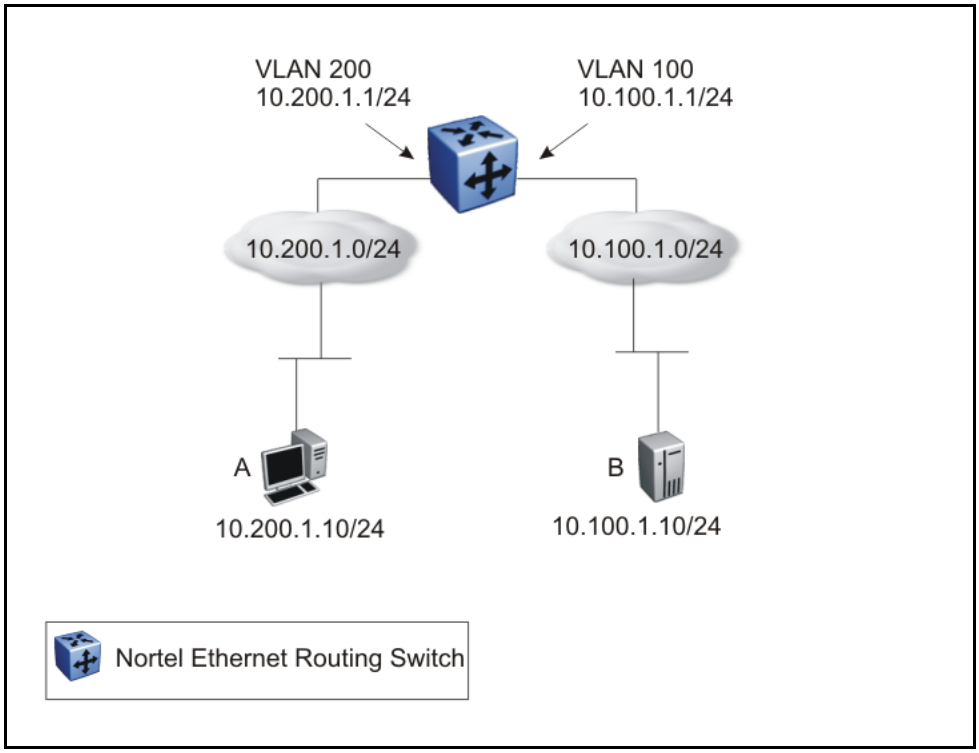

IP address 10.100.1.1/24 is assigned to VLAN 100, and IP address 10.200.1.1/24 is assigned to VLAN 200. As IP Routing is enabled, two local routes become active on the Nortel Ethernet Routing Switch as described in the following table.

	Network	Net-mask	Next-hop	Type
1	10.100.1.0	255.255.255.0	10.100.1.1	LOCAL
2	10.200.1.0	255.255.255.0	10.200.1.1	LOCAL

At this stage, both hosts A (10.200.1.10) and B (10.100.1.10) are reachable from the Ethernet Routing Switch. However, to achieve Layer 3 connectivity between A and B, additional configuration is required. Host A must know how to reach network 10.100.1.0/24, and host B must know how to reach network 10.200.1.0/24.

On host A, you must configure a route to network 10.100.1.0/24 through 10.200.1.1, or configure 10.200.1.1 as the default gateway for the host.

On host B, you must configure a route to network 10.200.1.0/24 through 10.100.1.1, or configure 10.100.1.1 as the default gateway for the host.

With these routes configured, the Ethernet Routing Switch can perform inter-VLAN routing, and packets can flow between hosts A and B.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Local and non-local static routes

After you create routable VLANs through IP address assignment, you can create static routes. With static routes, you can manually create specific routes to destination IP addresses. In this release, the Ethernet Routing Switch supports local and non-local static routes. Local routes have a next-hop that is on a directly connected network, while non-local routes have a next-hop that is not on a directly connected network. Non-local static routes are useful in situations where there are multiple paths to a network and the number of static routes can be reduced by using only one route with a remote gateway.

Static routes are not easily scalable. Thus, in a large or growing network this type of route management may not be optimal. Also, static routes do not have the capacity to determine the failure of paths. Thus, a router can still attempt to use a path after it has failed.

Default routes

Default routes specify a route to all networks for which there are no explicit routes in the Forwarding Information Base or the routing table. This static default route is a route to the network address 0.0.0.0 as defined by the Institute of Electrical and Electronics Engineers (IEEE) Request for Comment (RFC) 1812 standard.

The Ethernet Routing Switch uses the default route 0.0.0.0/0.0.0.0 for all Layer 3 traffic that does not match a specific route. This traffic is forwarded to the next-hop IP address specified in the default route.

Management VLAN

With IP routing enabled on the switch or stack, you can use any of the virtual router IP addresses for device management over IP. Any routable Layer 3 VLAN can carry the management traffic, including Telnet, Web, Simple Network Management Protocol (SNMP), BootP and Trivial File Transfer Protocol (TFTP). Without routing enabled, the management VLAN is reachable only through the switch or stack IP address, and only through ports that are members of the management VLAN. The management VLAN always exists on the switch and cannot be removed.

When routing is enabled on the Nortel Ethernet Routing Switch 5000 Series switches, the management VLAN behaves similar to other routable VLANs. The IP address is reachable through any virtual router interface, as long as a route is available.

Management route

On the Ethernet Routing Switch, you can configure a management route from the Management VLAN to a particular subnet. The management route is a static route that allows incoming management connections from the remote network to the management VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The management route transports traffic between the specified destination network and the Management VLAN only. It does not carry inter-VLAN routed traffic from the other Layer 3 VLANs to the destination network. This provides a management path to the router that is inaccessible from the other Layer 3 VLANs. While you can access the management VLAN from all static routes, other static routes cannot route traffic to the management route.

To allow connectivity through a management route, IP Routing must also be enabled globally and on the management VLAN interface.

The following figure shows an example of a management route allowing access to the management VLAN interface.

Figure 3
Management route

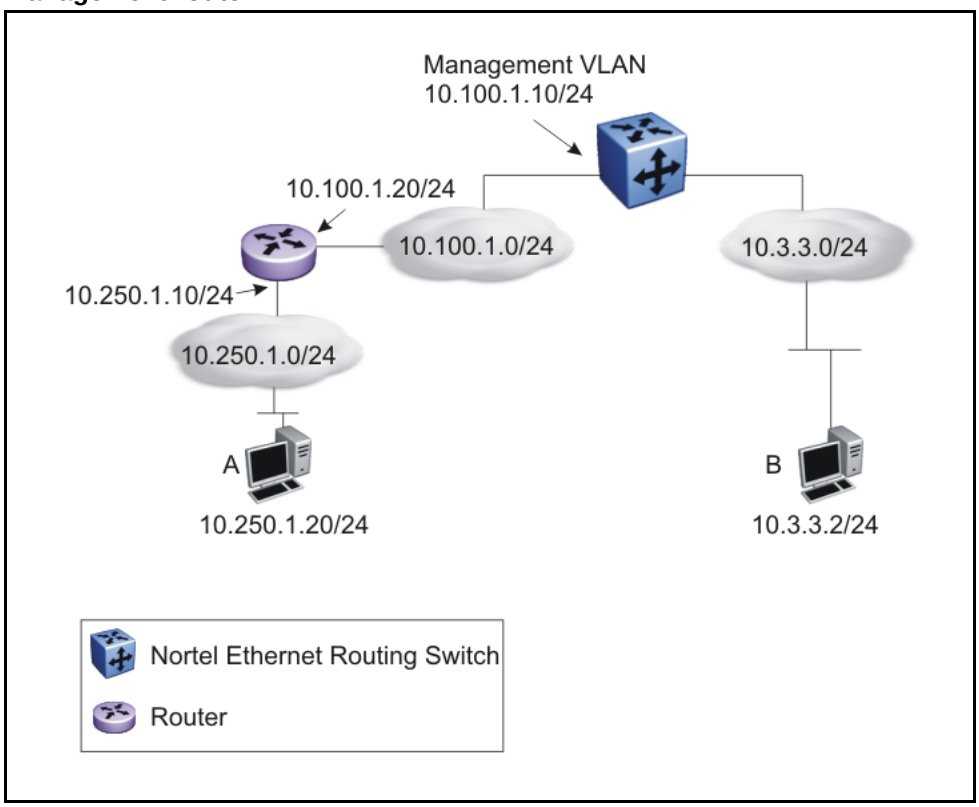

As network 10.250.1.0/24 is not directly connected to the Ethernet Routing Switch, to achieve connectivity from host 10.250.1.20 to the management VLAN, the Ethernet Routing Switch must know how to reach network 10.250.1.0/24. On the Ethernet Routing Switch, you can configure a management route to network 10.250.1.0/24 through 10.100.1.20. In this case, the following management route is active on the Ethernet Routing Switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

	Network	Net-mask	Next-hop	Type
1	10.250.1.0	255.255.255.0	10.100.1.20	MANAGEMENT

With this configured route, host A at 10.250.1.20 can perform management operations on the Ethernet Routing Switch. To do so, Host A also requires a route to 10.100.1.0/24 using 10.250.1.10 as the next hop, or with 10.250.1.10 as the default gateway.

If a Layer 3 VLAN is also configured for network 10.3.3.0/24, this provides a local route that host B at 10.3.3.2 can use to access the switch. However, host B cannot communicate with host A as the route to network 10.250.1.0/24 is a management route only. To provide connectivity between the two hosts, a static route must be configured to 10.250.1.0/24.

Multinetting

The Nortel Ethernet Routing Switch 5000 Series supports the definition and configuration of up to eight secondary interfaces on each VLAN (multinetting). With IP multinetting, you can associate multiple IP subnets with one VLAN. That is, connected hosts can belong to different IP subnets on the same VLAN.

Multinetting can be configured using NNCLI or Device Manager.

The following diagram illustrates a network with configured IP multinetting.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 4
Network with Multinetting

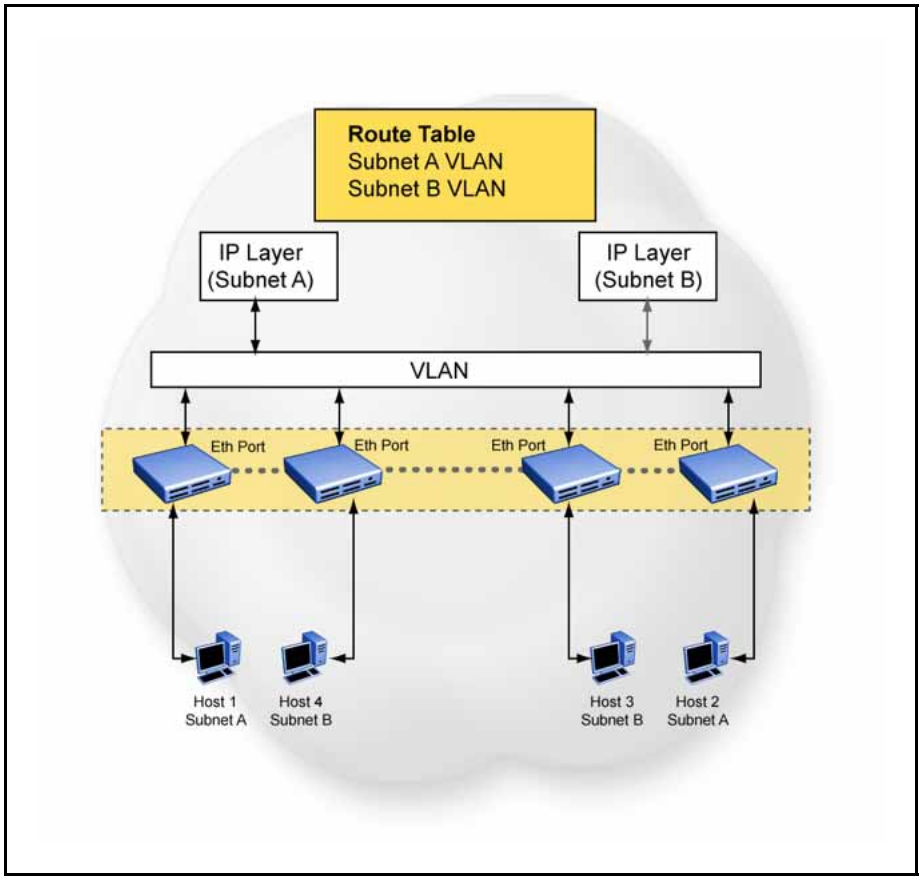

You can configure a static route with the next hop on the secondary interface. You can also add static ARP for a given IP address in the same subnet of a secondary interface.

Here are some limitations when you are working with secondary interfaces:

- you can have a maximum of eight secondary interfaces on each VLAN
- you can have a total maximum of 256 IP interfaces (including primary and secondary)
- all of the secondary interfaces on a VLAN are enabled or disabled together. There is no provision for configuring the administrative state of the secondary IP interfaces individually.
- dynamic routing is not available for secondary IP interfaces
- secondary interfaces are not supported on brouters
- a primary IP interface must be in place before secondary IP interfaces can be added; secondary interfaces must be deleted before you can delete the primary

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

If secondary interfaces are configured on the management VLAN, routing cannot be disabled globally or on the management VLAN. Secondary IP interfaces on the management VLAN are purged from NVRAM when

- a unit leaves the stack and the switch does not have a manually configured IP
- the switch fails to get the IP address through the BootP mode

The following are not supported on secondary interfaces:

- DHCP
- Proxy ARP
- UDP broadcast
- IPFIX
- VRRP, OSPF, RIP

Brouter port

The Nortel Ethernet Routing Switch 5000 Series supports the configuration of brouter ports. 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 configured for 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 non-routable 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.

When a brouter port is created, the following actions also take place on the switch:

- A port-based VLAN is created.
- The brouter port is added to the new port-based VLAN.
- The PVID of the brouter port is changed to the VLAN ID of the new VLAN.
- The STP participation of the brouter port is disabled.
- An IP address is assigned to the brouter VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Routing Information Protocol (RIP)

Routing Information Protocol (RIP) is a standards-based, dynamic routing protocol based on the Bellman-Ford (or distance vector) algorithm. It is used as an Interior Gateway Protocol (IGP). RIP allows routers to exchange information to compute the shortest routes through an IPv4-based network. The hop count is used as a metric to determine the best path to a remote network or host. The hop count cannot exceed 15 hops (the distance from one router to the next is one hop).

RIP is defined in RFC 1058 for RIP version 1 and RFC 2453 for RIP version 2. The most significant difference between the two versions is that, while RIP version 1 is classful, RIP version 2 is a classless routing protocol that supports variable length subnet masking (VLSM) by including subnet masks and next hop information in the RIP packet.

RIP operation

Each RIP router maintains a routing table, which lists the optimal route to every destination in the network. Each router advertises its routing information by sending routing information updates at regular intervals. Neighboring routers use this information to recalculate their routing tables and retransmit the routing information. For RIP version 1, no mask information is exchanged; the natural mask is always applied by the router receiving the update. For RIP version 2, mask information is always included.

RIP uses User Datagram Protocol (UDP) data packets to exchange routing information.

The sequence of processes governed by RIP is as follows:

1. When a router starts, it initializes the RIP data structures and then waits for indications from lower-level protocols that its interfaces are functional.
2. RIP advertisements are sent on all the interfaces that are configured to send routing information.
3. The neighbors send their routing tables and the new router updates its routing table based on the advertisements received.
4. From then on, each router in the network sends periodic updates to ensure a correct routing database.

RIP metrics

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. The distance from one router to the next is

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

considered to be one hop. This cost or hop count is known as the metric. The following figure shows the hop counts between various units in a network.

Figure 5
RIP hop counts

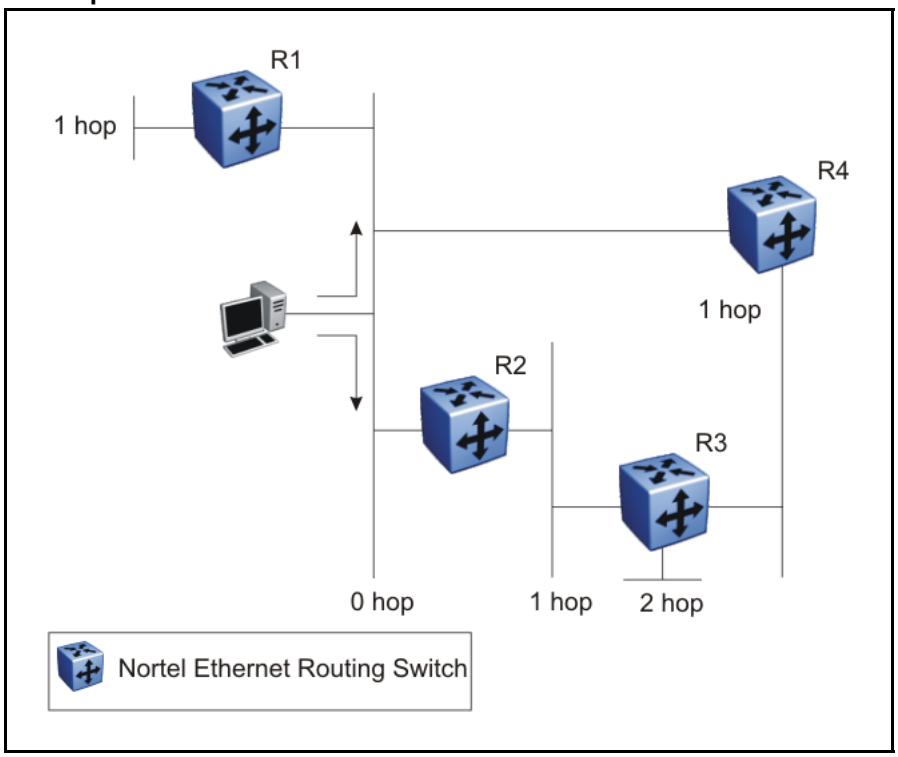

A directly connected network has a metric of zero. An unreachable network has a metric of 16. Therefore, 15 hops or 15 routers is the highest possible metric between any two networks.

Routing updates

Each RIP router advertises routing information updates out of all RIP-enabled interfaces at regular intervals (30 seconds by default). You can configure this interval using the update timer parameter. The routing updates contain information about known networks and the distances (hop count) associated with each. For RIP version 1, no mask information is exchanged; the natural mask is always applied by the router receiving the update. With RIP version 2, mask information is always included.

If a RIP router does not receive an update from another RIP router within a timeout period (180 seconds by default), it deletes the routes advertised by the nonupdating router from its routing table. You can configure this interval using the timeout interval parameter.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The router keeps aged routes from nonupdating routers temporarily in a garbage list and continues to advertise them with a metric of infinity (16) for a holddown period (120 seconds by default), so that neighbors know that the routes are unreachable. You can configure this interval using the holddown timer parameter. If a valid update for a garbage route is received within the holddown period, the router adds the route back into its routing table. If no update is received, the router completely deletes all garbage list entries for the nonupdating router.

Split horizon

To prevent routing loops, RIP uses the mechanism of split horizon, with or without poison reverse. Simple split horizon means that IP routes learned from a neighbor are not advertised back in updates to that neighbor. Split horizon with poison reverse means that these routes are advertised back to the neighbor, but they are “poisoned” with a metric of 16, which represents infinite hops in the network. The receiver neighbor therefore ignores this route.

By default, RIP split horizon is enabled without poison reverse.

Triggered updates

To promote fast convergence, RIP also supports a triggered updates option. With triggered updates enabled, a router sends update messages whenever it changes the metric for a route, even if it is not yet time for a regular update message.

RIP Send and Receive Modes

RIP can be configured to use a number of different send and receive modes depending on the specifics of the network configuration.

The following table lists the send modes supported.

Table 3
RIP send modes

Send Mode	Description	Result
rip1comp	This mode is used to broadcast RIP version 2 updates using RFC 1058 route consumption rules. This is the default send mode for the Nortel Ethernet Routing Switch 5000 Series .	• 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 version 2 update, including network mask • RIP version = 2

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Table 3
RIP send modes (cont'd.)

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 version 1 update, no network mask included • RIP version = 1
rip2	This mode is used to broadcast multicast RIP version 2 updates.	• Destination MAC is a multicast, 01-00-5e-00-00-09 • Destination IP is the RIP version 2 multicast address, 224.0.0.9 • RIP Update is formed as a RIP version 2 update including network mask • RIP version = 2
nosend	No RIP updates are sent on the interface.	None

The following table lists the receive modes supported.

Table 4
RIP receive modes

Receive Mode	Result
rip1OrRip2	RIP version 1 or RIP version 2 updates are accepted.
rip1	RIP version 1 and RIP version 1 compatible updates only are accepted.
rip2	RIP version 2 updates only are accepted.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Supported RIP capabilities on the 5000 switch

RIP supports the following standard behavior:

- periodic RIP updates about effective best routes
- garbage collection
- split horizon with or without poison reverse
- triggered update for changed RIP routes
- unicast to the specific query requestor
- broadcast/multicast of regular and triggered updates
- subnet mask (RIP version 2)
- routing table update based on the received RIP message
- global update timer
- holddown timer and timeout timer per device and per interface
- cost per device and per interface
- equal cost multipath (ECMP)

The Nortel Ethernet Routing Switch 5000 Series implementation of RIP also supports the following features:

- in and out routing policies
- auto-aggregation (also known as *auto-summarization*) of groups of adjacent routes into single entries

Many RIP features are configurable. The actual behavior of the protocol depends on the feature configurations.

RIP limitations

RIP has the following limitations:

- The protocol is limited to networks whose longest path is 15 hops.
- The protocol depends on counting to infinity to resolve certain unusual situations.
- The protocol uses fixed metrics (the hop number) to compare alternative routes, as opposed to real-time parameters such as measured delay, reliability, or load.
- RIP does not support address-less links.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Open Shortest Path First (OSPF) protocol

Open Shortest Path First (OSPF) is a classless Interior Gateway Protocol (IGP) that distributes routing information between routers belonging to a single autonomous system (AS). An OSPF AS is generally defined as a group of routers in a network that run OSPF and that operate under the same administration. Intended for use in large networks, OSPF is a link-state protocol that supports variable length subnet masking (VLSM) and tagging of externally-derived routing information.

ATTENTION
The Nortel Ethernet Routing Switch 5000 Series implementation of OSPF only supports broadcast and passive interfaces. Point-to-point and NBMA interfaces are not supported.

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 usable interfaces and reachable neighbors. Each router periodically checks for changes in its local state and shares detected changes by flooding link-state advertisements (LSA) 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.

In large networks, OSPF offers the following benefits:

- Fast convergence

  When the network topology 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.

- Scalable

  Because OSPF does not use hop count in its calculation, the routing domain is scalable.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

OSPF routes IP traffic based on the destination IP address, subnet mask, and IP TOS.

The Ethernet Routing Switch 5000 Series implementation of OSPF does not support TOS-based routing.

Autonomous system and areas

In large OSPF networks with many routers and networks, the link-state database (LSDB) and routing table on each router can become excessively large. Large route tables and LSDBs consume memory. In addition, the processing of additional LSAs puts added strain on the CPU to make forwarding decisions. To reduce these undesired effects, an OSPF network can be divided into subdomains called areas. Each area comprises a number of OSPF routers that have the same area ID. Subdividing the AS into areas significantly reduces the amount of routing protocol traffic compared to treating the entire AS as a single link-state domain.

When a network is divided into multiple areas, each router within an area maintains an LSDB only for the area to which it belongs. Each area is identified by a unique 32-bit area ID, expressed in IP address format (x.x.x.x). Area 0.0.0.0 is known as the backbone area and distributes routing information to all other areas.

Within the AS, packets are routed 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. Intra-area routing protects the area from bad routing information because no routing information obtained from outside the area can be used.

If the source and destination of a packet reside in different areas, inter-area routing is used. Inter-area routing must pass through the backbone area.

ABR

A router attached to two or more areas inside an OSPF network is identified as an Area Border Router (ABR). Each ABR maintains a separate topological database for each connected area. ABRs play an important role in OSPF networks by condensing the amount of disseminated OSPF information from one area to another. When the AS is divided into multiple areas, each nonbackbone area must be attached to the backbone area through an (ABR).

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

For routers that are internal to an area (identified as internal routers), the impact of a topology change is localized to the area in which it occurs. However, ABRs must maintain an LSDB for each area to which they belong. ABRs advertise changes in topology from one area to another by advertising summary LSAs.

Backbone area

The backbone area connects nonbackbone areas to each other. Traffic forwarded from one area to another must travel through the backbone. The backbone topology dictates the paths used between areas. The topology of the backbone area is invisible to other areas and the backbone has no knowledge of the topology of nonbackbone areas.

The area ID 0.0.0.0 is reserved for the backbone area.

Area border routers (ABR) cannot learn OSPF routes unless they have a connection to the backbone. Inter-area paths are selected by examining the routing table summaries for each connected ABR.

In inter-area routing, a packet travels along three contiguous paths:

1. First, the packet follows an intra-area path from the source to an ABR, which provides the link to the backbone.
2. From the source ABR, the packet travels through the backbone toward the destination area ABR.
3. At the destination area ABR, the packet takes another intra-area path to the destination.

The following figure shows an OSPF AS divide into three areas: a backbone area, a stub area, and a not-so-stubby area (NSSA). (Stub areas and NSSAs are described in subsequent sections.)

The figure also shows ABRs connecting the areas to one another and Autonomous System Border Routers (ASBR) connecting two areas to external networks. ASBRs redistribute external static or RIP routes into the OSPF network.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 6
OSPF network

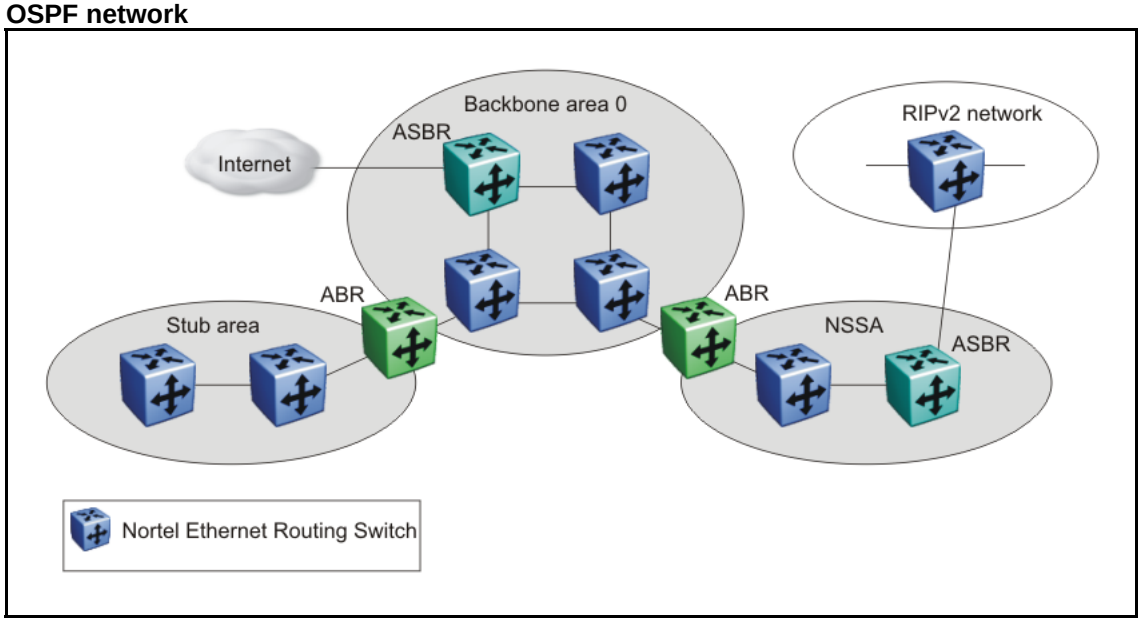

ASBR and external route advertisements

A router functions as an ASBR if one or more of its interfaces connect to a non-OSPF network (for example, RIP, or static routes). ASBRs advertise non-OSPF routes into OSPF domains.

An ASBR advertises external routes into the OSPF domain using autonomous system external (ASE) LSAs (LSA type 5). ASE LSAs flood across area borders. To conserve resources or control traffic flow, you can limit the number of ASBRs in your network.

OSPF considers the following routes to be 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

External route metrics

When an ASBR imports external routes, it imports OSPF route information using external type 1 or type 2 metrics. With external type 1 metrics, OSPF calculates the total cost by adding the external metric value and the internal cost to the ASBR. For external type 2 metrics, only the internal

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

OSPF cost to the ASBR is used in the routing decision. You can specify the metric type to use by configuring a route policy. For more information, see “Route policies” (page 54).

OSPF neighbors

In an OSPF broadcast network, any two routers that have an interface to the same network are neighbors. OSPF routers use the Hello Protocol to dynamically discover and maintain neighbor relationships.

Periodically, OSPF routers send Hello packets over all interfaces to the AllSPFRouters multicast address. These Hello packets include the following information:

- router priority
- router Hello Timer and Dead Timer values
- list of routers that sent the router Hello packets on this interface
- router choice for designated router (DR) and backup designated router (BDR)

Bidirectional communication is determined when a router discovers itself listed in its neighbor Hello packet.

Designated routers

To form an adjacency, two OSPF routers perform a database exchange process to synchronize their topological databases. When their databases are synchronized, the routers are said to be fully adjacent.

To limit the amount of routing protocol traffic, OSPF routers use the Hello Protocol to elect 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 (which on a large network can mean significant routing protocol traffic), all routers on the network form adjacencies with the DR and the BDR only, and send link-state information only to them. The DR redistributes this information to every other adjacent router.

The BDR receives link-state information from all routers on the network and listens for acknowledgements. If the DR fails, the BDR can transition quickly to the role of DR because its routing tables are up to date.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

OSPF operation

On broadcast multiaccess networks, the sequence of processes governed by OSPF is as follows:

1. When a router starts, it initializes the OSPF data structures and then waits for indications from lower-level protocols that the router interfaces are functional.
2. The router dynamically detects neighbors by sending and receiving Hello packets to the AllSPFRouters multicast address.
3. Using the Hello Protocol, a designated router (DR) and backup designated router (BDR) are elected for the network.
4. Each router forms an adjacency and exchanges database information only with the DR and the BDR.
5. The DR floods LSAs containing information about each router and its neighbors throughout the area to ensure that all routers in the area have an identical topological database.
6. From this database each router uses the OSPF routing algorithm (Dijkstra's algorithm) to calculate a shortest-path tree, with itself as root. This shortest-path tree in turn yields a routing table for the protocol.
7. After the network has converged, each OSPF router continues to periodically flood Hellos to maintain neighbor relationships. And at longer intervals, LSAs are retransmitted throughout the area. In addition, routers forwards LSAs to the DR if they detect a change in the state of a router or a link (that is, up or down). Upon receipt of an LSA, the DR can then flood the update to all routers in the area, enabling quick detection of dead routers on the network.

OSPF route advertisements

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

Because OSPF can specify a range of networks, it can 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.

Router types

As mentioned in preceding sections, routers in an OSPF network can have various roles depending on how you configure them. The following table describes the router types you can configure in an OSPF network.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Table 5
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 ASBR. Any router that distributes 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 ABR. ABRs play an important role in OSPF networks by condensing the amount of disseminated OSPF information.
Internal router (IR)	A router that has interfaces only within a single area inside an OSPF network is considered an IR. Unlike ABRs, IRs have topological information only about the area in which they are contained.
Designated router (DR)	In a broadcast network, a single router is elected to be the DR for that network. A DR ensures that all routers on the network are synchronized and advertises the network to the rest of the AS.
Backup designated router (BDR)	A BDR is elected in addition to the DR and, if the DR fails, can assume the DR role quickly.

LSA types

After the network has converged, OSPF does not require each router to keep sending its entire LSDB to its neighbors. Instead, each OSPF router floods only link-state change information in the form of LSAs throughout the area or AS. LSAs typically contain information about the router and its neighbors and are generated periodically to ensure connectivity or are generated by a change in state of the router or a link (that is, up or down).

The following table displays the seven LSA types exchanged between OSPF routers.

Table 6
OSPF LSA types

LSA type	LSA name	Description	Area of distribution
1	Router LSA	Type 1 LSAs are originated by every router to describe their set of active interfaces and neighboring routers. Type 1 LSAs are flooded only within the area. A backbone router can flood router link advertisements within the backbone area.	Only within the same area

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Table 6
OSPF LSA types (cont'd.)

LSA type	LSA name	Description	Area of distribution
2	Network LSA	Type 2 LSAs describe a network segment. In a broadcast network, the designated router (DR) originates network LSAs that list all routers on that LAN. Type 2 LSAs are flooded only within the area. A backbone DR can flood network links advertisements within the backbone area.	Only within the same area
3	Network-Summary LSA	Type 3 LSAs are originated by the area border router (ABR) to describe the networks that are reachable outside the area. An ABR attached to two areas generates a different network summary LSA for each area. ABRs also flood type 3 LSAs containing information about destinations within an area to the backbone area.	Passed between areas
4	ASBR-summary LSA	Type 4 LSAs are originated by the ABR to advertise the cost of the path to the closest ASBR from the router generating the advertisement.	Passed between areas
5	Autonomous System External [ASE] LSA	Type 5 LSAs are originated by the ASBR to describe the cost of the path to a destination outside the AS from the ASBR generating the advertisement. Type 5 LSAs are passed between areas. In stub and NSSA areas, type 5 LSA routes are replaced with a single default route.	Passed between areas
6	Group Membership LSA	Type 6 LSAs identify the location of multicast group members in multicast OSPF.	Passed between areas
7	NSSA External LSA	Type 7 LSAs are used in OSPF NSSAs to import external routes.	Translated between areas

Area types

OSPF supports multiple area types. The following sections describe the supported OSPF area types.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Stub area

As shown in the following figure, a stub area is configured at the edge of the OSPF routing domain and has only one ABR.

Figure 7
Stub area

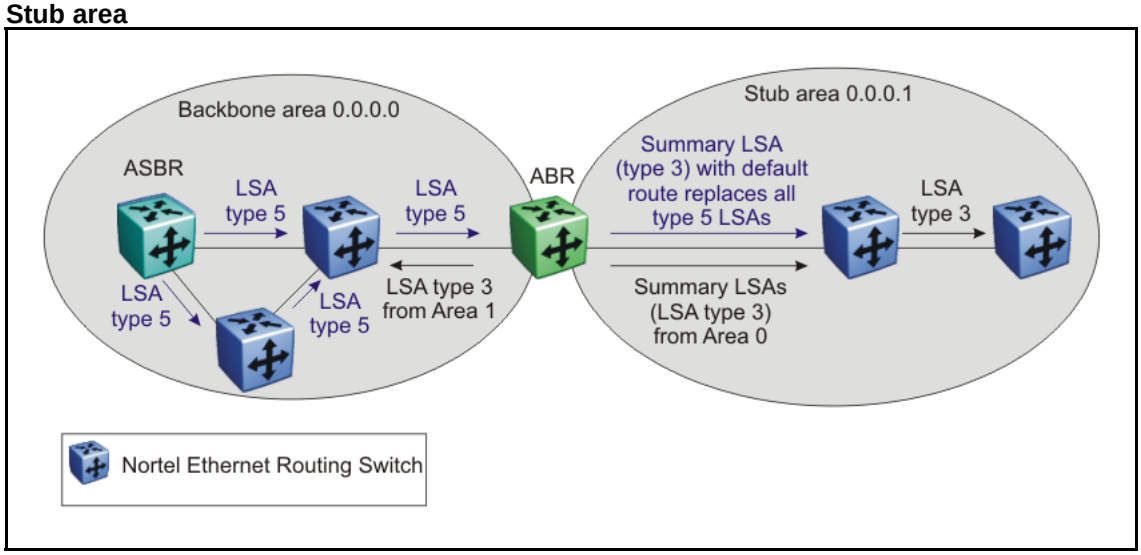

The ABR does not flood AS External LSAs (type 5) into a stub area. Instead, the ABR uses Summary LSAs (type 3) to advertise a default route (0.0.0.0) into the stub area for all external routes. As stub areas do not receive advertisements for external routes from the ABR, the size of the link state database in the stub area is reduced.

For internal routers in the stub area, any destinations that do not match intra-area or inter-area routes are passed to the ABR for routing to the external destinations.

Because stub areas do not support type 5 ASE LSAs, they cannot support ASBRs.

Totally stubby area

To further reduce the size of the stub area LSDB, you can configure a totally stubby area, which prevents redistribution of summary routes (Summary LSAs, type 3) from other areas into the stub area. As shown in the following figure, the totally stubby area ABR advertises a default route into the stub area not only for external routes, but for all destinations outside of the stub area.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 8
Totally stubby area

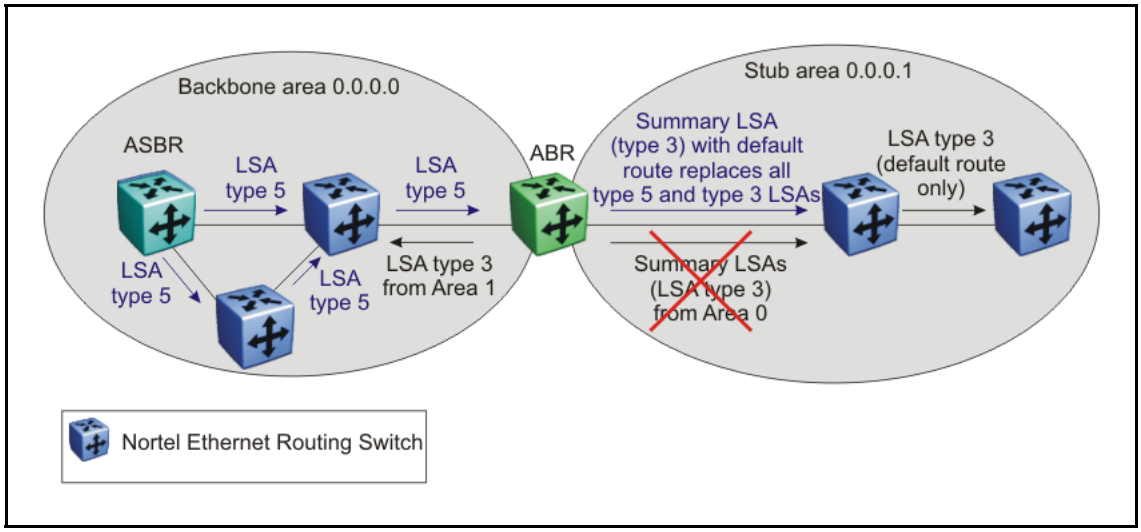

To configure a totally stubby area, you must disable import summaries on the stub area ABR.

Disabling import summaries is only allowed in the stub area.

Not so stubby area

Like a stub area, a not so stubby area (NSSA) is at the edge of an OSPF routing domain and it prevents the flooding of AS External LSAs into the NSSA by replacing them with a default route.

However, unlike a stub area, an NSSA can import small stub (non-OSPF) routing domains into OSPF. This allows the NSSA to import external routes, such as RIP routes, and advertise these routes throughout the network.

As shown in the following figure, a non-OSPF routing domain can connect to the NSSA to allow the external network to route traffic to the OSPF AS. One router in the NSSA must operate as an ASBR to provide a link to the non-OSPF domain.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 9
OSPF NSSA

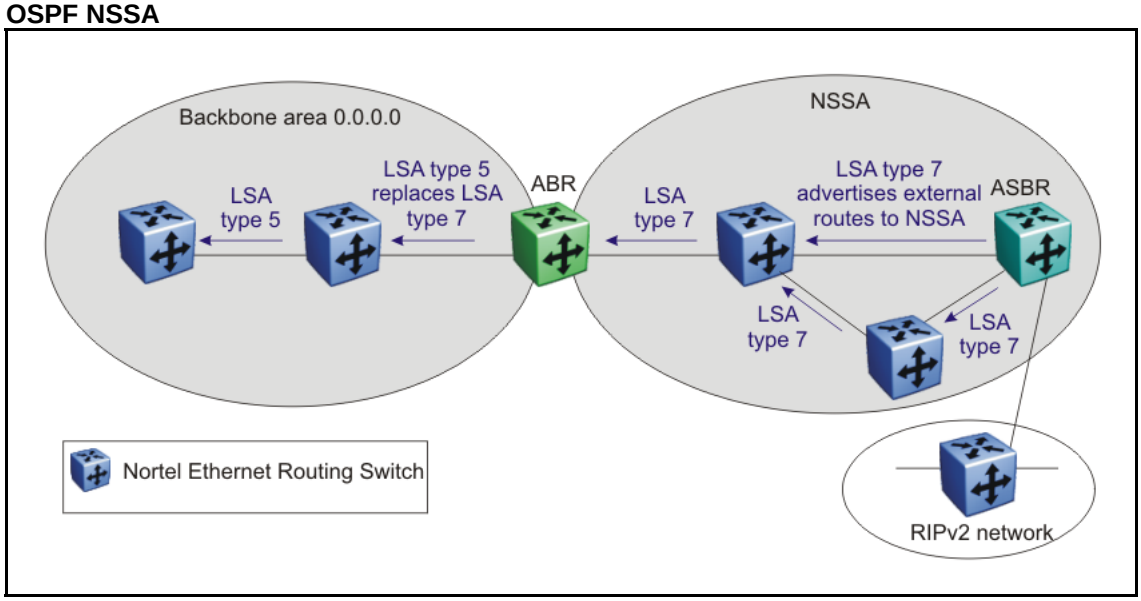

If the non-OSPF network is a small network, and the attached non-OSPF router has a default route to the OSPF network, this provides sufficient routing for any destinations that are outside the non-OSPF network.

Within the NSSA, the NSSA ASBR advertises route information imported from the external network using type 7 LSAs (NSSA External LSAs).

To propagate the external routes to other areas, the NSSA ABR translates these type 7 LSAs into type 5 LSAs (AS External LSAs). The ABR can flood the type 5 LSAs to the other areas so that the rest of the OSPF domain can learn about the non-OSPF destinations.

You can also configure the ABR to prevent the flooding of the external routes to other areas. To support this additional control over external route advertisement, the type 7 LSAs provide an Options field containing an N/P-bit that notifies the ABR which external routes can be advertised to other areas. When the NSSA N/P-bit is set to true (the default setting), the ABR exports the external route. When the NSSA N/P-bit is not set, the ABR drops the external route.

To manipulate the N/P-bit value for specific routes, you must configure a route policy on the Nortel Ethernet Routing Switch 5000 Series .

Normal area

A normal area is an area that is neither a bacbone nor a stub area that sends and receives LSA types 1 through 5. As illustrated in the following figure, a normal area supports Area Border Routers (ABRs) and Autonomous System Border Routers (ASBRs).

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 10
OSPF normal area

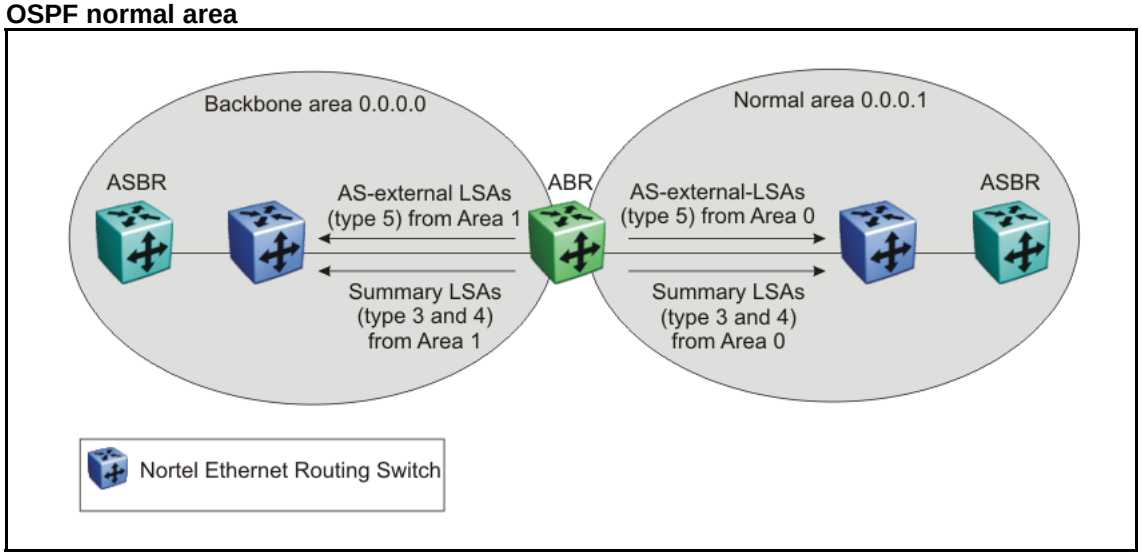

The Nortel Ethernet Routing Switch 5000 Series automatically becomes an ABR when it is connected to more than one area.

OSPF virtual link

The OSPF network can be partitioned into multiple areas. However, every non-backbone area must be connected to the backbone area through an ABR. If no physical connection to the backbone is available, you can create a virtual link.

A virtual link is established between two ABRs and is a logical connection to the backbone area through a non-backbone area called a transit area. Stub or NSSA areas cannot be transit areas.

In the following diagram, non-backbone ABR R4 establishes a virtual link with backbone ABR R1 across transit area 1.1.1.1. The virtual link connects area 2.2.2.2 to area 0.0.0.0.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 11
Virtual link between ABRs through a transit area

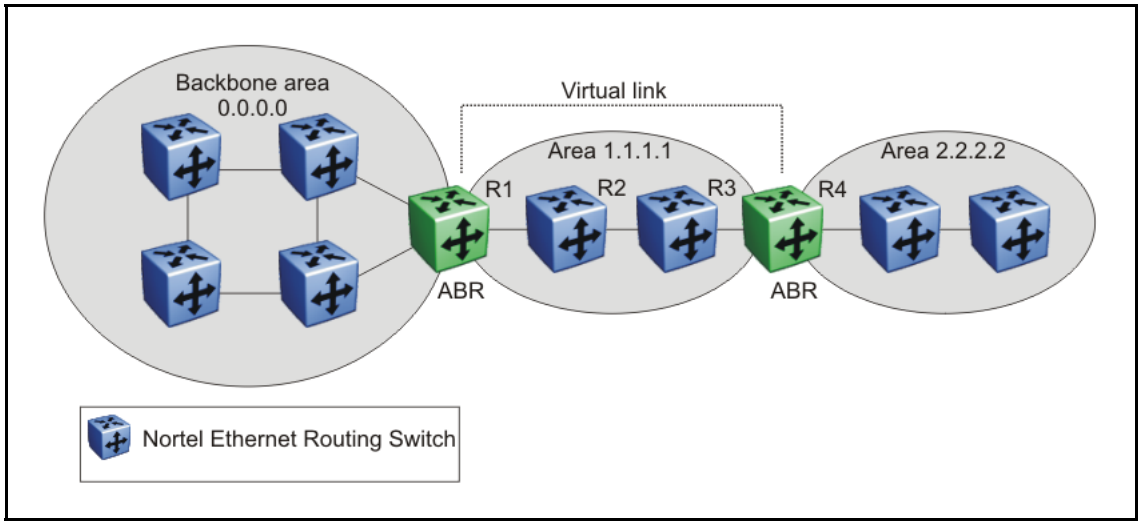

You can configure automatic or manual virtual links.

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. If a connection fails on the network, for example, 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.

Specify automatic virtual linking to ensure that a link is created to another router. When you specify automatic virtual linking, this feature is always ready to create a virtual link.

To configure automatic virtual link creation, enable automatic virtual link on both endpoint ABRs (the default value is disabled). Automatic virtual links are removed when the transit area is deleted, auto virtual link is disabled, or the router is no longer an ABR.

If automatic virtual linking uses more resources than you want to expend, a manual virtual link can be the better solution. Use this approach to conserve resources while maintaining specific control of where virtual links are placed in your OSPF network.

To add a virtual link manually, configure both endpoint ABRs with a neighbor router ID and transit area ID. You can configure up to 16 virtual links.

ATTENTION
Auto-created virtual links use default settings that cannot be modified. You can modify parameters for manually added virtual links.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

OSPF host route

An OSPF router with hosts directly attached to its interfaces can use host routes to advertise the attached hosts to its neighbors. You can configure up to 32 host routes.

Host routes are identified by the host IP address. You cannot configure the TOS for a host route as TOS-based routing is not supported. For each host directly connected to the router, configure the cost of the link to the host during host creation. You cannot modify this cost.

When a host is added to, or deleted from, a host route, the router updates the router LSAs and floods them to neighbors in each area where that router has an interface.

The following is an example of parameters for a host route advertised in the LSA.

Host route in LSA

- Type: 3 (stub network)
- LinkID: IP address of host directly connected to router
- Link Data: 0xFFFFFFFF
- Metric: configured cost of host

OSPF interfaces

An OSPF interface, or link, is configured on an IP interface. In the Ethernet Routing Switch 5000, an IP interface can be either a brouter port or a VLAN. The state information associated with the interface is obtained from the underlying lower level protocols and the routing protocol itself.

ATTENTION
To change the interface type of an enabled OSPF interface, you must first disable it, change the type, and then reenable it.

OSPF network types allow OSPF-neighboring between routers over various types of network infrastructures. You can configure each interface to support various network types. The following table describes the OSPF network interface types supported by the Ethernet Routing Switch 5000.

Broadcast interfaces

Broadcast interfaces automatically discover every OSPF router on the network by sending OSPF Hellos to the multicast group AllSPFRouters (224.0.0.5).

Neighboring is automatic and requires no configuration.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

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

Passive interfaces

A passive interface is an interfacing network in OSPF that does not generate LSAs or form adjacencies. Passive interfaces are typically used on an access network.

Using passive interfaces limits the amount of CPU cycles required to perform the OSPF routing algorithm.

Use a passive interface to enable an interface to advertise into an OSPF domain while limiting its adjacencies.

When you change the interface type to passive, the interface is advertised into the OSPF domain as an internal stub network with the following behaviors:

- does not send Hello packets to the OSPF domain
- does not receive Hello packets from the OSPF domain
- does not form adjacencies in the OSPF domain

The interface requires only that it be configured as passive to be advertised as an OSPF internal route. If the interface is not a passive interface, to advertise a network into OSPF and not form OSPF adjacencies, the interface must be configured as nonOSPF, and the local network must be redistributed as an autonomous system external (ASE) LSA.

The network behind a passive interface is treated as a stub network and does not form adjacencies. The network is advertised into the OSPF area as an internal route.

OSPF packets

OSPF runs over IP, which means that an OSPF packet is sent with an IP data packet header. The protocol field in the IP header is 89, which identifies it as an OSPF packet.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

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

- 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. A neighbor router that does not receive a Hello packet declares the other router dead.
- Database description (DD) packets are exchanged when a link is established between neighboring routers which synchronize their link-state databases.
- Link-state request packets describe one or more link-state advertisements that a router requests 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 contain one or more link-state advertisements and are sent following a change in network conditions.
- Link-state acknowledgement packets are sent to acknowledge receipt of link-state updates and contain the headers of the received link-state advertisements.

OSPF metrics

For OSPF, the best path to a destination is the path that offers the least-cost metric (least-cost delay). OSPF cost metrics are configurable, so you can specify preferred paths. You can configure metric speed globally or for specific 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, as shown in the following table.

Table 7
OSPF default metrics

Port type	Default OSPF metric
10 Mb/s	100
100 Mb/s	10
1000 Mb/s	1
10 000 Mb/s	1

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Automatic router ID change in a stack

If a unit leaves the stack and becomes standalone (when the stack disjoins), the router ID is automatically changed to its default value. This prevents router ID duplication in the OSPF routing domain.

To allow this feature to operate, IP blocking must be turned off (set to none) and OSPF must be globally enabled.

The new router ID value is temporary, that is, it is not saved in NVRAM. Therefore, upon reset, the old router ID is restored.

OSPF security mechanisms

The Ethernet Routing Switch 5000 implementation of OSPF includes security mechanisms to prevent unauthorized routers from attacking the OSPF routing domain. These security mechanisms prevent a malicious person from joining an OSPF domain and advertising false information in the OSPF LSAs. Likewise, security prevents a misconfigured router from joining an OSPF domain. Currently there are two security mechanisms supported: simple password security and Message Digest 5 (MD5) security.

Simple Password

The Simple Password security mechanism is a simple-text password that is transmitted in the OSPF headers. Only routers that contain the same authentication ID in their LSA headers can communicate with each other.

ATTENTION
Nortel recommends that you not 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.

Message Digest 5

Nortel 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 very difficult for a malicious user to compute or extrapolate the decrypting codes from the OSPF packets.

When you use MD5, each OSPF packet has a message digest appended to it. The digest must be matched between sending and receiving routers. The message digest is calculated at both the sending and receiving routers based on the MD5 key and any padding, and then compared. If the message digest computed at the sender and receiver does not match, the packet is rejected.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Each OSPF interface supports up to 2 keys, identifiable by key ID, to facilitate a smooth key transition during the rollover process. Only the selected primary key is used to encrypt the OSPF transmit packets.

Equal Cost MultiPath (ECMP)

The Equal Cost MultiPath (ECMP) feature allows routers to determine 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, links between routers can be used more efficiently when sending IP traffic. The ECMP feature supports the following protocols:

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

ECMP is not supported on the Nortel Ethernet Routing Switch 5510. ECMP works in a mixed stack but cannot run on any Nortel Ethernet Routing Switch 5510 units in the stack.

Route policies

Using standard routing schemes, a router forwards packets on routes that it has learned through routing protocols such as RIP and OSPF or through the introduction of static routes. With route policies, the router can forward packets based on rule sets created by the network administrator. These rule sets, or policies, are then applied to the learned or static routes.

On the Nortel Ethernet Routing Switch 5000 Series , you can configure route policies for RIP and OSPF. When used in conjunction with these protocols, route policies can be used to perform the following tasks that are not possible using traditional routing methods:

- Listen for routing updates from specific gateways.
- Listen for routing updates from specific networks.
- Assign a specific subnet mask to be included with a network in the routing table.
- Advertise routing updates from specific gateways.
- Advertise routing updates to specific networks.
- Assign a specific subnet mask to be included in the route summary packets.
- Advertise routes learned by one protocol to another.

The Ethernet Routing Switch 5000 Series supports the following types of policies:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- **Accept (In) Policies**

  Accept polices are applied to incoming routing updates before they are applied to the routing table. In the case of RIP, accept policies can be applied to all incoming packets. Only one policy can be created for each RIP interface. In the case of OSPF, accept policies are only applied to Type 5 External routes based on the advertising router ID. There can only be one OSPF accept policy per switch and the policy is applied before updates are added to the routing table from the link state database.

- **Announce (Out) Policies**

  Announce policies are applied to outgoing routing updates before the routing update packets are actually transmitted from the switch. In the case of RIP, announce policies can be applied to all outgoing packets. Only one policy can be created for each RIP interface. Announce policies are not supported for OSPF as OSPF requires routing information to be consistent throughout the OSPF domain.

- **Redistribution Policies**

  Redistribution policies are used to provide notification of addition or deletion of a route in the routing table by one protocol to another protocol. OSPF redistribution policies send redistributed routes as Type 5 External routes. To configure redistribution on a router, it must be an ASBR. There can be only one OSPF redistribution route per switch and redistribution must be enabled. The OSPF accept policy takes precedence over the redistribution policy. You cannot configure a redistribution policy for RIP.

Route policies consist of the following items:

- **Prefix Lists**
  - — List of IP addresses with subnet masks.
  - — Identified by a prefix list name and unique identifier.
  - — Prefix lists support the comparison of ranges of incoming masks.
- **Route Maps**
  - — Contain a set of match and set parameters.
  - — Match and set parameters can contain several prefix lists.
  - — A set of match and set parameters are identified by a sequence number.
  - — Accept and deny actions are associated with each sequenced parameter set.
  - — Sequence numbers act as a preference setting. Sets with a lower sequence number are preferred over those with a higher sequence number.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

To configure routing policies, create the appropriate prefix lists and then assign those prefix lists to route maps. Once all route maps have been created, assign them to the appropriate type of policy.

Route policies in a stack

In a stacked environment, the following rules apply to routing policies:

- The policy database is stored in all stack units.
- Policy configuration is supported from only the base unit. The base unit sends updates to non-base units to update the policy database in each stack unit.
- During database updates, only the database in the base unit is synchronized with the non-base unit. The database in the non-base units are deleted during the exchange.
- Only the policies stored in the base unit are used by RIP and OSPF for policy application.

DHCP relay

Dynamic Host Configuration Protocol (DHCP) is a mechanism to assign network IP addresses on a dynamic basis to clients who request an address. DHCP is an extension of the Bootstrap protocol (BootP). BootP/DHCP clients (workstations) generally use UDP broadcasts to determine their IP addresses and configuration information. If such a host is on a VLAN that does not include a DHCP server, the UDP broadcasts are by default not forwarded to servers located on different VLANs.

The Nortel Ethernet Routing Switch 5000 Series can resolve this issue using DHCP relay, which forwards the DHCP broadcasts to the IP address of the DHCP server. Network managers prefer to configure a small number of DHCP servers in a central location to lower administrative overhead. Routers must support DHCP relay so that hosts can access configuration information from servers several router hops away.

With DHCP relay enabled, the switch can relay client requests to DHCP servers on different Layer 3 VLANs or in remote networks. It also relays server replies back to the clients.

To relay DHCP messages, two Layer 3 VLANs must be created: one connected to the client and another providing a path to the DHCP server. DHCP relay can be enabled on a per-VLAN basis.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

ATTENTION
The DHCP Relay feature shares resources with QoS. If the DHCP Relay feature is enabled, a QoS policy with a precedence of 11 cannot be installed.

For further information on QoS policies refer to *Nortel Ethernet Routing Switch 5000 Series Configuration - Quality of Service* () (NN47200-504).

The following figure shows a DHCP relay example, with an end station connected to subnet 1, corresponding to VLAN 1. The Nortel Ethernet Routing Switch 5000 Series connects two subnets by means of the virtual routing function. When the end station generates a DHCP request as a limited UDP broadcast to the IP address of all 1s (that is, 255.255.255.255), with the DHCP relay function enabled, the Ethernet Routing Switch forwards DHCP requests to the host address of the DHCP server on VLAN 2.

Figure 12
DHCP relay operation

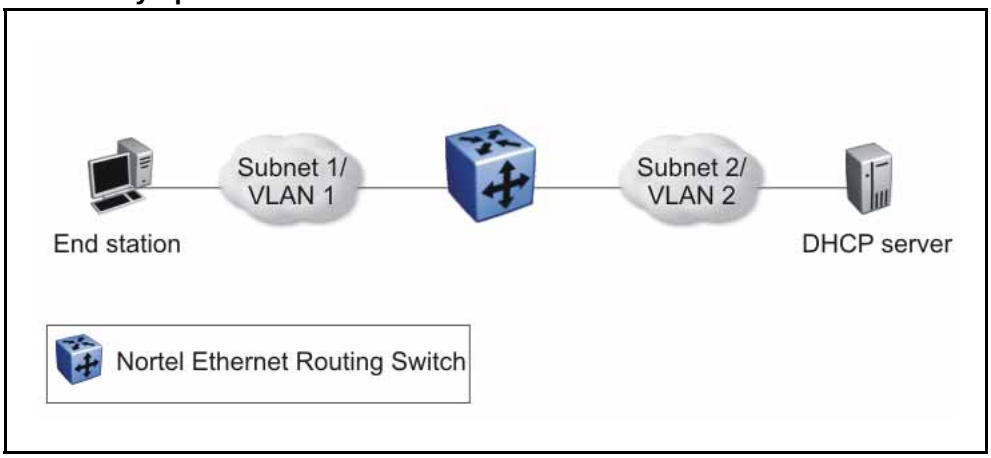

Forwarding DHCP packets

In the following figure, the DHCP relay agent address is 10.10.1.254. To configure the Nortel Ethernet Routing Switch 5000 Series to forward DHCP packets from the end station to the server, use 10.10.2.1 as the server address.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 13
Forwarding DHCP packets

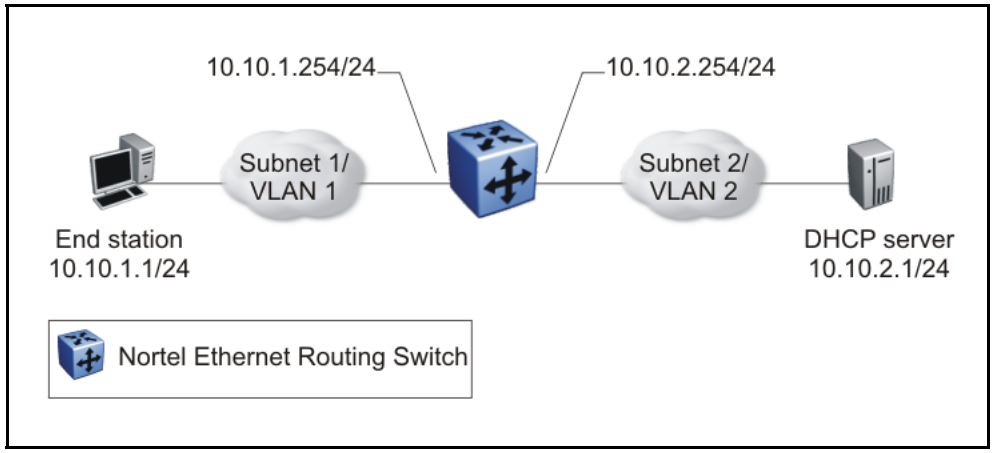

All BootP and DHCP broadcast packets that appear on the VLAN 1 router interface (10.10.1.254) are then forwarded to the DHCP server. In this case, the DHCP packets are forwarded as unicast to the DHCP server's IP address.

Multiple DHCP servers

Most enterprise networks use multiple DHCP servers for fault tolerance. The Nortel Ethernet Routing Switch 5000 Series can forward DHCP requests to multiple servers. You can configure up to 512 servers to receive copies of the forwarded DHCP messages.

To configure DHCP client requests to be forwarded to multiple different server IP addresses, specify the client VLAN as the DHCP relay agent for each of the destination server IP addresses.

In the following figure, two DHCP servers are located on two different VLANs. To configure the Nortel Ethernet Routing Switch 5000 Series to forward copies of the DHCP packets from the end station to both servers, specify the IP address of VLAN 1 (10.10.1.254) as the DHCP relay agent address and associate this relay agent with each of the DHCP server addresses, 10.10.2.1 and 10.10.3.1.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 14
Multiple BootP/DHCP servers

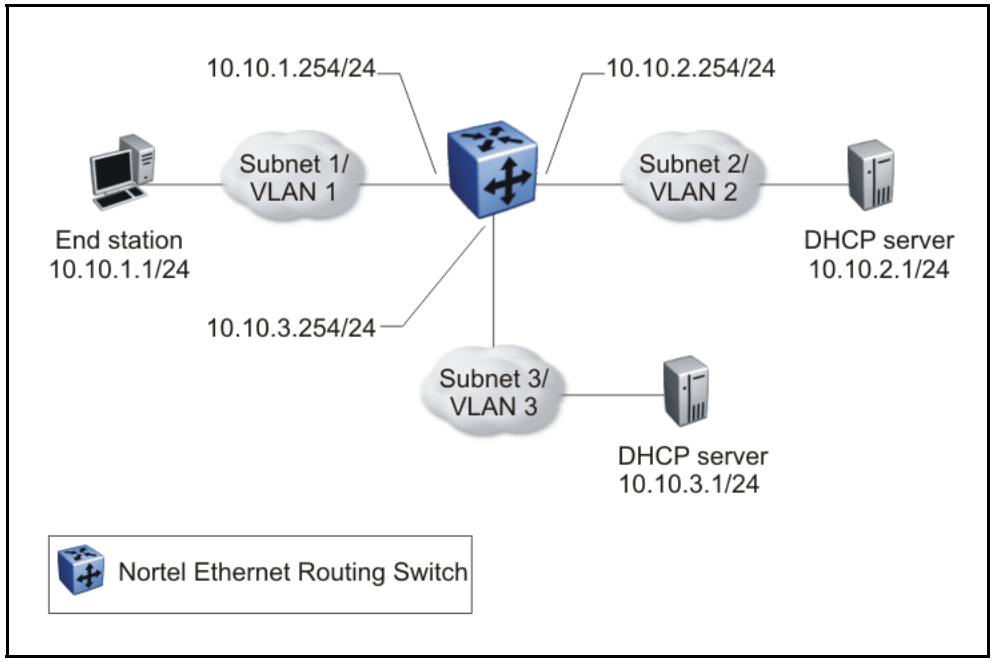

Differences between DHCP and BootP

With DHCP relay, the Nortel Ethernet Routing Switch 5000 Series supports the relay of DHCP and the Bootstrap protocol (BootP). The following differences between DHCP and BootP are specified in RFC 2131:

- BootP enables the retrieval of an ASCII configuration file name and configuration server address.
- A properly configured BootP server enables the switch to automatically learn its assigned IP address, subnet mask and the IP address of the default router (default gateway).
- DHCP defines mechanisms through which clients can be assigned a network address for a finite lease (allowing for reuse of IP addresses).
- DHCP provides the mechanism for clients to acquire all of the IP configuration parameters needed to operate.

DHCP uses the BootP message format defined in RFC 951. The remainder of the options field consists of a list of tagged parameters that are called *options* (RFC 2131).

UDP broadcast forwarding

By default, User Datagram Protocol (UDP) broadcast frames received on one VLAN are not routed to another VLAN. However, some network applications, such as the NetBIOS name service, rely on UDP broadcasts to request a service or locate a server. To allow UDP broadcasts to reach

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

a remote server, the Ethernet Routing Switch supports UDP broadcast forwarding, which forwards the broadcasts to the server through a Layer 3 VLAN interface.

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface to a configured IP address. The packet is sent as a unicast packet to the server.

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

- It must be a MAC-level broadcast.
- It must be an IP-limited broadcast.
- It must be for a configured UDP protocol.
- It must have a TTL value of at least 2.

For each ingress interface and protocol, the UDP broadcast packets are forwarded only to a unicast host address (the unicast IP address of the server for example).

When the UDP forwarding feature is enabled, a filter is installed that compares the UDP destination port of all packets against all the configured UDP forwarding entries. If a match occurs, the destination IP of the incoming packet is checked for consistency with the user-configured broadcast mask value for this source VLAN. If these conditions are met, the TTL field from the incoming packet is overwritten with the user-configured TTL value, the destination IP of the packet is overwritten with the configured destination IP, and the packet is routed to the destination as a unicast frame.

Directed broadcasts

With the directed broadcasts feature enabled, the Ethernet Routing Switch can determine if an incoming unicast frame is a directed broadcast for one of its interfaces. If so, the switch forwards the datagram onto the appropriate network using a link-layer broadcast.

With IP directed broadcasting enabled on a VLAN, the Ethernet Routing Switch forwards direct broadcast packets in two ways:

- through a connected VLAN subnet to another connected VLAN subnet
- through a remote VLAN subnet to the connected VLAN subnet

By default, this feature is disabled.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

ARP

The Address Resolution Protocol (ARP) allows the Ethernet Routing Switch to dynamically learn Layer 2 Media Access Control (MAC) addresses, and to build a table with corresponding Layer 3 IP addresses.

Network stations using the IP protocol need both a physical (MAC) address and an IP address to transmit a packet. If a network station knows only a network host's IP address, ARP enables the network station to determine a network host's physical address and bind the 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 message.
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 table for future use.
5. The network station's ARP table displays the association of the known MAC addresses to IP addresses.

The lifetime for the learned MAC addresses is a configurable parameter. The switch executes ARP lookups when this timer expires.

The default timeout value for ARP entries is 6 hours.

Static ARP

In addition to the dynamic ARP mechanism, the Ethernet Routing Switch supports a static mechanism that allows for static ARP entries to be added. With Static ARP, you can manually associate a device MAC address to an IP address. You can add and delete individual static ARP entries on the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Static ARP entries can be used to solve the following instances encountered on many networks:

- To communicate with a device that does not respond to an ARP request.
- To prevent an existing ARP entry from aging out.

When a static ARP entry is configured, both the IP address and MAC address of a device is assigned to a physical port. This includes the VLAN number if the physical port is associated with a VLAN.

Proxy ARP

Proxy ARP allows the Ethernet Routing Switch to respond to an ARP request from a locally attached host that is intended for a remote destination. It does so by sending an ARP response back to the local host with the MAC address of the switch interface that is connected to the host subnet. The reply is generated only if the switch has an active route to the destination network.

With Proxy ARP enabled, the connected host can reach remote subnets without the need to configure default gateways.

The following figure is an example of proxy ARP operation. In this example, host B wants to send traffic to host C, so host B sends an ARP request for host C. However, the Nortel Ethernet Routing Switch 5000 Series is between the two hosts, so the ARP message does not reach host C. To enable communication between the two hosts, the Nortel Ethernet Routing Switch 5000 Series intercepts the message, and responds to the ARP request with the IP address for host C but with the MAC address of the switch itself. Host B then updates its ARP table with the received information.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 15
Proxy ARP Operation

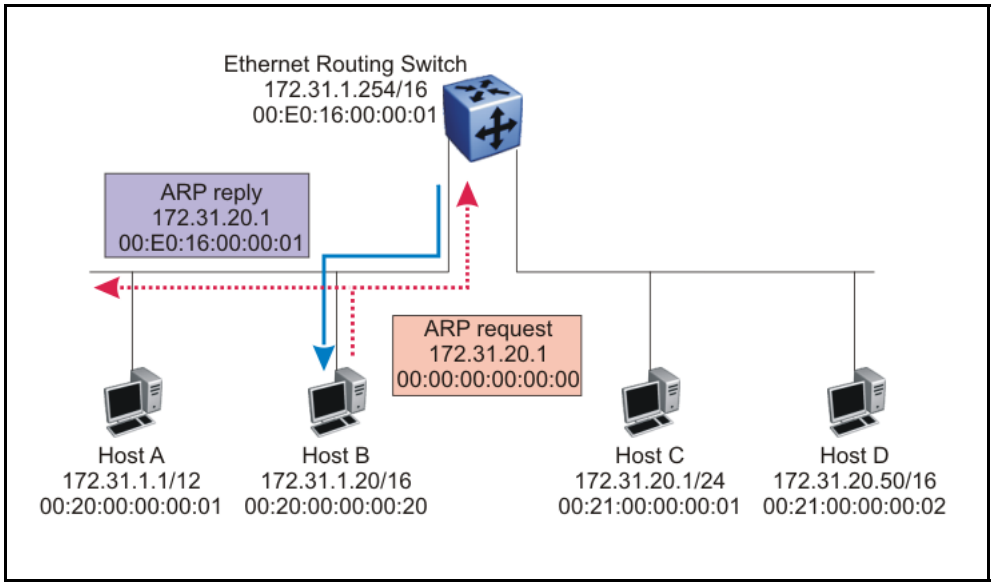

Nortel recommends Proxy ARP as a temporary fix only, for example, if you are gradually moving hosts from one addressing scheme to another and you still want to maintain connectivity between the disparately-addressed devices. You do not want Proxy ARP running as a general rule because it causes hosts to generate ARP messages for every address that they want to reach on the Internet.

IP blocking for stacks

IP Blocking is a Layer 3 feature of the Nortel Ethernet Routing Switch 5000 Series that provides safeguards for a stack where Layer 3 VLANs have port members across multiple stack units. IP Blocking is used whenever a unit leaves a stack or is rebooting inside the context of a stack. Depending on the setting in use, Layer 3 functionality is either continued or blocked by this feature.

You can set the IP Blocking mode on the base unit to either none or full.

When IP blocking is set to full, if any units leave the stack, those units run in Layer 2 mode. No Layer 3 settings remain on the units.

When IP blocking is set to none, if any units leave the stack, the Layer 3 configurations applied to the stack are still applied on the individual units.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

In a stack environment of 2 units, Nortel recommends that you use IP blocking mode none. In this case, you can expect the following functional characteristics:

- If either the stack base unit or nonbase unit becomes nonoperational, Layer 3 functionality continues to run on the remaining unit.

A disadvantage of this configuration is that if the nonoperational unit does not rejoin the stack, address duplication occurs.

In stack environments of more than 2 units, Nortel recommends that you use IP blocking mode full. In this case, you can expect the following functional characteristics:

- If the stack base unit becomes nonoperational, the following occurs:
  - The temporary base unit takes over base unit duties.
  - The temporary base unit takes over responsibility to manage Layer 3 functionality in the stack. When this occurs, the system updates the MAC addresses associated with each routing interface to be offset from the temporary base unit MAC address (rather than the base unit MAC address). During this period, some minor disruption may occur to routing traffic until end stations update their ARP cache with the new router MAC addresses. The Nortel Ethernet Routing Switch 5000 Series sends out gratuitous ARP messages on each routed VLAN for 5 minutes at 15 second intervals to facilitate quick failover in this instance.
  - If the nonoperational base unit does not rejoin the stack, no Layer 3 functionality runs on the unit.
- If a stack nonbase unit becomes nonoperational, the following occurs:
  - The stack continues to run normally with the base unit controlling Layer 3 functionality.
  - If the nonoperational nonbase unit does not rejoin the stack, no Layer 3 functionality runs on the unit.

By default, the IP blocking mode is none (disabled).

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

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

or more routers connecting a common subnet to the enterprise network. With end hosts using the virtual IP address as the default gateway, VRRP provides dynamic default gateway redundancy in the event of failure.

VRRP uses the following terms:

- **VRRP router**: a router running the VRRP protocol.
- **Virtual router**: the abstract object managed by VRRP that is assigned the virtual IP address and that acts as the default router for a set of IP addresses across a common network. Each virtual router is assigned a virtual router ID (VRID).
- **Virtual router master**: the VRRP router that assumes responsibility for forwarding packets sent to the IP address associated with the virtual router. The master router also responds to packets sent to the virtual router IP address and answers ARP requests for this IP address.
- **Virtual router backup**: the router or routers that can serve as the failover router if the master router becomes unavailable. If the master router fails, an election process provides a dynamic transition of forwarding responsibility to a new master router.
- **Priority**: an 8-bit value assigned to all VRRP routers. A higher value represents a higher priority for election to the master router. The priority can be a value from 1 to 255. 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. When a master router fails, an election process takes place among the backup routers to dynamically reassign the role of the master router. The host is unaware of the entire process.

VRRP operation

When you initialize a VRRP router, if there are no other VRRP routers enabled in the VLAN, then the initialized router assumes the role of the master router. When additional routers are enabled in the VLAN, an election process takes place among them to elect a master, based on their priority.

The master router functions as the forwarding router for the IP address associated with the virtual router. When a host sends traffic to a remote subnet, it sends an ARP request for the MAC address of the default gateway. In this case, the master router replies with the 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 on the Nortel Ethernet Routing Switch 5000 Series is automatically set as:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

00-00-5E-00-01-<VRID>

where **<VRID>** is an integer value in the range 1 to 255 that represents the virtual router identification.

The master router responds to ARP requests for the IP address, forwards packets with a destination MAC address equal to the virtual router MAC address, and accepts only packets addressed to the IP address associated with the virtual router. The master router also sends VRRP advertisements periodically (every 1 second by default) to all VRRP backup routers.

In the backup state, a VRRP router monitors the availability and state of the master 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 master router fails, the backup router with the highest priority assumes the role of the master router. It sends the VRRP advertisement and ARP request as described in the preceding paragraphs and transitions to the controlling state. The virtual router IP address and MAC address does not change, thereby providing transparent operation

VRRP topology example

The following figure shows a VRRP topology example.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 16
VRRP topology example

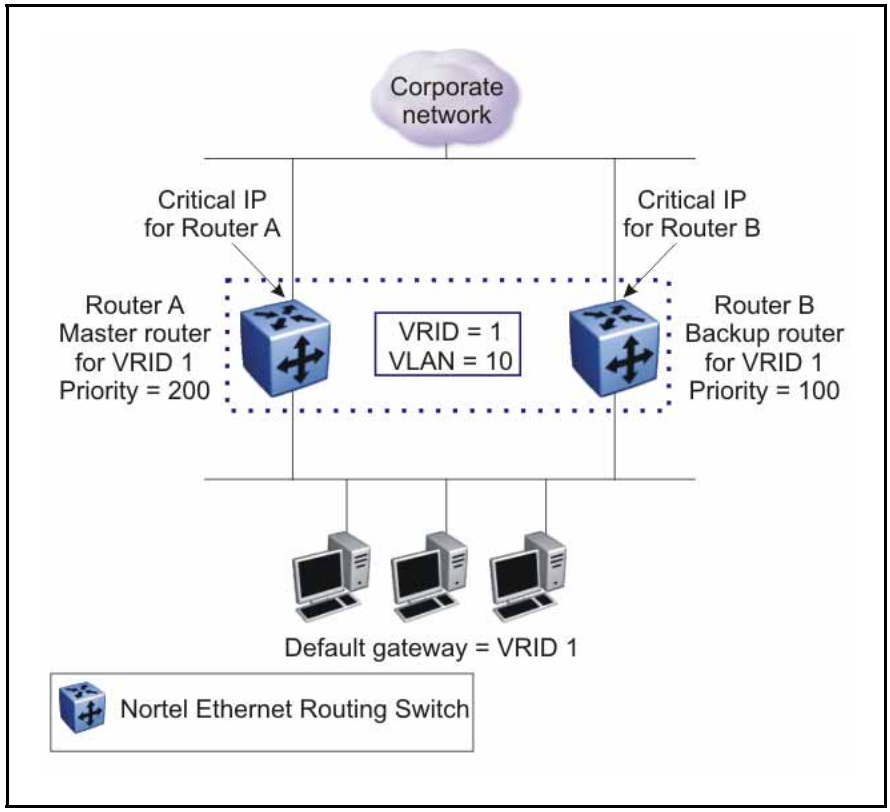

In this example, to configure router A as the master router and router B as the backup router, you can configure them for VRRP as follows:

1. On router A, create a VLAN, (in this case VLAN 10).
2. Assign an IP address to the VLAN for routing.
3. Configure VRRP properties for VLAN 10 on router A:
   - Assign a virtual router ID (in this case, VRID 1).
   - Set the virtual router IP address to a previously unassigned IP address.
   - Set the priority to a value above the priority of the Router B (in this case, 200).
4. On router B, create a matching VLAN (in this case, VLAN 10).
5. Assign an IP address to the VLAN for routing.
6. Configure VRRP properties for VLAN 10 on router B:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

— Assign the same virtual router ID as on router A (VRID 1).

— Configure the same virtual router IP address as on router A.

— Set the priority to a value below that on Router A (in this case, 100).

When you enable VRRP on both of these switches, an election process takes place, and because router A has the higher priority, it is elected the master router. It then assumes responsibility for the configured virtual router IP address.

Critical IP address

Within a VRRP VLAN, it is possible for one link to go down, while the remaining links in the VLAN remain operational. Because the VRRP VLAN continues to function, a virtual router associated with that VLAN does not register a master router failure.

As a result, if the local router IP interface connecting the virtual router to the external network fails, this does not automatically trigger a master router failover.

The critical IP address resolves this issue. If the critical IP address fails, it triggers a failover of the master router.

You can specify the local router IP interface uplink from the VRRP router to the network as the critical IP address. This ensures that, if the local uplink interface fails, VRRP initiates a master router failover to one of the backup routers.

In Figure 16 "VRRP topology example" (page 67), the local network uplink interface on router A is shown as the critical IP address for router A. As well, the similar network uplink is shown as the critical IP address for router B. Router B also requires a critical IP address for cases when it assumes the role of the master router.

VRRP and SMLT

The standard implementation of VRRP allows only one active master switch per IP subnet. All other VRRP interfaces in a network are in backup mode.

However, a deficiency occurs when VRRP-enabled switches use Split Multi-Link Trunking (SMLT).

Normally, if a switch is connected to two SMLT aggregation switches, the end host traffic is load-shared on all uplinks to the aggregation switches (based on the MLT traffic distribution algorithm). However, VRRP can only have one active routing interface enabled. All other VRRP routers are in

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

backup mode. Therefore, if the SMLT aggregation switches run VRRP, all traffic that reaches the backup VRRP router is forwarded over the Inter Switch Trunk (IST) link towards the master VRRP router. In this case, the IST link might not have enough bandwidth to carry all the aggregated traffic.

You can overcome this issue by assigning the backup router as the backup master router. The backup master router is a backup router that is permitted to actively load-share the routing traffic with a master router.

When the backup master router is enabled, the incoming host traffic can be load-shared over the SMLT links as normal. This configuration allows both switches to respond to ARP requests and forward traffic

The following figure shows a sample VRRP configuration with SMLT. As Router B is configured as the backup master, routing traffic is load-shared between the two devices.

Figure 17
VRRP with SMLT

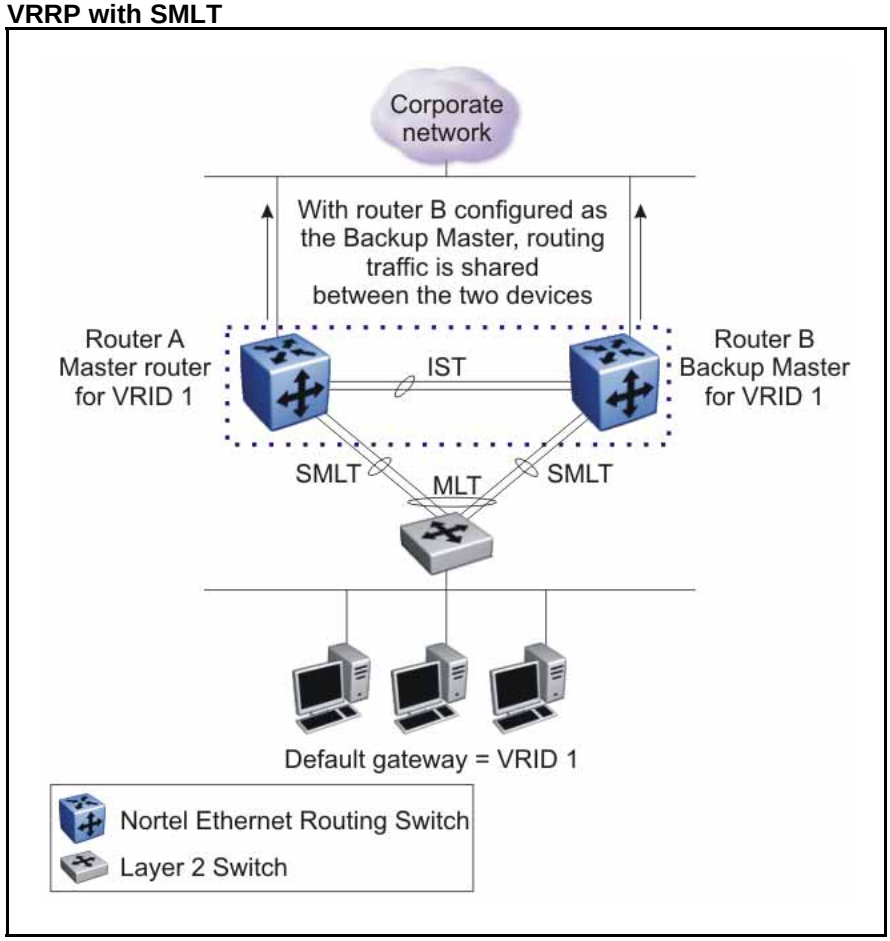

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

VRRP fast advertisement interval

With VRRP, you can set the advertisement interval between sending advertisement messages in seconds. This permits 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 network problems must be detected within hundreds of milliseconds.

To meet these requirements, Nortel Ethernet Routing Switch 5000 Series supports a fast advertisement interval parameter. The fast advertisement interval is similar to the advertisement interval except for the unit of measure and range. The fast advertisement interval is expressed in milliseconds and the range is from 200 to 1000 milliseconds. To use the fast advertisement interval, you must configure a value for the parameter and explicitly enable the feature.

When the fast advertisement interval is enabled, VRRP can only communicate with other Ethernet Routing Switch devices with the same settings.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IP multicast fundamentals

To manage multicast traffic, the Ethernet Routing Switch 5000 Series supports PIM-SM (for IGMPv1 and IGMPv2) and IGMP snooping (for IGMPv1, IGMPv2, and IGMPv3). You can enable IGMP snooping on a per-VLAN basis either on a Layer 2 or a Layer 3 VLAN. You can enable PIM-SM on Layer 3 VLANs only.

This chapter describes the fundamentals of IP multicast as they apply to the Ethernet Routing Switch 5000 Series.

Navigation

- "Overview of IP multicast" (page 71)
- "Internet Group Management Protocol " (page 78)
- "IGMP snooping" (page 85)
- "Protocol Independent Multicast-Sparse Mode" (page 95)
- "PIM passive interfaces" (page 104)

Overview of IP multicast

Most traditional network applications such as Web browsers and e-mail employ unicast connections in which each client sets up a separate connection to a server to access specific data. However, with certain applications such as audio and video streaming, more than one client accesses the same data at the same time. With these applications, if the server sends the same data to each individual client using unicast connections, the multiple connections waste both server and network capacity. For example, if a server offers a 1Mbit/sec live video stream for each client, a 100Mbit/sec NIC card on the server can be completely saturated after 90 client connections. The following figure shows an example of this waste of resources.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 18
Wasteful propagation of multiple copies of the same unicast stream

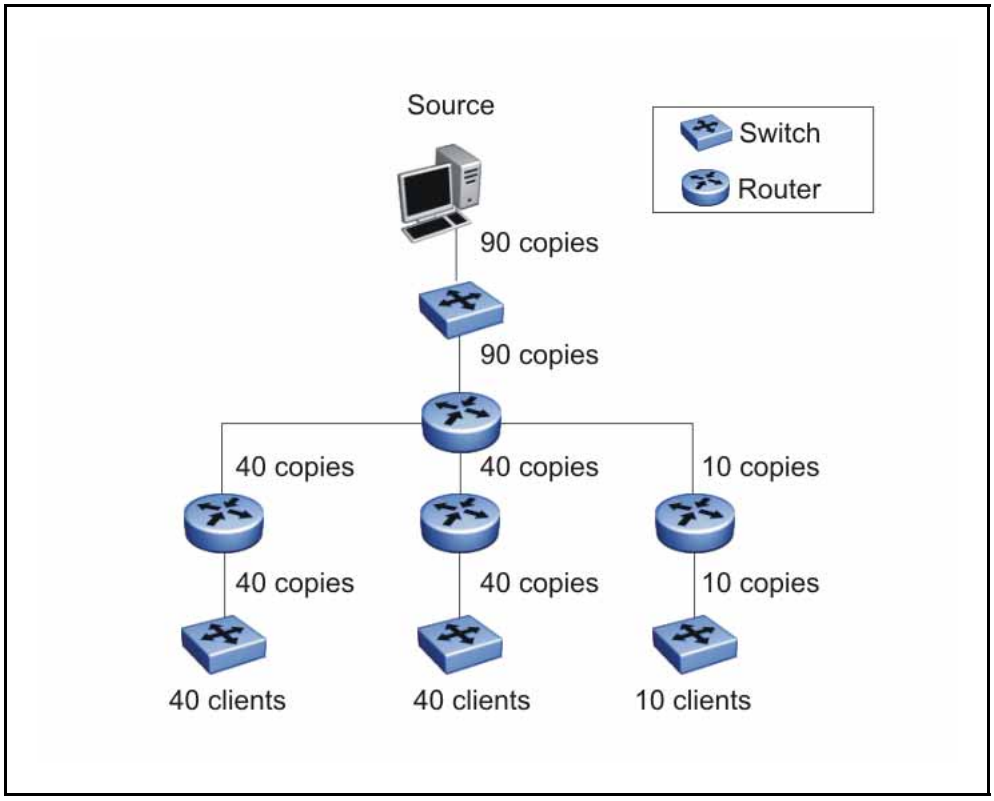

Multicasting provides the ability to transmit only one stream of data to all the interested clients at the same time. The following figure shows a simple example of how multicasting works. The source of the multicast data forwards only one stream to the nearest downstream router, and each subsequent downstream router forwards a copy of the same data stream to the recipients who are registered to receive it.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 19
1 stream replicated using multicasting

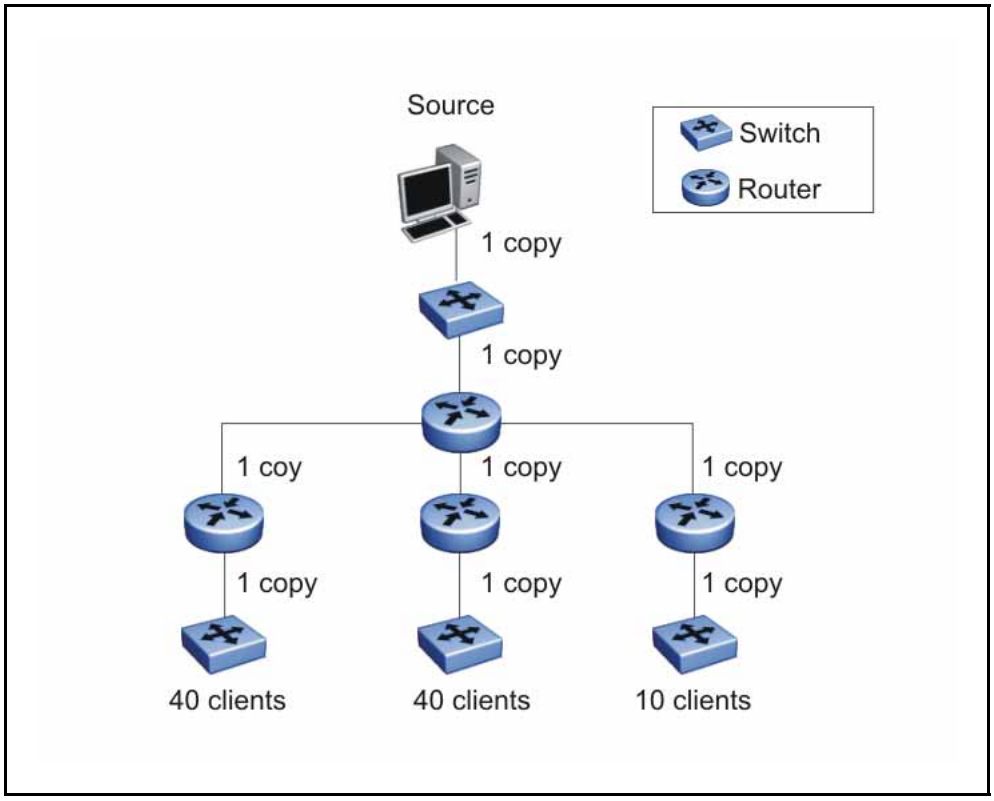

This one-to-many delivery mechanism is similar to broadcasting except that, while broadcasting transmits to all hosts in a network, multicasting transmits only to registered host groups. Because multicast applications transmit only one stream of data, which is then replicated to many receivers, multicasting saves a considerable amount of bandwidth.

Clients that want to receive the stream must register with the nearest multicast router to become a part of the receiving multicast group.

One downside to multicasting is that the multicast streams transmit data using UDP packets, which are not as reliable as TCP packets.

Applications that use multicasting to transmit data include:

- multimedia conferencing
- real-time data multicasts (such as stock tickers)
- gaming and simulations

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Multicast groups

To receive a multicast stream from a particular source, hosts must register with the nearest multicast router. The router adds all interested hosts to a multicast group, which is identified by a multicast IP address.

Multicast routers use Internet Group Membership Protocol (IGMP) to learn the existence of host group members on their directly attached subnets. To identify the hosts that want to be added to a group, a querier router sends out IGMP queries to each local network. A host that wants to belong to the group sends a response in the form of an IGMP membership report.

Each multicast router maintains a multicast routing table that lists each source, group (S,G) pair, which identifies the IP address of the source and the multicast address of the receiving group. For each (S,G) pair, the router maintains a list of downstream forwarding ports to which the multicast traffic is forwarded and the upstream port where the multicast traffic is received.

Multicast distribution trees

A multicast distribution tree consists of the routers that forward multicast data from a particular source to all registered multicast group members. When all routers in the multicast network have determined their upstream and downstream interfaces for a particular group, the multicast tree is formed. At the root of the tree is the source of the multicast stream, and the branches are the destination routers that want to receive the multicast stream. Different multicast protocols use different techniques to discover delivery paths.

The following figure is an example of a simple distribution tree, where the arrows indicate the multicast delivery path from the source to the group members along the multicast distribution tree.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 20
Multicast distribution tree

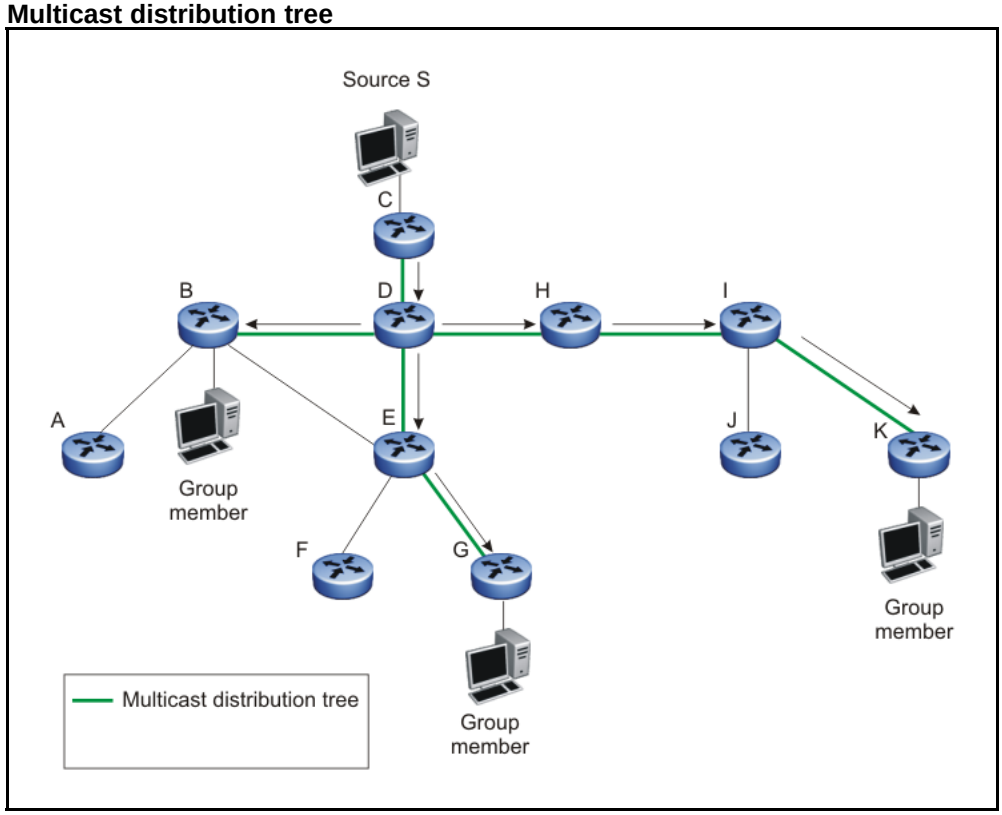

Reverse Path Forwarding

Reverse Path Forwarding is the means by which multicast routers ensure a loop-free topology. When a multicast packet arrives on an interface, the router compares the source address of the packet against the unicast routing table to determine whether the receiving interface is on the shortest path back to the source. If the receiving interface is the one the router would use to forward a unicast packet back to the source, the reverse path check passes, and the router forwards the multicast packet to its downstream neighbors. If the packet does not arrive on an upstream interface, the router discards the packet.

For example, in the preceding figure, if router E receives a packet from source S through router B, which is not on the optimal path back to the source, router E discards the packet. However, if router E receives the source packet from router D, which according to the unicast routing table is the next hop toward the source, router E forwards the incoming packet downstream to router G.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Without reverse path forwarding, loops can form in the network. For example, Router E can forward all packets coming from router D to router B, and router B can in turn forward the same traffic to router D, thereby causing a loop. With the RPF checks running on these routers, no loop can form.

Multicast addresses

Each multicast host group is assigned a unique multicast address. To reach all members of the group, a sender uses the multicast address as the destination address of the datagram.

An IP version 4 multicast address is a Class D address (the high-order bits are set to 1110) from 224.0.1.0 to 239.255.255.255. These addresses are assigned statically for use by permanent groups and dynamically for use by transient groups.

The multicast address range 224.0.0.0/24, 224.128.0.0/24, 225.0.0.0/24, 225.128.0.0/24 up to 239.0.0.0/24, 239.128.0.0/24 maps to reserved multicast MAC addresses. You cannot use these addresses for multicast data traffic.

IP multicast address ranges

IP multicast utilizes D class addresses, which range from 224.0.0.0 to 239.255.255.255. Although subnet masks are commonly used to configure IP multicast address ranges, the concept of subnets does not exist for multicast group addresses. Consequently, the usual unicast conventions where you reserve the all 0s subnets, all 1s subnets, all 0s host addresses, and all 1s host addresses do not apply when dealing with the IP multicast range of addresses.

Addresses from 224.0.0.0 through 224.0.0.255 are reserved by the Internet Assigned Numbers Authority (IANA) for link-local network applications. Packets with an address in this range are not forwarded by multicast capable routers by design. For example, Open Shortest Path First (OSPF) uses both 224.0.0.5 and 224.0.0.6 and Virtual Router Redundancy Protocol (VRRP) uses 224.0.0.18 to communicate across a local broadcast network segment.

IANA has also reserved the range of 224.0.1.0 through 224.0.1.255 for well-known applications. These addresses are also assigned by IANA to specific network applications. For example, the Network Time Protocol (NTP) uses 224.0.1.1 and Mtrace uses 224.0.1.32. RFC 1700 contains a complete list of these reserved numbers.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Multicast addresses in the 232.0.0.0/8 (232.0.0.0 to 232.255.255.255) range are reserved only for source-specific multicast (SSM) applications, such as one-to-many applications. (For more information, see RFC 4607). While this range is the publicly reserved range for SSM applications, private networks can use other address ranges for SSM.

Finally, addresses in the range 239.0.0.0/8 (239.0.0.0 to 239.255.255.255) are administratively scoped addresses, meaning they are reserved for use in private domains and cannot be advertised outside that domain. This multicast range is analogous to the 10.0.0.0/8, 172.16.0.0/20, and 192.168.0.0/16 private address ranges in the unicast IP space.

Technically, a private network can only assign multicast addresses from 224.0.2.0 through 238.255.255.255 to applications that are publicly accessible on the Internet. Multicast applications that are not publicly accessible can be assigned addresses in the 239.0.0.0/8 range.

IP to Ethernet multicast MAC mapping

Like IP, Ethernet has a range of MAC addresses that natively support Layer 2 multicast capabilities. However, while IP has a total of 28 addressing bits available for multicast addresses, Ethernet has only 23 addressing bits assigned to IP multicast. The multicast MAC address space for Ethernet is much larger than 23 bits, but only a subrange of that larger space is allocated to IP multicast by the Institute of Electrical and Electronics Engineers (IEEE). Because of this difference, 32 IP multicast addresses map to one Ethernet multicast MAC address.

IP multicast addresses map to Ethernet multicast MAC addresses by placing the low-order 23 bits of the IP address into the low-order 23 bits of the Ethernet multicast address 01:00:5E:00:00:00. Thus, more than one multicast address maps to the same Ethernet address. For example, all 32 addresses from 224.1.1.1, 224.129.1.1, 225.1.1.1, 225.129.1.1, and so on up to 239.1.1.1 and 239.129.1.1 map to the same 01:00:5E:01:01:01 multicast MAC address.

The following figure shows the mapping of multicast IP addresses to MAC addresses.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 21
Multicast IP address to MAC address mapping

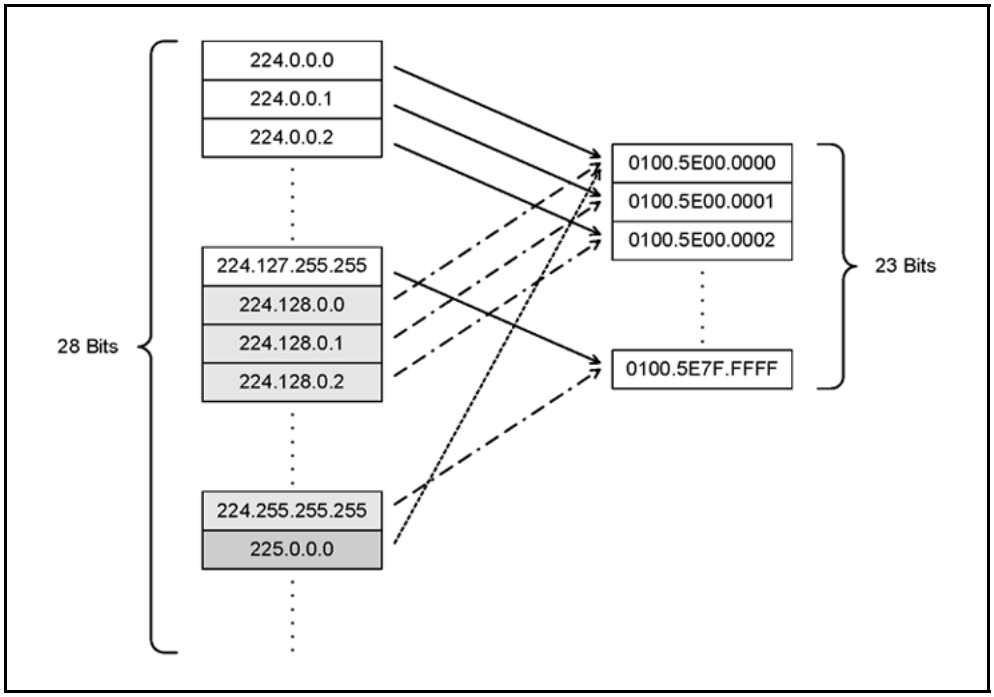

Most pure Layer 2 Ethernet switches handle Ethernet multicast by mapping a multicast MAC address to multiple switch ports in the MAC address table. However, the Ethernet Routing Switch 5000 Series switches IP multicast data based on the IP multicast address and not the MAC address. It internally maps IP multicast group addresses to the ports that contain group members. After an IP multicast packet is received, the lookup is based on IP group address, regardless of whether the VLAN is bridged or routed. This avoids the ambiguity in mapping 32 IP addresses to one MAC address.

If your network includes pure Layer 2 Ethernet switches that map each multicast MAC address to 32 IP addresses, the easiest way to avoid any potential issues is to use only a consecutive range of IP multicast addresses corresponding to the lower order 23 bits of that range. For example, use an address range from 239.0.2.0 through 239.127.255.255. A group address range of this size can accommodate the addressing needs of even the largest private enterprise.

Internet Group Management Protocol

IGMP is the Layer 3 protocol that IP multicast routers use to learn the existence of multicast group members on their directly attached subnets (for more information, see RFC 2236). With IGMP, hosts can register their desired group memberships to their local querier router.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

A multicast querier router communicates with hosts on a local network by sending IGMP queries. The router periodically sends a general query message to each local network of the router. A host that wants to join a multicast group sends a response in the form of a membership report requesting registration with a group. After the querier router registers hosts to a group, it forwards all incoming multicast group packets to the registered host networks. If any host on a subnet continues to participate in the group, all hosts, including nonparticipating end stations on that subnet, receive the IP Multicast stream.

IGMP versions are backward compatible and can all exist together on a multicast network.

The following sections provide more details on the differences between the different IGMP versions

IGMPv1 operation

IGMP version 1 is the simplest of the three IGMP versions and widely deployed.

IGMPv1 supports two different message types:

- 0x11—Membership Query message. Packets are sent to the all-systems multicast group (224.0.0.1).
- 0x12—Membership Report message. Packets are sent to the group that the host intends to join.

The IGMPv1 router periodically sends host membership queries (also known as general queries) to its attached local subnets to inquire if any hosts are interested in joining any multicast groups. The interval between queries is a configurable value on the router. A host that wants to join a multicast group sends a membership report message to the nearest router, one report for each joined multicast group. After receiving the report, the router adds the Multicast IP address and the host port to its forwarding table. The router then forwards any multicast traffic for that multicast IP address to the member ports.

The router keeps a list of multicast group memberships for each attached network, and a Group Membership Interval timer for each membership. Repeated IGMP membership reports refresh the timer. If no reports are received before the timer expires, the router sends a query message.

In some cases, the host does not wait for a query to send report messages to the router. Upon initialization, the host can immediately issue a report for each of the multicast groups that it supports. The router accepts and processes these asynchronous reports the same way it accepts requested reports.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IGMPv1 leave process

After hosts and routers are in a steady state, they communicate in a way that minimizes the exchange of queries and reports. The designated routers set up a path between the IP Multicast stream source and the end stations and periodically query the end stations to determine whether they want to continue participation. If any host on the subnet continues to participate, all hosts, including nonparticipating end stations on the subnet, receive the IP Multicast stream.

If all hosts on the subnet leave the group, the router continues to send general queries to the subnet. If no hosts send reports after three consecutive queries, the router determines that no group members are left on the subnet.

Host report suppression

A host that receives a query delays its reply by a random interval and listens for a reply from another host in the same host group. Consider a network that includes two host members—host A and host B—of the same multicast group. The router sends out a host membership query on the local network. Both host A and host B receive the query and listen on the network for a host membership report. The delay timer for host B expires first, so host B responds to the query with a membership report. Hearing the response, host A does not send a report of its own for the same group.

Figure 22
IGMP report suppression

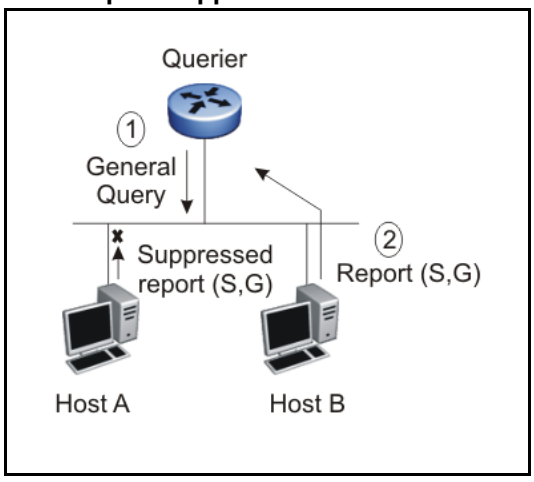

IGMPv2 operation

IGMPv2 extends the IGMPv1 features by implementing a host leave message to quickly report group membership termination to the routing protocol. Instead of routers sending multiple queries before determining

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

that hosts have left a group, the hosts can send a leave message. This feature is important for multicast groups with highly volatile group membership.

The IGMPv2 join process is similar to IGMPv1.

IGMPv2 also implements a querier election process.

IGMPv2 adds support for three new message types:

- 0x11—General Query and Group Specific Query message
- 0x16—Version 2 Membership Report (sent to the destination IP address of the group being reported)
- 0x17—Version 2 Membership Leave message (sent to all-router multicast address: 224.0.0.2)

IGMPv2 also supports IGMPv1 messages.

Host leave process

With IGMPv2, if the host that issued the most recent report leaves a group, it issues a leave message. The multicast router on the network then issues a group-specific query to determine whether other group members are present on the network. In the group-specific query message, the Group Address field is the group being queried (the Group Address field is 0 for the General Query message). If no host responds to the query, the router determines that no members belonging to that group exist on that interface.

The following figure shows an example of how IGMPv2 works.

Figure 23
IGMPv2

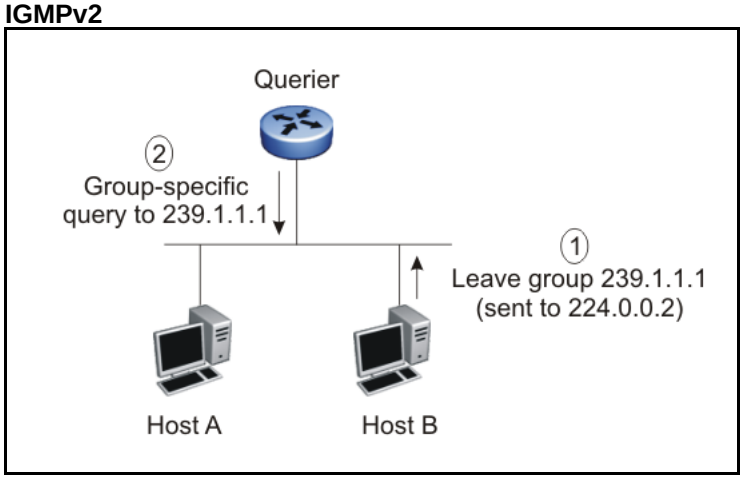

In this example:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- the host sends a leave message (to 224.0.0.2)
- the router sends a group-specific query to group 239.1.1.1
- no IGMP report is received
- group 239.1.1.1 times out

Maximum Response Time

Each IGMPv2 query from a router to a host includes a Maximum Response Time field, specifying the maximum time n in tenths of a second within which the host must issue a reply. The host uses this value to calculate a random value between 0 and n tenths of a second for the period that it waits before sending a response.

IGMPv1 queries do not specify a maximum response time. Instead, the maximum response time is a fixed value of 100, that is, 10 seconds.

Robustness value

As part of IGMP configuration, the robustness value lets you configure the switch to offset expected packet loss on a subnet. If you expect a network to lose IGMP query or membership report packets, you can increase the robustness value to offset the lost packets.

When the Ethernet Routing Switch receives an IGMP report from a host, the switch refreshes the expiration time for the group member. The timeout for a group member is a function of the query interval, robustness value, and the maximum response time. If the robustness value is increased, the group member lifetime increases as well, according to the following formula:

IGMP group member lifetime = (IGMP query interval * robustness value) + maximum response time.

If the network is congested, the switch is more likely to miss an IGMP report from a host, or the host can miss the query from the router. To offset the network loss, you can increase the robustness value to extend the life time for the group members.

Router alert

The router alert feature instructs the router to drop control packets that do not have the router-alert flag in the IP header. You can use this option to optimize the performance of multicast routers. When you enable router alert, a router needs to examine only the packets that have the router-alert option flagged, and can ignore all other multicast control packets destined to it. This optimizes the performance in packet processing. This feature is especially useful in routers-only networks. (It is difficult to force a host to send packets with the router-alert option.)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Querier election process

There is normally only one querier per subnet. When multiple IGMPv2 routers are present on a network, the router with the lowest IP address is elected to send queries. All multicast routers start up as a querier on each attached network. If a multicast router hears a query message from a router with a lower IP address, it becomes a nonquerier on that network.

IGMPv3 operation

IGMPv3 adds support for source filtering. The IGMPv3 host can report its interest in receiving multicast packets from only specific source addresses, or the host can report its interest in receiving multicast packets from all but specific source addresses.

IGMPv3 is mostly used in voice and video conferences where multiple people can be part of the same conference. The IGMPv3 packet format adds a v3 Report message type (0x22) and also includes Source-and-Group-specific Query messages.

The message type for Source-and-Group-specific Query message is 0x11, the same as IGMPv1 and IGMPv2. The different Query message versions are identified as follows:

- If the size of the IGMP message type is 8, then it is a v1 or v2 Query message.
- If the Group Address field is 0, then it is a General Query.
- If the Group Address field is a valid multicast IP address, then it is a Group-specific Query.
- If the Group Address field is a valid address and the Number of Sources field is nonzero, then it is a Group-and-Source specific Query message.

Each IGMPv3 Report contains a list of group records. The Group Record contains the multicast group address and the list of source addresses. The record type field specifies whether to INCLUDE or EXCLUDE the list of source addresses that are provided in the Source Address field. For example, to include packets from source 10.10.10.1, the report contains an INCLUDE(10.10.10.1) record.

The list of source addresses can be empty, which is represented by braces ({}), which means either to INCLUDE or EXCLUDE none. For example, the host that wants to receive packets from all group members can send a report with an EXCLUDE({}) record and a host that wants to leave a group can send a report with an INCLUDE({}) record, which is similar to a leave message.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

In the following figure, hosts A, B, C, D, E, and F are part of a conference group G1. All hosts except F send a report for group G1 with the mode as INCLUDE(A, B, C, D, E, F) containing all the source addresses. Host F, which is not interested in listening to C and D, sends a report to group G1 with the mode as EXCLUDE(C, D).

Figure 24
IGMPv3

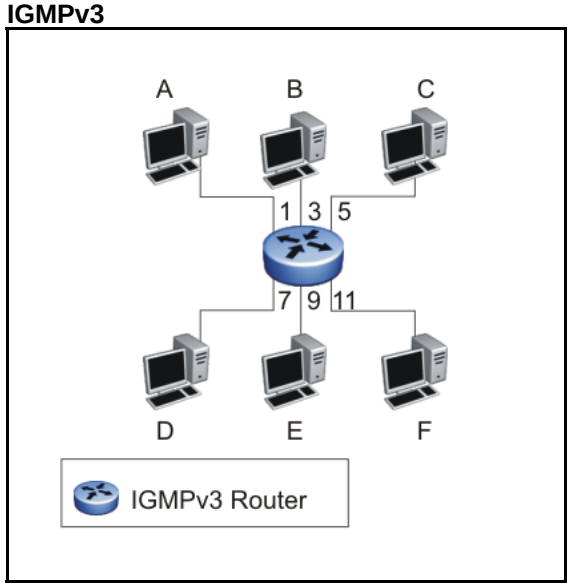

The router adds the multicast IP address and the list of sources in the forwarding table. The router forwards the packets from A, B, E, and F to all ports. If the packets are received from C and D, it is forwarded to all ports except port 11.

Multicast flow over Multi-Link Trunking

In the current release, the Ethernet Routing Switch 5000 Series supports multicast traffic and control packets only on the base link of the MLT.

IGMP requests for comment

For more information on IGMP, see the following requests for comment (RFC):

- For IGMPv1, see RFC 1112.
- For IGMPv2, see RFC 2236.
- For IGMPv3, see RFC 3376.
- For IGMP snooping, see RFC 4541.
- For IGMP management information bases (MIB), see RFC 2933

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IGMP snooping

If at least one host on a VLAN specifies that it is a member of a group, by default, the Ethernet Routing Switch 5000 Series forwards to that VLAN all datagrams bearing the multicast address of that group. All ports on the VLAN receive the traffic for that group.

The following figure shows an example of this scenario. In this example, the IGMP source provides an IP Multicast stream to a designated router. Because the local network contains receivers, the designated router forwards the IP Multicast stream to the network. Switches without IGMP snoop enabled flood the IP Multicast traffic to all segments on the local subnet. The receivers requesting the traffic receive the desired stream, but so do all other hosts on the network. Although the nonparticipating end stations can filter the IP Multicast traffic, the IP Multicast traffic still exists on the subnet and consumes bandwidth.

Figure 25
IP multicast propagation on a LAN without IGMP snooping

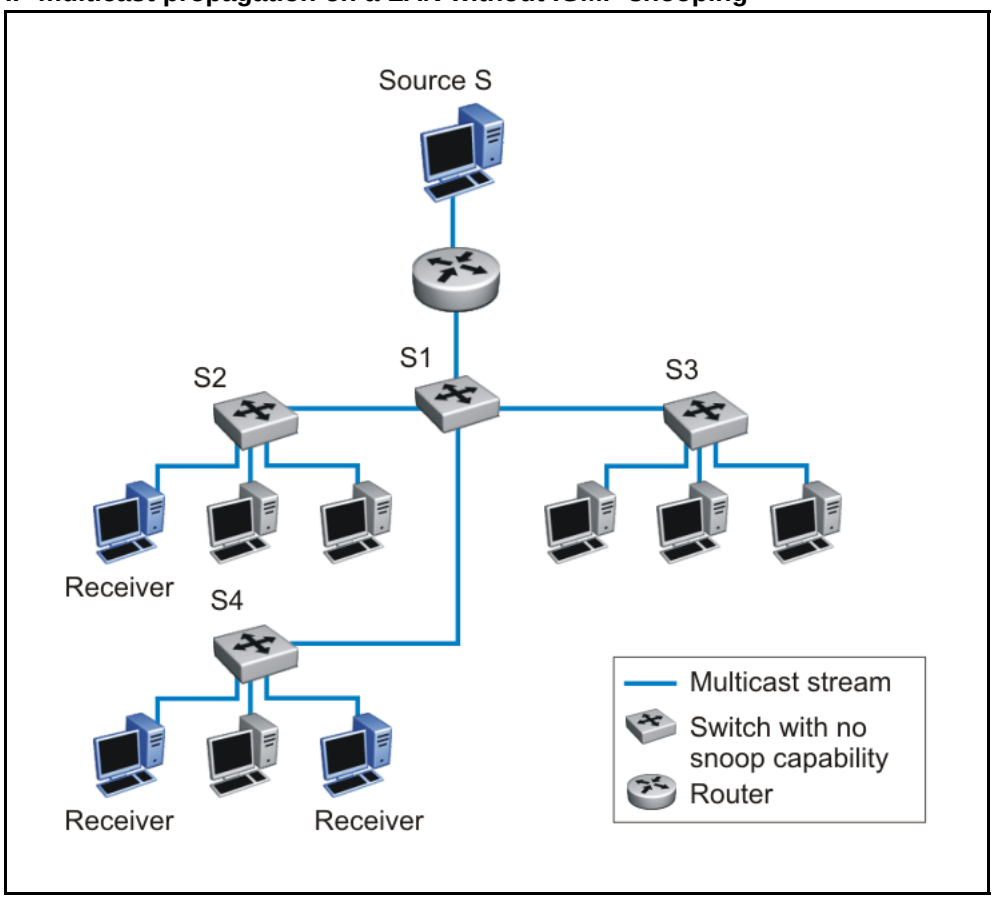

To prune ports that are not group members from receiving the group data, the Ethernet Routing Switch 5000 Series supports IGMP snoop for IGMPv1, IGMPv2, and IGMPv3. With IGMP snoop enabled on a VLAN,

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

the switch forwards the multicast group data to only those ports that are members of the group. Using IGMP snoop, VLANs can provide the same benefit as IP Multicast routers, but in the local area.

The Ethernet Routing Switch 5000 Series identifies multicast group members by listening to IGMP packets (IGMP reports, leaves, and queries) from each port. The switch suppresses the reports by not forwarding them out to other VLAN ports, forcing the members to continuously send their own reports. Using the information gathered from the reports, the switch builds a list of group members. After the group members are identified, the switch blocks the IP Multicast stream from exiting any port that does not connect to a group member, conserving bandwidth.

As shown in the following figure, after the switches learn which ports are requesting access to the IP Multicast stream, all other ports not responding to the queries are blocked from receiving the IP Multicast data.

Figure 26
5000 Series switch running IGMP snooping

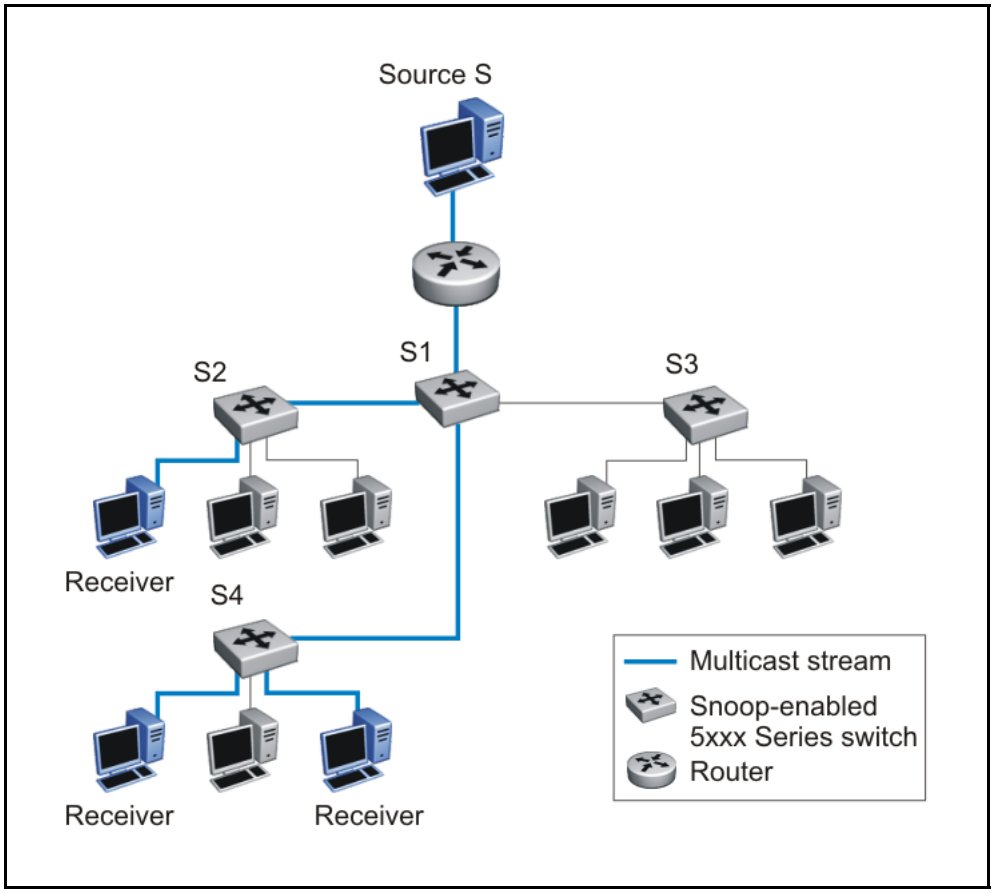

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The switch continues to forward the IGMP membership reports from the hosts to the multicast routers and forwards queries from multicast routers to all port members of the VLAN.

IGMPv3 snooping

In the IGMPv3 snooping mode, the Ethernet Routing Switch 5000 Series recognizes IGMPv3 reports and queries. While the Ethernet Routing Switch 5000 Series can understand and process all IGMPv3 record types, it does not support source filtering (INCLUDE, EXCLUDE, ALLOW, BLOCK of multicast sources) and it does not support SSM (Source Specific Multicast). The switch can recognize whether a source list is populated or blank, but cannot identify the specific sources to filter. As a result, it makes logical conclusions based on the group information available and the type of filter that is applied to a blank or populated source list.

The following table shows how IGMPv3 snooping handles different record types.

Table 8
IGMPv3 snooping with record types

IGMP v3 record type	Without multicast source ({ })	Action	With multicast source(s) ({S1, S2...})	Action
MODE_IS_INCLUDE (1)	This is INCLUDE NONE.	LEAVE the group.	This is INCLUDE multicast sources.	JOIN the group. Discard multicast source information.
MODE_IS_EXCLUDE (2)	This is EXCLUDE NONE.	JOIN the group.	This is EXCLUDE sources.	JOIN the group. Discard multicast source information.
CHANGE_TO_ INCLUDE_MODE (3)	This is include filter mode for multicast group.	LEAVE the group	This is include filter mode for multicast group.	JOIN the group. Discard multicast source information.
CHANGE_TO_ EXCLUDE_MODE (4)	This is exclude filter mode for multicast group.	JOIN the group.	This is exclude filter mode for multicast group.	JOIN the group. Discard multicast source information.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IGMP v3 record type	Without multicast source ({ })	Action	With multicast source(s) ({S1, S2...})	Action
ALLOW_NEW_ SOURCES (5)	This type is for allowing new sources. This record type comes with sources. (This case may not happen.)	JOIN the group.	This type is for allowing new sources.	JOIN the group. Discard multicast source information.
BLOCK_OLD_ SOURCES (6)	This type is for blocking existing sources.	JOIN the group.	This type is for blocking existing sources.	LEAVE the group. Discard multicast source information.

IGMP proxy

With IGMP snoop enabled, the switch can receive multiple reports for the same multicast group. Rather than forward each report upstream, the Ethernet Routing Switch 5000 Series can consolidate these multiple reports using the IGMP proxy feature. With IGMP proxy enabled, if the switch receives multiple reports for the same multicast group, it does not transmit each report to the upstream multicast router. Instead, the switch forwards the first report to the querier and suppresses the rest. If new information emerges, for example if the switch adds another multicast group or receives a query since the last report is transmitted upstream, then the switch forwards a new report to the multicast router ports.

To enable IGMP Proxy, you must first activate IGMP snooping.

In the following figure, switches S1 to S4 represent a LAN connected to an IP Multicast router. The router periodically sends Host Membership Queries to the LAN and listens for a response from end stations. All of the clients connected to switches S1 to S4 are aware of the queries from the router.

One client, connected to S2, responds with a host membership report. Switch S2 intercepts the report from that port, and generates a proxy report to its upstream neighbor, S1. Also, two clients connected to S4 respond with host membership reports, causing S4 to intercept the reports and to generate a consolidated proxy report to its upstream neighbor, S1.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 27
5000 Series switch running IGMP proxy

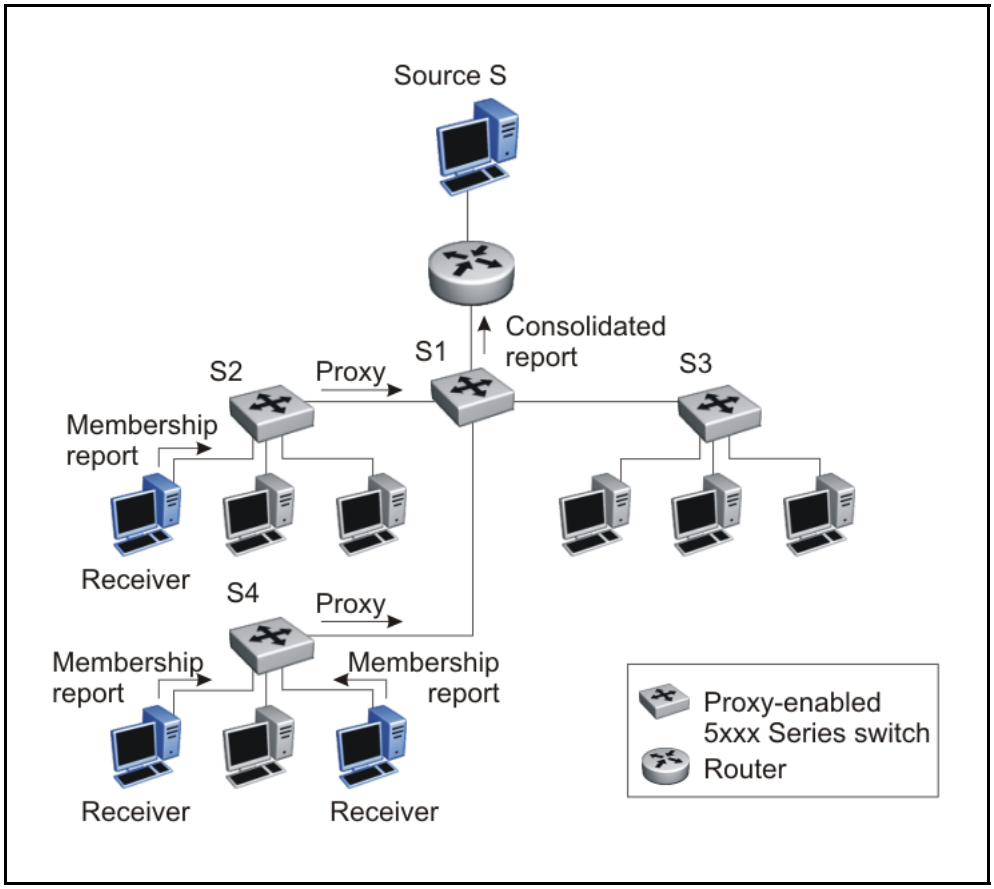

Switch S1 treats the consolidated proxy reports from S2 and S4 as if they were reports from any client connected to its ports, and generates a consolidated proxy report to the designated router. In this way, the router receives a single consolidated report from that entire subnet.

The consolidated proxy report generated by the switch remains transparent to Layer 3 of the International Standardization Organization, Open Systems Interconnection (ISO/OSI) model. (The switch IP address and MAC address are not part of proxy report generation.) The last reporting IGMP group member in each VLAN represents all of the hosts in that VLAN and IGMP group.

Report forwarding

When forwarding IGMP membership reports from group members, the Ethernet Routing Switch 5000 Series forwards the reports only to those ports where multicast routers are attached. For this the switch maintains a list of multicast querier routers and the multicast router (mrouter) ports on

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

which they are attached. The switch learns of the multicast querier routers by listening to the queries sent by the routers where source address is not 0.0.0.0., or to PIM ports.

Static mrouter port and nonquerier

If two IGMP routers are active on a VLAN, the router with the lower IP address is the querier, and the router with the higher IP address operates as a nonquerier. Only querier routers forward IGMP queries on the VLAN; nonqueriers do not forward IGMP queries. IGMP snoop considers the port on which the IGMP query is received as the active IGMP multicast router (mrouter) port. IGMP snoop is not aware of nonquerier IGMP routers.

By default, IGMP snoop forwards reports to the IGMP querier router only. To allow the switch to forward reports to the nonquerier router as well, you can configure the port connected to the nonquerier as a static mrouter port.

The following figure shows how static mrouter ports operate. Here, the Ethernet Routing Switch 5000 Series has port members 5/1 and 6/1 connected to IGMP routers in VLAN 10. In this case, router 1 is the IGMP querier because it has a lower IP address than router 2. Router 2 is then considered the nonquerier.

By default, the switch learns of the multicast querier routers by listening to the IGMP queries. In this case, port 6/1 connected to querier router 1 is identified as an mrouter port.

To forward reports to IGMP router 2 as well, you can configure port 5/1 on the switch as a static mrouter port. In this case, the IGMP reports are forwarded to both routers.

ATTENTION

Configure a static mrouter port only when multiple multicast routers are present that are not directly attached to one another but are directly attached to the VLAN (technically an invalid configuration). If multicast routers have an existing route between them (the valid configuration) and this field is configured, a multicast loop forms.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 28
Static mrouter port and nonquerier

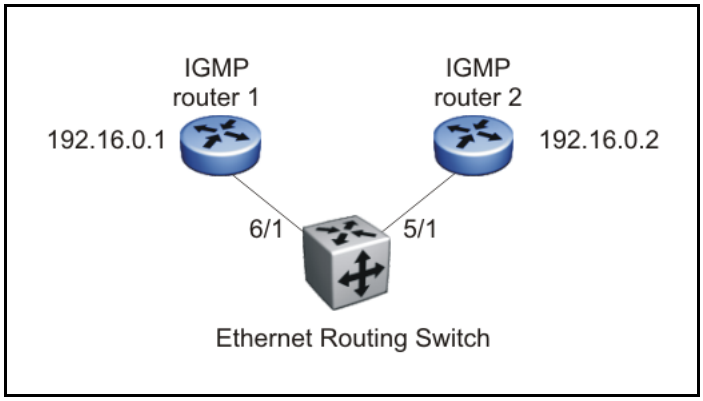

Unknown multicast packet filtering

With IGMP snoop enabled, if the switch receives multicast packets with destination addresses that it has not already registered using IGMP reports, the switch floods all such packets to all ports on the VLAN. All unknown multicast streams of a group are flooded on the VLAN until at least one port in the VLAN becomes a member of that group.

On the Ethernet Routing Switch 5000 Series, you can enable the unknown multicast filtering feature so that the unknown multicast packets are not flooded on the VLAN. To enable unknown multicast filtering, you can use the **`vlan igmp unknown-mcast-no-flood`** NNCLI command.

With this feature enabled, the switch forwards all unknown multicast traffic to IGMP static mrouter ports only. The traffic is not forwarded to dynamically discovered mrouter ports. If you require unknown multicast traffic to be forwarded to certain ports (for example, to forward Layer 3 multicast routing traffic), set the ports as static mrouter ports.

Nortel recommends that you enable this feature when IGMP snooping is enabled. User settings for the unknown multicast filtering feature are stored in NVRAM.

Allowing a multicast MAC address to flood all VLANs

The unknown multicast filtering feature introduces a potential problem when a Layer 2 VLAN is placed between two Layer 3 switches that are exchanging protocol packets such as OSPF. Since the protocols do not join a multicast group, the associated MAC addresses cannot be identified by the IGMP snooping process. These packets are dropped by the Layer 2 switch since the unknown multicast filtering feature is enabled. The two Layer 3 switches can never establish adjacencies and the OSPF protocol fails.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Using the **vlan igmp unknown-mcast-allow-flood** NNCLI command, you can specify MAC addresses or multicast IP addresses that need to be flooded on the switch even when the unknown multicast filtering feature is enabled. The specified MAC or IP addresses are added to the allow-flood table for all VLANs. Any matching packets are flooded on all ports of a VLAN.

Because multicast MAC addresses starting with 01:00:5E map to multiple IP addresses, you cannot specify 01:00:5E MAC addresses in the allow-flood table. Instead, you must specify the required multicast IP address to flood. For instance, you cannot add MAC address 01.00.5E.01.02.03 to the allow-flood table, but you can add IP address 224.1.2.3.

For all other types of MAC address, you can enter the MAC address directly to allow flooding. For example, to allow flooding of STP BPDUs, you can specify MAC address 01:80:c2:00:00:00.

IGMP snooping configuration rules

The IGMP snooping feature operates according to specific configuration rules. When configuring the switch for IGMP snooping, consider the following rules that determine how the configuration reacts in any network topology:

- IGMP snooping supports up to 240 multicast groups. (PIM-SM supports up to 1000 multicast groups.)

  If the multicast group table reaches its limit, a new entry cannot be added with a JOIN message or a new sender identifying a new group. The multicast stream from the new sender is discarded by the hardware. New entries can be added again when the table is not full.

- When you specify MAC addresses or IP addresses to be flooded on the switch, the specified MAC or IP addresses are flooded on all VLANs on the switch or stack. You cannot flood addresses for a specific VLAN only. In addition, if multicast join messages are received for IP addresses specified in the allow-flood table, these IP addresses are not displayed in the IGMP group membership table. In other words, the switch does not learn groups if they are specified in the allow-flood table.

- A port that is configured for port mirroring cannot be configured as a static mrouter port.

  The switch does support mirroring of IGMP control packets as well as multicast data packets.

- If a Multi-Link Trunk member is configured as a static mrouter port, all of the Multi-Link Trunk members are configured as static mrouter ports. Also, if you remove a static mrouter port, and it is a Multi-Link Trunk

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

member, all Multi-Link Trunk members are automatically removed as static mrouter port members.

- Static mrouter ports must be port members of at least one VLAN.
- The IGMP snooping feature is not STP dependent.
- The IGMP snooping feature is not Rate Limiting dependent.
- The snooping feature must be enabled for the proxy feature to have any valid meaning.
- Static mrouter ports are configured per VLAN.
- The Ethernet Routing Switch 5510 does not support IGMP with Multiple Spanning Tree.

ATTENTION
Because IGMP snooping is set up per VLAN, all IGMP changes are implemented according to the VLAN configuration for the specified ports.

IGMP and stacking

All IGMP features that are supported in the standalone mode are also supported in stacking mode. The configuration of IGMP from the NNCLI is supported only from the base unit. This behavior s similar to all the other Layer 3 functions and routing protocols. However, it is different from the IGMP implementation in previous releases.

Default IGMP values

Parameters	Range	Default Value
Snooping	Enable/Disable	Disable
Version	1–3	2
Proxy	Enable/Disable	Disable
Query Interval	0–65535	125
Last Member Query Interval	0–255	10
Query Max. Response Time	0–255	100
Robust Value	2–255	2
Router Alert	Enable/Disable	Disable
Multicast Router ports	Port masks	Disable port masks

IGMP snooping interworking with Windows clients

This section describes an interworking issue between Windows clients and the Ethernet Routing Switch 5000 Series when you enable IGMP snooping for multicast traffic.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Under normal IGMP snooping operation, as soon as a client joins a specific multicast group, the group is no longer unknown to the switch and the switch sends the multicast stream only to the ports which request it.

Windows clients, in response to IGMPv2 queries from the switch, initially reply with IGMPv2 reports. However, after a period of time, the Windows clients switch to IGMPv3 reports, which the Ethernet Routing Switch 5000 Series does not recognize. In this case, the switch prunes the Windows client from the group and only forwards traffic to any non-Microsoft clients that are left in the group. If no other group members are left, the switch can revert to flooding all ports (in which case, the Windows client still receives the stream). Alternatively, the switch may be pruned altogether from the multicast group (in which case, the Windows client no longer receives the stream.)

To force a Windows client to only use IGMPv1 or IGMPv2 reports so that these symptoms do not occur, change the TCP/IP settings in the Windows Registry located under the following registry key:

```
HKEY_LOCAL_MACHINE
\SYSTEM
\CurrentControlSet
\Services
\Tcpip
\Parameters
```

The specific parameter which controls the IGMP Version is:

```
IGMPVersion
Key:  Tcpip\Parameters
Value Type:  REG_DWORD—Number
Valid Range:  2, 3, 4
Default:  4
```

To set the Windows client to only utilize IGMPv2 change the IGMPVersion parameter to 3 (2 specifies IGMPv1, 3 specifies IGMPv2, and 4 specifies IGMPv3).

The IGMPVersion parameter may not be present in the list of the TCP/IP parameters. By default, the system assumes the IGMPv3 value (4). To configure it for IGMPv2, you must create the parameter as a DWORD key in the registry and specify Decimal 3.

ATTENTION
If you edit the Windows registry incorrectly, you can severely damage your system. As a minimal safeguard, back up your system data before undertaking changes to the registry.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IGMP Send Query

A multicast query router communicates with hosts on a local network by sending IGMP queries. This router periodically sends a general query message to each local network of the router. This is standard multicast behavior.

Nortel recommends that each VLAN using IGMP multicast have a router performing multicast queries. This router typically has PIM or DVMRP enabled. Currently, PIM is only available for standalone devices. Networks with no standalone devices currently have no capability for implementing the pruning of IGMP traffic. The IGMP Send Query functionality allows a switch or stack to be configured as an active query router without the need for dedicating a standalone switch in each network to the task.

There are several behavioral differences between a traditional query router and a switch or stack using the IGMP Send Query functionality. The following differences should be noted:

- There is no election process. When a switch or stack restarts, the code will send some queries as part of IGMP start up. This process will stop other devices sending queries while they detect the new device starting up. The last active device sending queries on the network is the active one. This is not the case with Layer 3 IGMP behavior.
- If the current active device stops sending queries, a timeout period must elapse before another device takes over. This may result in an ageout of groups, and subsequent flooding, before a new query is sent and the pruning process restarts. This occurs only during the transition between active query devices. Once the new device is established, queries will be sent as configured in the Query Interval and Robust Values fields.
- Multiple active query devices are not supported. Enabling multiple devices establishes one active device and other devices listening to take over should the active device fail.

IGMP Send Query functionality can only be enabled when IGMP snooping is active on the switch or stack.

Successful deployment of this feature is dependant on the addition of IP addresses from all devices in the IGMP domain. This is true even when non-management VLANs are used.

Protocol Independent Multicast-Sparse Mode

Protocol Independent Multicast-Sparse Mode (PIM-SM), as defined in RFC 2362, supports multicast groups spread out across large areas of a company or the Internet. Unlike dense-mode protocols, such as Distance Vector Multicast Routing Protocol (DVMRP), that initially flood multicast

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

traffic to all routers over an entire internetwork, PIM-SM sends multicast traffic only to routers that belong to a specific multicast group and that choose to receive the traffic. This technique reduces traffic flow over wide area network (WAN) links and minimizes the overhead costs of processing unwanted multicast packets.

Dense-mode protocols that use the flood-and-prune technique are efficient when receivers are densely populated; however, for sparsely populated networks, PIM-SM is more efficient.

PIM-SM is independent of any specific unicast routing protocol, but it does require the presence of a unicast routing protocol, such as RIP or OSPF. PIM-SM uses the information from the unicast routing table to create and maintain multicast trees that allow PIM-enabled routers to communicate.

A PIM-SM network consists of several multipoint data streams, each targeted to a small number of LANs in the internetwork. For example, customers whose networks consist of multiple hosts on different LANs in many dispersed locations can use PIM-SM to simultaneously access a video data stream, such as a video teleconference.

In some cases, PIM-SM stream initialization can take several seconds.

PIM-SM concepts and terminology

The following sections describe PIM-SM concepts and terminology.

PIM-SM sources and receivers

With PIM-SM, a host can be a source, a receiver, or both:

- A source, also known as a sender, sends multicast data to a multicast group.
- A receiver receives multicast data from one or several sources that send data to a multicast group.

PIM neighbor discovery

To discover neighbors, PIM routers exchange PIM hello packets. When PIM is enabled on a router interface, the interface forwards PIM hello packets to the all-PIM-Routers multicast address (224.0.0.13).

Each PIM hello packet contains a holdtime that specifies the period that the receiving router must wait before declaring the neighbor unreachable. This holdtime is configurable as the query interval for each interface. Each PIM interface continues to send hello messages at the configured query interval.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Required elements for PIM-SM operation

PIM-SM operates in a domain of contiguous routers that have PIM-SM enabled. Each router must run an underlying unicast routing protocol to provide routing table information to PIM-SM.

Each PIM-SM domain requires the following routers:

- designated routers (DR)
- rendezvous-point (RP) router
- bootstrap router (BSR)

Within the PIM-SM domain, each group can have only one active RP router and one active BSR. The active BSR is chosen among a list of candidate-BSRs, and the active RP is chosen among a list of candidate-RPs. You can configure the Ethernet Routing Switch 5000 Series to be a candidate-BSR, a candidate-RP, or both.

Designated router

The designated router (DR) serves as the link from sources and receivers to the other routers in the PIM-SM domain. There are typically multiple DRs in a PIM-SM domain.

On any subnet, the DR is the PIM-SM router with the highest IP address. The DR performs the following tasks:

- sends register messages to the RP router on behalf of directly connected sources
- sends join/prune messages to the upstream router on behalf of directly connected receivers
- maintains information about the status of the active RP router

ATTENTION
You cannot manually configure a router as a DR. If a router is enabled with PIM-SM and it is the PIM-SM router with the highest IP address on the subnet, it automatically acts as the DR for any directly attached sources and receivers, as required.

Rendezvous-point router

A multicast group has only one active rendezvous-point (RP) router. The RP performs the following tasks:

- manages one or several IP Multicast groups
- becomes the root for the shared tree to these groups
- accepts join messages from receivers

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- registers sources that want to send data to group members
- forwards data to the group

At the RP router, receivers meet new sources. Sources register with the RP to identify themselves to other routers on the network; receivers join the RP-based multicast distribution tree to learn about new sources.

For each multicast group, PIM-SM builds a multicast distribution tree, known as the shared tree, with the RP at the root and all receivers downstream from the RP. Although you can physically locate the RP anywhere on the network, the RP must be as close to the source as possible.

Active RP selection

The active RP is calculated among a list of candidate RPs (C-RP). Within each group, you can configure multiple PIM-SM routers as C-RPs.

Each C-RP sends unicast advertisement messages to the BSR. The BSR creates a list of C-RPs, which is referred to as the RP set. The BSR periodically sends bootstrap messages that contain the complete RP set to all routers in the group. Each router uses the same hash function to determine which router in the set is going to be the RP (given the same RP set, each router points to the same RP). If the active RP fails, routers can recalculate the active RP using the reduced set of C-RPs.

Static RP

You can use the static RP feature to configure a static entry for an RP. Static RP-enabled routers do not learn about C-RPs through the BSR. With static RP enabled, the router ignores BSR messages and loses all dynamically learned BSR information. When you configure static RP entries, the router adds them to the RP set as though they are learned through the BSR.

You can use the static RP feature when dynamic learning is not needed, typically in small networks or for security reasons. You can also enable static RP to allow communication with routers from other vendors that do not use the BSR mechanism. Some vendors use early implementations of PIM-SMv1 that do not support the BSR or proprietary mechanisms like the Cisco Auto-RP. For a network to work properly with static RP, all the routers in the network (including routers from other vendors) must be configured with the same RP or RPs, if several RPs are present in the network.

To configure static RP on a router, the next hop of the unicast route toward the static RP must be a PIM-SM neighbor. If a route change causes the next hop toward an already configured static RP to become

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

a non-PIM neighbor, the PIM-SM protocol fails on the router. The state of the configured RP on the router remains invalid until it can be reached through a PIM neighbor.

To avoid a single point of failure, you can also configure redundant static RPs.

When you configure a static RP, take into account the following considerations:

- You cannot configure a static RP-enabled router as a BSR or as a C-RP.
- All dynamically learned BSR information is lost. However, if you disable static RP, the router clears the static RP information and regains the BSR functionality.
- Static RPs do not age; that is, they cannot time out.
- Routers do not advertise static RPs; therefore, if a new PIM-SM neighbor joins the network, this new neighbor does not know about the static RP unless you configure the neighbor with that static RP.
- All the routers in the network (including routers from other vendors) must map to the same RP.
- In a PIM-SM domain with both static and dynamic RP routers, you cannot configure one of the (local) interfaces of the static RP routers as RP.
- To avoid a single point of failure, you can configure redundant static RPs for the same group prefix. If a mix of Nortel and other vendor routers exist across the network, ensure that all routers use the same active RP because other vendors can use different algorithms to elect the active RP. The Ethernet Routing Switch 5000 Series uses the hash function defined in the PIM-SM standard to elect the active RP, with the highest C-RP address selected to break a tie. Other vendors can use the lowest IP address to break the tie.
- You cannot assign a priority to static RP entries, although the Ethernet Routing Switch accepts priority values from non-Nortel routers for interoperability.
- A static RP that you configure on the router is alive as long as the router has a unicast route to the network for the static RP. If the router loses this route, it invalidates the static RP and uses the hash algorithm to remap all affected groups. If the router regains this route, it validates the static RP and uses the hash algorithm to remap the affected groups.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Bootstrap router

The bootstrap router (BSR) receives advertisement messages from the C-RPs. The BSR adds the C-RPs and their group prefixes to the RP set. The BSR sends bootstrap messages that contain the complete RP set to all routers in the domain to allow them to learn group-to-RP mappings.

Only one BSR exists for each PIM-SM domain.

Active BSR selection

Within a PIM-SM domain, you can configure a set of routers as candidate BSRs (C-BSR). The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs have equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with a higher priority to the domain, it automatically becomes the new BSR.

PIM-SM shared trees and shortest-path trees

PIM-SM uses two types of multicast distribution trees to deliver data packets to group members: shared trees and shortest-path trees (SPT).

Shared tree

The shared tree connects all members of the multicast group to a central core router, the active RP, which is at the root of the shared tree.

The construction of the shared tree begins when a host sends an IGMP membership report to a local DR to join a multicast group. The DR in turn signals join messages toward the RP. The intermediate routers toward the RP add the group entry when forwarding the join messages. When the join messages reach the RP, the RP adds the tree branch to the shared tree for the group.

Although a shared tree is less efficient than a source-rooted tree, PIM-SM shared tree reduces the network bandwidth during tree construction and maintenance, as flood-and-prune messages are not required.

The following figure shows an example of an RP-based shared tree.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 29
RP-based shared tree

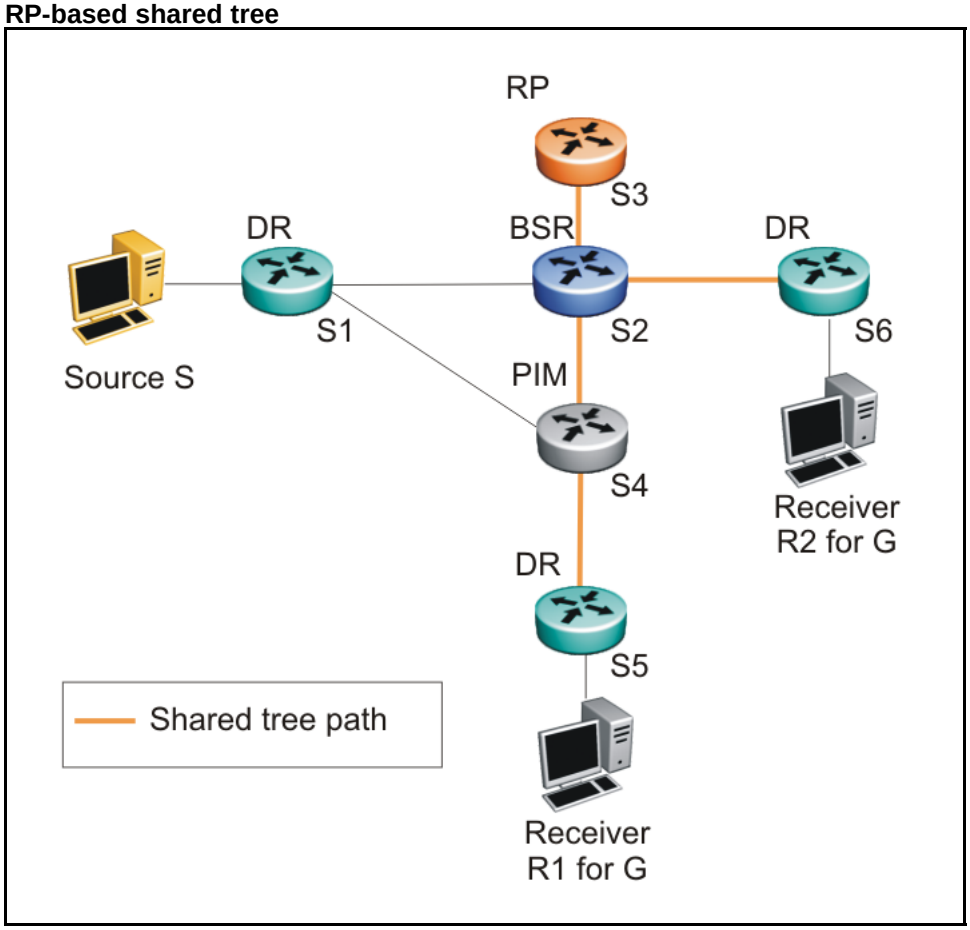

Traffic forwarding with the shared tree

All group traffic initially flows from the RP downstream through the shared tree to the receivers. To forward multicast data from a source to group members, the source DR encapsulates the multicast packets in Register messages that it then unicasts to the RP. The RP decapsulates the Register messages, and then forwards the multicast data to any existing group members downstream using the shared tree.

In the shared tree, the RP router represents a potential bottleneck and a single point of failure. As a result, PIM-SM allows local DRs to bypass the share tree and switch to a source-rooted shortest path tree.

Shortest path tree

When multicast packets arrive at the receiver DR, the DR can identify the IP address of the source. If the DR determines that the shared tree is not the optimal path back to the source, it sends a join message directly to the source DR. This new direct path from the source to the receiver DR is

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

the source-based shortest-path tree (SPT). When the receiver DR starts receiving traffic directly from the source, it sends a prune message to the RP to stop sending messages over the shared tree.

With the 5000 Series switches, the DR switches to the SPT after it receives the first packet from the RP.

The following figure shows an example of a source-based SPT.

Figure 30
Source-based SPT

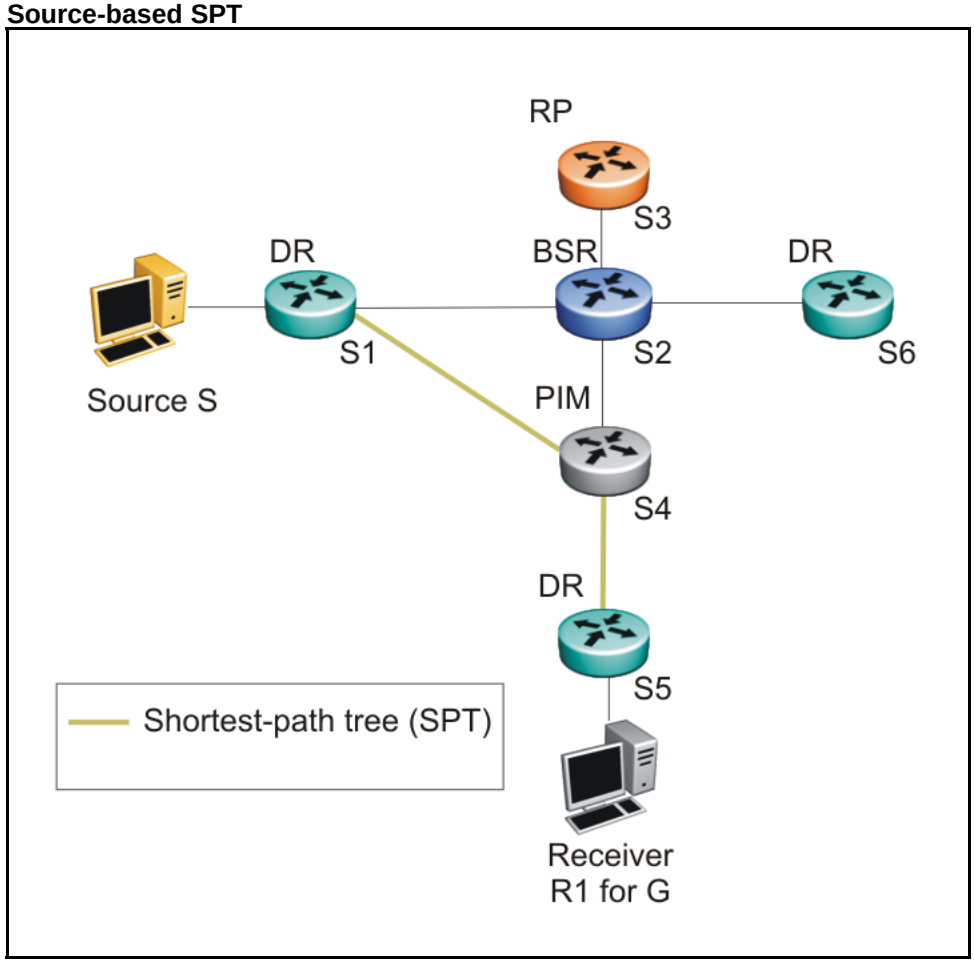

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Receiver joining a group and receiving data from a source

The following steps describe how the receiver R1 in Figure 29 "RP-based shared tree" (page 101) and Figure 30 "Source-based SPT" (page 102) joins multicast group G:

1. The BSR distributes RP information to all switches in the network. In this example, based on the RP hash function, S3 is the RP for group G.
2. Receiver R1 multicasts an IGMP host membership report for group G, which the DR (S5) receives.
3. Acting on this report, S5 creates a (*,G) route entry in the multicast forwarding table and sends a (*,G) join to the RP.
4. The intermediate routers toward the RP (S4 and S2) add the (*,G) route entry when forwarding the join message to the RP.
5. The RP adds the port that receives the join as a downstream port for the (*,G) group.
6. The source S starts multicasting data to group G.
7. The source DR (S1) encapsulates the data in a Register message that it unicasts to the RP (S3).
8. S3 decapsulates the multicast data and forwards it down the shared tree. Group member S5 receives the data and forwards it to receiver R1.
9. After S5 receives the first packet, it knows the IP address for the source. S5 creates an (S,G) entry in the multicast forwarding table, and sends a (S,G) join to the source. All intermediate routers along the path to the source create the (S,G) entry. S5 also prunes itself from the RP shared tree.
10. S1 forwards multicast packets to S5 over the SPT.

ATTENTION
The PIM-SM topology shown in this example is simplified and is not the best design for a network if the source and receiver are placed as shown. In general, RPs are placed as close as possible to sources.

Source-to-RP SPT

Rather than continue to receive multicast traffic from the source through unicast Register messages, the RP also switches to a source-based SPT. After it receives the first source Register message, it sends a join message to the source DR to receive the data through a multicast rather than unicast stream. After it receives the first multicast packet over the SPT, the RP sends a register-stop message to the source to stop sending the data in register messages.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

On the Ethernet Routing Switch 5000 Series, the DR only forwards the first multicast packet as a Register packet to the RP, and immediately goes into discard mode until it receives a join message from the RP. During this time, there is brief data loss of the multicast stream.

After the source DR processes the join message, the DR forwards native multicast packets to the RP over the SPT path.

Register suppression timeout

If a source registers with an RP, but no receivers are registered to receive the traffic, the RP sends a register-stop to the source.

After receiving a register-stop message from the RP, the source DR starts a register suppression timer (the default value is 60 seconds).

Shortly before the register suppression timer expires, the source DR sends a register message with no encapsulated packets to the RP router. This null-register message prompts the RP router to determine whether new downstream receivers joined the group. If no new members have joined the group, the RP router sends another register-stop message to the DR for the source, and the register suppression timer restarts. In this way, the DR can regularly poll the RP to determine whether any new members have joined the group without forwarding larger traffic packets to the RP unnecessarily.

A lower register suppression timeout produces traffic bursts from the DR more frequently, whereas with a higher value, new receivers face a longer join latency.

Receivers leaving a group

If all directly connected members of a multicast group leave or time out, and no downstream members remain, the DR sends a prune message upstream and PIM-SM deletes the route entry after that entry times out.

PIM assert

When a PIM router connects a source to a LAN segment and it detects a second PIM router with a route to the same source on the same segment, the routers exchange Assert messages to determine which router is to forward the multicast stream on the segment.

PIM passive interfaces

You can specify whether you want a PIM interface to be active or passive. The default is active. Active interfaces can transmit and receive PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. This feature is useful when you have a high number of PIM interfaces and these interfaces are connected to end users, not to other routers.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

A PIM passive interface drops any messages of the following type:

- Hello
- Join/Prune
- Register
- Register-Stop

ATTENTION
A device can send Register and Register-Stop messages to a PIM passive interface, but that interface cannot send out these messages.

- Assert
- Candidate-RP-Advertisement
- Bootstrap

If a PIM passive interface receives any of these types of messages, it drops them, and the switch logs a message, detailing the type of protocol message received and the IP address of the sending device. These log messages help to identify the device that is performing routing on the interface, which is useful if you must disable a device that is not operating correctly.

The PIM passive interface maintains information through the IGMP protocol about hosts that are related to senders and receivers, but the interface does not maintain information about any PIM neighbors.

You can also use the PIM passive interface feature as a security measure to prevent routing devices from becoming attached and participating in the multicast routing of the network.

You can configure a PIM passive interface as a BSR or an RP, although Nortel does not recommend these options.

ATTENTION
Before you change the state (active or passive) of a PIM interface, disable PIM on that interface. Disabling PIM prevents instability in the PIM operations, especially when neighbors are present or streams are received.

PIM-SM capabilities and limitations

The following list describes the capabilities and limitations of PIM-SM on the Ethernet Routing Switch 5000 Series.

- PIM-SM is not supported on the Ethernet Routing Switch 5510 platform.
- PIM-SM is not supported in a stack in the current release.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- You cannot allow the PIM-SM shared path tree or SPT to span across any Layer 2 switches. Be sure to implement your topology such that the unicast routes from any DR to the PIM RP and to all multicast sources travel through directly-connected PIM neighbors only. Otherwise, network issues may arise.
- PIM-SSM is not supported in the current release.
- PIM-SM cannot be enabled on brouter ports.
- PIM-SM is not supported over SMLT or IST.
- PIM-SM is not supported on a secondary IP of a Layer 3 VLAN.
- A maximum of 32 PIM-SM active interfaces are supported.
- A maximum of 64 total PIM-SM interfaces are supported.
- You can configure only one Candidate-RP per switch for any number of groups (up to 50 groups).
- You can configure static RP for up to 50 groups.
- You can configure every PIM-enabled interface as Candidate-BSR.
- PIM-SM supports forwarding of the multicast stream on ECMP, but traffic balancing is not supported. PIM-SM picks one route for its RPF check and uses it for all streams when joining a source on this route.
- On the Ethernet Routing Switch 5600, a maximum of 1000 (S,G) entries is supported.
- On the Ethernet Routing Switch 5500, a maximum of 490 (S,G) entries is supported.
- A Layer 2 IGMP snooping-enabled switch can learn a maximum of 240 groups from clients. However, a PIM router can learn more that 240 groups if it is connected to more than one snooping-enabled switch. If each Layer 2 switch learns 240 groups, the number of groups the PIM router learns is: 240 * number of Layer 2 switches. However, the number of (*,G) entries on the PIM router is limited to 960 for a 5600 switch and 490 for a 5500 switch.
- If a PIM server and IGMP receiver are in the same VLAN, you cannot connect them to the same port. To have a PIM server and IGMP receiver on the same port, the server and receiver must be in different tagged VLANs.
- With static RP, priority is not supported in a pure Nortel-only solution. If the 5000 Series switch is connected to a non-Nortel router that is running static RP, then the 5000 Series switch can learn the priority as advertised by the non-Nortel router.
- Passive interfaces are supported on the edge only (where the port only has connections to either clients or servers). Make sure that any

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

passive interfaces are not in the path of any PIM RPF paths, otherwise the network may not work.

Enabling or disabling routing with IGMP enabled

You cannot enable PIM-SM and IGMP snooping on the same VLAN. As a result, when enabling or disabling routing on a VLAN, the IGMP functionality operates as follows:

- When a VLAN is changed from Layer 3 to Layer 2, IGMP reduces the functionality to snooping.
- When an IGMP snooping-enabled VLAN is changed from Layer 2 to Layer 3, the switch continues to support snooping until you enable PIM-SM.
- When a Layer 3 VLAN has v1 and v2 snooping enabled and you try to enable PIM-SM, the switch displays a warning message to disable snooping before enabling PIM-SM.

Nonsupported IGMP features

The following list describes nonsupported IGMP features on the Ethernet Routing Switch 5000 Series.

- IGMPv3 (only snooping is supported)
- IGMPv3 snooping – Source Filtering (INCLUDE and EXCLUDE)
- Multicast Router Discovery
- Fast Leave
- Channel Limit—Limit the number of groups a host (port) can join at a time
- Access Control feature—Block/Allow range of addresses
- IGMP Static Address configuration

Default PIM-SM values

The following table describes the PIM-SM default values.

Parameter	Definition	Range	Default Value
Global PIM-SM status	Indicates the status of PIM-SM on the switch.	Enabled/Disabled	Disabled
PIM mode	Specifies the global PIM mode on the switch.	Sparse mode (SSM is not supported in the current release)	Sparse mode

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Parameter	Definition	Range	Default Value
Bootstrap Period	At the elected BSR, this is the interval between originating bootstrap messages.	5–32 757 seconds	60 seconds
C-RP Advertise Timeout	Indicates the frequency with which candidate RPs periodically send C-RP-Adv messages.	5–26 214 seconds	60 seconds
Unicast Route Change Timeout	Specifies how often the routing information that PIM uses is updated from the routing table manager (RTM).	2–65 535 seconds	5 seconds
Join/Prune Interval	Indicates how long the switch waits between sending out join/prune messages to the upstream neighbors.	1–18 724 seconds	60 seconds
Register Suppress Timeout	Specifies how often the source DR polls the RP using data packets encapsulated in Register messages.	5–65 535 seconds	60 seconds
Data Discard Timer	After the router forwards the first source packet to the RP, this value specifies how long (in seconds) the router discards subsequent source data while waiting for a join from the RP.	5–65 535 seconds	60 seconds
Static RP	Indicates the status of static RP on the switch.	Enabled/Disabled	Disabled
Forward Cache Timeout	Indicates the PIM-SM forward cache expiry value. This value is used in aging PIM-SM mroutes.	10–86 400 seconds	210 seconds
VLAN PIM-SM status	Indicates the status of PIM-SM on the VLAN.	Enabled/Disabled	Disabled
Hello Interval	Sets the hello interval for the VLAN.	0–18 724 seconds	30 seconds

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Parameter	Definition	Range	Default Value
Interface Type	Sets the interface type on a particular VLAN.	• active: allows PIM-SM control traffic to be transmitted and received. • passive: prev ents PIM-SM control traffic from being transmitted or received.	Active
Candidate-BSR priority	Indicates whether the router is acting as a C-BSR on a particular VLAN, and if so, the priority associated with it.	0 to 255	–1 (indicates that the interface is not a Candidate BSR)
Candidate-RP	Indicates whether the VLAN interface is configured as a C-RP. With the Ethernet Routing Switch 5000 Series, you can configure only one local interface as a C-RP for any number of groups.	IP address of the C-RP interface and the associated group and mask.	None defined (disabled)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IP routing configuration using NNCLI

This chapter describes the procedures you can use to configure routable VLANs using the NNCLI.

The Nortel Ethernet Routing Switch 5000 Series are Layer 3 (L3) switches. This means that a regular Layer 2 VLAN becomes a routable Layer 3 VLAN if an IP address and MAC address are attached to the VLAN. When routing is enabled in Layer 3 mode, every Layer 3 VLAN is capable of routing as well as carrying the management traffic. You can use any Layer 3 VLAN instead of the Management VLAN to manage the switch.

For more information on creating and configuring VLANs, see *Configuration — VLANs, Spanning Tree, and Link Aggregation* (NN47200-502).

IP routing configuration procedures

To configure inter-VLAN routing on the switch, perform the following steps:

Step	Action
1	Enable IP routing globally.
2	Assign an IP address to a specific VLAN or brouter port. Routing is automatically enabled on the VLAN or brouter port when you assign an IP address to it.
	--End--

In the above procedure, you are not required to enable IP routing as the first step. All IP routing parameters can be configured on the Nortel Ethernet Routing Switch 5000 Series before routing is actually enabled on the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IP routing configuration navigation

- "Configuring global IP routing status" (page 112)
- "Displaying global IP routing status" (page 112)
- "Configuring an IP address for a VLAN" (page 113)
- "Configuring IP routing status on a VLAN" (page 114)
- "Configuring a secondary IP address for a VLAN" (page 114)
- "Displaying the IP address configuration and routing status for a VLAN" (page 115)
- "Displaying IP routes" (page 116)
- "Performing a traceroute" (page 117)
- "Entering Router Configuration mode" (page 118)

Configuring global IP routing status

Use this procedure to enable and disable global routing at the switch level. By default, routing is disabled.

Procedure steps

Step	Action
1	To configure the status of IP routing on the switch, enter the following from the Global Configuration mode: `[no] ip routing`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
no	Disables IP routing on the switch.

Displaying global IP routing status

Use this command to display the status of IP blocking on the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To display the status of IP blocking on the switch, enter the following from the User EXEC mode: `show ip routing`

--End--

Configuring an IP address for a VLAN

To enable routing an a VLAN, you must first configure an IP address on the VLAN.

Procedure steps

Step	Action
1	To configure an IP address on a VLAN, enter the following from the VLAN Interface Configuration mode: `[no] ip address <ipaddr> <mask> [<MAC-offset>]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the configured IP address and disables routing on the VLAN.
<ipaddr>	Specifies the IP address to attach to the VLAN.
<mask>	Specifies the subnet mask to attach to the VLAN
[<MAC-offset>]	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address. The valid range is 1-256. Specify the value 1 for the Management VLAN only. If no MAC offset is specified, the switch applies one automatically.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring IP routing status on a VLAN

Use this procedure to enable and disable routing for a particular VLAN.

Procedure steps

Step	Action
1	To configure the status of IP routing on a VLAN, enter the following from the VLAN Interface Configuration mode: `[default] [no] ip routing`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
default	Disables IP routing on the VLAN.
no	Disables IP routing on the VLAN.

Configuring a secondary IP address for a VLAN

Use this procedure to configure a secondary IP interface to a VLAN (also known as multinetting). You can have a maximum of eight secondary IP addresses for every primary address, and you must configure the primary address before configuring any secondary addresses.

Primary and secondary interfaces must reside on different subnets.

To remove a primary IP address from a VLAN, you must first remove all secondary addresses from the VLAN.

Prerequisites

- Configure a primary IP address on the VLAN.

Procedure steps

Step	Action
1	To configure the secondary IP interface on the VLAN, enter the following from the VLAN Interface Configuration mode.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
[no] ip address <ip address> <mask> [<mac offset>]
secondary
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
no	Removes the configured IP address. To remove a primary IP address from a VLAN, you must first remove all secondary addresses from the VLAN.
<ipaddr>	Specifies the IP address to attach to the VLAN.
<mask>	Specifies the subnet mask to attach to the VLAN
[<MAC-offset>]	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address. The valid range is 1-256. Specify the value 1 for the Management VLAN only. If no MAC offset is specified, the switch applies one automatically.

Job aid: Example of adding a secondary IP interface to a VLAN

Primary and secondary interfaces must reside on different subnets. In the following example, 4.1.0.10 is the primary IP and 4.1.1.10 is the secondary IP.

```
(config)# interface vlan 4
(config-if)# ip address 4.1.0.10 255.255.255.0 6
(config-if)# ip address 4.1.1.10 255.255.255.0 7 secondary
```

Displaying the IP address configuration and routing status for a VLAN

Use this procedure to display the IP address configuration and the status of routing on a VLAN.

Procedure steps

Step	Action
1	To display the IP address configuration on a VLAN, enter the following from the VLAN Privileged Exec mode:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
show vlan ip [vid <vid>]
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[vid <vid>]	Specifies the VLAN ID of the VLAN to be displayed. Range is 1-4094.

Job aid

The following table shows the field descriptions for the **show vlan ip** command.

Field	Description
Vid	Specifies the VLAN ID.
ifIndex	Specifies an index entry for the interface.
Address	Specifies the IP address associated with the VLAN.
Mask	Specifies the mask.
MacAddress	Specifies the MAC address associated with the VLAN.
Offset	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address.
Routing	Specifies the status of routing on the VLAN: enabled or disabled.

Displaying IP routes

Use this procedure to display all active routes in the routing table.

Route entries appear in ascending order of the destination IP addresses.

Procedure steps

Step	Action
1	To display all active routes in the routing table, enter the following from the User EXEC command mode: **show ip route [<dest-ip>] [-s <subnet> <mask>] [summary]**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
[<dest-ip>]	Specifies the destination IP address of the route to display.
[-s <subnet> <mask>]	Specifies the destination subnet of the routes to display.
[summary]	Displays a summary of IP route information.

Job aid

The following show sample outputs for the **`show ip route`** command.

show ip route command output

Ip Route

DST	MASK	NEXT	COST	VLAN	PORT	PROT	TYPE	PRF
0.0.0.0	0.0.0.0	10.3.2.137	1	1	1/21	S	IB	5
2.2.2.0	255.255.255.0	2.2.2.2	1	2	----	C	DB	0
10.3.2.0	255.255.255.0	10.3.2.199	1	1	----	C	DB	0

Total Routes: 3

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

show ip route summary command output

```
------------------------
Connected routes :  65
Static routes :  2
RIP routes :  512
OSPF routes :  512
------------------------
Total routes :  1091
------------------------
```

Performing a traceroute

Use this procedure to display the route taken by IP packets to a specified host.

When applied to a stack, this procedure can be executed only on the base unit.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To perform a traceroute, enter the following from the Global Configuration mode: **traceroute <Hostname|A.B.C.D.|ipv6-addr> <-m> <-p> <-q> <-v> <-w> <1-1464>**
2	Type **CTRL+C** to interrupt the command.
	--End--

Variable definitions

The following table describes the command variables.

Variable	Value
Hostname	Specifies the name of the remote host.
A.B.C.D.	Specifies the IP address of the remote host.
ipv6-addr	Specifies the IPv6 address of the remote host.
-m	Specifies the maximum time to live (ttl). The value for this parameter is in the rage from 1-255. The default value is 10. Example: **traceroute 10.3.2.134 -m 10**
-p	Specifies the base UDP port number. The value for this parameter is in the range from 0-65535. Example: **traceroute 1.2.3.4 -p 87**
-q	Specifies the number of probes per time to live. The value for this parameter is in the range from 1-255. The default value is 3. Example: **traceroute 10.3.2.134 -q 3**
-v	Specifies verbose mode. Example: **traceroute 10.3.2.134 -v**
-w	Specifies the wait time per probe. The value for this parameter is in the range from 1-255. The default value is 5 seconds. Example: **traceroute 10.3.2.134 -w 15**
<1-1464>	Specifies the UDP probe packet size. **TIP**: probe packet size is 40 plus specified data length in bytes. Example: **traceroute 10.3.2.134 -w 60**

Entering Router Configuration mode

Use this procedure to enter Router Configuration mode and configure parameters related to routing protocols. Router Configuration mode is used to configure RIP, OSPF, and VRRP.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To enter router mode, enter the following from the Global Configuration mode: `router {rip | ospf | vrrp}`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Brouter port configuration using NNCLI

A brouter port is a single-port VLAN that can route IP packets as well as bridge all non-routable traffic. The following section describes the NNCLI commands used to configure and manage brouter ports on the Nortel Ethernet Routing Switch 5000 Series .

Brouter port configuration navigation

- "Configuring a brouter port" (page 121)
- "Displaying the brouter port configuration" (page 122)

Configuring a brouter port

Use this procedure to create and manage a brouter port on the switch.

Procedure steps

Step	Action
1	To configure a brouter port, enter the following from the FastEthernet Interface Configuration mode: **brouter [port <brouter_port>] vlan <vid> subnet <ip_address/mask> [routing enable]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the brouter configuration from the switch

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<brouter_port>	Specifies the port to configure as a brouter port. When modifying a brouter this is the port of the existing brouter to modify.
<vid>	Specifies the VLAN ID of the brouter. When creating a new brouter port, this is the VLAN ID assigned to the brouter port.
<ip_address/mask>	Specifies the IP address and subnet mask of the brouter. When creating a new brouter, this is the IP address and subnet mask assigned. When modifying a brouter port, this is the new IP address and subnet mask to assign to the port. The subnet mask portion is expressed as a value between 0 and 32.
[routing enable]	Enables Layer 3 routing on the brouter port.

Displaying the brouter port configuration

Use this procedure to display the brouter port configuration on the switch.

Procedure steps

Step	Action
1	To display the brouter port configuration, enter the following from the User Exec mode: **`show brouter [port <brouter_port>]`**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[port <brouter_port>]	Narrows the scope the command to the specified port. Omission of this parameter displays all ports.

Job aid

The following table shows the field descriptions for the **`show brouter`** command.

Field	Description
Port	Specifies the brouter port number.
Brouter VID	Specifies the brouter VLAN ID.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
IP/Mask	Specifies the IP address and subnet mask of the brouter.
IP Routing	Specifies whether IP routing is enabled or disabled on the brouter.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Static route configuration using NNCLI

This chapter describes the procedures you can use to configure static routes using the NNCLI.

Static route configuration navigation

- "Configuring a static route" (page 125)
- "Displaying static routes" (page 126)
- "Configuring a management route " (page 127)
- "Displaying the management routes " (page 128)

Configuring a static route

Use this procedure to configure a static route. Create static routes to manually configure a path to destination IP address prefixes.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLANs to be routed.

Procedure steps

Step	Action
1	To configure a static route, enter the following from the Global Configuration command mode: `[no] ip route <dest-ip> <mask> <next-hop> [<cost>] [disable] [enable] [weight <cost>]`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the specified static route.
<dest-ip>	Specifies the destination IP address for the route being added. 0.0.0.0 is considered the default route.
<mask>	Specifies the destination subnet mask for the route being added.
<next-hop>	Specifies the next hop IP address for the route being added.
[<cost>]	Specifies the weight, or cost, of the route being added. Range is 1-65535.
[disable]	Disables the specified static route.
[enable]	Enables the specified static route.
[weight <cost>]	Changes the weight, or cost, of an existing static route. Range is 1-65535.

Displaying static routes

Use this procedure to display all static routes, whether these routes are active or inactive.

Procedure steps

Step	Action
1	To display a static route, enter the following from the User EXEC command mode: **`show ip route static [<dest-ip>] [-s <subnet> <mask>]`**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<dest-ip>	Specifies the destination IP address of the static routes to display.
[-s <subnet> <mask>]	Specifies the destination subnet of the routes to display.

Job aid

The following table shows the field descriptions for the **`show ip route static`** command.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
DST	Identifies the route destination.
MASK	Identifies the route mask.
NEXT	Identifies the next hop in the route.
COST	Identifies the route cost.
VLAN	Identifies the VLAN ID on the route.
PORT	Specifies the ports.
PROT	Specifies the routing protocols. For static routes, options are LOC (local route) or STAT (static route).
TYPE	Indicates the type of route as described by the Type Legend on the NNCLI screen.
PRF	Specifies the route preference.

Configuring a management route

Use this procedure to create a management route to the far end network, with a next-hop IP address from the management VLAN's subnet. A maximum of 4 management routes can be configured on the switch.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the management VLAN interface.

Procedure steps

Step	Action
1	To configure a static management route, enter the following from the Global Configuration command mode: `[no] ip mgmt route <dest-ip> <mask> <next-hop>`

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
[no]	Removes the specified management route.
<dest-ip>	Specifies the destination IP address for the route being added.
<mask>	Specifies the destination subnet mask for the route being added.
<next-hop>	Specifies the next hop IP address for the route being added.

Displaying the management routes

Use this procedure to display the static routes configured for the management VLAN.

Procedure steps

Step	Action
1	To display the static routes configured for the management VLAN, enter the following from the User EXEC mode: **`show ip mgmt route`**

--End--

Job aid

The following table shows the shows the field descriptions for the **`show ip mgmt route`** command.

Field	Description
Destination IP	Identifies the route destination.
Subnet Mask	Identifies the route mask.
Gateway IP	Identifies the next hop in the route.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

OSPF configuration using NNCLI

This chapter describes the procedures you can use to configure OSPF using NNCLI.

The 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 and the tagging of externally-derived routing information.

OSPF commands used during the configuration and management of VLANs in the Interface Configuration mode can be used to configure *any* VLAN regardless of the one used to log into the command mode. Insert the keyword **vlan** with the number of the VLAN to be configured after the command keywords **ip ospf**. The current VLAN remains the one used to log into the Interface Configuration command mode after the command execution.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Assign an IP address to the VLAN that you want to enable with OSPF.

  Routing is automatically enabled on the VLAN when you assign an IP address to it.

OSPF configuration navigation

- “Configuring the OSPF hardware mode” (page 131)
- “Configuring the router ID” (page 131)
- “Configuring global OSPF status” (page 131)
- “Configuring global OSPF parameters ” (page 132)
- “Displaying global OSPF parameters” (page 133)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- “Configuring OSPF area parameters” (page 133)
- “Displaying OSPF area configuration” (page 135)
- “Displaying OSPF area range information” (page 135)
- “Enabling OSPF routing on an interface” (page 136)
- “Assigning an interface to an OSPF area” (page 136)
- “Configuring the OSPF properties for an interface” (page 137)
- “Displaying OSPF interface timers” (page 138)
- “Displaying OSPF interface configurations” (page 139)
- “Displaying OSPF neighbors” (page 139)
- “Specifying a router as an ASBR” (page 140)
- “Configuring the OSPF authentication type for an interface” (page 140)
- “Defining simple authentication keys for OSPF interfaces” (page 141)
- “Defining MD5 keys for OSPF interfaces” (page 141)
- “Displaying OSPF MD5 keys” (page 142)
- “Applying an MD5 key to an OSPF interface” (page 142)
- “Displaying OSPF interface authentication configuration” (page 143)
- “Configuring a virtual link” (page 144)
- “Creating a virtual interface message digest key ” (page 145)
- “Configuring automatic virtual links” (page 146)
- “Displaying OSPF virtual links” (page 149)
- “Displaying OSPF virtual neighbors” (page 149)
- “Configuring an OSPF host route” (page 149)
- “Displaying OSPF host routes” (page 150)
- “Displaying the OSPF link state database” (page 150)
- “Displaying the external link state database” (page 151)
- “Initiating an SPF run to update the OSPF LSDB” (page 151)
- “Displaying OSPF default port metrics” (page 152)
- “Displaying OSPF statistics” (page 152)
- “Displaying OSPF interface statistics” (page 152)
- “Clearing OSPF statistics counters” (page 153)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring the OSPF hardware mode

Use this procedure to configure the OSPF hardware-compatibility mode.

Procedure steps

Step	Action
1	To configure the OSPF hardware-compatibility mode, enter the following from the Global Configuration command mode: `ip ospf op-mode {5510 | non-5510}`

--End--

Configuring the router ID

Use this procedure to configure the router ID, which is expressed in the form of an IP address.

Procedure steps

Step	Action
1	To configure the router ID, enter the following from the OSPF Router Configuration command mode: `[no] router-id <router_id>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Resets the router ID to 0.0.0.0.
<router_id>	Specifies the unique identifier for the router.

Configuring global OSPF status

Use this procedure to configure the status of OSPF globally on the switch.

By default, OSPF is disabled.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To configure OSPF globally on the switch, enter the following from the Global Configuration command mode: `[default] [no] router ospf enable`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Disables OSPF globally on the switch.
[no]	Disables OSPF globally on the switch.
enable	Enables OSPF globally on the switch. If omitted, enters OSPF Router configuration mode without enabling OSPF.

Configuring global OSPF parameters

Use this procedure to define the global OSPF parameters, including default cost metric, RFC 1583 compatibility, OSPF holddown timer, and OSPF system traps.

Procedure steps

Step	Action

1	To define the global OSPF parameters, enter the following from the OSPF Router Configuration command mode:
	`[default] [no] default-cost {ethernet | fast-ethernet | gig-ethernet | ten-gig-ethernet} <metric_value>`
	`rfc1583-compatibility enable`
	`timers basic holddown <timer_value>`
	`trap enable`

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
[default]	Sets the specified parameter to the default value.
[no]	Disables the specified feature (applicable only to rcf-1583-compatibilit y enable and trap enable).
{ethernet | fast-ethernet | gig-ethernet | ten-gig-ethernet} <metric_value>	Specifies the default cost metric to assign to the specified port type. The metric value is an integer between 1 and 65535. The default values are as follows: • ethernet (10 Mb/s): 100 • fast-ethernet (100 Mb/s): 10 • gig-ethernet (1000 Mb/s): 1 • ten-gig-ethernet (10000 Mb/s): 1
rfc1583-compatibility enable	Enables RFC 1583 compatibility on the switch.
timers basic holddown <timer_value>	Specifies a holddown timer value between 3 and 60 seconds.
trap enable	Enables OSPF system traps.

Displaying global OSPF parameters

Use this procedure to display global OSPF parameters.

Procedure steps

Step	Action
1	To display global OSPF parameters, enter the following from the User EXEC command mode: **show ip ospf**

--End--

Configuring OSPF area parameters

Use this procedure to configure OSPF area parameters.

Procedure steps

Step	Action
1	To configure the OSPF area parameters, enter the following from the OSPF Router Configuration command mode:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
[default] [no] area <area-id>
[default-cost {0-16777215}]
[import {external | noexternal | nssa}]
[import-summaries {enable}]
[range {subnet_mask} [{nssa-entlink | summary-link}]
[advertise-mode {no-summarize | summarize | suppress}]
[advertise-metric {0-65535}]]
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Sets the specified parameter to the default value (applicable only for default-cost, import, import-summaries, and range).
[no]	Removes the specified OSPF configuration (applicable only for import-summaries [disables] and range [removes the specified range]).
<area-id>	Specifies the Area ID in dotted decimal notation (A.B.C.D).
default-cost {0-16777215}	Specifies the default cost associated with an OSPF stub area.
import {external | noexternal | nssa}	Specifies the area type by defining the area's support for importing Autonomous System external link state advertisements: • external: specifies a normal area • noexternal: specifies a stub area • nssa: specifies an NSSA ***Note:*** The configuration of a totally stubby area (no summary advertising) is a two step process. First, define an area with the import flag set to *noexternal*. Second, disable import summaries in the same area with the command **`no area <area-id> import-summaries enable`**.
import-summaries {enable}	Controls the import of summary link state advertisements into stub areas. This setting has no effect on other areas.
range {subnet_mask} [{nssa-entlink | summary-link}] [advertise-mode {no-summarize | summarize | suppress}] [advertise-metric {0-65535}]	Specifies range parameters for the OSPF area.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying OSPF area configuration

Use this procedure to display OSPF area configuration.

Procedure steps

Step	Action
1	To display OSPF area configuration, enter the following from the User EXEC command mode: `show ip ospf area [<area-id>]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<area-id>	Displays configuration information about the specified OSPF area. Omitting this parameter displays information for all OSPF areas.

Displaying OSPF area range information

Use this procedure to display OSPF area range information.

Procedure steps

Step	Action
1	To display OSPF area range information, enter the following from the User EXEC command mode: `show ip ospf area-range <range>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<range>	Displays configuration information about the specified OSPF area range. Omitting this parameter displays information for all OSPF area ranges.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Enabling OSPF routing on an interface

Use this procedure to enable OSPF routing on an interface and assign the interface to an OSPF area.

Procedure steps

Step	Action
1	To enable OSPF on an interface, enter the following from the OSPF Router Configuration command mode: **network <ip_address> [area <area_id>]** OR enter the following from the Interface Configuration command mode **[no] ip ospf vlan <vid> area <area_id>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables OSPF routing on an interface.
<ip_address>	Specifies the IP address of interface to be enabled for OSPF routing.
area <area_id>	Specifies the ID of the area assigned to the interface in dotted decimal notation (A.B.C.D).
<vid>	The VLAN ID to be enabled for OSPF routing.

Assigning an interface to an OSPF area

Use this procedure to assign an interface to an OSPF area.

Procedure steps

Step	Action
1	To assign an interface to an OSPF area, enter the following from the Interface Configuration command mode: **ip ospf area <area-id>**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
<area-id>	Specifies the unique ID of the area to which the interface connects. An area ID of 0.0.0.0 indicates the OSPF area backbone and is created automatically by the switch.

Configuring the OSPF properties for an interface

Use this command to used to configure OSPF properties for an interface.

Procedure steps

Step	Action
1	To configure OSPF interface properties, enter the following from the Interface Configuration command mode: **ip ospf** **advertise-when-down enable** **cost <interface_cost>** **dead-interval <interval>** **hello-interval <interval>** **mtu-ignore enable** **network {broadcast | passive}** **priority <0-255>** **retransmit-interval <interval>** **transit-delay <interval>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
advertise-when-down enable	Enables advertisement of the OSPF interface even when the interface is operationally unavailable.
cost <interface_cost>	Specifies the cost assigned to the interface. This is an integer value between 1 and 65535.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
dead-interval <interval>	Specifies a dead interval for the interface. This is the interval of time that a neighbor waits for a Hello packet from this interface before the neighbor declares it down. This is an integer value between 0 and 2147483647.
hello-interval <interval>	Specifies the amount of time between transmission of hello packets from this interface. This is an integer value between 1 and 65535.
mtu-ignore enable	Instructs the interface to ignore the packet MTU size specified in Database Descriptors.
network {broadcast | passive}	Defines the type of OSPF interface this interface is.
priority <priority_value>	Assigns a priority to the interface for the purposes of Designated Router election. This is an integer value between 0 and 255.
retransmit-interval <interval>	Defines the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. This is an integer value between 0 and 3600.
transit-delay <interval>	Defines the transit delay for this OSPF interface in seconds. The transit delay is the estimated number of seconds it takes to transmit a link-state update over the interface. This is an integer value between 0 and 3600.

Displaying OSPF interface timers

Use this procedure to display OSPF interface timers

Procedure steps

Step	Action
1	To display OSPF timers, enter the following from the User EXEC command mode: **show ip ospf timer {interface [FastEthernet <portlist> | vlan <vid>] | virtual-links}** OR

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
show ip ospf int-timers
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
interface [FastEthernet <portlist> | vlan <vid>]	Displays configured timers for the specified interface. If no interface is specified, all interface timers are displayed.
virtual-links	Displays configured timers for virtual links.

Displaying OSPF interface configurations

Use this procedure to display OSPF interface configurations.

Procedure steps

Step	Action
1	To display OSPF interface configurations, enter the following from the User EXEC command mode: `show ip ospf interface [FastEthernet <portlist> | vlan <vid>]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[FastEthernet <portlist> | vlan <vid>]	Displays OSPF configuration for the specified interface. If no interface is specified, all interface configurations are displayed.

Displaying OSPF neighbors

Use this procedure to display information about OSPF neighbors for the router.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To display OSPF neighbors, enter the following from the User EXEC command mode: `show ip ospf neighbor`

--End--

Specifying a router as an ASBR

Use this procedure to identify a router as an Autonomous System Boundary Router.

Procedure steps

Step	Action
1	To configure a router as an ASBR, enter the following from the OSPF Router Configuration command mode: `[no] as-boundary-router [enable]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
no	Removes the ASBR configuration for the router.

Configuring the OSPF authentication type for an interface

Use this procedure to configure the interface authentication type.

Procedure steps

Step	Action
1	To configure the interface authentication type, enter the following from the Interface Configuration mode:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip ospf authentication-type {message-digest | simple |
none}
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
{message-digest | simple | none}	Specifies the authentication type.

Defining simple authentication keys for OSPF interfaces

Use this procedure to configure an interface authentication password.

Procedure steps

Step	Action
1	To configure an interface authentication password, enter the following from the Interface Configuration mode: `ip opsf authentication-key <password>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<password>	Specifies the password to be configured. This password can be up to 8 characters in length.

Defining MD5 keys for OSPF interfaces

Use this procedure to define the MD5 keys.

Procedure steps

Step	Action
1	To define an MD5 key, enter the following from the Interface Configuration command mode:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip ospf message-digest-key <key_number> md5 <key_value>
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<key_number>	Specifies an index value for the MD5 key being configured. This is an integer value between 1 and 255.
<key_value>	Specifies the value of the MD5 key. This is a string value of up to 16 characters in length.

Displaying OSPF MD5 keys

Use this procedure to display OSPF MD5 key configuration.

Procedure steps

Step	Action
1	To display OSPF MD5 keys, enter the following from the User EXEC command mode: `show ip ospf authentication {interface [FastEthernet <portlist> | vlan <vid>] | virtual-links}`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
interface [FastEthernet <portlist> | vlan <vid>]	Displays configured MD5 authentication keys for the specified interface. If no interface is specified, all interface MD5 keys are displayed.
virtual-links	Displays configured MD5 authentication keys for virtual links.

Applying an MD5 key to an OSPF interface

Use this procedure to specify the primary MD5 key (configured using the **ip ospf message-digest-key** command) to use for authentication in instances where interface authentication uses an MD5 key.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Each OSPF interface supports up to 2 keys, identifiable by key ID, to facilitate a smooth key transition during the rollover process. Only the selected primary key is used to encrypt the OSPF transmit packets.

Assuming that all routers already use the same key for authentication and a new key is required, the process of key change is as follows:

1. Add the second key to all routers. The routers will continue to send OSPF packets encrypted with the old key.
2. Activate the second key on all routers by setting it as the primary key. Routers will send OSPF packets encrypted with the new key while still accepting packets using the old key. This is necessary as some routers will not have activated the new key.
3. After all routers activate the new key, remove the old key.

Procedure steps

Step	Action
1	To specify the primary MD5 key, enter the following from the Interface Configuration command mode: **ip ospf primary-md5-key <key_id>**
	--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<key_id>	Specifies the index value for the MD5 key to apply. This is an integer value between 1 and 255.

Displaying OSPF interface authentication configuration

Use this procedure to display the authentication type and key applied to interfaces.

Procedure steps

Step	Action
1	To display OSPF authentication configuration for interfaces, enter the following from the User EXEC command mode:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
show ip ospf int-auth
```

--End--

Configuring a virtual link

Use this procedure to create a virtual interface.

Procedure steps

Step	Action
1	To create a virtual interface, enter the following from the OSPF Router Configuration command mode: `[no] area virtual-link <area-id> <nghbr-router-id> {[authentication-key <WORD>] [authentication-type {none|simple|message-digest}] [primary-md5-key <1-255>] [dead-interval <1-2147483647>] [hello-interval <1-65535>] [retransmit-interval <1-3600>] [transit-delay <1-3600>]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Deletes a virtual interface.
<area_id>	Specifies the transit area ID in dotted decimal notation (A.B.C.D).
<nghbr-router-id>	Specifies the neighbor router ID expressed as an IP address.
authentication-key <WORD>	Specifies the unique identifier assigned to the authentication key.
authentication-type	Specifies one of the following authentication types: • none • simple • password • message digest MD5 TIP: Up to 2 MD5 keys are allowed for message digest. The default authentication type is none.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
primary-md5-key	Specifies the user-selected key used to encrypt OSPF protocol packets for transmission.
dead-interval	Specifies the time interval, in seconds, that a Hello packet has not been transmitted from the virtual interface before its neighbors declare it down. Expressed as an integer from 1-2147483647, the default dead interval value is 60 seconds.
hello-interval	Specifies the time interval, in seconds, between transmission of Hello packets from the virtual interface. Expressed as an integer from 1-65535, the hello-interval default value is 10 seconds.
retransmit-interval	Specifies the time interval, in seconds, between link stage advertisement retransmissions for adjacencies belonging to the virtual interface. Expressed as an integer from 1-3600, the default value is 5 seconds.
transit-delay	Specifies the estimated number of seconds required to transmit a link state update packet over the virtual interface. Expressed as an integer from 1-3600, the default value is 1 second.

Creating a virtual interface message digest key

Use this procedure to create a virtual interface message digest key.

Procedure steps

Step	Action
1	To create a virtual interface message digest key, enter the following from the OSPF Router Configuration command mode: **area virtual-link message-digest-key <A.B.C.D.> <A.B.C.D./0-32> <1-255> md5-key <WORD>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Deletes a virtual interface message digest key.
<A.B.C.D.>	Specifies the transit area Id expressed as an IP address.
<W.X.Y.Z>	Specifies the neighbor router ID expressed as an IP address.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<1-255>	Specifies the primary MD5 key value, expressed as an integer from 1-255.
md5-key <WORD>	Specifies the user-selected key used to encrypt OSPF protocol packets for transmission.

Configuring automatic virtual links

Use this command to enable global automatic Virtual Link creation.

Procedure steps

Step	Action
1	To enable global automatic Virtual Link creation, enter the following from the OSPF Router Configuration command mode: `[no] auto-vlink`
	--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables global automatic Virtual Link creation.

Job aid: example of configuring automatic virtual links

Consider the following situation:

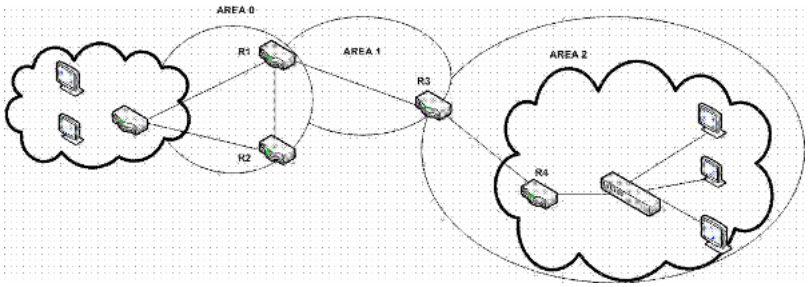

In this case, R4 in Area2 cannot be physically connected to Area0 (for some reason) and it will be connected to R3 which is NOT a backbone ABR (like R1 is for instance). As Area2 is not directly connected to

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

backbone Area0 or directly connected to a backbone ABR router, clients from Area2 will not be able to access anything outside Area2. Also, router R3 is an ABR router connected to two non-backbone areas.

In order to solve these problems, virtual-link must be configured between router R3 and R1 which are both ABRs. Virtual-link cannot be configured on non-ABR routers.

Consider the following Router IDs:

- R1 : 1.0.0.0
- R3 : 3.0.1.0
- R4 : 4.0.2.0

The virtual-link can be configured in two ways on ABR routers :

- Configuring the virtual link manually
- Configuring the virtual link automatically

The following is an example for creating an auto virtual link:

Table 9
Creating auto virtual link

```
R1 (config-router)#auto-vlink
```

Example : 1

```
R1(config)#show ip ospf
Router ID: 1.0.0.0
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  0
External Link-State Checksum:  0(0x0)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  67
New Link-State Advertisements Received:  722
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Enabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

```
R3 (config-router)#auto-vlink
```

Example : 2

```
R3(config)#show ip ospf
Router ID: 3.0.1.0
Admin Status:  Enabled
Version Number:  2
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
    Area Border Router Oper Status:  True
    AS Boundary Router Config Status:  False
    External Link-State Advertisements:  0
    External Link-State Checksum:  0(0x0)
    Type-of-Service (TOS) Routing Supported:  False
    Originated Link-State Advertisements:  67
    New Link-State Advertisements Received:  722
    OSPF Traps:  Disabled
    Auto Virtual Link Creation:  Enabled
    SPF Hold-Down Time:  10
    RFC 1583 Compatibility:  Enabled
```

The following is an example for deleting an auto virtual link:

Table 10
Deleting auto virtual link

```
R1 (config-router)#no auto-vlink

    Example : 1
    R1(config)#show ip ospf
    Router ID: 1.0.0.0
    Admin Status:  Enabled
    Version Number:  2
    Area Border Router Oper Status:  True
    AS Boundary Router Config Status:  False
    External Link-State Advertisements:  0
    External Link-State Checksum:  0(0x0)
    Type-of-Service (TOS) Routing Supported:  False
    Originated Link-State Advertisements:  67
    New Link-State Advertisements Received:  722
    OSPF Traps:  Disabled
    Auto Virtual Link Creation:  Disabled
    SPF Hold-Down Time:  10
    RFC 1583 Compatibility:  Enabled

R3 (config-router)#no auto-vlink

    Example : 2
    R3(config)#show ip ospf
    Router ID: 3.0.1.0
    Admin Status:  Enabled
    Version Number:  2
    Area Border Router Oper Status:  True
    AS Boundary Router Config Status:  False
    External Link-State Advertisements:  0
    External Link-State Checksum:  0(0x0)
    Type-of-Service (TOS) Routing Supported:  False
    Originated Link-State Advertisements:  67
    New Link-State Advertisements Received:  722
    OSPF Traps:  Disabled
    Auto Virtual Link Creation:  Disabled
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

Displaying OSPF virtual links

Use this procedure to display the configuration of OSPF virtual links.

Procedure steps

Step	Action
1	To display OSPF virtual links, enter the following from the User EXEC command mode: **show ip ospf virtual-links**
	--End--

Displaying OSPF virtual neighbors

Use this procedure to display OSPF virtual neighbors.

Procedure steps

Step	Action
1	To display OSPF virtual neighbors, enter the following from the User EXEC command mode: **show ip ospf virtual-neighbors**
	--End--

Configuring an OSPF host route

Use this procedure to add a host to a router.

Procedure steps

Step	Action
1	To add a host to a router, enter the following from the OSPF Router Configuration command mode: **[no] host-route <A.B.C.D.> metric <0-65535>**
	--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Deletes a host route from ther router.
<A.B.C.D.>	Specifies the host IP address.
metric <0-65535>	Specifies an integer between 0 and 65535 representing the configured cost of the host route.

Job aid: example of configuring an OSPF host route

The following is an example for creating a host route:

```
R3(config)#router ospf
R3(config-router)#host-route 11.11.11.111 metric 10
R3(config-router)#show ip ospf host-route
```

Host IP	Metric
11.11.11.111	10

Displaying OSPF host routes

Use this procedure to display OSPF host routes.

Procedure steps

Step	Action
1	To display OSPF host routes, enter the following from the User EXEC command mode: `show ip ospf host-route`

--End--

Displaying the OSPF link state database

Use this procedure to display the OSPF link state database.

Procedure steps

Step	Action
1	To display the OSPF link state database, enter the following from the User EXEC command mode:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
show ip ospf lsdb
{[area <area-id>]
[lsa-type <type>]
[lsid <ip_address>]
[adv-rtr <router_id>] |
detail [<router_id>] }
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[area <area-id>]	Displays OSPF LSDB information related to the specified area.
[lsa-type <type>]	Displays OSPF LSDB information for the specified LSA type.
[lsid <ip_address>]	Displays OSPF LSDB information for the specified link state ID.
[adv-rtr <router_id>]	Displays OSPF LSDB information related to the specified advertisement router.
detail [<router_id>]	Display detailed OSPF LSDB information related to the specified advertisement router. If no router is specified, all detailed LSDB information is displayed.

Displaying the external link state database

Use this procedure to display the external link state database.

Procedure steps

Step	Action
1	To display OSPF ASE LSAs, enter the following from the User EXEC command mode: `show ip ospf ase`

--End--

Initiating an SPF run to update the OSPF LSDB

Manually initiate an SPF run to immediately update the link state database. Use this procedure, in the following situations:

- when you need to immediately restore a deleted OSPF-learned route
- as a debug mechanism when the routing table entries and the link-state database are not synchronized

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To immediately initiate an SPF run, enter the following from the Global Configuration command mode: **ip ospf spf-run**
	--End--

Displaying OSPF default port metrics

Use this procedure to display OSPF default metrics for different port types.

Procedure steps

Step	Action
1	To display OSPF default metrics, enter the following from the User EXEC command mode: **show ip ospf default-cost**
	--End--

Displaying OSPF statistics

Use this procedure to display OSPF statistics.

To clear OSPF statistics counters, use the **clear ip ospf counters** command.

Procedure steps

Step	Action
1	To display OSPF statistics, enter the following from the User EXEC command mode: **show ip ospf stats**
	--End--

Displaying OSPF interface statistics

Use this procedure to display OSPF interface statistics.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To display OSPF interface statistics, enter the following from the User EXEC command mode: `show ip ospf ifstats <if-ip> [mismatch] [detail]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<if-ip>	Displays OSPF statistics for the specified interface IP address. Omitting this parameter displays statistics for the backbone area.
mismatch	Displays statistics where the area ID not matched.
detail	Display detailed statistics.

Clearing OSPF statistics counters

Use this procedure to clear OSPF statistics counters, including mismatch counters.

This procedure is applicable only to the base unit in a stack.

Procedure steps

Step	Action
1	To clear OSPF statistics counters, enter the following from the Global Configuration command mode: `clear ip ospf counters <1-4094>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<1-4094>	Specifies the VLAN ID. Range is 1-4094. If no VLAN is specified, the command clears OSPF global counters.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

OSPF configuration examples using NNCLI

The following sections provide OSPF configuration examples using NNCLI.

Navigation

- "Basic OSPF configuration examples" (page 155)
- "Advanced OSPF configuration examples" (page 158)

Basic OSPF configuration examples

This section contains the steps necessary for the initial configuration of OSPF on the switch. More advanced configuration examples can be found in "Advanced OSPF configuration examples" (page 158).

Note: In many of the following configuration examples, a brouter port is used to create a connection to the network core. This practice does not imply that a brouter port is the only means through which a core connection can be established. The use of a brouter port is only one of many ways to create such a connection.

Basic OSPF configuration

A basic OSPF configuration will learn OSPF routes from other OSPF devices and propagate routes to other OSPF devices. The following procedure describes the creation of a basic OSPF configuration:

Step	Action
1	Log into User EXEC mode. `5530-24TFD> enable`
2	Log into Global Configuration mode. `5530-24TFD# config terminal`
3	Enable IP routing globally.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config)# ip routing
```

4 Enable OSPF globally.

```
5530-24TFD(config)# router ospf en
```

5 Log into the OSPF router configuration mode. It is not necessary to make any changes at this time but entering the router configuration mode is a good way to verify that the mode has been activated.

```
5530-24TFD(config)# router ospf
```

Note: The remainder of this procedure refers to VLAN 35. Although VLAN 35 is used for this example, any port type VLAN could be used.

6 Create a port type VLAN as VLAN number 35 in spanning tree protocol group 1.

```
5530-24TFD(config)# vlan create 35 type port 1
```

7 Log into the Interface Configuration mode for VLAN 35.

```
5530-24TFD(config)# interface vlan 35
```

8 Enable IP routing on VLAN 35.

```
5530-24TFD(config-if)# ip routing
```

9 Assign an IP address to VLAN 35.

```
5530-24TFD(config-if)# ip address 1.1.2.25
255.255.255.0
```

10 Enable OSPF in VLAN 35.

```
5530-24TFD(config-if)# ip ospf en
```

11 Return to Global Configuration mode.

```
5530-24TFD(config-if)# exit
```

12 By default all ports belong to a newly created VLAN. This command removes all of the ports from VLAN 35.

```
5530-24TFD(config)# vlan members remove 35 all
```

13 Add ports 1 through 10 to VLAN 35.

```
5530-24TFD(config)# vlan members add 35 1-10
```

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Basic ASBR configuration

The Autonomous System Boundary Router (ASBR) is used in OSPF to import routes that come from non-OSPF sources such as:

- Local interfaces that are not part of OSPF.
- RIP interfaces.
- RIP learned routes.
- Static routes.

This quick reference will help in the configuration of OSPF to import these types of routes. This will allow the rest of the OSPF network to learn them as OSPF routes. To create a basic ASBR configuration, follow this procedure:

Step	Action
1	Log into User EXEC mode. `5530-24TFD> enable`
2	Log into Global Configuration mode. `5530-24TFD# config terminal`
3	Log into the OSPF router configuration mode. `5530-24TFD(config)# router ospf`
4	Enable ASBR functionality. `5530-24TFD(config-router)# as-boundary-router en`

5	Use the following commands to select the type of routes that OSPF will distribute to other OSPF devices. RIP, direct, and static routes are supported.
	`5530-24TFD(config-router)# redistribute rip en`
	`5530-24TFD(config-router)# redistribute direct en`
	`5530-24TFD(config-router)# redistribute static en`

6	Return to Global Configuration mode. `5530-24TFD(config-router)# exit`

7	Once the commands in step 5 have be used to select the types of routes to redistribute, apply the changes globally with the following commands.
	`5530-24TFD(config)#ip ospf apply redistribute rip`
	`5530-24TFD(config)#ip ospf apply redistribute direct`
	`5530-24TFD(config)#ip ospf apply redistribute static`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring ECMP for OSPF

To configure ECMP with OSPF, use the following procedure.

Usage of ECMP with OSPF is not supported on the 5510 models.

Step	Action
1	Log into User EXEC mode. `5530-24TFD> enable`
2	Log into Global Configuration mode. `5530-24TFD# config terminal`
3	Set the number of ECMP paths to use with OSPF. Up to four paths can be used. `5530-24TFD(config)# ospf maximum-path 2` This command tells the router to use up to two equal-cost paths to get to any OSPF network destination.
4	The configuration can be verified using the following command. `5530-24TFD(config)# show ecmp`

--End--

Advanced OSPF configuration examples

This section contains examples of common OSPF-related configuration tasks.

The Nortel Ethernet Routing Switch 5000 Series 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 NNCLI commands used to create the configuration.

Configuring an IP OSPF interface

An OSPF interface can be configured on a brouter port or on a VLAN. The following section demonstrates the creation of the example OSPF interface illustrated below.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 31
OSPF interface example topology

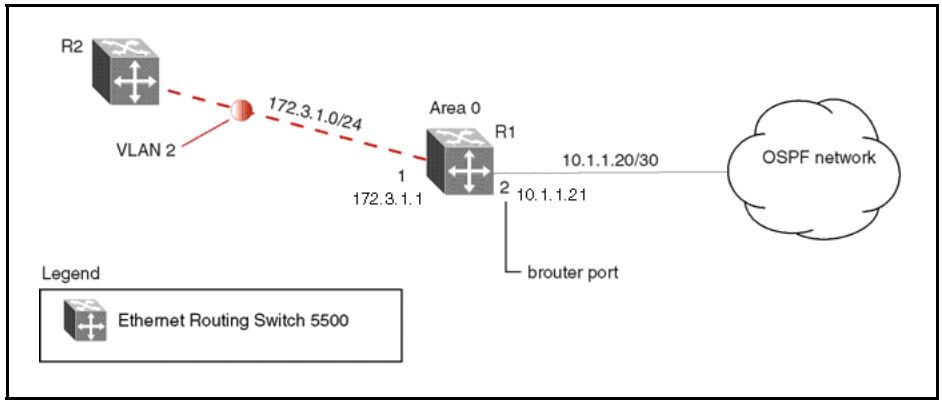

To create the OSPF interface illustrated above for router R1, follow this procedure:

Step	Action

1	Configure brouter port OSPF interface.
	
	Configure port 2 as a brouter port with VLAN ID of 2134 and enable OSPF on this interface.
	
	`5530-24TFD# config terminal`
	`5530-24TFD(config)# interface fast 2`
	`5530-24TFD(config-if)# brouter port 2 vlan 2134 subnet 10.1.1.21/30`
	`5530-24TFD(config-if)# router ospf`
	`5530-24TFD(config-router)# network 10.1.1.21`

2	Configure the VLAN OSPF interface.
	
	Create a port-based VLAN (VLAN 2) using spanning tree group 1, assign IP address 172.3.1.1 to VLAN 2 and enable OSPF on this interface.
	
	`5530-24TFD(config)# vlan create 2 type port`
	`5530-24TFD(config)# spanning-tree stp 1 add-vlan 2`
	`5530-24TFD(config)# vlan member add 2 1`
	`5530-24TFD(config)# interface vlan 2`
	`5530-24TFD(config-if)# ip address 172.3.1.1 255.255.255.0`
	`5530-24TFD(config-if)# router ospf`
	`5530-24TFD(config-router)# network 172.3.1.1`

3	Assign a router ID to the new interface and enable OSPF globally.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# router-id 1.1.1.1
5530-24TFD(config-router)# exit
5530-24TFD(config)# router ospf enable
```

--End--

OSPF security configuration example using Message Digest 5

In the configuration example illustrated below, MD5 is configured between router R1 and R2.

Figure 32
MD5 configuration example

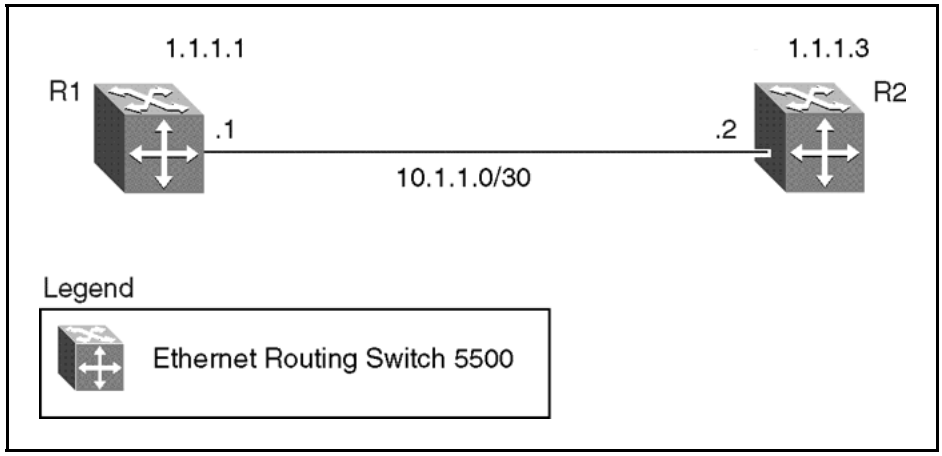

To replicate the above configuration example using the key ID 2 and key value **qwsdf89**, perform the following steps:

Step	Action
1	Configure MD5 authentication on R1.

```
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip ospf message-digest-key 2 md5
qwsdf89
5530-24TFD(config-if)# ip ospf primary-md5-key 2
5530-24TFD(config-if)# ip ospf authentication-type
message-digest
```

2	Configure MD5 authentication on R2.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip ospf message-digest-key 2 md5
qwsdf89
5530-24TFD(config-if)# ip ospf primary-md5-key 2
5530-24TFD(config-if)# ip ospf authentication-type
message-digest
```

--End--

Configuring OSPF network types

OSPF network types were created to allow OSPF-neighboring between routers over different types of network infrastructures. With this feature, each interface can be configured to support the various network types.

In the example configuration illustrated below, VLAN 2 on Nortel Ethernet Routing Switch 5000 Series 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 33
OSPF network example

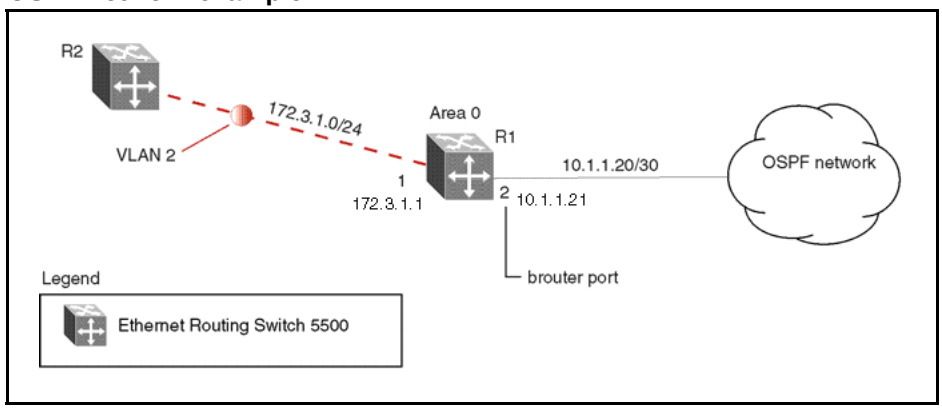

To create the configuration illustrated above for router R1, use the following commands:

```
5530-24TFD(config)# vlan create 2 type port
5530-24TFD(config)# vlan mem add 2 1
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 172.3.1.1 255.255.255.0
5530-24TFD(config-if)# ip ospf network passive
```

The Nortel Ethernet Routing Switch 5000 Series supports the following types of networks:

- **Broadcast** - Automatically discovers every OSPF router on the network by sending OSPF hellos to the multicast group

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

AllSPFRouters (224.0.0.5). Neighboring is automatic and requires no configuration. This interface type is typically used in an Ethernet environment.

- **Passive** - Allows interface network to be included in OSPF without generating LSAs or forming adjacencies. Typically used on an access network. This also limits the amount of CPU cycles required to process the OSPF routing algorithm.

Configuring Area Border Routers (ABR)

Configuration of an OSPF ABR is an automatic process on the Nortel Ethernet Routing Switch 5000 Series ; no user intervention is required. The Nortel Ethernet Routing Switch 5000 Series automatically becomes an OSPF ABR when it has operational OSPF interfaces belonging to more than one area.

In the configuration example below, the Nortel Ethernet Routing Switch 5000 Series R1 is automatically configured as an OSPF ABR after it is configured with an OSPF interface for area 0.0.0.0 and 0.0.0.2.

Figure 34
ABR configuration example

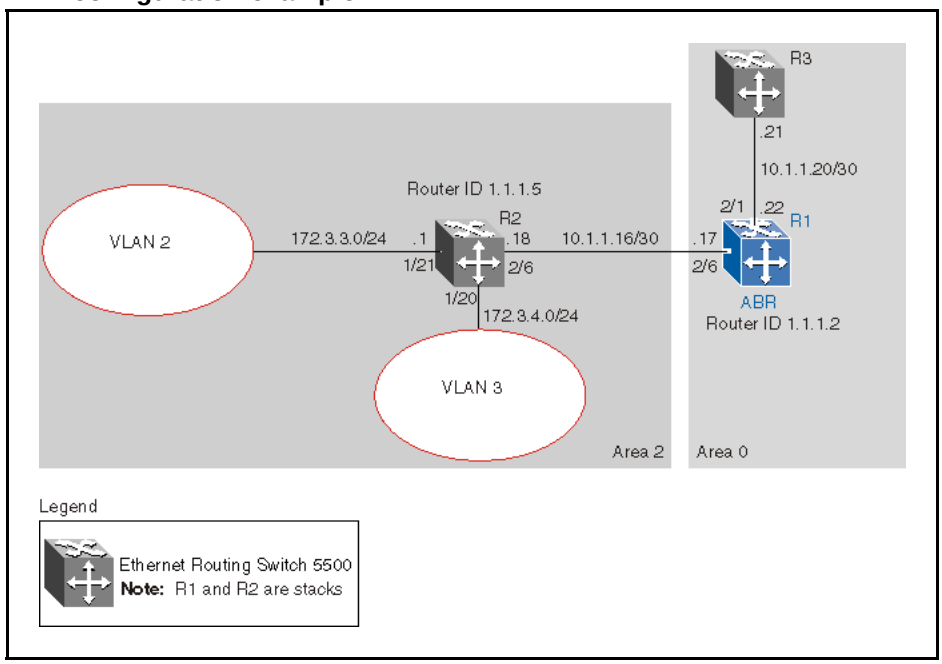

To recreate the illustrated ABR configuration, use the following procedure:

Step	Action
1	Configure an OSPF interface on port 2/6.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configure port 2/6 as a brouter port in VLAN 100.

```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(config-if)# brouter port 2/6 vlan 100 subnet
10.1.1.17/30
```

2 Configure an OSPF interface on port 2/1.

Configure port 2/1 as a brouter port in VLAN 200 and enable OSPF on this interface.

```
5530-24TFD(config)# interface fast 2/1
5530-24TFD(config-if)# brouter port 2/1 vlan 200 subnet
10.1.1.22/30
5530-24TFD(config-if)# ip ospf enable
```

3 Enable OSPF.

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 created and OSPF interface 10.1.1.17 is added to area 0.0.0.2, R1 automatically becomes an ABR.

```
5530-24TFD(config-router)# router-id 1.1.1.2
5530-24TFD(config-router)# area 0.0.0.2
5530-24TFD(config-router)# network 10.1.1.17 area
0.0.0.2
5530-24TFD(config)# router ospf enable
```

4 Configure area range.

Configure R1 to enclose the two networks (172.3.3.0 and 172.3.4.0) into an address range entry 172.3.0.0 in area 0.0.0.2. R1 will generate a single summary advertisement into the backbone for 172.3.0.0 with metric 100.

```
5530-24TFD(config-router)# area 0.0.0.2 range
172.3.0.0/16 summary-link advertise-mode summarize
advertise-metric 100
```

--End--

To display the created areas, use the **show ip ospf area** command. Usage of this command on the example configuration would yield the following output:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Area ID: 0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  2
Reachable Area Border Routers:  0
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
Area ID: 0.0.0.2
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  2
Reachable Area Border Routers:  1
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
```

To display area ranges, use the **show ip ospf area-range** command. Usage of this command on the example configuration would yield the following output:

```
Area ID Range Subnet/Mask Range Type Advertise Mode Metric
---------- ----------------- ------------ -------------- ------
0.0.0.2 172.3.0.0/16 Summary Link Summarize 100
```

To display ABR status, use the **show ip ospf** command. Usage of this command on the example configuration would yield the following output:

```
Router ID: 1.1.1.2
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  2
External Link-State Checksum:  45698(0xb282)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  5
New Link-State Advertisements Received:  34
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

Configuring Autonomous System Border Routers (ASBR)

An ASBR is a router that has a connection to another Autonomous System to distribute any external routes that originated from a protocol into OSPF. A Nortel Ethernet Routing Switch 5000 Series configured as an ASBR can:

- Distribute all OSPF routes to RIP.
- Distribute RIP, direct, or static routes to OSPF.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

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

The following configuration example displays a Nortel Ethernet Routing Switch 5000 Series configured as an ASBR between an OSPF and RIP version 2 network. In this example, the router distributes all OSPF routes to the RIP network and all RIP routes to the OSPF network.

Figure 35
ASBR distribution example

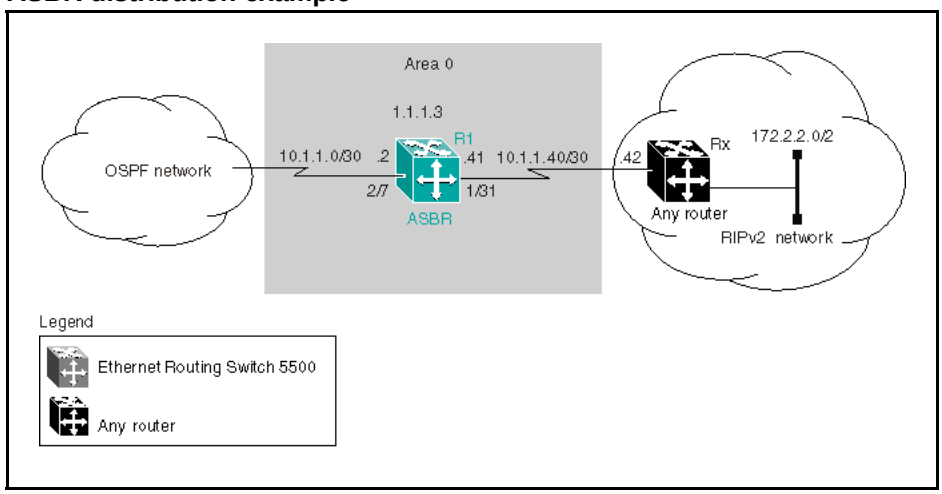

Use the following procedure to replicate the ASBR distribution example:

Step	Action

1	Configure RIP.
	Configure the RIP interface on R1 by configuring port 1/31 as a brouter port in VLAN 100 and enabling RIP on this interface.
	`5530-24TFD(config)# interface fast 1/31`
	`5530-24TFD(config-if)# brouter port 1/31 vlan 100 subnet 10.1.1.41/30`
	`5530-24TFD(config)# router rip`
	`5530-24TFD(config-router)# network 10.1.1.41`

2	Configure the RIP interface for RIP version 2 mode only.
	`5530-24TFD(config)# router rip enable`
	`5530-24TFD(config)# interface vlan 100`
	`5530-24TFD(config-if)# ip rip receive version rip2 send version rip2`

3	Configure the OSPF interface. Configure port 2/7 as a brouter port in VLAN 200 and enable OSPF on this interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config)# interface fast 2/7
5530-24TFD(config-if)# brouter port 2/7 vlan 200 subnet
10.1.1.2/30
5530-24TFD(config-if)# router ospf
5530-24TFD(config-router)# network 10.1.1.2
```

4 Make R1 the ASBR.

Configure R1 as an ASBR and assign the OSPF Router-ID.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# as-boundary-router enable
5530-24TFD(config-router)# router-id 1.1.1.3
5530-24TFD(config)# router ospf enable
```

5 Configure OSPF route distribution.

Configure OSPF route distribution to import RIP into OSPF. The Nortel Ethernet Routing Switch 5000 Series distributes the RIP routes as AS-external LSA (LSA type 5), using external metric type 1.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# redistribute rip enable
metric 10 metric-type type1
5530-24TFD(config)# ip ospf apply redistribute rip
```

6 Configure a route policy.

A route policy is required for OSPF to RIP route redistribution. After the route policy is created, apply it to the RIP interface. The following command creates a route policy named **allow** which distributes both direct and OSPF interfaces.

```
5530-24TFD(config)# route-map allow permit 1 enable
match protocol direct,ospf
```

7 Apply the route policy to the RIP Out Policy.

The following commands apply the route policy created in step 6 to RIP interface 10.1.1.41.

```
5530-24TFD(config)# interface vlan 100
5530-24TFD(config-if)# ip rip out-policy allow
```

--End--

The configuration steps described in the above example distributes all OSPF routes to RIP. However, there are times when it can be more advantageous 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:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Step	Action
1	Configure an IP prefix list with a default route. The following command creates an IP prefix list named **default** with an IP address of 0.0.0.0. `5530-24TFD(config)# ip prefix-list default 0.0.0.0/0`

2	Configure a route policy.
	
	Create a route policy named **Policy_Default** which distributes the IP prefix list created in step 1. Note that **ospf** is selected as the **match-protocol** value. This causes the default route to be advertised through RIP only if OSPF is operational.
	
	`5530-24TFD(config)# route-map Policy_Default permit 1 enable match protocol ospf set injectlist default`
	`5530-24TFD(config)# route-map Policy_Default 1 set metric-type type1`

3	Apply the route policy to the RIP Out Policy.
	
	Apply the route policy created in step 2 to RIP interface 10.1.1.41.
	
	`5530-24TFD(config)# interface vlan 100`
	`5530-24TFD(config-if)# ip rip out-policy Policy_Default`

	--End--

Stub area configuration example

In the configuration example illustrated below, the Nortel Ethernet Routing Switch 5000 Series R1 is configured in Stub Area 2, and R2 is configured as a Stub ABR for Area 2.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 36
OSPF stub area example

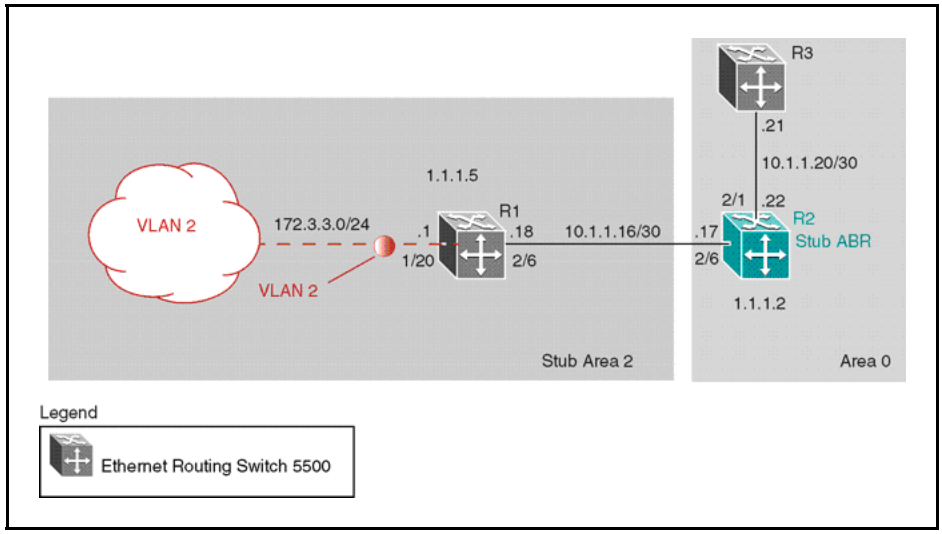

Note: AS-external LSAs are not flooded into a stub area. Instead, only one default route to external destinations is distributed into the stub area by the stub ABR router. The area default cost specifies the cost for advertising the default route into stub area by the ABR.

Use the procedure outlined below to perform the stub area configuration illustrated above:

Note: This example assumes that global IP routing has been enabled on the switch. Global IP routing is enabled on the switch in Global Configuration mode using the **`ip routing`** command.

Step	Action
1	Configure router R1.

a Configure the OSPF interface on R1.

Configure port 2/6 as a brouter port in VLAN 100.

```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(config-if)# brouter port 2/6 vlan 100
subnet 10.1.1.18/30
```

b Configure VLAN 2 on R1.

Create VLAN 2 and assign an IP address to it.

```
5530-24TFD(config)# vlan create 2 type port
5530-24TFD(config)# vlan mem add 2 1/20
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 172.3.3.1
255.255.255.0
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

c Enable OSPF on R1.

Configure R1 in stub area 2 with the Router-ID 1.1.1.5. Add the OSPF interfaces to area 2 and enable OSPF on these interfaces.

```
5530-24TFD(config-router)# router-id 1.1.1.5
5530-24TFD(config-router)# area 0.0.0.2 import
noexternal
5530-24TFD(config-router)# network 10.1.1.18 area
0.0.0.2
5530-24TFD(config-router)# network 172.3.3.1 area
0.0.0.2
5530-24TFD(config)# router ospf enable
```

2 Configure router R2.

a Configure the OSPF interface on R2.

Configure port 2/6 as a brouter port in VLAN 100.

```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(config-if)# brouter port 2/6 vlan 100
subnet 10.1.1.17/30
```

b Configure the second OSPF interface on R2.

Configure port 2/1 as a brouter port in VLAN 300. Enable OSPF on this interface.

```
5530-24TFD(config)# interface fast 2/1
5530-24TFD(config-if)# brouter port 2/1 vlan 300
subnet 10.1.1.22/30
5530-24TFD(config-if)# ip ospf enable
```

c Enable OSPF on R2.

Configure R2 in stub area 2 with an area default cost of 10. Disable import summary to prevent R2 from sending summary LSAs of area 0 into area 2. R2 will originate only summary LSA for default route into area 2. 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 and OSPF interface 10.1.1.17 is added to area 0.0.0.2, R2 automatically becomes a stub ABR.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config-router)# router-id 1.1.1.2
5530-24TFD(config-router)# area 0.0.0.2 import
noexternal
5530-24TFD(config-router)# no area 0.0.0.2
import-summary enable
5530-24TFD(config-router)# area 0.0.0.2 default-cost
10
5530-24TFD(config-router)# network 10.1.1.17 area
0.0.0.2
5530-24TFD(config)# router ospf enable
```

--End--

NSSA configuration example

The NSSA configuration example illustrated below demonstrates a Nortel Ethernet Routing Switch 5000 Series configured as a NSSA ASBR router.

Figure 37
NSSA configuration example

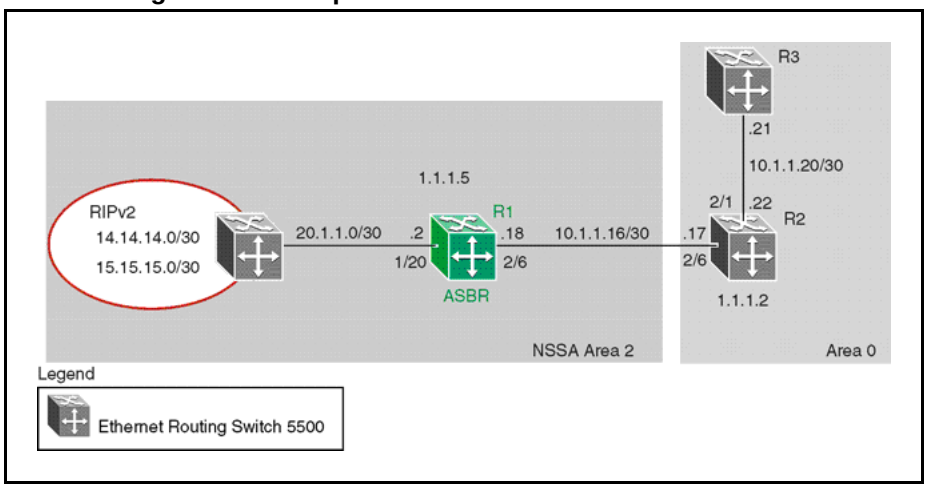

To configure the example illustrated above, follow this procedure:

Step	Action
1	Configure router R1.
	a Configure the RIP interface on R1. Configure port 1/20 as a brouter port in VLAN 100 and enable RIP on this interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config)# interface fast 1/20
5530-24TFD(config-if)# brouter port 1/20 vlan 100
subnet 20.1.1.2/30
5530-24TFD(config)# router rip
5530-24TFD(config-router)# network 20.1.1.2
```

b Enable RIP globally and configure the RIP version 2 interface.

```
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 100
5530-24TFD(config-if)# ip rip receive version rip2
send version rip2
```

c Configure the OSPF interface on R1.

Configure port 2/6 as a brouter port in VLAN 200.

```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(config-if)# brouter port 2/6 vlan 200
subnet 10.1.1.18/30
```

d Enable OSPF on R1.

Configure R1 as an ASBR, assign OSPF Router-ID 1.1.1.5, create OSPF NSSA area 2, add the OSPF interface 10.1.1.18 to area 2, and enable OSPF on the interface.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# as-boundary-router
enable
5530-24TFD(config-router)# router-id 1.1.1.5
5530-24TFD(config-router)# area 0.0.0.2 import nssa
5530-24TFD(config-router)# network 10.1.1.18 area
0.0.0.2
5530-24TFD(config)# router ospf enable
```

e Configure a route policy to distribute Direct and OSPF to RIP.

Create a route policy named **Rip_Dist** that distributes directly connected and OSPF routes into RIP.

```
5530-24TFD(config)# route-map Rip_Dist permit 1
enable match protocol direct,ospf set metric-type
type1
```

f Apply the **Rip_Dist** route policy to RIP Out Policy.

```
5530-24TFD(config)# interface vlan 100
5530-24TFD(config-if)# ip rip out-policy Rip_Dist
```

g Configure OSPF route distribution to distribute RIP routes as AS-external LSA type 1.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# redistribute rip enable
metric-type type1
5530-24TFD(config)# ip ospf apply redistribute rip
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Controlling NSSA external route advertisements

In an OSPF NSSA, the NSSA N/P-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 Nortel Ethernet Routing Switch 5000 Series . When the NSSA N/P-bit is not set true, the ABR drops the external route. A route policy can be created on the Nortel Ethernet Routing Switch 5000 Series to manipulate the N/ p-bit value.

For example, the illustration below shows a RIP network located in NSSA 2. If advertising the 15.15.15.0/24 network to area 0 is the only desired action, perform the following tasks:

- Enable R1 as an OSPF ASBR.
- Create NSSA area 0.0.0.2.
- 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 38
External route advertisement example

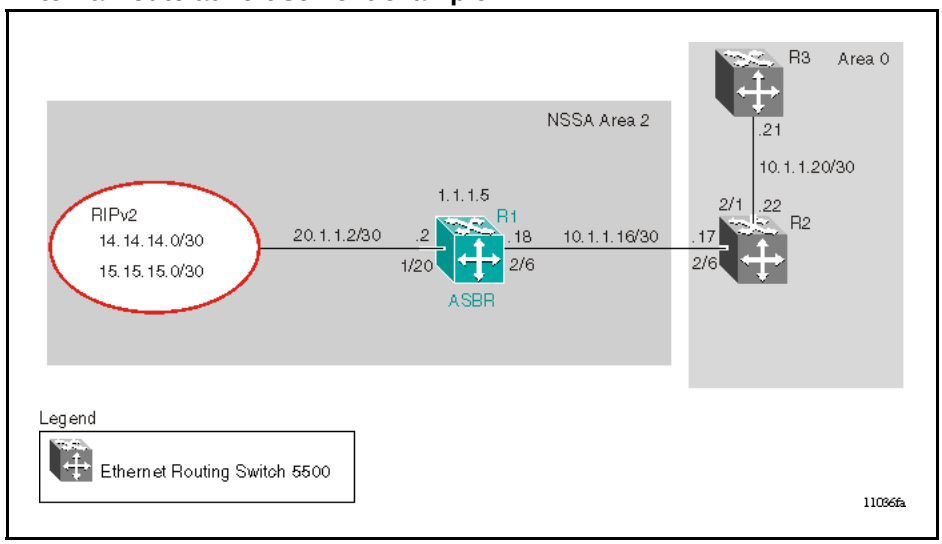

The following procedure describes the commands used to replicate the above configuration example:

Step	Action
1	Configure the RIP interface. Configure port 1/20 as a brouter port in VLAN 200 and enables RIP on this interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config)# interface fast 1/20
5530-24TFD(config-if)# brouter port 1/20 vlan 200 subnet
20.1.1.2/30
5530-24TFD(config)# router rip
5530-24TFD(config-router)# network 20.1.1.2
```

2 Globally enable RIP and configure a RIP interface for RIP version 2.

```
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 200
5530-24TFD(config-if)# ip rip receive version rip2 send
version rip2
```

3 Configure the OSPF interface.

Configure port 2/6 as a brouter port.

```
5530-24TFD(config)# interface fast 2/6
5530-24TFD(config-if)# brouter port 2/6 vlan 100 subnet
10.1.1.18/30
```

4 Enable OSPF.

Configure R1 as an ASBR, assign the OSPF Router-ID 1.1.1.5, create OSPF NSSA area 2, add the OSPF interface 10.1.1.18 to area 2, and enable OSPF on the interface. Enable ASBR and OSPF globally.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# router-id 1.1.1.5
5530-24TFD(config-router)# as-boundary-router enable
5530-24TFD(config-router)# area 0.0.0.2 import nssa
5530-24TFD(config-router)# network 10.1.1.18 area
0.0.0.2
5530-24TFD(config)# router ospf enable
```

5 Create a route policy named **Rip_Dist** that distributes directly connected and OSPF routes into RIP.

```
5530-24TFD(config)# route-map Rip_Dist permit 1 enable
match protocol direct,ospf set metric-type type1
```

6 Apply route policy to RIP Out Policy.

```
5530-24TFD(config)# interface vlan 200
5530-24TFD(config-if)# ip rip out-policy Rip_Dist
```

7 Add two prefix lists (**15net** and **14net**) that are associated with the network addresses from the RIP version 2 network.

```
5530-24TFD(config)# ip prefix-list 15net 15.15.15.0/24
5530-24TFD(config)# ip prefix-list 14net 14.14.14.0/24
```

8 Create a route policy named **P_bit** that sets the NSSA N/P-bit only for the prefix list named **15net**.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config)# route-map P_bit permit 1 enable
match network 15net set nssa-pbit enable
5530-24TFD(config)# route-map P_bit permit 2 enable
match network 14net
5530-24TFD(config)# no route-map P_bit 2 set nssa-pbit
enable
```

9 Configure OSPF route distribution to distribute RIP routes as AS-external LSA Type 1.

```
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# redistribute rip enable
metric-type type1 route-policy P_bit
5530-24TFD(config)# ip ospf apply redistribute rip
```

--End--

Configuring a multi-area complex

The multi-area complex configuration example described in this section uses five Nortel Ethernet Routing Switch 5000 Series devices (R1 to R5) in a multi-area configuration.

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 a real world topology example with command descriptions.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 39
Multi-area complex example

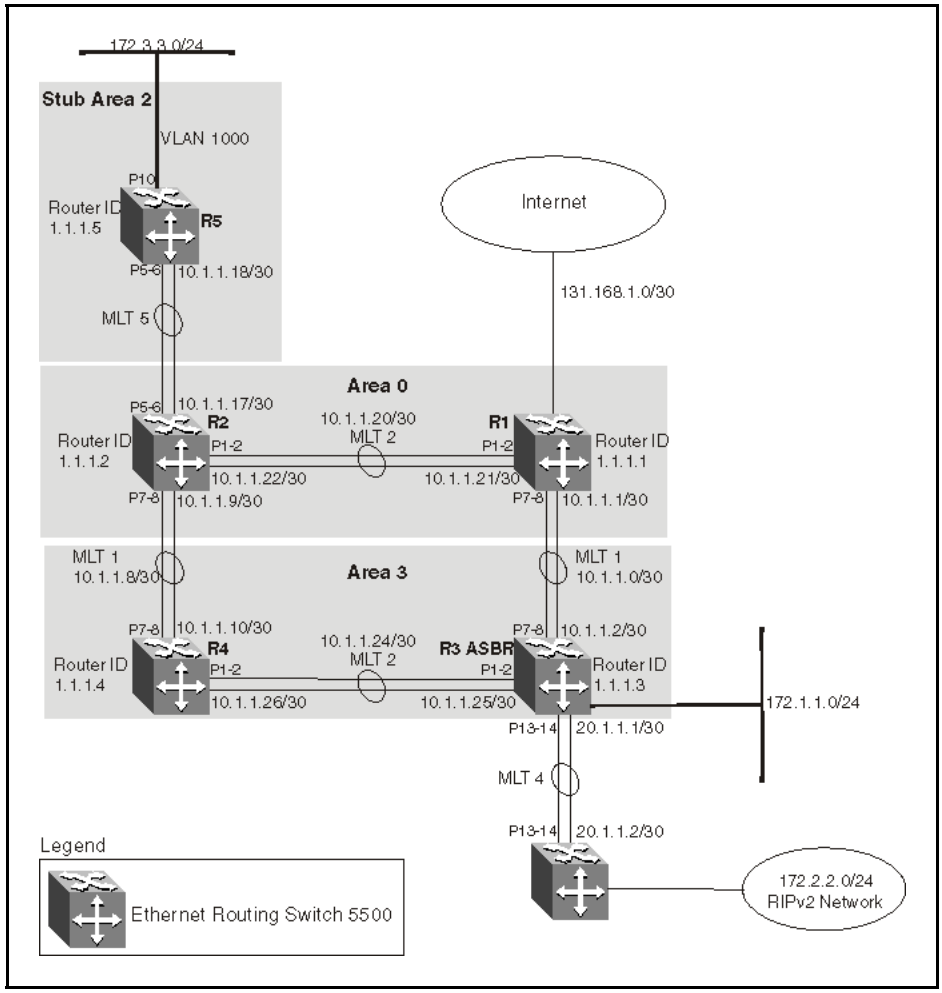

For this configuration example, the Nortel Ethernet Routing Switch 5000 Series devices R1 through R5 are configured as follows:

- R1 is an OSPF ABR that is associated with OSPF Area 0 and 3.
- 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 stub router in Area 2.
- All interfaces used for this configuration are ethernet, therefore the OSPF interfaces are broadcast.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- 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 of 100) and R2 becomes backup designated router (BDR) with a priority value of 50.

Stub and NSSA areas are used to reduce the LSDB size by excluding external LSAs. The stub ABR advertises a default route into the stub area for all external routes.

The following list describes the commands used to create the illustrated configuration. A similar listing could be provided by using the **show running-config** command.

The following commands illustrate the status of the routers in the configuration example. Accompanying each command is the output matching to the configuration example.

R1 configuration commands

```
! *** STP (Phase 1) *** !
spanning-tree stp 2 create
spanning-tree stp 3 create
spanning-tree cost-calc-mode dot1d
spanning-tree port-mode normal
spanning-tree stp 1 priority 8000
spanning-tree stp 1 hello-time 2
spanning-tree stp 1 max-age 20
spanning-tree stp 1 forward-time 15
spanning-tree stp 1 tagged-bpdu disable tagged-bpdu-vid 4001
spanning-tree stp 1 multicast-address 01:80:c2:00:00:00
spanning-tree stp 2 priority 8000
spanning-tree stp 2 hello-time 2
spanning-tree stp 2 max-age 20
spanning-tree stp 2 forward-time 15
spanning-tree stp 2 tagged-bpdu enable tagged-bpdu-vid 4002
spanning-tree stp 2 multicast-address 01:80:c2:00:00:00
spanning-tree stp 3 priority 8000
spanning-tree stp 3 hello-time 2
spanning-tree stp 3 max-age 20
spanning-tree stp 3 forward-time 15
spanning-tree stp 3 tagged-bpdu enable tagged-bpdu-vid 4003
spanning-tree stp 3 multicast-address 01:80:c2:00:00:00
! *** VLAN *** !
vlan configcontrol autopvid
auto-pvid
vlan name 1 "VLAN #1"
vlan create 102 name "VLAN #102" type port
vlan create 103 name "VLAN #103" type port
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
vlan ports 1-24 tagging unTagAll filter-untagged-frame disable
filter-unregistered-frames enable priority 0
vlan ports 25-26 tagging tagAll filter-untagged-frame disable
filter-unregistered-frames enable priority 0
vlan members 1 24-26
vlan members 102 1-2
vlan members 103 7-8
vlan ports 1-2 pvid 102
vlan ports 3-6 pvid 1
vlan ports 7-8 pvid 103
vlan ports 9-26 pvid 1
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 102 snooping disable
vlan igmp 102 proxy disable robust-value 2 query-interval 125
vlan igmp 103 snooping disable
vlan igmp 103 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
no mlt
mlt 1 name "Trunk #1" enable member 7-8 learning normal
mlt 1 learning normal
mlt 1 bpdu all-ports
mlt 1 loadbalance basic
mlt 2 name "Trunk #2" enable member 1-2 learning normal
mlt 2 learning normal
mlt 2 bpdu all-ports
mlt 2 loadbalance basic
! *** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
spanning-tree stp 2 add-vlan 102
spanning-tree stp 3 add-vlan 103
spanning-tree stp 2 enable
spanning-tree stp 3 enable
interface FastEthernet ALL
spanning-tree port 24-26 learning normal
spanning-tree port 1-2 stp 2 learning normal
spanning-tree port 7-8 stp 3 learning normal
spanning-tree port 24-26 cost 1 priority 80
spanning-tree port 1-2 stp 2 cost 1 priority 80
spanning-tree port 7-8 stp 3 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26
enable
exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 1 stp 3 learning normal
mlt spanning-tree 2 stp 2 learning normal
! *** L3 *** !
no ip directed-broadcast enable
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip routing
interface vlan 102
ip address 10.1.1.21 255.255.255.252 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 103
ip address 10.1.1.1 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
ip arp timeout 360
ip dhcp-relay
ip blocking-mode none
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.1
no as-boundary-router enable
no trap enable
timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.3 import external
area 0.0.0.3 import-summaries enable
exit
enable
configure terminal
interface vlan 103
ip ospf area 0.0.0.3
ip ospf network broadcast
ip ospf priority 100
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf transmit-delay 1
ip ospf retransmit-interval 5
ip ospf hello-interval 10
ip ospf dead-interval 40
ip ospf enable
exit
interface vlan 102
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 100
ip ospf authentication-type none
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
exit
```

R2 configuration commands

```
! !  *** STP (Phase 1) *** !
spanning-tree stp 2 create
spanning-tree cost-calc-mode dot1d
spanning-tree port-mode normal
spanning-tree stp 1 priority 8000
spanning-tree stp 1 hello-time 2
spanning-tree stp 1 max-age 20
spanning-tree stp 1 forward-time 15
spanning-tree stp 1 tagged-bpdu disable tagged-bpdu-vid 4001
spanning-tree stp 1 multicast-address 01:80:c2:00:00:00
spanning-tree stp 2 priority 8000
spanning-tree stp 2 hello-time 2
spanning-tree stp 2 max-age 20
spanning-tree stp 2 forward-time 15
spanning-tree stp 2 tagged-bpdu enable tagged-bpdu-vid 4002
spanning-tree stp 2 multicast-address 01:80:c2:00:00:00
! *** VLAN *** !
vlan configcontrol autopvid
auto-pvid
vlan name 1 "VLAN #1"
vlan create 100 name "VLAN #100" type port
vlan create 101 name "VLAN #101" type port
vlan create 102 name "VLAN #102" type port
vlan ports 1-2 tagging tagAll filter-untagged-frame disable filter-unregister
ed-frames enable priority 0
vlan ports 3-6 tagging unTagAll filter-untagged-frame disable
filter-unregistered-frames enable priority 0
vlan ports 7-8 tagging tagAll filter-untagged-frame disable filter-unregister
ed-frames enable priority 0
vlan ports 9-26 tagging unTagAll filter-untagged-frame disable
filter-unregistered-frames enable priority 0
vlan members 1 1-26
vlan members 100 5-6
vlan members 101 7-8
vlan members 102 1-2
vlan ports 1-2 pvid 102
vlan ports 3-4 pvid 1
vlan ports 5-6 pvid 100
vlan ports 7-8 pvid 101
vlan ports 9-26 pvid 1
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 100 snooping disable
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
vlan igmp 100 proxy disable robust-value 2 query-interval 125
vlan igmp 101 snooping disable
vlan igmp 101 proxy disable robust-value 2 query-interval 125
vlan igmp 102 snooping disable
vlan igmp 102 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
no mlt
mlt 1 name "Trunk #1" enable member 7-8 learning normal
mlt 1 learning normal
mlt 1 bpdu all-ports
mlt 1 loadbalance basic
mlt 2 name "Trunk #2" enable member 1-2 learning normal
mlt 2 learning normal
mlt 2 bpdu all-ports
mlt 2 loadbalance basic
mlt 5 name "Trunk #5" enable member 5-6 learning normal
mlt 5 learning normal
mlt 5 bpdu all-ports
mlt 5 loadbalance basic
! *** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
spanning-tree stp 1 add-vlan 100
spanning-tree stp 2 add-vlan 101
spanning-tree stp 2 add-vlan 102
spanning-tree stp 2 enable
interface FastEthernet ALL
spanning-tree port 1-26 learning normal
spanning-tree port 1-2,7-8 stp 2 learning normal
spanning-tree port 1-26 cost 1 priority 80
spanning-tree port 1-2,7-8 stp 2 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26 enable
exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 1 stp 1 learning normal
mlt spanning-tree 1 stp 2 learning normal
mlt spanning-tree 2 stp 1 learning normal
mlt spanning-tree 2 stp 2 learning normal
mlt spanning-tree 5 stp 1 learning normal
! *** L3 *** !
no ip directed-broadcast enable
ip routing
interface vlan 1
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 100
ip address 10.1.1.17 255.255.255.252 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 101
ip address 10.1.1.9 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 102
ip address 10.1.1.22 255.255.255.252 4
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
ip arp timeout 360
ip dhcp-relay
ip blocking-mode none
! *** ECMP *** !
maximum-path 1 rip
maximum-path 1 ospf
maximum-path 1
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.2
no as-boundary-router enable
no trap enable
timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.2 import noexternal
default-cost 1
area 0.0.0.2 import-summaries enable
area 0.0.0.3 import external
area 0.0.0.3 import-summaries enable
exit
enable
configure terminal
interface vlan 101
ip ospf area 0.0.0.3
ip ospf network broadcast
ip ospf priority 50
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf transmit-delay 1
ip ospf retransmit-interval 5
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip ospf hello-interval 10
ip ospf dead-interval 40
ip ospf enable
exit
interface vlan 100
ip ospf area 0.0.0.2
ip ospf network broadcast
ip ospf priority 50
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
exit
interface vlan 102
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 50
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
exit
interface vlan 1
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
no ip ospf enable
exit
```

R3 configuration commands

```
! *** STP (Phase 1) *** !
spanning-tree stp 3 create
spanning-tree cost-calc-mode dot1d
spanning-tree port-mode normal
spanning-tree stp 1 priority 8000
spanning-tree stp 1 hello-time 2
spanning-tree stp 1 max-age 20
spanning-tree stp 1 forward-time 15
spanning-tree stp 1 tagged-bpdu disable tagged-bpdu-vid 4001
spanning-tree stp 1 multicast-address 01:80:c2:00:00:00
spanning-tree stp 3 priority 8000
spanning-tree stp 3 hello-time 2
spanning-tree stp 3 max-age 20
spanning-tree stp 3 forward-time 15
spanning-tree stp 3 tagged-bpdu enable tagged-bpdu-vid 4003
spanning-tree stp 3 multicast-address 01:80:c2:00:00:00
! *** VLAN *** !
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
vlan configcontrol automatic
auto-pvid
vlan name 1 "VLAN #1"
vlan create 103 name "VLAN #103" type port
vlan create 104 name "VLAN #104" type port
vlan create 105 name "VLAN #105" type port
vlan create 1001 name "VLAN #1001" type port
vlan ports 1-2 tagging tagAll filter-untagged-frame disable filter-unregister
ed-frames enable priority 0
vlan ports 3-6 tagging unTagAll filter-untagged-frame disable
filter-unregistered-frames enable priority 0
vlan ports 7-8 tagging tagAll filter-untagged-frame disable filter-unregister
ed-frames enable priority 0
vlan ports 9-26 tagging unTagAll filter-untagged-frame disable
filter-unregistered-frames enable priority 0
vlan members 1 4-6,9,12,15-26
vlan members 103 7-8
vlan members 104 1-2
vlan members 105 13-14
vlan members 1001 10
vlan ports 1-2 pvid 104
vlan ports 3-6 pvid 1
vlan ports 7-8 pvid 103
vlan ports 9 pvid 1
vlan ports 10 pvid 1001
vlan ports 11-12 pvid 1
vlan ports 13-14 pvid 105
vlan ports 15-26 pvid 1
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 103 snooping disable
vlan igmp 103 proxy disable robust-value 2 query-interval 125
vlan igmp 104 snooping disable
vlan igmp 104 proxy disable robust-value 2 query-interval 125
vlan igmp 105 snooping disable
vlan igmp 105 proxy disable robust-value 2 query-interval 125
vlan igmp 1001 snooping disable
vlan igmp 1001 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
no mlt
mlt 1 name "Trunk #1" enable member 7-8 learning normal
mlt 1 learning normal
mlt 1 bpdu all-ports
mlt 1 loadbalance basic
mlt 2 name "Trunk #2" enable member 1-2 learning normal
mlt 2 learning normal
mlt 2 bpdu all-ports
mlt 2 loadbalance basic
mlt 4 name "Trunk #4" enable member 13-14 learning normal
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
mlt 4 learning normal
mlt 4 bpdu all-ports
mlt 4 loadbalance basic
! *** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
spanning-tree stp 3 add-vlan 103
spanning-tree stp 3 add-vlan 104
spanning-tree stp 1 add-vlan 105
spanning-tree stp 1 add-vlan 1001
spanning-tree stp 3 enable
interface FastEthernet ALL
spanning-tree port 4-6,9,12-26 learning normal
spanning-tree port 1-2,7-8 stp 3 learning normal
spanning-tree port 4-6,9,12-26 cost 1 priority 80
spanning-tree port 1-2,7-8 stp 3 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26 enable
exit
interface FastEthernet ALL
spanning-tree port 10 learning disable
exit
interface FastEthernet ALL
exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 1 stp 3 learning normal
mlt spanning-tree 2 stp 3 learning normal
mlt spanning-tree 4 stp 1 learning normal
! *** L3 *** !
no ip directed-broadcast enable
ip routing
interface vlan 1
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 103
ip address 10.1.1.2 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 104
ip address 10.1.1.25 255.255.255.252 4
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 105
ip address 20.1.1.1 255.255.255.0 5
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip dhcp-relay
exit
interface vlan 1001
ip address 172.1.1.1 255.255.255.0 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
ip arp timeout 360
ip dhcp-relay
ip blocking-mode none
! *** Route Policies *** !
route-map Allow permit 1
route-map Allow 1 enable
route-map Allow 1 match protocol direct,ospf
no route-map Allow 1 match interface
route-map Allow 1 match metric 0
no route-map Allow 1 match network
no route-map Allow 1 match next-hop
route-map Allow 1 match route-type any
no route-map Allow 1 match route-source
no route-map Allow 1 set injectlist
route-map Allow 1 set mask 0.0.0.0
route-map Allow 1 set metric 5
route-map Allow 1 set nssa-pbit enable
route-map Allow 1 set ip-preference 0
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.3
as-boundary-router enable
no trap enable
timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.0 import external
area 0.0.0.0 import-summaries enable
area 0.0.0.3 import external
area 0.0.0.3 import-summaries enable
redistribute direct metric 10 metric-type type2 subnets allow
redistribute direct enable
redistribute rip metric 10 metric-type type2 subnets allow
redistribute rip enable
exit
enable
configure terminal
interface vlan 103
ip ospf area 0.0.0.3
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf transmit-delay 1
ip ospf retransmit-interval 5
ip ospf hello-interval 10
ip ospf dead-interval 40
ip ospf enable
exit
interface vlan 104
ip ospf area 0.0.0.3
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
exit
interface vlan 105
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
no ip ospf enable
exit
interface vlan 1001
ip ospf area 0.0.0.3
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
exit
interface vlan 1
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
no ip ospf enable
exit
! *** RIP *** !
router rip
router rip enable
timers basic holddown 120
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
timers basic timeout 180 update 30
default-metric 8
no network 10.1.1.2
no network 10.1.1.25
network 20.1.1.1
no network 172.1.1.1
no network 203.203.100.52
exit
enable
configure terminal
interface vlan 103
no ip rip advertise-when-down enable
no ip rip auto-aggregation enable
no ip rip default-listen enable
no ip rip default-supply enable
ip rip cost 1
ip rip holddown 120
ip rip listen enable
no ip rip poison enable
no ip rip proxy-announce enable
ip rip receive version rip1OrRip2
ip rip send version rip1Comp
ip rip timeout 180
no ip rip triggered enable
ip rip supply enable
exit
interface vlan 104
no ip rip advertise-when-down enable
no ip rip auto-aggregation enable
no ip rip default-listen enable
no ip rip default-supply enable
ip rip cost 1
ip rip holddown 120
ip rip listen enable
no ip rip poison enable
no ip rip proxy-announce enable
ip rip receive version rip1OrRip2
ip rip send version rip1Comp
ip rip timeout 180
no ip rip triggered enable
ip rip supply enable
exit
interface vlan 105
no ip rip advertise-when-down enable
no ip rip auto-aggregation enable
no ip rip default-listen enable
no ip rip default-supply enable
ip rip cost 1
ip rip holddown 120
ip rip listen enable
ip rip out-policy Allow
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
no ip rip poison enable
no ip rip proxy-announce enable
ip rip receive version rip1OrRip2
ip rip send version rip1Comp
ip rip timeout 180
no ip rip triggered enable
ip rip supply enable
exit
interface vlan 1001
no ip rip advertise-when-down enable
no ip rip auto-aggregation enable
no ip rip default-listen enable
no ip rip default-supply enable
ip rip cost 1
ip rip holddown 120
ip rip listen enable
no ip rip poison enable
no ip rip proxy-announce enable
ip rip receive version rip1OrRip2
ip rip send version rip1Comp
ip rip timeout 180
no ip rip triggered enable
ip rip supply enable
exit
interface vlan 1
no ip rip advertise-when-down enable
no ip rip auto-aggregation enable
no ip rip default-listen enable
no ip rip default-supply enable
ip rip cost 1
ip rip holddown 120
ip rip listen enable
no ip rip poison enable
no ip rip proxy-announce enable
ip rip receive version rip1OrRip2
ip rip send version rip1Comp
ip rip timeout 180
no ip rip triggered enable
ip rip supply enable
exit
```

R4 configuration commands

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
! *** STP (Phase 1) *** !
spanning-tree stp 3 create
spanning-tree cost-calc-mode dot1d
spanning-tree port-mode normal
spanning-tree stp 1 priority 8000
spanning-tree stp 1 hello-time 2
spanning-tree stp 1 max-age 20
spanning-tree stp 1 forward-time 15
spanning-tree stp 1 tagged-bpdu disable tagged-bpdu-vid 4001
spanning-tree stp 1 multicast-address 01:80:c2:00:00:00
spanning-tree stp 3 priority 8000
spanning-tree stp 3 hello-time 2
spanning-tree stp 3 max-age 20
spanning-tree stp 3 forward-time 15
spanning-tree stp 3 tagged-bpdu enable tagged-bpdu-vid 4003
spanning-tree stp 3 multicast-address 01:80:c2:00:00:00
! *** VLAN *** !
vlan configcontrol automatic
auto-pvid
vlan name 1 "VLAN #1"
vlan create 101 name "VLAN #101" type port
vlan create 104 name "VLAN #104" type port
vlan ports 1-26 tagging unTagAll filter-untagged-frame disable
filter-unregistered-frames enable priority 0
vlan members 1 3-6,9-26
vlan members 101 7-8
vlan members 104 1-2
vlan ports 1-2 pvid 104
vlan ports 3-6 pvid 1
vlan ports 7-8 pvid 101
vlan ports 9-26 pvid 1
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 101 snooping disable
vlan igmp 101 proxy disable robust-value 2 query-interval 125
vlan igmp 104 snooping disable
vlan igmp 104 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
no mlt
mlt 1 name "Trunk #1" enable member 7-8 learning normal
mlt 1 learning normal
mlt 1 bpdu all-ports
mlt 1 loadbalance basic
mlt 2 name "Trunk #2" enable member 1-2 learning normal
mlt 2 learning normal
mlt 2 bpdu all-ports
mlt 2 loadbalance basic
! *** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
spanning-tree stp 3 add-vlan 101
spanning-tree stp 3 add-vlan 104
spanning-tree stp 3 enable
interface FastEthernet ALL
spanning-tree port 3-6,9-26 learning normal
spanning-tree port 1-2,7-8 stp 3 learning normal
spanning-tree port 3-6,9-26 cost 1 priority 80
spanning-tree port 1-2,7-8 stp 3 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26 enable
exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 1 stp 3 learning normal
mlt spanning-tree 2 stp 3 learning normal
! *** L3 *** !
no ip directed-broadcast enable
ip routing interface vlan 1
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 101
ip address 10.1.1.10 255.255.255.252 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 104
ip address 10.1.1.26 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
ip arp timeout 360
ip dhcp-relay ip blocking-mode none
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.4
no as-boundary-router enable
no trap enable
timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.0 import external
area 0.0.0.0 import-summaries enable
area 0.0.0.3 import external
area 0.0.0.3 import-summaries enable
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
exit
enable
configure terminal
interface vlan 101
ip ospf area 0.0.0.3
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf transmit-delay 1
ip ospf retransmit-interval 5
ip ospf hello-interval 10
ip ospf dead-interval 40
ip ospf enable
exit
interface vlan 104
ip ospf area 0.0.0.3
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
exit
interface vlan 1
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
no ip ospf enable
exit
```

R5 configuration commands

```
! *** STP (Phase 1) *** !
spanning-tree cost-calc-mode dot1d
spanning-tree port-mode normal
spanning-tree stp 1 priority 8000
spanning-tree stp 1 hello-time 2
spanning-tree stp 1 max-age 20
spanning-tree stp 1 forward-time 15
spanning-tree stp 1 tagged-bpdu disable tagged-bpdu-vid 4001
spanning-tree stp 1 multicast-address 01:80:c2:00:00:00
! *** VLAN *** !
vlan configcontrol autopvid
auto-pvid
vlan name 1 "VLAN #1"
vlan create 100 name "VLAN #100" type port
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
vlan create 1000 name "VLAN #1000" type port
vlan ports 1-26 tagging unTagAll filter-untagged-frame disable
filter-unregistered-frames enable priority 0
vlan members 1 24-26
vlan members 100 5-6
vlan members 1000 10
vlan ports 1-4 pvid 1
vlan ports 5-6 pvid 100
vlan ports 7-9 pvid 1
vlan ports 10 pvid 1000
vlan ports 11-26 pvid 1
vlan igmp unknown-mcast-no-flood disable
vlan igmp 1 snooping disable
vlan igmp 1 proxy disable robust-value 2 query-interval 125
vlan igmp 100 snooping disable
vlan igmp 100 proxy disable robust-value 2 query-interval 125
vlan igmp 1000 snooping disable
vlan igmp 1000 proxy disable robust-value 2 query-interval 125
vlan mgmt 1
! *** MLT (Phase 1) *** !
mlt 5 name "Trunk #5" enable member 5-6 learning normal
mlt 5 learning normal
mlt 5 bpdu all-ports
mlt 5 loadbalance basic
! *** STP (Phase 2) *** !
spanning-tree stp 1 add-vlan 1
spanning-tree stp 1 add-vlan 100
spanning-tree stp 1 add-vlan 1000
interface FastEthernet ALL
spanning-tree port 5-6,24-26 learning normal
spanning-tree port 5-6,24-26 cost 1 priority 80
spanning-tree bpdu-filtering port 1-26 timeout 120
no spanning-tree bpdu-filtering port 1-26 enable
exit
interface FastEthernet ALL
spanning-tree port 10 learning disable
exit
! *** MLT (Phase 2) *** !
mlt spanning-tree 5 stp 1 learning normal
! *** L3 *** !
no ip directed-broadcast enable
ip routing interface vlan 1
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
interface vlan 100
ip address 10.1.1.18 255.255.255.252 2
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
exit
interface vlan 1000
ip address 172.3.3.1 255.255.255.252 3
ip dhcp-relay min-sec 0 mode bootp_dhcp
no ip dhcp-relay broadcast
ip dhcp-relay
exit
ip arp timeout 360
ip dhcp-relay
ip blocking-mode none
! *** OSPF *** !
router ospf enable
router ospf
router-id 1.1.1.5
no as-boundary-router enable
no trap enable
timers basic holddown 10
rfc1583-compatibility enable
default-cost ethernet 100
default-cost fast-ethernet 10
default-cost gig-ethernet 1
default-cost ten-gig-ethernet 1
area 0.0.0.0 import external
area 0.0.0.0 import-summaries enable
area 0.0.0.2 import noexternal
default-cost 1
area 0.0.0.2 import-summaries enable
exit
enable
configure terminal
interface vlan 100
ip ospf area 0.0.0.2
ip ospf network broadcast
ip ospf priority 0
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf transmit-delay 1
ip ospf retransmit-interval 5
ip ospf hello-interval 10
ip ospf dead-interval 40
ip ospf enable
exit
interface vlan 1000
ip ospf area 0.0.0.2
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
ip ospf enable
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
exit
interface vlan 1
ip ospf area 0.0.0.0
ip ospf network broadcast
ip ospf priority 1
ip ospf authentication-type none
ip ospf mtu-ignore enable
no ip ospf advertise-when-down enable
no ip ospf enable
exit
```

Router R1 Status

show vlan

```
Id Name Type Protocol User PID Active IVL/SVL Mgmt
--- --------- ---- -------- -------- ------ ------- ----
1 VLAN #1 Port None 0x0000 Yes IVL Yes
Port Members:  1-2,5-7,9-14,16-17,19-26

2 VLAN #2 Port None 0x0000 Yes IVL No
Port Members:  3-4,8,18

5 VLAN #5 Port None 0x0000 Yes IVL No
Port Members:  15

Total VLANs:3
```

show vlan ip

```
================================================================================
Id ifIndex Address Mask MacAddress Offset Routing
================================================================================
Primary Interfaces ------------------------------------------------------------
------------------- 1 10001 10.100.111.200 255.255.255.0 00:11:F9:35:84:40 1
Enabled
2 10002 3.3.3.1 255.255.255.0 00:11:F9:35:84:41 2 Enabled
5 10005 10.10.10.1 255.255.255.0 00:11:F9:35:84:44 5 Enabled
--------------------------------------------------------------------------------
- Secondary Interfaces ---------------------------------------------
----------------------- 2 14096 4.4.4.1 255.255.255.0 00:11:F9:35:84:42 3
Enabled
2 18190 5.5.5.1 255.255.255.0 00:11:F9:35:84:43 4 Enabled
```

show ip ospf

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Router ID: 1.1.1.1
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  2
External Link-State Checksum:  49786(0xc27a)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  427
New Link-State Advertisements Received:  811
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

show ip ospf area

```
Area ID: 0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  35
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  15
Link-State Advertisements Checksum:  551120(0x868d0)
Area ID: 0.0.0.3
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  37
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  1
Link-State Advertisements:  13
Link-State Advertisements Checksum:  454461(0x6ef3d)
```

show ip ospf interface

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Interface:  10.1.1.1
Area ID: 0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  100
Designated Router:  10.1.1.1
Backup Designated Router:  10.1.1.2
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  10.1.1.21
Area ID: 0.0.0.0
Admin State:  Enabled
Type:  Broadcast
Priority:  100
Designated Router:  10.1.1.21
Backup Designated Router:  10.1.1.22
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
```

show ip ospf neighbor

```
Interface Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
--------- ------------ -------------- --- ----- ----------- ----
10.1.1.1 1.1.1.3 10.1.1.2 1 Full 0 Dyn
10.1.1.21 1.1.1.2 10.1.1.22 50 Full 0 Dyn
Total OSPF Neighbors:  2
```

show ip route

```
=============================================================================
Ip Route
=============================================================================
DST MASK NEXT COST VLAN PORT PROT TYPE PRF
-----------------------------------------------------------------------------
172.2.2.0 255.255.255.0 10.1.1.2 10 103 T#1 O IB 120
172.1.1.0 255.255.255.0 10.1.1.2 20 103 T#1 O IB 20
172.3.3.0 255.255.255.252 10.1.1.22 30 102 T#2 O IB 25
20.1.1.0 255.255.255.0 10.1.1.2 10 103 T#1 O IB 120
10.1.1.24 255.255.255.252 10.1.1.2 20 103 T#1 O IB 20
10.1.1.20 255.255.255.252 10.1.1.21 1 102 ---- C DB 0
10.1.1.16 255.255.255.252 10.1.1.22 20 102 T#2 O IB 25
10.1.1.0 255.255.255.252 10.1.1.1 1 103 ---- C DB 0
10.1.1.8 255.255.255.252 10.1.1.2 30 103 T#1 O IB 20
Total Routes:  9
-----------------------------------------------------------------------------
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

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

Router R2 Status

show vlan

```
Id Name Type Protocol User PID Active IVL/SVL Mgmt
--- --------- ---- -------- -------- ------ ------- ----
1 VLAN #1 Port None 0x0000 Yes IVL Yes
Port Members:  1-26
100 VLAN #100 Port None 0x0000 Yes IVL No
Port Members:  5-6
101 VLAN #101 Port None 0x0000 Yes IVL No
Port Members:  7-8
102 VLAN #102 Port None 0x0000 Yes IVL No
Port Members:  1-2
```

show vlan ip

```
Id ifIndex Address Mask MacAddress Offset Routing
1 10001 203.203.100.53 255.255.255.0 00:15:9B:F3:70:40 1 Enabled
100 10100 10.1.1.17 255.255.255.252 00:15:9B:F3:70:41 2 Enabled
101 10101 10.1.1.9 255.255.255.252 00:15:9B:F3:70:42 3 Enabled
102 10102 10.1.1.22 255.255.255.252 00:15:9B:F3:70:43 4 Enabled
```

show ip ospf

```
Router ID: 1.1.1.2
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  True
AS Boundary Router Config Status:  False
External Link-State Advertisements:  2
External Link-State Checksum:  49786(0xc27a)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  99
New Link-State Advertisements Received:  66
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

show ip ospf area

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Area ID: 0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  8
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  15
Link-State Advertisements Checksum:  551120(0x868d0)
Area ID: 0.0.0.2
Import Summaries:  Yes
Import Type:  No External
Intra-Area SPF Runs:  10
Reachable Area Border Routers:  1
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  9
Link-State Advertisements Checksum:  274851(0x431a3)
Stub Metric:  1
Stub Metric Type:  OSPF Metric
Area ID: 0.0.0.3
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  13
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  1
Link-State Advertisements:  13
Link-State Advertisements Checksum:  454461(0x6ef3d)
```

show ip ospf interface

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Interface:  10.1.1.9
Area ID: 0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  50
Designated Router:  10.1.1.9
Backup Designated Router:  10.1.1.10
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  10.1.1.17
Area ID: 0.0.0.2
Admin State:  Enabled
Type:  Broadcast
Priority:  50
Designated Router:  10.1.1.17
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  10.1.1.22
Area ID: 0.0.0.0
Admin State:  Enabled
Type:  Broadcast
Priority:  50
Designated Router:  10.1.1.21
Backup Designated Router:  10.1.1.22
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  203.203.100.53
Area ID: 0.0.0.0
Admin State:  Disabled
Type:  Broadcast
Priority:  1
Designated Router:  0.0.0.0
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
```

show ip ospf neighbor

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Interface Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
--------- ------------ -------------- --- ----- ----------- ----
10.1.1.9 1.1.1.4 10.1.1.10 1 Full 0 Dyn
10.1.1.17 1.1.1.5 10.1.1.18 0 Full 0 Dyn
10.1.1.22 1.1.1.1 10.1.1.21 100 Full 0 Dyn
Total OSPF Neighbors:  3
```

show ip route

```
==========================================================================
Ip Route
==========================================================================
DST MASK NEXT COST VLAN PORT PROT TYPE PRF
--------------------------------------------------------------------------
172.3.3.0 255.255.255.252 10.1.1.18 20 100 T#5 O IB 20
172.2.2.0 255.255.255.0 10.1.1.10 10 101 T#1 O IB 120
172.1.1.0 255.255.255.0 10.1.1.10 30 101 T#1 O IB 20
203.203.100.0 255.255.255.0 203.203.100.53 1 1 ---- C DB 0
20.1.1.0 255.255.255.0 10.1.1.10 10 101 T#1 O IB 120
10.1.1.24 255.255.255.252 10.1.1.10 20 101 T#1 O IB 20
10.1.1.20 255.255.255.252 10.1.1.22 1 102 ---- C DB 0
10.1.1.16 255.255.255.252 10.1.1.17 1 100 ---- C DB 0
10.1.1.8 255.255.255.252 10.1.1.9 1 101 ---- C DB 0
10.1.1.0 255.255.255.252 10.1.1.10 30 101 T#1 O IB 20
Total Routes:  10
--------------------------------------------------------------------------
TYPE Legend:  I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best
Route,E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

Router R3 Status

show vlan

```
Id Name Type Protocol User PID Active IVL/SVL Mgmt
--- --------- ---- -------- -------- ------ ------- ----
1 VLAN #1 Port None 0x0000 Yes IVL Yes
Port Members:  4-6,9,12,15-26
103 VLAN #103 Port None 0x0000 Yes IVL No
Port Members:  7-8
104 VLAN #104 Port None 0x0000 Yes IVL No
Port Members:  1-2
105 VLAN #105 Port None 0x0000 Yes IVL No
Port Members:  13-14
1001 VLAN #1001 Port None 0x0000 Yes IVL No
Port Members:  10
```

show vlan ip

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Id ifIndex Address Mask MacAddress Offset Routing
1 10001 203.203.100.52 255.255.255.0 00:15:9B:F1:FC:40 1 Enabled
103 10103 10.1.1.2 255.255.255.252 00:15:9B:F1:FC:42 3 Enabled
104 10104 10.1.1.25 255.255.255.252 00:15:9B:F1:FC:43 4 Enabled
105 10105 20.1.1.1 255.255.255.0 00:15:9B:F1:FC:44 5 Enabled
1001 11001 172.1.1.1 255.255.255.0 00:15:9B:F1:FC:41 2 Enabled
```

show ip rip

```
Default Import Metric:  8
Domain:
HoldDown Time:  120
Queries:  0 Rip:  Enabled
Route Changes:  1
Timeout Interval:  180
Update Time:  30
```

show ip rip interface

```
IP Address Enable Send Receive Advertise When Down
-------------- ------ ------------- ----------- ------------------
10.1.1.2 false rip1Compatible rip1OrRip2 false
10.1.1.25 false rip1Compatible rip1OrRip2 false
20.1.1.1 true rip1Compatible rip1OrRip2 false
172.1.1.1 false rip1Compatible rip1OrRip2 false
203.203.100.52 false rip1Compatible rip1OrRip2 false

RIP Dflt Dflt Trigger AutoAgg
IP Address Cost Supply Listen Update Enable Supply Listen Poison Proxy
---------- ---- ------ ------ ------- ------ ------ ------ ------ -----
10.1.1.2 1 false false false false true true false false
10.1.1.25 1 false false false false true true false false
20.1.1.1 1 false false false false true true false false
172.1.1.1 1 false false false false true true false false
203.203.100.52 1 false false false false true true false false
IP Address RIP In Policy
-------------- -------------
10.1.1.2
10.1.1.25
20.1.1.1
172.1.1.1
203.203.100.52

IP Address RIP Out Policy
-------------- --------------
10.1.1.2
10.1.1.25
20.1.1.1 Allow
172.1.1.1
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
203.203.100.52

IP Address Holddown Timeout
-------------- -------- -------
10.1.1.2 120 180
10.1.1.25 120 180
20.1.1.1 120 180
172.1.1.1 120 180
203.203.100.52 120 180
```

show route-map detail

```
================================================================
Route Policy ================================================================
====
Name Allow, Id 1, Seq 1 -------------------------------------------------
-------------
Match:
enable :  enable
mode :  permit
match-protocol :  direct,ospf
match-interface :
match-metric :  0
match-network :
match-next-hop :
match-route-type :  any
match-route-src :
Set:
set-injectlist :
set-mask :  0.0.0.0
set-metric :  5
set-metric-type :  type2
set-nssa-pbit :  enable
set-metric-type-internal :  0
set-preference :  0
----------------------------------------------------------------
```

show ip ospf redistribute

```
Source Metric Metric Type Subnet Enabled Route Policy
------ ------ ----------- -------- ------- ------------
Direct 10 Type 2 Allow True
RIP 10 Type 2 Allow True
```

show ip ospf

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Router ID: 1.1.1.3
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  False
AS Boundary Router Config Status:  True
External Link-State Advertisements:  2
External Link-State Checksum:  49786(0xc27a)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  9
New Link-State Advertisements Received:  39
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

show ip ospf area

```
Area ID: 0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  1
Reachable Area Border Routers:  0
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
Area ID: 0.0.0.3
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  4
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  1
Link-State Advertisements:  13
Link-State Advertisements Checksum:  448840(0x6d948)
```

show ip ospf

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Interface:  10.1.1.2
Area ID: 0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  10.1.1.1
Backup Designated Router:  10.1.1.2
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  10.1.1.25
Area ID: 0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  10.1.1.26
Backup Designated Router:  10.1.1.25
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  20.1.1.1
Area ID: 0.0.0.0
Admin State:  Disabled
Type:  Broadcast
Priority:  1
Designated Router:  0.0.0.0
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  172.1.1.1
Area ID: 0.0.0.3
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  172.1.1.1
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  203.203.100.52
Area ID: 0.0.0.0
Admin State:  Disabled
Type:  Broadcast
Priority:  1
Designated Router:  0.0.0.0
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

show ip ospf neighbor

```
Interface Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
--------- ------------ -------------- --- ----- ----------- ----
10.1.1.2 1.1.1.1 10.1.1.1 100 Full 0 Dyn
10.1.1.25 1.1.1.4 10.1.1.26 1 Full 0 Dyn
Total OSPF Neighbors:  2
```

show ip route

```
==============================================================================
Ip Route
==============================================================================
DST MASK NEXT COST VLAN PORT PROT TYPE PRF
------------------------------------------------------------------------------
172.2.2.0 255.255.255.0 20.1.1.2 2 105 T#4 R IB 100
172.3.3.0 255.255.255.252 10.1.1.1 40 103 T#1 O IB 25
172.1.1.0 255.255.255.0 172.1.1.1 1 1001 ---- C DB 0
20.1.1.0 255.255.255.0 20.1.1.1 1 105 ---- C DB 0
10.1.1.16 255.255.255.252 10.1.1.1 30 103 T#1 O IB 25
10.1.1.20 255.255.255.252 10.1.1.1 20 103 T#1 O IB 25
10.1.1.24 255.255.255.252 10.1.1.25 1 104 ---- C DB 0
10.1.1.8 255.255.255.252 10.1.1.26 20 104 T#2 O IB 20
10.1.1.0 255.255.255.252 10.1.1.2 1 103 ---- C DB 0
Total Routes:  9
------------------------------------------------------------------------------
TYPE Legend:  I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best
Route,E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

Router R4 Status

show vlan

```
Id Name Type Protocol User PID Active IVL/SVL Mgmt
--- --------- ---- -------- -------- ------ ------- ----
1 VLAN #1 Port None 0x0000 Yes IVL Yes
Port Members:  3-6,9-26
101 VLAN #101 Port None 0x0000 Yes IVL No
Port Members:  7-8
104 VLAN #104 Port None 0x0000 Yes IVL No
Port Members:  1-2
```

show vlan ip

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Id ifIndex Address Mask MacAddress Offset Routing
1 10001 203.203.100.54 255.255.255.0 00:15:9B:F2:2C:40 1 Enabled
101 10101 10.1.1.10 255.255.255.252 00:15:9B:F2:2C:41 2 Enabled
104 10104 10.1.1.26 255.255.255.252 00:15:9B:F2:2C:42 3 Enabled
```

show ip ospf

```
Router ID: 1.1.1.4
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  False
AS Boundary Router Config Status:  False
External Link-State Advertisements:  2
External Link-State Checksum:  45698(0xb282)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  5
New Link-State Advertisements Received:  34
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

show ip ospf area

```
Area ID: 0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  1
Reachable Area Border Routers:  0
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
Area ID: 0.0.0.3
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  3
Reachable Area Border Routers:  2
Reachable Autonomous System Border Routers:  1
Link-State Advertisements:  13
Link-State Advertisements Checksum:  409758(0x6409e)
```

show ip ospf interface

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Interface: 10.1.1.10
Area ID: 0.0.0.3
Admin State: Enabled
Type: Broadcast
Priority: 1
Designated Router: 10.1.1.9
Backup Designated Router: 10.1.1.10
Authentication Type: None
MTU Ignore: Yes
Advertise When Down: No
Metric Value: 10
Interface: 10.1.1.26
Area ID: 0.0.0.3
Admin State: Enabled
Type: Broadcast
Priority: 1
Designated Router: 10.1.1.25
Backup Designated Router: 10.1.1.26
Authentication Type: None
MTU Ignore: Yes
Advertise When Down: No
Metric Value: 10
Interface: 203.203.100.54
Area ID: 0.0.0.0
Admin State: Disabled
Type: Broadcast
Priority: 1
Designated Router: 0.0.0.0
Backup Designated Router: 0.0.0.0
Authentication Type: None
MTU Ignore: Yes
Advertise When Down: No
Metric Value: 10
```

show ip ospf neighbor

```
Interface Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
--------- ------------- --------------- --- --------- ----------- ----
10.1.1.10 1.1.1.2 10.1.1.9 50 Full 0 Dyn
10.1.1.26 1.1.1.3 10.1.1.25 1 Full 0 Dyn
Total OSPF Neighbors: 2
```

show ip route

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
==========================================================================
Ip Route
==========================================================================
DST MASK NEXT COST VLAN PORT PROT TYPE PRF
--------------------------------------------------------------------------
172.2.2.0 255.255.255.0 10.1.1.25 10 104 T#2 O IB 120
172.3.3.0 255.255.255.252 10.1.1.9 30 101 T#1 O IB 25
172.1.1.0 255.255.255.0 10.1.1.25 20 104 T#2 O IB 20
20.1.1.0 255.255.255.0 10.1.1.25 10 104 T#2 O IB 120
10.1.1.16 255.255.255.252 10.1.1.9 20 101 T#1 O IB 25
10.1.1.20 255.255.255.252 10.1.1.9 20 101 T#1 O IB 25
10.1.1.24 255.255.255.252 10.1.1.26 1 104 ---- C DB 0
10.1.1.8 255.255.255.252 10.1.1.10 1 101 ---- C DB 0
10.1.1.0 255.255.255.252 10.1.1.25 20 104 T#2 O IB 20
Total Routes:  9
--------------------------------------------------------------------------
TYPE Legend:  I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best
Route,E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

Router R5 Status

show vlan

```
Id Name Type Protocol User PID Active IVL/SVL Mgmt
--- ---------- ---- -------- -------- ------ ------- ----
1 VLAN #1 Port None 0x0000 Yes IVL Yes
Port Members:  24-26
100 VLAN #100 Port None 0x0000 Yes IVL No
Port Members:  5-6
1000 VLAN #1000 Port None 0x0000 Yes IVL No
Port Members:  10
```

show vlan ip

```
Id ifIndex Address Mask MacAddress Offset Routing
1 10001 203.203.100.51 255.255.255.0 00:15:9B:F8:1C:40 1 Enabled
100 10100 10.1.1.18 255.255.255.252 00:15:9B:F8:1C:41 2 Enabled
1000 11000 172.3.3.1 255.255.255.252 00:15:9B:F8:1C:42 3 Enabled
```

show ip ospf

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Router ID: 1.1.1.5
Admin Status:  Enabled
Version Number:  2
Area Border Router Oper Status:  False
AS Boundary Router Config Status:  False
External Link-State Advertisements:  0
External Link-State Checksum:  0(0x0)
Type-of-Service (TOS) Routing Supported:  False
Originated Link-State Advertisements:  48
New Link-State Advertisements Received:  387
OSPF Traps:  Disabled
Auto Virtual Link Creation:  Disabled
SPF Hold-Down Time:  10
RFC 1583 Compatibility:  Enabled
```

show ip ospf area

```
Area ID: 0.0.0.0
Import Summaries:  Yes
Import Type:  External
Intra-Area SPF Runs:  3
Reachable Area Border Routers:  0
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  0
Link-State Advertisements Checksum:  0(0x0)
Area ID: 0.0.0.2
Import Summaries:  Yes
Import Type:  No External
Intra-Area SPF Runs:  11
Reachable Area Border Routers:  1
Reachable Autonomous System Border Routers:  0
Link-State Advertisements:  9
Link-State Advertisements Checksum:  274851(0x431a3)
Stub Metric:  1
Stub Metric Type:  OSPF Metric
```

show ip ospf interface

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Interface:  10.1.1.18
Area ID: 0.0.0.2
Admin State:  Enabled
Type:  Broadcast
Priority:  0
Designated Router:  10.1.1.17
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  172.3.3.1
Area ID: 0.0.0.2
Admin State:  Enabled
Type:  Broadcast
Priority:  1
Designated Router:  172.3.3.1
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
Interface:  203.203.100.51
Area ID: 0.0.0.0
Admin State:  Disabled
Type:  Broadcast
Priority:  1
Designated Router:  0.0.0.0
Backup Designated Router:  0.0.0.0
Authentication Type:  None
MTU Ignore:  Yes
Advertise When Down:  No
Metric Value:  10
```

show ip ospf

```
Interface Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
--------- -------------- -------------- --- --------- ----------- ----
10.1.1.18 1.1.1.2 10.1.1.17 50 Full 0 Dyn
Total OSPF Neighbors:  1
```

show ip route

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
==========================================================================
Ip Route
==========================================================================
DST MASK NEXT COST VLAN PORT PROT TYPE PRF
--------------------------------------------------------------------------
172.3.3.0 255.255.255.252 172.3.3.1 1 1000 ---- C DB 0
172.1.1.0 255.255.255.0 10.1.1.17 40 100 T#5 O IB 25
10.1.1.16 255.255.255.252 10.1.1.18 1 100 ---- C DB 0
10.1.1.24 255.255.255.252 10.1.1.17 30 100 T#5 O IB 25
10.1.1.20 255.255.255.252 10.1.1.17 20 100 T#5 O IB 25
10.1.1.8 255.255.255.252 10.1.1.17 20 100 T#5 O IB 25
10.1.1.0 255.255.255.252 10.1.1.17 40 100 T#5 O IB 25
0.0.0.0 0.0.0.0 10.1.1.17 11 100 T#5 O IB 25
Total Routes:  8
--------------------------------------------------------------------------
TYPE Legend:  I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best
Route,E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

Diagnosing neighbor state problems

At initial startup, routers transmit hello packets in an attempt to find other OSPF routers with which form adjacencies. After 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. The following table lists the states a router can go through during the process of forming an adjacency.

Table 11
OSPF neighbor states

Step	State	Description
1	Down	Indicates that a neighbor was configured manually, but the router did 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. The Nortel Ethernet Routing Switch 5000 Series does not support NBMA type.
3	Init	The router 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 (LSDB) is exchanged

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Table 11
OSPF neighbor states (cont'd.)

7	Loading	Indicates the processing state of the LSDB for input into the routing table. The router can request LSA for missing or corrupt routes.
8	Full	Indicates the normal full adjacency state.

OSPF neighbor state information

Neighbor state information can be accessed by using the **show ip ospf neighbor** command.

```
5530-24TFD#show ip ospf neighbor
Interface Nbr Router ID Nbr IP Address Pri State RetransQLen Perm
--------- ------------- ------------- --- ----- ----------- ----
10.1.1.22 1.1.1.1 10.1.1.21 100 Full 0 Dyn
10.1.1.17 1.1.1.5 10.1.1.18 0 Full 0 Dyn
10.1.1.9 1.1.1.4 10.1.1.10 1 Full 0 Dyn
```

Problems with OSPF occur most often 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 can become stuck in an **Init** state and not form adjacencies. There are several possible causes for this problem:

- Authentication mismatch or configuration problem
- Area mismatch for Stub or NSSA
- Area ID mismatch
- Hello Interval or Dead Interval mismatch

To determine any mismatches in OSPF configuration, use the **show ip ospf ifstats mismatch** command.

ExStart/Exchange problems

Even though routers can recognize each other and have moved beyond two way communications, routers can become stuck in the **ExStart/Exchange** state.

A mismatch in maximum transmission unit (MTU) sizes between the routers usually causes this type of problem. For example, one router could be set for a high MTU size and the other router a smaller value. Depending on the size of the link state database, the router with the smaller value may not be able to process the larger packets and thus

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

be stuck in this 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.

Note: The Nortel Ethernet Routing Switch 5000 Series automatically checks for OSPF MTU mismatches.

In the Nortel Ethernet Routing Switch 5000 Series , the supported MTU size for OSPF is 1500 bytes by default. Incoming OSPF database description (DBD) packets are dropped if their MTU size is greater than this value.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

RIP configuration using NNCLI

This section describes how to configure RIP using NNCLI.

RIP is a distance vector protocol used to dynamically discover network routes based on information passed between routers in the network.

Prerequisites

- Enable IP routing globally.
- Assign an IP address to the VLAN or brouter port that you want to enable with RIP.

  Routing is automatically enabled on the VLAN when you assign an IP address to it.

RIP configuration procedures

To configure RIP routing on the Ethernet Routing Switch, perform the following steps:

Step	Action
1	Enable RIP globally.
2	Configure global RIP properties as required.
3	Enable RIP on the desired VLAN or brouter interfaces.
4	Configure interface RIP properties as required.

--End--

RIP configuration navigation

- “Configuring the global RIP status” (page 216)
- “Configuring the RIP global timeout, holddown timer, and update timer” (page 216)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- “Configuring the default RIP metric value” (page 218)
- “Displaying global RIP information” (page 218)
- “Configuring the RIP status on an interface” (page 219)
- “Configuring RIP parameters for an interface ” (page 220)
- “Manually triggering a RIP update” (page 224)
- “Displaying RIP interface configuration” (page 222)

Configuring the global RIP status

Use this procedure to globally enable RIP on the switch.

Procedure steps

Step	Action
1	To configure RIP on the switch, enter the following from the Global Configuration command mode: `[default] [no] router rip enable`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
default	Globally disables RIP on the switch.
no	Globally disables RIP on the switch.

Configuring the RIP global timeout, holddown timer, and update timer

Use this procedure to set the RIP global timeout, holddown timer, and update timer.

Procedure steps

Step	Action
1	To configure the global RIP timers, enter the following from the RIP Router Configuration command mode:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
[default] timers basic
holddown <holdown-timer>
timeout <global-timeout>
update <update-timer>
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Returns the parameters to the factory default timer values: • holddown timer: 120 seconds • global timeout: 180 seconds • update timer: 30 seconds
<holdown-timer>	Specifies the global holddown timer, which is the length of time (in seconds) that RIP maintains a route in the garbage list after determining that it is unreachable. During this period, RIP continues to advertise the garbage route with a metric of infinity (16). If a valid update for a garbage route is received within the holddown period, the router adds the route back into its routing table. If no update is received, the router deletes the garbage list entry. Range is 0–360 seconds. Default is 120 seconds.
<global-timeout>	Specifies the global timeout interval parameter. If a RIP router does not receive an update from another RIP router within the configured timeout period, it moves the routes advertised by the nonupdating router to the garbage list. The timeout interval must be greater than the update timer. Range is 15–259200 seconds. Default is 180 seconds.
<update-timer>	Specifies a value for the RIP update timer, which is the time interval (in seconds) between regular RIP updates. The update timer value must be less than the timeout interval. Range is 0–360 seconds. Default is 30 seconds.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring the default RIP metric value

Use this procedure to configure a default metric to apply to routes not learned through RIP but imported into the RIP domain. The Ethernet Routing Switch applies this default metric to redistributed routes if the associated route policy does not specify a metric for the redistributed protocol, such as OSPF. The range is 0 to 15, and the default is 8.

Procedure steps

Step	Action
1	To configure the default RIP metric value, enter the following from the RIP Router Configuration mode: **[default] default-metric <metric_value>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<metric_value>	Specifies a metric value between 0 and 15.
default	Returns the switch to the factory default RIP default import metric value: 8.

Displaying global RIP information

Use this procedure to display the global RIP configuration.

Procedure steps

Step	Action
1	To display the global RIP configuration, enter the following from the User EXEC mode: **show ip rip**

--End--

Job aid

The following table shows the field descriptions for the **show ip rip** command.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Default Import Metric	Indicates the value of the default import metric.
Domain	Indicates the value inserted into the Routing Domain field of all RIP packets sent on this device. This value is not configurable.
HoldDown Time	Indicates the value of the holddown timer.
Queries	Indicates the number of responses the router has sent in response to RIP queries from other systems.
Rip	Indicates whether RIP is enabled.
Route Changes	Indicates the number of route changes the RIP process has made to the routing database.
Timeout Interval	Indicates the RIP timeout interval.
Update Time	Indicates the value of the RIP update timer.

Configuring the RIP status on an interface

Use this procedure to configure the RIP status on a VLAN interface or brouter port.

Procedure steps

Step	Action
1	To configure RIP status on an interface, enter the following from the Interface Configuration command mode: **[default] [no] ip rip enable** OR enter the following from the Router Configuration command mode: **[no] network <ip_address>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Disables RIP on the interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
[no]	Disables RIP on the IP interface.
<ip_address>	The IP address of the interface to be configured.

Configuring RIP parameters for an interface

Use this procedure to configure RIP parameters for an interface.

Procedure steps

Step	Action
1	To configure RIP parameters for an interface, enter the following from the Interface Configuration mode: **[default] [no] ip rip** **[advertise-when-down enable]** **[auto-aggregation enable]** **[cost <cost>]** **[default-listen enable]** **[default-supply enable]** **[holddown {<holddown> | global}]** **[listen enable]** **[poison enable]** **[proxy-announce enable]** **[receive version {rip1 | rip1orrip2 | rip 2}]** **[send version {rip1 | rip1orrip2 | rip 2}]** **[supply enable]** **[timeout {<timeout> | global}]** **[triggered enable]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
default	Sets the specified parameter to the default value.
no	Removes the specified configuration from the switch.
advertise-when-down enable	Enables RIP advertisements for an interface even when the link to the network fails. The router continues to advertise the subnet even if that particular network is no longer connected (no link in the enabled VLAN). This feature does not advertise the route until the VLAN is first enabled. After the VLAN is enabled, the route is advertised even when the link fails. By default, advertise when down functionality is disabled.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
auto-aggregation enable	Enables auto aggregation on the RIP interface. After you enable auto aggregation, the Ethernet Routing Switch automatically aggregates routes to their natural net mask when they are advertised on an interface in a network of a different class. Automatic route aggregation can be enabled only in RIP2 mode or RIP1 compatibility mode. By default, auto aggregation is disabled.
cost <cost>	Specifies the RIP cost (metric) for this interface in a range from 1 to 15. The default cost is 1.
default-listen enable	Enables the interface to accept default routes learned through RIP updates. The default setting is disabled.
default-supply enable	Enables the interface to send default route information in RIP updates. This setting takes effect only if a default route exists in the routing table. The default setting is disabled.
holddown	Specifies the interface holddown timer, which is the length of time (in seconds) that RIP maintains a route in the garbage list after determining that it is unreachable. During this period, RIP continues to advertise the garbage route with a metric of infinity (16). If a valid update for a garbage route is received within the holddown period, the router adds the route back into its routing table. If no update is received, the router deletes the garbage list entry. If the interface timer is configured, this setting overrides the global parameter and does not change if the global parameter is modified. Range is 0–360 seconds. The default value is set by the global holddown parameter, which has a default of 120 seconds.
listen enable	Enables this interface to listen for RIP advertisements. The default value is enabled.
poison enable	Specifies whether RIP routes on the interface learned from a neighbor are advertised back to the neighbor. If poison reverse is disabled, split horizon is invoked and 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 which a route is learned are "poisoned" with a metric of 16. The receiving neighbor ignores this route because the metric 16 indicates infinite hops in the network. By default, poison reverse is disabled.
proxy-announce enable	Enables proxy announcements on a RIP interface. When proxy announcements are enabled, the source of a route and its next hop are treated as the same when processing received updates. So, instead of the advertising router being used as the source, the next hop is. Proxy announcements are disabled by default.
receive {rip1 | rip1orrip2 | rip 2}	Specifies the RIP version received on this interface. Default is rip1orrip2.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
send {notsend | rip1 | rip1comp | rip 2}	Specifies the RIP version sent on an interface. Default is rip1compatible.
supply enable	Enables RIP route advertisements on this interface. The default value is enabled.
timeout <timeout>	Specifies the RIP timeout value on this interface. If a RIP interface does not receive an update from another RIP router within the configured timeout period, it moves the routes advertised by the nonupdating router to the garbage list. The timeout interval must be greater than the update timer. Range is 15–259200 seconds. The default value is set by the global timeout parameter, which has a default of 180 seconds. The interface timer setting overrides the global parameter and does not change if the global parameter is changed.
triggered enable	Enables automatic triggered updates on this RIP interface. Default is disabled.

Displaying RIP interface configuration

Use this procedure to display configuration for a RIP interface.

Procedure steps

Step	Action
1	To display RIP interface configuration, enter the following from the User EXEC mode: **show ip rip** **[interface** **[<vid>]** **[FastEthernet [<portlist>]]** **[VLAN [<vid>]]]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[interface]	Displays RIP statistics by interface. Omission of this key word displays general RIP information
[<vid>]	Displays RIP information for the specified VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
[FastEthernet [<portlist>]]	Displays RIP information for the specified ports. If no ports are specified, all port information is displayed.
[VLAN [<vid>]]	Displays RIP information for the specified VLAN. If no VLAN ID is specified, all VLAN information is displayed.

Job aid

The following table shows the field descriptions for the **`show ip rip interface`** command.

Field	Description
unit/port	Indicates the unit and port of the RIP interface.
IP Address	Indicates the IP address of the RIP interface.
Enable	Indicates whether RIP is enabled or disabled on the interface.
Send	Indicates which send mode is enabled.
Receive	Indicates which receive mode is enabled.
Advertise When Down	Indicates whether the advertise when down feature is enabled.
RIP Cost	Indicates the RIP cost (metric) for this interface.
Dflt Supply	Indicates whether the interface sends the default route in RIP updates, if a default route exists in the routing table.
Dflt Listen	Indicates whether the interface listens for default routes in RIP updates.
Trigger Update	Indicates whether triggered updates are enabled.
AutoAgg Enable	Indicates whether auto aggregation is enabled.
Supply	Indicates whether the interface is enabled to supply updates for RIP.
Listen	Indicates whether the interface is enabled to listen for RIP routes.
Poison	Indicates whether RIP routes on the interface learned from a neighbor are advertised back to the neighbor.
Proxy	Indicates whether proxy announcements are enabled.
RIP IN Policy	Indicates the RIP policy for inbound filtering on the interface.
RIP Out Policy	Indicates the RIP policy for outbound filtering on the interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Holddown	Indicates the value of the RIP holddown timer for the interface.
Timeout	Indicate the RIP timeout interval for the interface.

Manually triggering a RIP update

Use this procedure to manually trigger a RIP update on an interface.

Procedure steps

Step	Action
1	To manually trigger a RIP update, enter the following from the Privileged EXEC command mode: **manualtrigger ip rip interface vlan <vid>**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

RIP configuration examples using NNCLI

This section provides configuration examples for common IP routing tasks using NNCLI.

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

RIP is configured on a VLAN or brouter port basis.

> ***Note:*** In many of the following configuration examples, a brouter port is used to create a connection to the network core. This practice does not imply that a brouter port is the only means through which a core connection can be established. The use of a brouter port is only one of many ways to create such a connection.

Navigation

- "RIP configuration tasks" (page 225)
- "Configuring RIP" (page 227)
- "Configuring RIP version 2" (page 230)
- "Using RIP accept policies" (page 232)
- "Using RIP announce policies" (page 234)

RIP configuration tasks

To perform a basic RIP configuration on a VLAN, perform the following steps.

Step	Action
1	Configure the interface, assign an IP address and add ports.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD# enable
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 51 name "VLAN-51" type
port
5530-24TFD(config)# interface vlan 51
5530-24TFD(config-if)# ip address 10.10.1.1
255.255.255.0
5530-24TFD(config-if)# exit
5530-24TFD(config)# vlan members add 51 8-9
```

2 Enable RIP using one of the following command sequences.

```
5530-24TFD(config)# interface vlan 51
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# exit
```

OR

```
5530-24TFD(config)# router rip
5530-24TFD(config-router)# network 10.10.1.1
5530-24TFD(config-router)# exit
```

3 Select the VLAN to configure RIP interface properties.

```
5530-24TFD(config)# interface vlan 51
```

4 Disable Supply RIP Updates on the VLAN, if required.

```
5530-24TFD(config-if)# ip rip supply disable
```

5 Disable Listen for RIP Updates on the VLAN, if required.

```
5530-24TFD(config-if)# ip rip listen disable
```

6 Enable Default Route Supply on the VLAN, if a default route exists in the route table.

```
5530-24TFD(config-if)# ip rip default-supply enable
```

7 Enable Default Route Listen on the VLAN to add a default route to the route table, if advertised from another router.

```
5530-24TFD(config-if)# ip rip default-listen enable
```

8 Add the Out Route Policy to the VLAN (this step assumes that you have previously configured the route policy).

```
5530-24TFD(config-if)# ip rip out-policy map1
```

9 Enable Triggered Updates on the VLAN, if required.

```
5530-24TFD(config-if)# ip rip triggered enable
```

10 Configure the cost of the VLAN link by entering a value of 1 to 15; where 1 is the default.

```
5530-24TFD(config-if)# ip rip cost 2
```

11 Configure send mode parameters on the VLAN.

```
5530-24TFD(config-if)# ip rip send version rip2
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

12 Configure receive mode parameters on the VLAN.

```
5530-24TFD(config-if)# ip rip receive version rip2
```

13 Enable poison reverse on the VLAN.

```
5530-24TFD(config-if)# ip rip poison enable
```

--End--

Configuring RIP

This section describes the set up of a basic RIP configuration between two Nortel Ethernet Routing Switch 5000 Series routers. As shown in the following diagram, router ERS2 is configured between router ERS1 and the edge of the network core. Two VLANs (VLAN 2 and 3) are associated with ERS1.

Figure 40
RIP configuration example

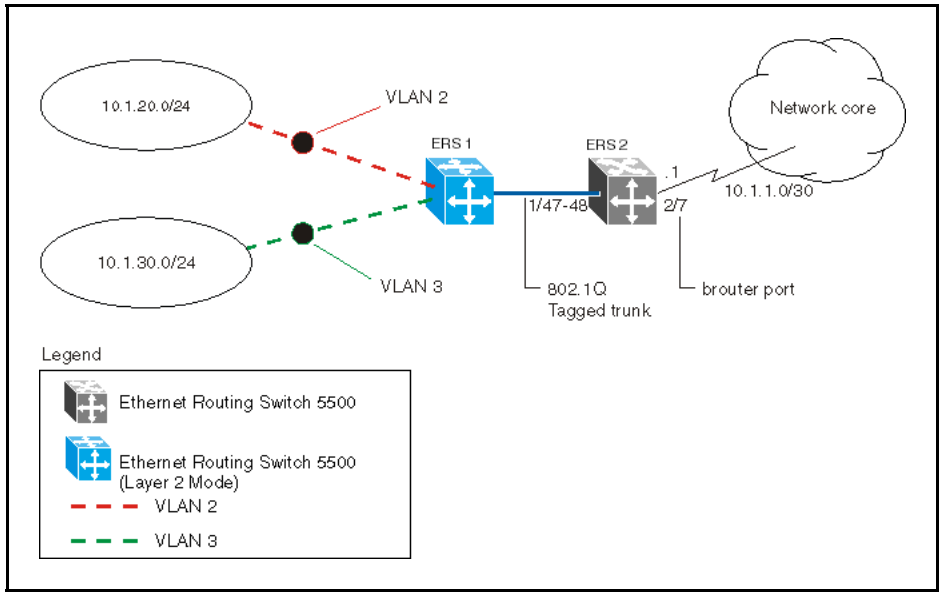

For the purposes of this example:

- ERS1 is an edge switch with two configured VLANs, VLAN 2 and 3. It is connected to aggregation switch ERS2 on ports 1/47 and 1/48.
- Port 2/7 of ERS2 is configured as a brouter port with RIP to connect to the network core.

Use the following procedure to configure router ERS 2 and reproduce the illustrated RIP configuration:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Step	Action

1 Configure tagging on ports 1/47 and 1/48. Tagging is required to support multiple VLANs on the same interface.

```
5530-24TFD# enable
5530-24TFD# config terminal
5530-24TFD(config)# vlan ports 1/47-48 tagging tagAll
```

2 Configure ERS2 for VLAN 2 access.

a Create a port-based VLAN (VLAN 2) using spanning tree group 1 and include ports 1/47 and 1/48 in VLAN 2.

```
5530-24TFD(config)# vlan create 2 name "VLAN-2" type
port
5530-24TFD(config)# vlan member add 2 port 1/47-48
```

b Assign the IP address 10.1.20.2/24 to VLAN 2.

```
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.2
255.255.255.0
```

c Enable RIP for VLAN 2 and disable RIP supply and listen. RIP supply and listen are not required because no router is attached to VLAN 2.

```
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip supply disable
5530-24TFD(config-if)# ip rip listen disable
```

3 Configure ERS2 for VLAN 3 access.

a Create a port-based VLAN (VLAN 3) using spanning tree group 1 and include ports 1/47 and 1/48 in VLAN3.

```
5530-24TFD(config)# vlan create 3 name "VLAN-3" type
port
5530-24TFD(config)# vlan member add 3 port 1/47-48
```

b Assign the IP address 10.1.30.2/24 to VLAN 3.

```
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip address 10.1.30.2
255.255.255.0
```

c Enable RIP for VLAN 3 and disable RIP supply and listen. RIP supply and listen are not required because no router is attached to VLAN 3.

```
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip supply disable
5530-24TFD(config-if)# ip rip listen disable
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

4 Configure brouter port 2/7 on ERS2.

a Assign the IP address 10.1.1.1/30 to port 2/7 using brouter VLAN 2090.

```
5530-24TFD(config)# interface FastEthernet 2/7
5530-24TFD(config-if)# brouter vlan 2090 subnet
10.1.1.1/30
```

Note: Usage of the **brouter** command above requires the usage of Variable Length Subnetting. Usage of a dotted decimal subnet mask is not allowed.

b Enable RIP on the interface.

```
5530-24TFD(config)# interface FastEthernet 2/7
5530-24TFD(config-if)# ip rip enable
```

5 Enable IP routing and RIP globally.

```
5530-24TFD(config)# ip routing
5530-24TFD(config)# router rip enable
```

--End--

A list of the commands used to create this configuration can be displayed using the **show running-config** command. Using this command on ERS2 would list the following commands:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
! *** VLAN *** !
vlan igmp unknown-mcast-no-flood disable
vlan configcontrol strict
auto-pvid
vlan name 1 "VLAN #1"
vlan create 2 name "VLAN-2" type port
vlan create 3 name "VLAN-3" type port
vlan members 2 1/47-48
vlan members 3 1/47-48
! *** RIP *** !
router rip
router rip enable
timers basic holddown 120
timers basic timeout 180 update 30 default-metric 8
network 10.1.20.2
network 10.1.30.2
network 10.1.1.1
interface vlan 2
no ip rip listen enable
no ip rip supply enable
interface vlan 3
no ip rip listen enable
no ip rip supply enable
! *** Brouter Port *** !
interface fastEthernet ALL
brouter port 2/7 vlan 3 subnet 10.1.1.1/30
```

The following commands can be used to confirm the configuration of RIP parameters:

Command	Description
`show vlan`	This command is used to display information about the currently configured switch VLANs.
`show vlan ip`	This command is used to display IP address information about VLANs that have been assigned addresses on the switch.
`show ip rip`	This command displays information on the global switch RIP configuration.
`show ip route`	This command displays the switch routing table.
`show ip rip interface`	This command displays information about the RIP interfaces present on the switch.

Configuring RIP version 2

When RIP is enabled on an interface, it operates by default in **rip1compatible** send mode and **rip1orRip2** receive mode. Depending on configuration requirements, the Nortel Ethernet Routing Switch 5000 Series can be configured to operate using RIP version 1 or 2. The

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

configuration illustrated below demonstrates a Nortel Ethernet Routing Switch 5000 Series switch that has been configured to operate use RIP version 2 only.

Note: This example builds on the previous RIP configuration.

Figure 41
RIPv2 configuration example

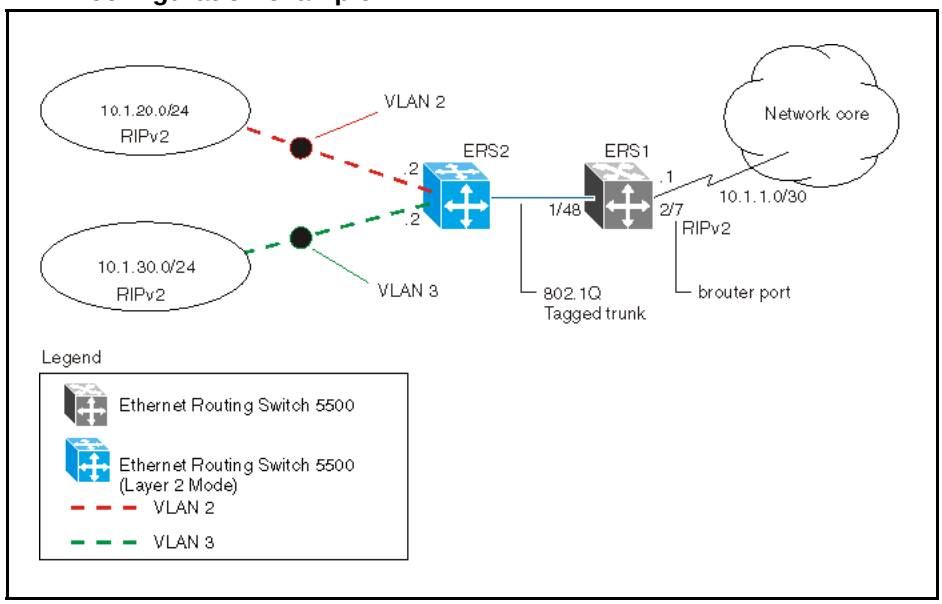

Use the following procedure to configure ERS2 to add RIP version 2 to VLAN 2, VLAN 3, and the brouter port:

Step	Action
1	Configure RIP version 2 on VLAN 2. Enable RIP version 2 mode on the IP address used for VLAN 2.

```
5530-24TFD# enable
5530-24TFD# config terminal
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip rip send version rip2
5530-24TFD(config-if)# ip rip receive version rip2
```

2	Configure RIP version 2 on VLAN 3. Enable RIP version 2 mode on the IP address used for VLAN 3.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD# enable
5530-24TFD# config terminal
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip rip send version rip2
5530-24TFD(config)# router rip enable
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip rip receive version rip2
```

3 Configure RIP version 2 on the brouter port. Enable RIP version 2 mode on the IP address used for the brouter port.

```
5530-24TFD(config)# interface FastEthernet 2/7
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip send version rip2
5530-24TFD(config-if)# ip rip receive version rip2
```

--End--

Using RIP accept policies

RIP accept policies are used on the Nortel Ethernet Routing Switch 5000 Series to selectively accept routes from RIP updates. If no policies are defined, the default behavior is applied. This default behavior is to add all learned routes to the route table. RIP accept policies are 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 illustrated below, the Nortel Ethernet Routing Switch 5000 Series (ERS1) is configured with a RIP accept policy. This creates a single route directed to ERS3 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 ERS1.

A summary route is calculated 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

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 42
Accept policy configuration

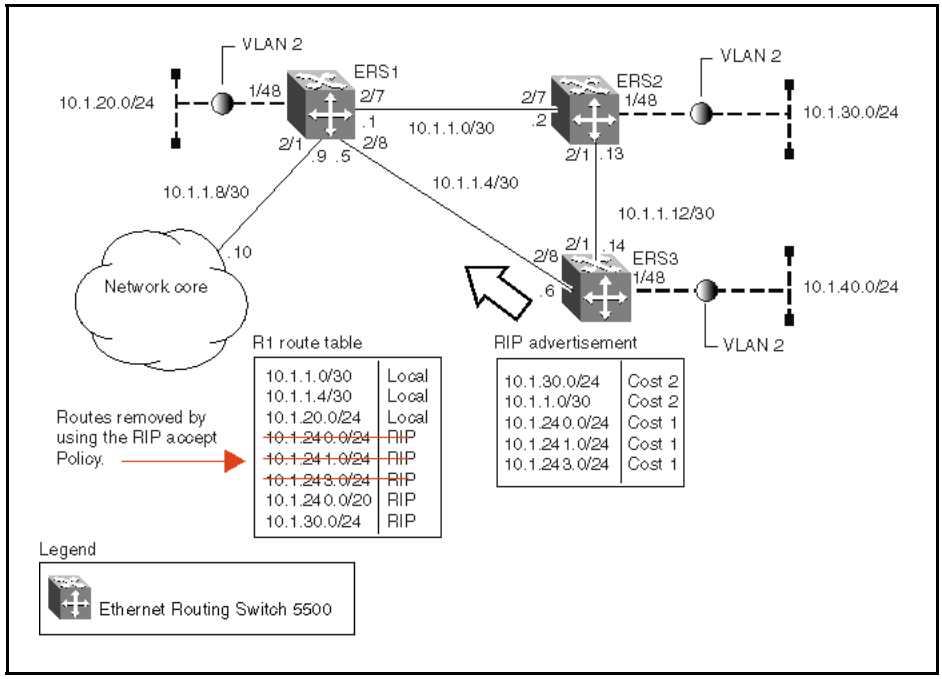

Use the following steps to recreate the above configuration example:

Step	Action
1	Configure the IP prefix list on ERS1. Create a prefix list named **Prefix_1** with an IP range from 10.1.240.0 to 10.1.255.0. `5530-24TFD(config)# ip prefix-list Prefix_1 10.1.240.0/20 ge 20 le 32`

2	Configure the route policy named **rip_pol_1** with match criteria using the IP prefix configured in step 1. This injects one route of 10.1.240.0/20 into the route table.
	
	`5530-24TFD(config)# route-map rip_pol_1 1`
	`5530-24TFD(config)# route-map rip_pol_1 1 enable`
	`5530-24TFD(config)# route-map rip_pol_1 permit 1 enable`
	`5530-24TFD(config)# route-map rip_pol_1 permit 1 match network Prefix_1`
	`5530-24TFD(config)# route-map rip_pol_1 permit 1 set-injectlist Prefix_1`

3	Add the route policy created in step 2 to both RIP core ports.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config)# interface FastEthernet 2/7
5530-24TFD(config-if)# brouter vlan 2090 subnet
10.1.1.1/30
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip in-policy rip_pol_1

5530-24TFD(config)# interface FastEthernet 2/8
5530-24TFD(config-if)# brouter vlan 2091 subnet
10.1.1.5/30
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip in-policy rip_pol_1
```

--End--

The **show running-config** command is used to display the current configuration of a switch. Using this command on the above configuration would yield the following results:

```
rip_pol_1
! *** Route Policies *** !
ip prefix-list Prefix_1 10.1.240.0/20 ge 20 le 32
route-map rip_pol_1
route-map rip_pol_1 1 enable
no route-map rip_pol_1 1 match interface
route-map rip_pol_1 1 match metric 0
route-map rip_pol_1 1 match network Prefix_1
no route-map rip_pol_1 1 match next-hop
route-map rip_pol_1 1 match route-type any
no route-map rip_pol_1 match route-source
route-map rip_pol_1 1 set injectlist Prefix_1
route-map rip_pol_1 set mask 0.0.0.0
route-map rip_pol_1 set metric 0
route-map rip_pol_1 set nssa-pbit enable
route-map rip_pol_1 set ip-preference 0
! *** Brouter Port *** !
interface fastEthernet ALL
brouter port 2/7
vlan 2090 subnet 10.1.1.1/30
ip rip in-policy rip_pol_1
brouter port 2/8 vlan 2091 subnet 10.1.1.5/30
ip rip in-policy rip_pol_1
```

Using RIP announce policies

In the previous configuration example, a RIP accept policy is used on ERS1 to insert a single route into its route table for all networks from ERS3. Instead of using an accept policy on ERS1, a RIP announce policy on ERS3 could be used to announce a single route to both ERS1 and ERS2 for the local network range.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

To configure the RIP announce policy on ERS3, use the following configuration steps:

Step	Action
1	Configure the IP prefix list on ERS3 named **Prefix_1** with the IP address 10.1.240.0. `5530-24TFD(config)# ip prefix-list Prefix_1 10.1.240.0/20 ge 20 le 32`

2	Configure the route policy named **Policy_Rip** with match criteria using the IP prefix configured in step 1.
	`5530-24TFD(config)# route-map rip_pol_1 1`
	`5530-24TFD(config)# route-map rip_pol_1 1 enable`
	`5530-24TFD(config)# route-map rip_pol_1 permit 1 enable`
	`5530-24TFD(config)# route-map rip_pol_1 permit 1 set-injectlist Prefix_1`

3	Add the route policy created in step 2 to both RIP core ports.
	`5530-24TFD(config)# interface FastEthernet 2/1`
	`5530-24TFD(config-if)# brouter vlan 2091 subnet 10.1.1.14/30`
	`5530-24TFD(config-if)# ip rip enable`
	`5530-24TFD(config-if)# ip rip out-policy rip_pol_1`
	
	`5530-24TFD(config)# interface FastEthernet 2/8`
	`5530-24TFD(config-if)# brouter vlan 2090 subnet 10.1.1.6/30`
	`5530-24TFD(config-if)# ip rip enable`
	`5530-24TFD(config-if)# ip rip out-policy rip_pol_1`

	--End--

To limit the advertising of routes using the announce policy from the routing table, a route policy should be created to deny the route. To configure the RIP announce policy with a limited announce policy on ERS3, use the following configuration steps:

Step	Action
1	Configure the IP prefix list named **Prefix_2** with the IP address 10.1.240.0. `5530-24TFD(config)# ip prefix-list Prefix_2 10.1.240.0/20 ge 20 le 20`
2	Configure the IP route policy named **rip_pol_2** with match criteria using the IP prefix configured in Step 1.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD(config)# route-map rip_pol_2 deny 1 enable
match network Prefix_2
5530-24TFD(config)# route-map rip_pol_2 1 match network
Prefix_2
```

3 Add the Route Policy created in step 2 to both RIP core ports.

```
5530-24TFD(config)# interface FastEthernet 2/1
5530-24TFD(config-if)# brouter vlan 2091 subnet
10.1.1.14/30
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)#ip rip out-policy rip_pol_2
5530-24TFD(config)# interface FastEthernet 2/8
5530-24TFD(config-if)# brouter vlan 2090 subnet
10.1.1.6/30
5530-24TFD(config-if)# ip rip enable
5530-24TFD(config-if)# ip rip out-policy rip_pol_2
```

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

ECMP configuration using NNCLI

This section describes the procedures you can use to configure ECMP using NNCLI.

The Equal Cost MultiPath (ECMP) feature allows routers to determine 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.

ATTENTION
ECMP is not supported on the Nortel Ethernet Routing Switch 5510. ECMP works in a mixed stack but cannot run on any Nortel Ethernet Routing Switch 5510 units in the stack.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Configure routing (RIP, OSPF, or static routes) on the switch.

ECMP configuration navigation

- "Configuring the number of ECMP paths allotted for RIP" (page 237)
- "Configuring the number of ECMP paths allotted for OSPF" (page 238)
- "Configuring the number of ECMP paths allotted for static routes" (page 239)
- "Displaying ECMP path information" (page 239)

Configuring the number of ECMP paths allotted for RIP

Use this procedure to configure the number of ECMP paths allotted for the Routing Information Protocol (RIP).

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To configure the number of ECMP paths allotted for RIP, enter the following from the Global Configuration mode: `[default] [no] rip maximum-path <path_count>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Sets the maximum ECMP paths allowed to the default value, 1.
[no]	Sets the maximum ECMP paths allowed to the default value, 1.
<path_count>	Represents the number of ECMP paths allotted. This is a value between 1 and 4. The default is 1.

Configuring the number of ECMP paths allotted for OSPF

Use this procedure to configure the number of ECMP paths allotted for the Open Shortest Path First (OSPF) protocol.

Procedure steps

Step	Action
1	To configure the number of ECMP paths allotted for OSPF, enter the following from the Global Configuration mode: `[default] [no] ospf maximum-path <path_count>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
default	Sets the maximum ECMP paths allowed to the default value, 1.
[no]	Sets the maximum ECMP paths allowed to the default value, 1.
<path_count>	Represents the number of ECMP paths allotted. This is a value between 1 and 4. The default is 1.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring the number of ECMP paths allotted for static routes

Use this procedure to configure the number of ECMP paths allotted to static routes.

Procedure steps

Step	Action
1	To configure the number of ECMP paths allotted to static routes enter the following from the Global Configuration mode: **[default] [no] maximum-path <path_count>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
default	Sets the maximum ECMP paths allowed to the default value, 1.
no	Sets the maximum ECMP paths allowed to the default value, 1.
<path_count>	Represents the number of ECMP paths allotted. This is a value between 1 and 4. The default is 1.

Displaying ECMP path information

Use this procedure to display ECMP path information.

Procedure steps

Step	Action
1	To display ECMP path information, enter the following from the User Exec mode: **show ecmp**

--End--

Job aid

The following table shows the field descriptions for the **show ecmp** command.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Protocol	Indicates the protocol.
MAX-PATH	Indicates the maximum number of equal-cost paths supported for the listed protocol.

ECMP configuration examples

Equal Cost Multipath (ECMP) is an IP feature for load-balancing routed IP traffic across up to four equal-cost paths for each supported protocol. ECMP supports OSPF, RIP, and static routes. Some benefits of using ECMP:

- Supported protocols will rerun when an ECMP path fails, and the other configured paths will automatically take the load.
- Load sharing implies better use of network facilities.

ECMP is selected based on the source and destination IP address in the packet. The *hash_control* register has a *HASH_SELECT* field which is set to 5 (lower CRC-32).

R1 = CRC32 (SIP, DIP)

R2 = R1 & 0x1F(The Least Significant 5 bits are selected)

ecmp_index = R2 % (ecmp_count + 1)

> ***Note:*** The value *ecmp_count* above is zero-based in the hardware so if four paths are present then the value is three. This is why the value is *ecmp_count + 1*.

The ECMP traffic distribution algorithm is demonstrated in the following example:

Consider two network devices, Device 1 at the IP address 192.1.1.3 and Device 2 at 192.1.1.4. Device 1 send to Device 2 so that 192.1.1.3 is the source IP address (SIP) and 192.1.1.4 is the destination IP address (Device 2).

To calculate the CRC32 for the example source and destination IP address noted above, the following calculations would be made:

- CRC32 polynomial : x ^ 32 + x ^ 28 + x ^ 23 + x ^ 22 + x ^ 16 + x ^ 12 + x ^ 11 + x ^ 10 + x ^ 8 + x ^ 7 + x ^ 5 + x ^ 4 + x ^ 2 + x ^ 1 + 1
- R1 = CRC32 (0xc0010103, 0xc0010104) = 0xf474b549
- R2 = (0xf474b549 & 0x1f) = 9

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

If, for the purposes of this example, it is assumed that the ECMP count is 4 (hardware entries 0 though 3), the following calculation is then made:

- ecmp_index = 9 % (4+1) = 1

This means that in this example, the second path at hardware index 1 in the ECMP table will be used.

In the configuration example below, the following command would enable two OSPF ECMP paths on router R1:

```
5530-24TFD(config)#ospf maximum-path 2
```

Figure 43
ECMP configuration example

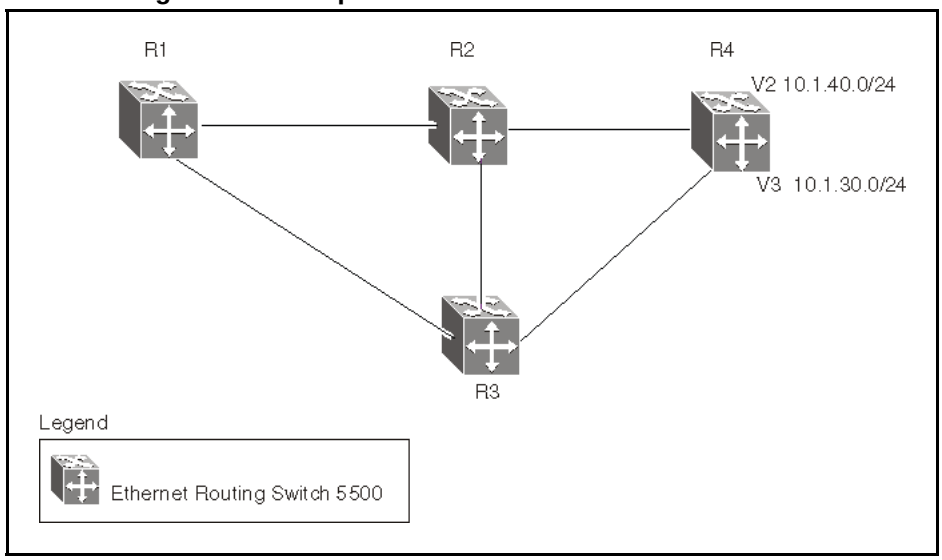

Use the following commands to enable ECMP on each of the supported protocols:

- **OSPF**
  ospf maximum-path <path_count>
- **RIP**
  rip maximum-path <path_count>
- **Static Routes**
  maximum-path <path_count>

In all commands above, the **<path_count>** parameter represents the number of ECMP paths allotted. This is a value between 1 and 4. The default is 1.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying the IP routing table

After ECMP configuration is complete, verify the ECMP paths in the routing table using the **show ip route** command. The following example displays the output for this command:

```
==========================================================================
Ip Route
==========================================================================
DST MASK NEXT COST VLAN PORT PROT TYPE PRF
--------------------------------------------------------------------------
0.0.0.0 0.0.0.0 10.100.111.1 10 1 19 S IB 5
3.3.3.0 255.255.255.0 3.3.3.1 1 2 - C DB 0
4.4.4.0 255.255.255.0 4.4.4.1 1 2 - C DB 0
5.5.5.0 255.255.255.0 5.5.5.1 1 2 - C DB 0
10.10.10.0 255.255.255.0 10.10.10.1 1 5 - C DB 0
10.100.111.0 255.255.255.0 10.100.111.200 1 1 - C DB 0
Total Routes:  6
--------------------------------------------------------------------------
TYPE Legend:  I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best
Route,E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

Paths shown with the letter **E** in the **TYPE** column are designated equal-cost paths. In this example, two routes to IP address 10.1.40.0 and two routes to IP address 10.1.30.0 are displayed.

Displaying global ECMP configuration

To confirm global ECMP configuration, use the show ecmp command. A sample output from this command is displayed below:

```
5530-24TFD# show ecmp
Protocol MAX-PATH
-------- --------
static:  1
rip:  2
ospf:  4
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Route policies configuration using NNCLI

This section describes the procedures you can use to configure route policies using NNCLI.

Using standard routing schemes, packets are forwarded based on routes that have been learned by the router through routing protocols such as RIP and OSPF or through the introduction of static routes. Route policies provide the ability to forward packets based on rule sets created by the network administrator. These rule sets, or policies, are then applied to the learned or static routes.

Route policies configuration procedures

To configure routing policies, perform the following steps:

Step	Action
1	Create the appropriate prefix lists.
2	Assign those prefix lists to route maps.
3	Apply the route maps to the appropriate type of policy.

--End--

Route policies configuration navigation

- “Configuring prefix lists” (page 244)
- “Configuring route maps” (page 244)
- “Applying a RIP accept (in) policy” (page 247)
- “Applying a RIP announce (out) policy ” (page 248)
- “Configuring an OSPF accept policy” (page 249)
- “Applying the OSPF accept policy” (page 250)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- "Displaying the OSPF accept policy" (page 250)
- "Configuring an OSPF redistribution policy" (page 251)
- "Applying the OSPF redistribution policy" (page 252)
- "Displaying the OSPF redistribution policy" (page 252)

Configuring prefix lists

Use this procedure configure up to four prefix lists for use in route policies.

Procedure steps

Step	Action
1	To configure a prefix list, enter the following from the Global Configuration mode: `[no] ip prefix-list <prefix_name> {<ip_address/mask> [ge <mask_from>] [le <mask_to>]} [name <new_prefix_name>]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes a prefix list or a prefix from a list.
<prefix_name>	Specifies the name assigned to the prefix list.
<ip_address/mask>	Specifies the IP address and subnet mask of the prefix list. The subnet mask is expressed as a value between 0 and 32.
ge <mask_from>	Specifies the lower bound of the mask length. This value, when combined with the higher bound mask length (**le**), specifies a subnet range covered by the prefix list.
le <mask_to>	Specifies the higher bound of the mask length. This value, when combined with the lower bound mask length (**ge**), specifies a subnet range covered by the prefix list.
name <new_prefix_name>	Assigns a new name to previously configured prefix list.

Configuring route maps

Use this procedure to define route maps used in the configuration of route policies.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To configure a route map, enter the following from the Global Configuration mode: **[no] route-map <map_name> [permit | deny] <sequence_number> [enable] [match {interface <prefix_list> | metric <metric_value> | network <prefix_list> | next-hop <prefix_list> | protocol <protocol_name> | route-source <prefix_list> | route-type <route_type>}] [name <new_map_name>] [set {injectlist <prefix_list> | ip-preference <pref> | mask <ip_address> | metric <metric_value> | metric-type <metric_type> | nssa-pbit enable}]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
no	Removes the specified route map.
<map_name>	Specifies the name associated with this route map.
[permit | deny]	Specifies the action to be taken when this policy is selected for a specific route. A value of **permit** indicates that the route is used while **deny** indicates that the route is ignored.
<sequence_number>	Specifies the secondary index value assigned to individual policies inside a larger policy group.
enable	Specifies whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
[match {interface <prefix_list> | metric <metric_value> | network <prefix_list> | next-hop <prefix_list> | protocol <protocol_name> | route-source <prefix_list> | route-type <route_type>}]	If configured, the switch matches the specified criterion: • interface <prefix_list>: matches the IP address of the received interface against the contents of the specified prefix list. • metric <metric_value>: matches the metric of the incoming advertisement or existing route against the specified value, an integer value from 0 to 65535. If 0, then this field is ignored. The default is 0. • network <prefix_list>: matches the destination network against the contents of the specified prefix list. • next-hop <prefix_list>: matches the next hop IP address of the route against the contents of the specified prefix list. • protocol <protocol_name>: matches the protocol through which a route is learned. Options are direct, static, rip, ospf, and any. Multiple protocols can be specified by using a comma-separated list. • route-source <prefix_list>: matches the source IP address for RIP routes against the contents of the specified prefix list. • route-type <route_type>: Specifies the route type to be matched. Options are any, external, external-1, external-2, internal, and local.
[name <new_map_name>]	Specifies a new name to be assigned to a previously configured route map.
[set {injectlist <prefix_list> | ip-preference <pref> | mask <mask> | metric <metric_value> | metric-type <metric_type> | nssa-pbit enable}]	If configured, the switch sets the specified parameter: • injectlist <prefix_list>: replaces the destination network of the route that matches this policy with the contents of the specified prefix list. • ip-preference <pref>: specifies the route preference value to be assigned to the route that matches this policy. Valid range is 0–255. If 0 (the default value), the global preference value is used. Used for accept policies only. • mask <mask>: sets the mask of the route that matches this policy. Used for RIP accept policies only. • metric <metric_value>: sets the value of the metric to be assigned to matching routes. This is an integer value between 0 and 65535. • metric-type <metric_type>: sets the metric type for routes to be imported into the OSPF routing protocol. Options are type1 and type2. • nssa-pbit enable: enables the NSSA N/P-bit, which notifies the ABR to export the matching external route. Used for OSPF policies only.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying route maps

Use this procedure to display configured route maps.

Procedure steps

Step	Action
1	To display route maps, enter the following from the Global Configuration mode: **show route-map [detail] <map_name>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[detail]	Provides detailed information on the route maps.
<map_name>	Specifies the name of the route map to display.

Applying a RIP accept (in) policy

Use this procedure to specify a RIP Accept (In) policy for an interface. This policy takes the form of a previously configured route map. Only one policy can be created for each RIP interface.

Procedure steps

Step	Action
1	To specify a RIP Accept policy for an interface, enter the following from the Interface Configuration mode: **[default] [no] ip rip in-policy <rmap_name>**
2	To display RIP interface configuration, enter the following from the User EXEC mode: **show ip rip** **[interface** **[<vid>]** **[FastEthernet [<portlist>]]** **[VLAN [<vid>]]]**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Removes the in policy associated with this interface.
[no]	Removes the in policy associated with this interface.
<rmap_name>	Applies the previously configured route map as the RIP accept policy.
[interface]	Displays RIP statistics by interface. Omission of this key word displays general RIP information
[<vid>]	Displays RIP information for the specified VLAN.
[FastEthernet [<portlist>]]	Displays RIP information for the specified ports. If no ports are specified, all port information is displayed.
[VLAN [<vid>]]	Displays RIP information for the specified VLAN. If no VLAN ID is specified, all VLAN information is displayed.

Applying a RIP announce (out) policy

Use this procedure to specify a RIP Announce (Out) policy for an interface. This policy takes the form of a previously configured route map. Only one policy can be created for each RIP interface.

Procedure steps

Step	Action
1	To apply a RIP Announce (Out) policy to an interface, enter the following from the Interface Configuration mode: **[default] [no] ip rip out-policy <rmap_name>**
2	To display RIP interface configuration, enter the following from the User EXEC mode: **show ip rip** **[interface** **[<vid>]** **[FastEthernet [<portlist>]]** **[VLAN [<vid>]]]**

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
default	Removes the out policy associated with this interface.
no	Removes the out policy associated with this interface.
<rmap_name>	Applies the previously configured route map as the RIP announce policy.
[interface]	Displays RIP statistics by interface. Omission of this key word displays general RIP information
[<vid>]	Displays RIP information for the specified VLAN.
[FastEthernet [<portlist>]]	Displays RIP information for the specified ports. If no ports are specified, all port information is displayed.
[VLAN [<vid>]]	Displays RIP information for the specified VLAN. If no VLAN ID is specified, all VLAN information is displayed.

Configuring an OSPF accept policy

Use this procedure to configure the router to accept advertisements from another router in the system. The referenced policy takes the form of a previously configured route map.

Accept policies are only applied to Type 5 External routes based on the advertising router ID. There can only be one OSPF accept policy on the switch and the policy is applied before updates are added to the routing table from the link state database.

Procedure steps

Step	Action
1	To configure the OSPF accept-advertisements router policy, enter the following from the OSPF Router Configuration mode: `[no] accept adv-rtr <router_ip_address> [enable] [metric-type {any | type1 | type2}] [route-policy <rmap_name>]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Configures the router to not accept advertisements from another router in the system.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
router_ip_address	Represents the IP address of the router from which advertisements are to be accepted. The value *0.0.0.0* denotes that advertisements from all routers are accepted.
enable	Enables the accept entry for the router specified in the **<ip_address>** parameter.
metric-type {any | type1 | type2}	Indicates the type of OSPF external routes that will be accepted from this router.
route-policy <rmap_name >	Specifies the name of a previously configured route map to be used for filtering external routes advertised by the specified advertising router before accepting them into the routing table.

Applying the OSPF accept policy

Use this procedure to apply the configured OSPF accept policy to the switch.

Procedure steps

Step	Action
1	To apply the OSPF accept policy to the switch, enter the following from the Global Configuration mode: **ip ospf apply accept**

--End--

Displaying the OSPF accept policy

Use this procedure to display the OSPF accept policy.

Procedure steps

Step	Action
1	To display the OSPF accept policy, enter the following from the Global Configuration mode: **show ip ospf accept**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
[detail]	Provides detailed information on the route maps.
<map_name>	Specifies the name of the route map to display.

Configuring an OSPF redistribution policy

Use this procedure to configure OSPF route redistribution. Redistribution of direct, RIP, and static routes is currently supported.

OSPF redistribution policies send redistributed routes as Type 5 External routes. There can be only one OSPF redistribution policy on the switch. The OSPF accept policy takes precedence over the redistribution policy.

Procedure steps

Step	Action
1	To configure OSPF route redistribution, enter the following from the Router Configuration mode: **`redistribute <route_type> [enable] [route-policy <rmap_name>] [metric <metric_value>] [metric-type <metric_type>] [subnets <subnet_setting>`**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables an OSPF route policy or OSPF route redistribution completely.
<route_type>	Specifies the source protocol to be redistributed. Valid options are **direct**, **rip**, and **static**.
<rmap_name>	Specifies the route policy to associate with route redistribution. This is the name of a previously configured route map.
<metric_value>	Specifies the metric value to associate with the route redistribution. This is an integer value between 0 and 65535.
<metric_type>	Specifies the metric type to associate with the route redistribution. Valid options are **type1** and **type2**.
<subnet_setting>	Specifies the subnet advertisement setting of this route redistribution. This determines whether individual subnets are advertised. Valid options are **allow** and **suppress**.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Applying the OSPF redistribution policy

Use this procedure to apply the configured OSPF route redistribution policy to the switch.

Procedure steps

Step	Action
1	To apply route redistribution to OSPF, enter the following from the Global Configuration mode: `ip ospf apply redistribute {direct | rip | static}`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
direct	Applies only direct OSPF redistribution configuration to the switch
rip	Applies only RIP OSPF redistribution configuration on the switch.
static	Applies only static OSPF redistribution configuration on the switch.

Displaying the OSPF redistribution policy

Use this procedure to display the OSPF redistribution policy configuration and status.

Procedure steps

Step	Action
1	To display the OSPF redistribution policy, enter the following from the Global Configuration mode: `show ip ospf redistribute`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

DHCP relay configuration using NNCLI

This chapter describes the procedures you can use to configure DHCP relay using the NNCLI.

ATTENTION
DHCP relay uses a hardware resource that is shared by switch Quality of Service applications. When DHCP relay is enabled globally, the Quality of Service filter manager will not be able to use precedence 11 for configurations. For the filter manager to be able to use this resource, DHCP relay must be disabled for the entire unit or stack.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be set as the DHCP relay agent.
- Ensure that a route to the destination DHCP server is available on the switch.

DHCP relay configuration procedures

To configure DHCP relay, perform the following steps:

Step	Action
1	Ensure that DHCP relay is enabled globally. (DHCP relay is enabled by default.)
2	Configure the DHCP relay forwarding path, specifying the VLAN IP as the DHCP relay agent and the remote DHCP server as the destination.
3	Enable DHCP for the specific VLAN.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

DHCP relay configuration navigation

- "Configuring global DHCP relay status" (page 254)
- "Displaying the global DHCP relay status" (page 254)
- "Specifying a local DHCP relay agent and remote DHCP server" (page 255)
- "Displaying the DHCP relay configuration" (page 256)
- "Configuring DHCP relay status and parameters on a VLAN" (page 256)
- "Displaying the DHCP relay configuration for a VLAN" (page 257)
- "Displaying DHCP relay counters" (page 258)
- "Clearing DHCP relay counters for a VLAN" (page 259)

Configuring global DHCP relay status

Use this procedure to configure the global DHCP relay status. DHCP relay is enabled by default.

Procedure steps

Step	Action
1	To configure the global DHCP relay status, enter the following from the Global Configuration mode: `[no] ip dhcp-relay`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables DHCP relay.

Displaying the global DHCP relay status

Use this procedure to display the current DHCP relay status for the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To display the global DHCP relay status, enter the following from the User EXEC command mode: **show ip dhcp-relay**

--End--

Specifying a local DHCP relay agent and remote DHCP server

Use this procedure to specify a VLAN as a DHCP relay agent on the forwarding path to a remote DHCP server. The DHCP relay agent can forward DHCP client requests from the local network to the DHCP server in the remote network.

The DHCP relay feature is enabled by default, and the default mode is BootP-DHCP.

Prerequisites

- Enable IP routing and configure an IP address on the VLAN to configure as a DHCP relay agent.

Procedure steps

Step	Action
1	To configure a VLAN as a DHCP relay agent, enter the following from the Global Configuration mode: **[no] ip dhcp-relay fwd-path <relay-agent-ip> <DHCP-server> [enable] [disable] [mode {bootp | bootp-dhcp | dhcp}]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the specified DHCP forwarding path.
<relay-agent-ip>	Specifies the IP address of the VLAN that serves as the local DHCP relay agent.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<DHCP-server>	Specifies the address of the remote DHCP server to which DHCP packets are to be relayed.
[enable]	Enables the specified DHCP relay forwarding path.
[disable]	Disables the specified DHCP relay forwarding path.
[mode {bootp | bootp-dhcp | dhcp}]	Specifies the mode for DHCP relay. • BootP only • BootP and DHCP • DHCP only If you do not specify a mode, the default DHCP and BootP is used.

Displaying the DHCP relay configuration

Use this procedure to display the current DHCP relay agent configuration.

Procedure steps

Step	Action
1	To display the DHCP relay configuration, enter the following from the User EXEC command mode: **`show ip dhcp-relay fwd-path`**

--End--

Job aid

The following table shows the field descriptions for the **`show ip dhcp-relay fwd-path`** command.

Field	Description
INTERFACE	Specifies the interface IP address of the DHCP relay agent.
SERVER	Specifies the IP address of the DHCP server.
ENABLE	Specifies whether DHCP is enabled.
MODE	Specifies the DHCP mode.

Configuring DHCP relay status and parameters on a VLAN

Use this procedure to configure the DHCP relay parameters on a VLAN. To enable DHCP relay on the VLAN, enter the command with no optional parameters.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To configure DHCP relay on a VLAN, enter the following from the VLAN Interface Configuration mode: **[no] ip dhcp-relay [broadcast] [min-sec <min-sec>] [mode {bootp | dhcp | bootp_dhcp}]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables DHCP relay on the specified VLAN.
[broadcast]	Enables the broadcast of DHCP reply packets to the DHCP clients on this VLAN interface.
min-sec <min-sec>	The switch immediately forwards a BootP/DHCP packet if the 'secs' field in the BootP/DHCP packet header is greater than the configured min-sec value; otherwise, the packet is dropped. Range is 0-65535. The default is 0.
mode {bootp | dhcp | bootp_dhcp}	Specifies the type of DHCP packets this VLAN supports: • bootp - Supports BootP only • dhcp - Supports DHCP only • bootp_dhcp - Supports both BootP and DHCP

Displaying the DHCP relay configuration for a VLAN

Use this procedure to display the current DHCP relay parameters configured for a VLAN.

Procedure steps

Step	Action
1	To display the DHCP relay VLAN parameters, enter the following from the Privileged EXEC command mode: **show vlan dhcp-relay [<vid>]**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
[<vid>]	Specifies the VLAN ID of the VLAN to be displayed. Range is 1-4094.

Job aid

The following table shows the field descriptions for the **show ip dhcp-relay** command.

Field	Description
IfIndex	Indicates the VLAN interface index.
MIN_SEC	Indicates the minimum time, in seconds, to wait between receiving a DHCP packet and forwarding the DHCP packet to the destination device. A value of zero indicates forwarding is done immediately without delay.
ENABLED	Indicates whether DHCP relay is enabled on the VLAN.
MODE	Indicates the type of DHCP packets this interface supports. Options include none, BootP, DHCP, and both.
ALWAYS_BROADCAST	Indicates whether DHCP reply packets are broadcast to the DHCP client on this VLAN interface.

Displaying DHCP relay counters

Use this procedure to display the current DHCP relay counters. This includes the number of requests and the number of replies.

Procedure steps

Step	Action
1	To display the DHCP relay counters, enter the following from the User EXEC command mode: **show ip dhcp-relay counters**

--End--

Job aid

The following table shows the field descriptions for the **show ip dhcp-relay counters** command.

Field	Description
INTERFACE	Indicates the interface IP address of the DHCP relay agent.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
REQUESTS	Indicates the number of DHCP requests.
REPLIES	Indicates the number of DHCP replies.

Clearing DHCP relay counters for a VLAN

Use this procedure to clear the DHCP relay counters for a VLAN.

Procedure steps

Step	Action
1	To clear the DHCP relay counters, enter the following from the VLAN Interface Configuration command mode: `ip dhcp-relay clear-counters`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

UDP broadcast forwarding configuration using NNCLI

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface to a configured IP address. This section describes how to configure UDP broadcast forwarding using NNCLI.

The UDP broadcast forwarding feature cannot be enabled or disabled on a global level. When you attach the first UDP forwarding list to a VLAN interface, the feature is enabled. When you remove the last UDP forwarding list from a VLAN, the feature is disabled.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a UDP forwarding interface.
- Ensure that a route to the destination address is available on the switch.

UDP broadcast forwarding configuration procedures

To configure UDP broadcast forwarding, perform the following steps:

Step	Action
1	Create UDP protocol entries that specify the protocol associated with each UDP port that you want to forward.
2	Create a UDP forwarding list that specifies the destination IP addresses for each forwarding UDP port. (You can create up to 128 UDP forwarding lists.)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

3	Apply UDP forwarding lists to local VLAN interfaces.

--End--

UDP broadcast forwarding configuration navigation

- "Configuring UDP protocol table entries" (page 262)
- "Displaying the UDP protocol table" (page 262)
- "Configuring a UDP forwarding list" (page 263)
- "Applying a UDP forwarding list to a VLAN" (page 264)
- "Displaying the UDP broadcast forwarding configuration" (page 265)
- "Clearing UDP broadcast counters on an interface" (page 266)

Configuring UDP protocol table entries

Use this procedure to create UDP table entries that identify the protocols associated with specific UDP ports that you want to forward.

Procedure steps

Step	Action
1	To configure a UDP table entry, enter the following from the Global Configuration mode: **ip forward-protocol udp [<forwarding_port> <protocol_name>]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<forwarding_port>	Specifies the UDP port number. Range is 1-65535.
<protocol_name>	Specifies the UDP protocol name.

Displaying the UDP protocol table

Use this procedure to display the configured UDP protocol table entries.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To display the UDP protocol table, enter the following from the User Exec mode: **show ip forward-protocol udp**

--End--

Job aid

The following table shows the field descriptions for the **show ip forward-protocol udp** command.

Field	Description
UDP_PORT	Indicates the UDP ports.
PROTOCOL_NAME	Indicates the name of the associated protocol.

Configuring a UDP forwarding list

Use this procedure to configure a UDP forwarding list, which associates UDP forwarding ports with destination IP addresses. Each forwarding list can contain multiple port/destination entries. You can configure a maximum of 16 port/destination entries in one forwarding list.

You can configure up to 128 forwarding lists.

Procedure steps

Step	Action
1	To configure a UDP port forwarding list, enter the following from the Global Configuration mode: **ip forward-protocol udp portfwdlist <forward_list> <udp_port> <dest_ip> [name <list_name>]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<forward_list>	Specifies the ID of the UDP forwarding list. Range is 1-128.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<udp_port>	Specifies the port on which the UDP forwarding originates.
<dest_ip>	Specifies the destination IP address of the UDP forwarding.
<list_name>	Specifies the name of the UDP forwarding list being created. (maximum 15 characters).

Applying a UDP forwarding list to a VLAN

Use this procedure to associate a UDP forwarding list to a VLAN interface (you can attach only one list at a time to a VLAN interface).

You can bind the same UDP forwarding list to a maximum of 16 different VLANs.

Procedure steps

Step	Action
1	To associate a UDP forwarding list to a VLAN, enter the following from the VLAN Interface Configuration mode: **`ip forward-protocol udp [vlan <vid>] [portfwdlist <forward_list>] [broadcastmask <bcast_mask>] [maxttl <max_ttl>]`**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<vid>	Specifies the VLAN ID on which to attach the UDP forwarding list. This parameter is optional, and if not specified, the UDP forwarding list is applied to the interface specified in the **interface vlan** command.
<forward_list>	Specifies the ID of the UDP forwarding list to attach to the selected VLAN interface.
<bcast_mask>	Specifies the 32 bit mask used by the selected VLAN interface to take forwarding decisions based on the destination IP address of the incoming UDP broadcast traffic. If you do not specify a broadcast mask value, the mask of the interface to which the list is attached is used. (See Note 1.)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<max_ttl>	Specifies the time to live (TTL) value inserted in the IP headers of the forwarded UDP packets coming out of the selected VLAN interface. If you do not specify a TTL value, the default value (4) is used. (See Note 1.)
Note 1: If you specify maxttl and/or broadcastmask values with no portfwdlist specified, then the switch saves the settings for this interface. When a portfwdlist is subsequently attached to this interface, and the maxttl and/or broadcastmask value are not defined, the saved parameters are automatically attached to the list. But if when specifying the portfwdlist, you also specify the maxttl and/or broadcastmask, your specified properties are used, regardless of any previous configurations.	

Displaying the UDP broadcast forwarding configuration

Use this procedure to display the UDP broadcast forwarding configuration.

Procedure steps

Step	Action
1	To display the UDP broadcast forwarding configuration, enter the following from the User Exec mode: **show ip forward-protocol udp [interface [vlan <1-4094>]] [portfwdlist [<portlist>]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[interface [vlan <1-4094>]]	Displays the configuration and statistics for a VLAN interface. If no VLAN is specified, the configuration for all UDP forwarding-enabled VLANs is displayed.
[portfwdlist [<forward_list>]	Displays the specified UDP forwarding list. If no list is specified, a summary of all forwarding lists is displayed.

Job aids

The following table shows the field descriptions for the **show ip forward-protocol udp** command.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
UDP_PORT	Indicates the UDP ports.
PROTOCOL_NAME	Indicates the name of the protocol.

The following table shows the field descriptions for the **show ip forward-protocol udp interfaces** command.

Field	Description
INTF_ADDR	Indicates the IP address of the interface.
FWD LISTID	Identifies the UDP forwarding policy.
MAXTTL	Indicates the maximum TTL.
RXPKTS	Indicates the number of received packets.
FWDPKTS	Indicates the number of forwarded packets.
DRPDEST UNREACH	Indicates the number of dropped packets that cannot reach the destination.
DRP_UNKNOWN PROTOCOL	Indicates the number of packets dropped with an unknown protocol.
BDCASTMASK	Indicates the value of the broadcast mask.

The following table shows the field descriptions for the **show ip forward-protocol udp portfwdlist** command.

Field	Description
LIST_ID	Specifies the specific UDP forwarding policy number.
NAME	Specifies the name of the UDP forwarding policy.

Clearing UDP broadcast counters on an interface

Use this procedure to clear the UDP broadcast counters on an interface.

Procedure steps

Step	Action
1	To clear the UDP broadcast counters, enter the following from the Privileged Exec Configuration mode: **clear ip forward-protocol udp counters <1-4094>**

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<1-4094>	Specifies the VLAN ID.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Directed broadcasts configuration using NNCLI

This chapter describes procedures you can use to configure and display the status of directed broadcasts using NNCLI.

Directed broadcasts configuration navigation

- "Configuring directed broadcasts" (page 269)
- "Displaying the directed broadcast configuration" (page 270)

Configuring directed broadcasts

Use this procedure to enable directed broadcasts on the switch. By default, directed broadcasts are disabled.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a broadcast interface.
- Ensure that a route (local or static) to the destination address is available on the switch.

Procedure steps

Step	Action
1	To enable directed broadcasts, enter the following from the Global Configuration mode: `ip directed-broadcast enable`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying the directed broadcast configuration

Use this procedure to display the status of directed broadcasts on the switch. By default, directed broadcasts are disabled.

Procedure steps

Step	Action
1	To display directed broadcast status, enter the following from the User EXEC mode: **show ip directed-broadcast**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Static ARP and Proxy ARP configuration using NNCLI

This chapter describes the procedures you can use to configure Static ARP, Proxy ARP, and display ARP entries using the NNCLI.

Static ARP and Proxy ARP configuration navigation

- "Static ARP configuration" (page 271)
- "Displaying the ARP table" (page 272)
- "Proxy ARP configuration" (page 274)

Static ARP configuration

This section describes how to configure Static ARP using the NNCLI.

Configuring a static ARP entry

Use this procedure to create and enable a static ARP entry.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN.

Procedure steps

Step	Action
1	To configure a static ARP entry, enter the following from the Global Configuration mode: `[no] ip arp <A.B.C.D> <aa:bb:cc:dd:ee:ff> <unit / port> [vid <1-4094>]`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the specified ARP entry.
<A.B.C.D>	Specifies the IP address of the device being set as a static ARP entry.
<aa:bb:cc:dd:ee:ff>	Specifies the MAC address of the device being set as a static ARP entry.
<unit / port>	Specifies the unit and port number to which the static ARP entry is being added.
vid <1 - 4094>	Specifies the VLAN ID to which the static ARP entry is being added.

Configuration example: Adding a static ARP entry to a VLAN

The following is an example of adding a static ARP entry to a VLAN or brouter port:

```
5530-24TFD(config)# ip arp 10.1.1.23 00:00:11:43:54:23 1/48 vid
1
```

Configuration example: Deleting a static ARP entry

The following is an example of deleting a static ARP entry:

```
5530-24TFD(config)# no ip arp 172.2.2.13
```

Displaying the ARP table

Use the following procedures to display the ARP table, configure a global timeout for ARP entries, and clear the ARP cache.

Navigation

- "Displaying ARP entries" (page 272)
- "Configuring a global timeout for ARP entries" (page 273)
- "Clearing the ARP cache" (page 274)

Displaying ARP entries

Use this procedure to display ARP entries.

Procedure steps

Step	Action
1	To display ARP entries, enter the following from the User Exec mode: `show arp-table`

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

OR

```
show ip arp [static | dynamic]
[<ip-addr> | {-s <subnet> <mask>}]
[summary]
```

The **show ip arp** command is invalid if the switch is not in Layer 3 mode.

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<ip-addr>	Specifies the IP address of the ARP entry to be displayed.
-s <subnet> <mask>	Displays ARP entries for the specified subnet only.
static	Displays all configured static entries, including those without a valid route.

Job aid

The following table shows the field descriptions for the **show ip arp** command.

Field	Description
IP Address	Specifies the IP address of the ARP entry.
Age (min)	Displays the ARP age time.
MAC Address	Specifies the MAC address of the ARP entry.
VLAN-Unit/Port/Trunk	Specifies the VLAN/port of the ARP entry.
Flags	Specifies the type of ARP entry. S=Static, D=Dynamic, L=Local, B=Broadcast.

Configuring a global timeout for ARP entries

Use this procedure to configure an aging time for the ARP entries.

Procedure steps

Step	Action
1	To configure a global timeout for ARP entries, enter the following from the Global Configuration mode:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip arp timeout <timeout>
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<timeout>	Specifies the amount of time in minutes before an ARP entry ages out. Range is 5-360. The default value is 360 minutes.

Configuration example: Changing the global ARP timeout

The following is an example of setting a new default aging time:

```
5530-24TFD(config)# ip arp timeout 180
```

The new setting can be confirmed by using the **show ip routing** command.

Clearing the ARP cache

Use this procedure to clear the cache of ARP entries.

Procedure steps

Step	Action
1	To clear the ARP cache, enter the following from the Global Configuration mode: `clear arp-cache`

--End--

Proxy ARP configuration

This section describes how to configure Proxy ARP using the NNCLI.

Navigation

- "Configuring proxy ARP status" (page 275)
- "Displaying proxy ARP status on a VLAN" (page 275)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring proxy ARP status

Use this procedure to enable proxy ARP functionality on a VLAN. By default, proxy ARP is disabled.

Prerequisite

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a Proxy ARP interface.

Procedure steps

Step	Action
1	To configure proxy ARP status, enter the following from the VLAN Interface Configuration mode: `[default] [no] ip arp-proxy enable`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[default]	Disables proxy ARP functionality on the VLAN.
[no]	Disables proxy ARP functionality on the VLAN.

Displaying proxy ARP status on a VLAN

Use this procedure to display the status of proxy ARP on a VLAN.

Procedure steps

Step	Action
1	To display proxy ARP status for a VLAN, enter the following from the User EXEC mode: `show ip arp-proxy interface [vlan <vid>]`

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<vid>	Specifies the ID of the VLAN to display. Range is 1-4094.

Job aid

The following table shows the field descriptions for the **show ip arp-proxy interfaces** command.

Field	Description
Vlan	Identifies a VLAN.
Proxy ARP status	Specifies the status of Proxy ARP on the VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IP blocking configuration using NNCLI

This chapter describes the procedures you can use to configure and display the status of IP blocking in a stack using NNCLI.

IP blocking configuration navigation

- "Configuring IP blocking for a stack" (page 277)
- "Displaying IP blocking status" (page 277)

Configuring IP blocking for a stack

Use this procedure to set the IP blocking mode in the stack.

Procedure steps

Step	Action
1	To configure IP blocking, enter the following from the Global Configuration mode: `ip blocking-mode {full | none}`
	--End--

Variable definitions

The following table describes the command variables.

Variable	Value
full	Select this parameter to set IP blocking to full. This never allows a duplicate IP address in a stack.
none	Select this parameter to set IP blocking to none. This allows duplicate IP addresses unconditionally.

Displaying IP blocking status

Use this command to display the status of IP blocking on the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To display the status of IP blocking on the switch, enter the following from the User EXEC mode: **show ip blocking**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

VRRP configuration using NNCLI

The Virtual Router Redundancy Protocol (VRRP) 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.

This section describes the procedures you can use to configure VRRP on a VLAN using NNCLI.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Assign an IP address to the VLAN that you want to enable with VRRP.

  Routing is automatically enabled on the VLAN when you assign an IP address to it.

VRRP configuration procedures

To enable VRRP on a VLAN, perform the following steps:

Step	Action
1	Enable VRRP globally on the switch.
2	Assign a virtual router IP address to a virtual router ID.
3	Configure the priority for this router as required.
4	Enable the virtual router.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

VRRP configuration navigation

- "Configuring global VRRP status" (page 280)
- "Assigning an IP address to a virtual router ID " (page 281)
- "Configuring the router priority for a virtual router ID" (page 281)
- "Configuring the status of the virtual router" (page 282)
- "Configuring the critical IP address" (page 283)
- "Configuring the VRRP critical IP status" (page 283)
- "Configuring a backup master" (page 282)
- "Configuring the VRRP holddown timer" (page 284)
- "Configuring VRRP holddown action" (page 284)
- "Configuring the VRRP advertisement interval" (page 285)
- "Configuring the fast advertisement interval" (page 285)
- "Configuring fast advertisement status" (page 286)
- "Configuring ICMP echo replies" (page 287)
- "Enabling VRRP traps" (page 287)
- "Displaying VRRP configuration and statistics" (page 288)

Configuring global VRRP status

Use this procedure to configure the global VRRP status on the switch.

Procedure steps

Step	Action
1	To configure the global VRRP status on the switch, enter the following from the Global Configuration mode: `[no] router vrrp enable`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Globally disables VRRP on the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Assigning an IP address to a virtual router ID

Use this procedure to associate an IP address with a virtual router ID.

Procedure steps

Step	Action
1	To assign an IP address to a virtual router ID, enter the following from the Interface Configuration mode: `[no] ip vrrp address <vr_id> <ip_address>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Removes the IP address from the virtual router ID.
<vr_id>	Specifies the virtual router being configured. This is a value between 1 and 255.
<ip_address>	Represents the address to be associated with the virtual router ID.

Configuring the router priority for a virtual router ID

Use this procedure to assign a priority to the router for a specified virtual router ID.

Procedure steps

Step	Action
1	To assign a priority to the router for a particular virtual router ID, enter the following from the Interface Configuration mode: `ip vrrp <vr_id> priority <priority_value>`

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<vr_id>	Specifies the virtual router ID for which the router priority is being configured.
<priority_value>	Specifies the priority assigned to the router for the specified virtual router ID. This is a value between 1 and 255.

Configuring the status of the virtual router

Use this procedure to enable the virtual router.

Procedure steps

Step	Action
1	To enable or disable the virtual router, enter the following from the Interface Configuration mode: `[no] ip vrrp <vr_id> enable`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router being configured.
[no]	Disables the virtual router.

Configuring a backup master

Use this procedure to configure the VRRP backup master functionality. Enable backup master on both the master and backup routers.

Procedure steps

Step	Action
1	To configure the backup master, enter the following from the Interface Configuration mode: `[no] ip vrrp <vr_id> backup-master enable`

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
[no]	Disables the VRRP backup master functionality.
<vr_id>	Specifies the virtual router being configured.

Configuring the critical IP address

Use this procedure to set the critical IP address on the router.

Procedure steps

Step	Action
1	To configure the critical IP address, enter the following from the Interface Configuration mode: **[no] ip vrrp <vr_id> critical-ip-addr <ip_address>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router being configured.
<ip_address>	Specifies the critical IP address to be used.
[no]	Removes the configured critical IP.

Configuring the VRRP critical IP status

Use this procedure to configure the status of the VRRP critical IP functionality.

Procedure steps

Step	Action
1	To configure the critical IP status, enter the following from the Interface Configuration mode: **[no] ip vrrp <vr_id> critical-ip enable**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router being configured.
[no]	Disables the VRRP critical IP functionality

Configuring the VRRP holddown timer

Use this procedure to set the VRRP holddown timer.

Procedure steps

Step	Action
1	To set the VRRP holddown timer, enter the following from the Interface Configuration mode: **ip vrrp <vr_id> holddown-timer <timer_value>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router being configured.
<timer_value>	Specifies the holddown timer value. This is a value in seconds between 0 and 21600.

Configuring VRRP holddown action

Use this procedure to define the action this interface takes when a holddown timer threshold has been reached.

Procedure steps

Step	Action
1	To define the action for the interface to perform in the holddown state, enter the following from the Interface Configuration mode:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip vrrp <vr_id> action {none | preempt}
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router being configured.
{none | preempt}	If **none** is selected as the action, no action is taken when a holddown timer threshold is reached. If **preempt** is selected as the action, the holddown timer is cancelled.

Configuring the VRRP advertisement interval

Use this procedure to configure the VRRP advertisement interval.

Procedure steps

Step	Action
1	To configure the VRRP advertisement interval enter the following from the Interface Configuration mode: **ip vrrp <vr_id> adver-int <interval>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router being configured.
<interval>	Specifies the advertisement interval in seconds. This is an integer value between 1 and 255

Configuring the fast advertisement interval

Use this procedure to set the interval used in the VRRP fast advertisement functionality.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To enable or disable fast advertisements, enter the following from the Interface Configuration mode: `ip vrrp <vr_id> fast-adv-int <interval>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router being configured.
<interval>	Specifies the fast advertisement interval. This is a value in milliseconds between 200 and 1000.

Configuring fast advertisement status

Use this procedure to enable the VRRP fast advertisement functionality.

Procedure steps

Step	Action
1	To enable or disable fast advertisements, enter the following from the Interface Configuration mode: `[no] ip vrrp <vr_id> fast-adv enable`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<vr_id>	Specifies the virtual router being configured.
[no]	Disables the VRRP fast advertisement functionality.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring ICMP echo replies

Use this procedure to configure ICMP echo replies from VRRP associated addresses.

Procedure steps

Step	Action
1	To configure ICMP echo replies, enter the following from the VRRP Router Configuration mode: `[no] ping-virtual-address enable`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables ICMP echo replies from VRRP associated addresses.

Enabling VRRP traps

Use this procedure to enable the sending of SNMP notifications after virtual router state changes.

Procedure steps

Step	Action
1	To enable VRRP traps, enter the following from the VRRP Router Configuration mode: `send-trap enable`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables the sending of SNMP notifications after virtual router state changes.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying VRRP configuration and statistics

Use this procedure to display VRRP configuration information and statistics. If no parameters are specified, basic global configuration information is displayed.

Procedure steps

Step	Action
1	To display global VRRP properties, enter the following from the User EXEC mode: **show ip vrrp**
2	To display the VRRP virtual IP address configuration for a VLAN, enter the following from the User EXEC mode: **show ip vrrp address** **[addr <A.B.C.D>]** **[vrid <1-255>]** **[interface [addr <A.B.C.D>]** **[vlan <1-4094>]** **[vrid <1-255>]]**
3	To display detailed VRRP interface configuration information, enter the following from the User EXEC mode: **show ip vrrp interface** **[verbose]** **[vlan <1-4094>]** **[vrid <1-255>]**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
addr <A.B.C.D>	Displays VRRP configurations associated with the specified IP address.
interface [addr <A.B.C.D>] [vlan <1-4094>] [vrid <1-255>]	Displays VRRP configurations associated with the specified interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
vrid <1-255>	Displays VRRP configurations associated with the specified virtual router ID.
vlan <1-4094>	Displays VRRP configurations associated with the specified VLAN.
verbose	Displays additional VRRP configuration information.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

VRRP configuration examples using NNCLI

This section provides configuration examples showing how to configure VRRP on a Nortel Ethernet Routing Switch 5000 Series .

Navigation

- “Configuring normal VRRP operation” (page 291)
- “Configuring VRRP with SMLT” (page 297)
- “Configuring VRRP with SLT” (page 301)

Configuring normal VRRP operation

The following configuration example shows how to provide VRRP service for two edge host locations.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 44
VRRP example topology

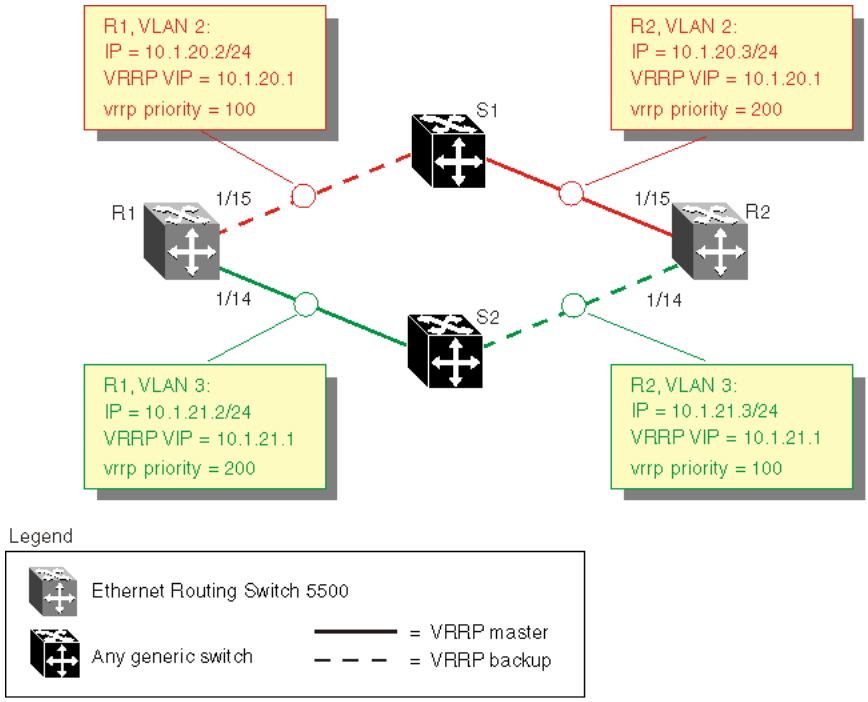

In this example, the switches have the following duties:

- R1 is the VRRP master for S2
- R2 is the VRRP master for S1

In this example, VRRP is enabled with OSPF as the routing protocol on R1 and R2.

The VRRP priority setting is used to determine which router will become the VRRP master and which will become the VRRP backup. In instances where the priority setting is the same for two routers, the higher IP address becomes the tie breaker. Therefore, it is very important to set the correct VRRP priority. VRRP fast advertisement will also be enabled in this example to allow for fast failover detection.

The following procedure describes the steps necessary to reproduce the example described above:

Step	Action
1	Configure VLAN 2 on router R1.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

a Create VLAN 2 on router R1.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 2 type port
```

b Configure the ports for VLAN 2 on R1.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan members add 2 1/15
```

c Configure an IP address for VLAN 2.

Add IP address 10.1.20.2 / 255.255.255.0 to VLAN 2.

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.2
255.255.255.0
```

d Configure an OSPF interface for VLAN 2.

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# network 10.1.20.2
```

e Configure VRRP on VLAN 2.

The VRRP VIP address of 10.1.20.1 is added to VLAN 2 using a VRID of 1.

Note 1: The VRRP priority is not configured here; it is left at factory default of 100. Instead, the priority setting on router R2 will be set to a higher value when R2 is configured.

Note 2: Fast advertisement is disabled by default. Fast advertisement is proprietary to Nortel to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. If fast VRRP advertisement is desired, enable fast advertisement.

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip vrrp address 1 10.1.20.1
5530-24TFD(config-if)# ip vrrp 1 enable
```

2 Configure VLAN 3 on router R1.

a Configure VLAN 3 on router R1 using spanning tree group 1.

```
5530-24TFD# config terminal
5530-24TFD# vlan create 3 type port
```

b Configure the ports for VLAN 3 on R1.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan members add 3 1/14
```

c Configure an IP address for VLAN 3.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Add IP address 10.1.21.2 / 255.255.255.0 to VLAN 3.

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 3
5530-24TFD(config)# ip address 10.1.21.2
255.255.255.0
```

d Configure an OSPF interface for VLAN 3.

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# network 10.1.21.2
```

e Configure VRRP on VLAN 3.

The VRRP VIP address of 10.1.21.1 is added to VLAN 2 using a VRID of 2.

> ***Note:*** Fast advertisement is disabled by default. Fast advertisement is proprietary to Nortel to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. If fast VRRP advertisement is desired, enable fast advertisement.

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip vrrp address 2 10.1.21.1
5530-24TFD(config-if)# ip vrrp 2 priority 200
5530-24TFD(config-if)# ip vrrp 2 enable
```

3 Configure VLAN 2 on router R2.

a Create VLAN 2 on router R2.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 2 type port
```

b Configure the ports for VLAN 2 on R2.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan members add 2 1/15
```

c Configure an IP address for VLAN 2.

Add IP address 10.1.20.3 / 255.255.255.0 to VLAN 2.

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.3
255.255.255.0
```

d Configure an OSPF interface for VLAN 2.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# network 10.1.20.3
```

e Configure VRRP on VLAN 2.

The VRRP VIP address of 10.1.21.1 is added to VLAN 2 using a VRID of 1.

Note 1: For this example the VRRP priority value is set to 200. This allows router R2 to be elected as the VRRP master router.

Note 2: Fast advertisement is disabled by default. Fast advertisement is proprietary to Nortel to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. If fast VRRP advertisement is desired, enable fast advertisement.

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip vrrp address 1 10.1.20.1
5530-24TFD(config-if)# ip vrrp 1 enable
5530-24TFD(config-if)# ip vrrp 1 priority 200
```

4 Configure VLAN 3 on router R2.

a Configure VLAN 3 on router R2.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 3 type port
```

b Configure the ports for VLAN 3 on R1.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan members add 3 1/14
```

c Configure an IP address for VLAN 3.

Add IP address 10.1.21.3 / 255.255.255.0 to VLAN 3.

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip address 10.1.21.3
255.255.255.0
```

d Configure an OSPF interface for VLAN 3.

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# network 10.1.21.3
```

e Configure VRRP on VLAN 3.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The VRRP VIP address of 10.1.21.1 is added to VLAN 2 using a VRID of 2.

Note: Fast advertisement is disabled by default. Fast advertisement is proprietary to Nortel to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. If fast VRRP advertisement is desired, enable fast advertisement.

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 3
5530-24TFD(config-if)# ip vrrp address 2 10.1.21.1
5530-24TFD(config-if)# ip vrrp 2 enable
```

--End--

After the VRRP configuration has been completed, use the **show ip vrrp** and **show ip vrrp interface verbose** commands to display VRRP configuration information and statistics.

Configuration command listing

This following list is a complete sequence of the commands used in this configuration:

1. VLAN Configuration for Router R1

   ```
   config tvlan create 2 type port vlan members remove 2
   1/1-1/14,2/1-2/8,3/1-3/8vlan members add 2 1/15 interface
   vlan 2 ip address 10.1.20.2 255.255.255.0 router ospf enable
   router ospf network 10.1.20.2 router vrrp ena interface vlan
   2ip vrrp address 1 10.1.20.1 ip vrrp 1 enablevlan create 3 type
   portinterface vlan 3ip address 10.1.21.2 255.255.255.0router
   ospf enablerouter ospfnetwork 10.1.21.2router vrrp
   enainterface vlan 3ip vrrp address 2 10.1.21.1ip vrrp 2
   priority 200ip vrrp 2 enable
   ```

2. VLAN Configuration for Router R2

   ```
   config t vlan create 2 type port vlan members remove 2
   1/1-1/14,2/1-2/8,3/1-3/8 vlan members add 2 1/15 interface
   vlan 2 ip address 10.1.20.3 255.255.255.0 router ospf enable
   router ospf network 10.1.20.3 router vrrp ena interface
   vlan 2 ip vrrp address 1 10.1.20.1 ip vrrp 1 priority 200 ip
   vrrp 1 enable vlan create 3 type port vlan members remove
   3 1/1-1/14,1/15,2/1-2/8,3/1-3/8 vlan members add 3 1/14
   interface vlan 3 ip address 10.1.21.3 255.255.255.0 router
   ospf enable router ospf network 10.1.21.3 router vrrp ena
   interface vlan 3 ip vrrp address 2 10.1.21.1 ip vrrp 2 enable
   ```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring VRRP 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 demonstrated below, both R1 and R2 switches are configured with a port-based VLAN (VLAN 2) with SMLT and VRRP set to enabled. This topology provides failover protection and load-balancing.

The Nortel Ethernet Routing Switch 5000 Series (E1), running in Layer 2 mode, is configured with one port-based VLAN and one Multi-Link Trunking (MLT) group for the aggregate uplink ports. The Nortel Ethernet Routing Switch 5000 Series switches (R1 and R2) are configured with backup master enabled so that both switches can reply to ARP.

Figure 45
VRRP with SMLT configuration

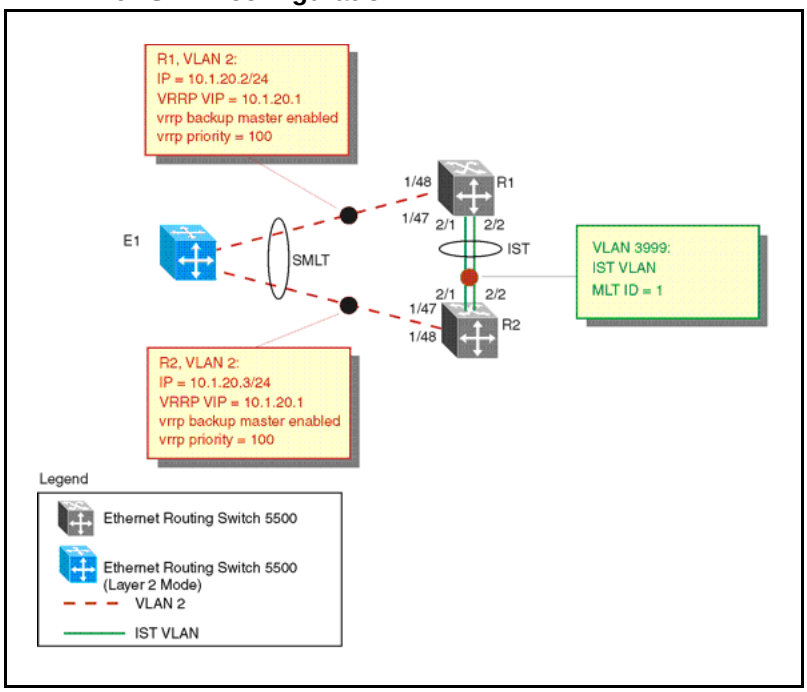

The following procedure would be used to recreate the illustrated topology:

Step	Action
1	Configure the IST VLAN on R1

a Configure IST VLAN 3999 on R1

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 3999 type port
5530-24TFD(config)# interface vlan 3999
5530-24TFD(config-if)# ip address 2.1.1.1
255.255.255.0
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

b Configure IST MLT on R1

```
5530-24TFD# config terminal
5530-24TFD(config)# mlt 1 member 2/1-2/2
5530-24TFD(config)# vlan port 2/1-2/2 tagging enable
5530-24TFD(config)# mlt 1 enable
```

c Configure the IST and add the IST to VLAN 3999

```
5530-24TFD# config terminal
5530-24TFD(config)# interface mlt 1
5530-24TFD(config-if)# ist enable peer-ip 2.1.1.2
vlan 3999
```

2 Configure VRRP and SMLT for access VLAN to E1

a Configure VLAN 2 on R1

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 2 type port
```

b Create IP address for VLAN 2

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.2
255.255.255.0
```

c Configure the access port for VLAN 2 on R1 and add VLAN 2 to the IST and SMLT groups

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan members add 2 1/48,1/47,2/
1,2/2
```

Note: 2/1 and 2/2 are IST ports. 1/48,1/47 are SMLT ports.

d Create SMLT on R1

```
5530-24TFD# config terminal
5530-24TFD(config)# mlt 2 member 1/47,1/48
5530-24TFD(config)# mlt 2 enable
5530-24TFD(config)# interface mlt 2
5530-24TFD(config-if)# smlt 1
```

e Enable OSPF interface on VLAN 2 of R1

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# network 10.1.20.2
```

f Configure VRRP VIP address for VLAN2 of R1

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 2
5530-24TFD(config)# ip vrrp address 1 10.1.20.1
5530-24TFD(config)# ip vrrp 1 enable
5530-24TFD(config)# ip vrrp 1 backup-master enable
```

Note: Fast advertisement is disabled by default. Fast advertisement is proprietary to Nortel to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. If fast VRRP advertisement is desired, enable fast advertisement.

3 Configure the IST VLAN for router R2

a Configure IST VLAN 3999 on R2

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 3999 type port
5530-24TFD(config)# interface vlan 3999
5530-24TFD(config-if)# ip address 2.1.1.2
255.255.255.0
```

b Configure IST MLT on R2

```
5530-24TFD# config terminal
5530-24TFD(config)# mlt 1 member 2/1-2/2
5530-24TFD(config)# mlt 1 enable
5530-24TFD(config)# vlan port 2/1-2/2 tagging enable
```

c Configure an IST peer for R2 and add the IST to VLAN 3999

```
5530-24TFD# config terminal
5530-24TFD(config)# interface mlt 1
5530-24TFD(config-if)# ist enable peer-ip 2.1.1.2
vlan 3999
```

4 Configure VRRP and SMLT for VLAN access to E1

a Configure VLAN 2 on R2

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan create 2 type port
```

b Create an IP address for VLAN 2

```
5530-24TFD# config terminal
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip address 10.1.20.3
255.255.255.0
```

c Configure the access port for VLAN 2 on R2 and add VLAN 2 to the IST and SMLT groups

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
5530-24TFD# config terminal
5530-24TFD(config)# vlan members add 2 1/47, 1/48,
2/1.  2/2
```

Note: 1/47 and 1/48 are SMLT ports. 2/1 and 2/2 are IST ports.

d Create SMLT on R2

```
5530-24TFD# config terminal
5530-24TFD(config)# mlt 2 member 1/47, 1/48
5530-24TFD(config)# mlt 2 enable
5530-24TFD(config)# interface mlt 2
5530-24TFD(config-if)# smlt 1
```

e Enable OSPF interface for VLAN 2 on R2

```
5530-24TFD# config terminal
5530-24TFD(config)# router ospf enable
5530-24TFD(config)# router ospf
5530-24TFD(config-router)# network 10.1.20.3
```

f Configure VRRP VIP address for VLAN 2 on R2

```
5530-24TFD# config terminal
5530-24TFD(config)# router vrrp ena
5530-24TFD(config)# interface vlan 2
5530-24TFD(config-if)# ip vrrp address 1 10.1.20.1
5530-24TFD(config-if)# ip vrrp 1 enable
5530-24TFD(config-if)# ip vrrp 1 backup-master enable
```

Note: Fast Advertisement is disabled by default. Fast advertisement is proprietary to Nortel to support an advertisement interval from 200 to 1000 milliseconds (ms) with a default of 200. If fast VRRP advertisement is desired, enable fast advertisement.

--End--

Configuration command listing

This following list is a complete sequence of the commands used in this configuration:

1. Configuration for R1

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
#MLT CONFIGURATION #config tmlt 1 member 2/1-2/2mlt 1
enavlan port 2/2-2/2 tagging enableinterface mlt 1ist
enable peer-ip 2.1.1.2 vlan 3999mlt 2 member 1/48,1/47mlt
2 enableinterface mlt 2smlt 1#VLAN CONFIGURATION #config
tvlan members remove 1 1/47-1/48,2/1-2/2vlan create 2 type
portvlan members remove 2 1/1-1/46,2/3-2/8,3/1-3/8vlan
members add 2 1/47-1/48,2/1-2/2interface vlan 2ip
address 10.1.20.2 255.255.255.0router ospf enablerouter
ospfnetwork 10.1.20.2router vrrp enainterface vlan
2ip vrrp address 1 10.1.20.1ip vrrp 1 enableip vrrp 1
backup-master enablevlan create 3999 type portvlan
members remove 3999 1/1-1/47,2/3-2/8,3/1-3/8vlan members
add 3999 2/1-2/2interface vlan 3999ip address 2.1.1.1
255.255.255.0#PORT CONFIGURATION – PHASE II #config tmlt
spanning-tree 1 stp all learning disablemlt spanning-tree 2
stp all learning disable
```

2. Configuration for R2

```
#MLT CONFIGURATION # config tmlt 1 member 2/1-2/2mlt 1
enablevlan port 2/2-2/2 tagging enableinterface mlt 1ist
enable peer-ip 2.1.1.1 vlan 3999mlt 2 member 1/48,1/47mlt
2 enableinterface mlt 2smlt 1#VLAN CONFIGURATION #config
tvlan members remove 1 1/47-1/48,2/1-2/2vlan create 2 type
portvlan members remove 2 1/1-1/46,2/3-2/8,3/1-3/8vlan
members add 2 1/47-1/48,2/1-2/2interface vlan 2ip
address 10.1.20.3 255.255.255.0router ospf enablerouter
ospfnetwork 10.1.20.2router vrrp enainterface vlan
2ip vrrp address 1 10.1.20.1ip vrrp 1 enableip vrrp 1
backup-master enablevlan create 3999 type portvlan
members remove 3999 1/1-1/47,2/3-2/8,3/1-3/8vlan members
add 3999 2/1-2/2interface vlan 3999ip address 2.1.1.2
255.255.255.0#PORT CONFIGURATION – PHASE II #config tmlt
spanning-tree 1 stp all learning disablemlt spanning-tree 2
stp all learning disable
```

Configuring VRRP with SLT

The following illustration and configuration file examples demonstrate a VRRP configuration with SLT.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 46
VRRP with SLT configuration

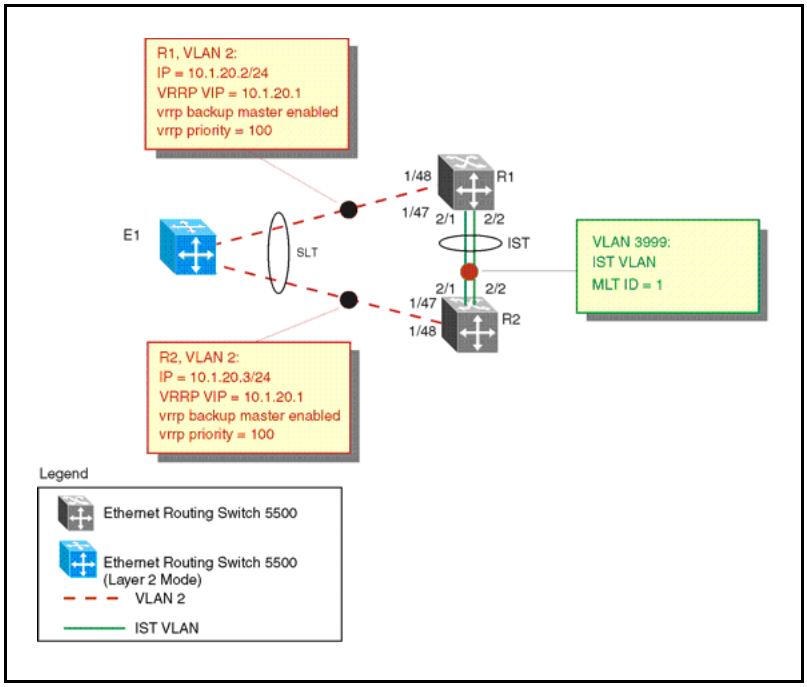

The following commands would recreate the above configuration:

1. Configuration for R1

```
#MLT CONFIGURATION #config tmlt 1 member 2/1-2/2mlt
1 enablevlan port 2/2-2/2 tagging enableinterface
mlt 1ist enable peer-ip 2.1.1.2 vlan 3999interface
fast-Ethernet 1/48smlt 1#VLAN CONFIGURATION #config
tvlan members remove 1 1/48,2/1-2/2vlan create 2 type
portvlan members remove 2 1/1-1/47,2/3-2/8,3/1-3/8vlan
members add 2 1/47,2/1-2/2interface vlan 2ip address
10.1.20.2 255.255.255.0router ospf enablerouter
ospfnetwork 10.1.20.2router vrrp enainterface vlan
2ip vrrp address 1 10.1.20.1ip vrrp 1 enableip vrrp 1
backup-master enablevlan create 3999 type portvlan
members remove 3999 1/1-1/47,2/3-2/8,3/1-3/8vlan members
add 3999 2/1-2/2interface vlan 3999ip address 2.1.1.1
255.255.255.0#PORT CONFIGURATION - PHASE II #config
tmlt spanning-tree 1 stp all learning disableinterface
fast-Ethernet 1/48spanning-tree stp 1 learning disable
```

2. Configuration for R2

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
#MLT CONFIGURATION #config tmlt 1 member 2/1-2/2mlt
1 enablevlan port 2/2-2/2 tagging enableinterface
mlt 1ist enable peer-ip 2.1.1.1 vlan 3999interface
fast-Ethernet 1/48smlt 1#VLAN CONFIGURATION #config
tvlan members remove 1 1/48,2/1-2/2vlan create 2 type
portvlan members remove 2 1/1-1/47,2/3-2/8,3/1-3/8vlan
members add 2 1/48,2/1-2/2interface vlan 2ip address
10.1.20.3 255.255.255.0router ospf enablerouter
ospfnetwork 10.1.20.2router vrrp enainterface vlan
2ip vrrp address 1 10.1.20.1ip vrrp 1 enableip vrrp 1
backup-master enablevlan create 3999 type portvlan
members remove 3999 1/1-1/47,2/3-2/8,3/1-3/8vlan members
add 3999 2/1-2/2interface vlan 3999ip address 2.1.1.2
255.255.255.0#PORT CONFIGURATION - PHASE II #config
tmlt spanning-tree 1 stp all learning disableinterface
fast-Ethernet 1/48spanning-tree stp 1 learning disable
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IGMP snooping configuration using NNCLI

This chapter describes the procedures you can use to configure IGMP snooping on a VLAN using NNCLI.

IGMP snooping configuration procedures

Step	Action
1	To configure IGMP snooping, the only required configuration is to enable snooping on the VLAN. All related configurations, listed below, are optional and can be configured to suit the requirements of your network.
	--End--

IGMP snooping configuration navigation

- "Job aid: Roadmap of IGMP NNCLI commands" (page 306)
- "Configuring IGMP snooping on a VLAN" (page 307)
- "Configuring IGMP send query on a VLAN" (page 308)
- "Configuring IGMP proxy on a VLAN" (page 308)
- "Configuring the IGMP version on a VLAN" (page 309)
- "Configuring static mrouter ports on a VLAN" (page 310)
- "Displaying IGMP snoop, proxy, and mrouter configuration" (page 311)
- "Configuring IGMP parameters on a VLAN" (page 312)
- "Configuring the router alert option on a VLAN" (page 313)
- "Displaying IGMP interface information" (page 314)
- "Displaying IGMP group information" (page 437)
- "Configuring unknown multicast packet filtering" (page 317)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- "Displaying the status of unknown multicast packet filtering" (page 318)
- "Displaying the multicast MAC addresses for which flooding is allowed" (page 319)
- "Displaying IGMP cache information" (page 319)
- "Flushing the router table" (page 320)
- "Configuring IGMP selective channel block" (page 321)

Job aid: Roadmap of IGMP NNCLI commands

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
VLAN Interface Configuration mode	
ip igmp	last-member-query-interval <last-mbr-query-int>
	mrouter <portlist>
	proxy
	query-interval <query-int>
	query-max-response <query-max-resp>
	robust-value <robust-val>
	router-alert
	snooping
	version <1-3>
Global Configuration mode	
ip igmp	flush vlan <vid> <grp-member|mrouter|sender>
vlan igmp <vid>	[snooping {enable | disable}]
	[proxy {enable | disable}]
	[query-interval <query-int>] [robust-value <robust-val>]
	unknown-mcast-no-flood {enable | disable}
	unknown-mcast-allow-flood <H.H.H> <mcast_ip_address>

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Command	Parameter
	{v1-members | v2-members} [add | remove] <portlist>
Privileged EXEC mode	
show ip igmp	cache
	group [count] [group <A.B.C.D>] [member-subnet <A.B.C.D>/<0-32>]
	interface [vlan <vid>]
	snooping
show vlan igmp	unknown-mcast-allow-flood
	unknown-mcast-no-flood
	<vid>
show vlan multicast membership	<vid>

Configuring IGMP snooping on a VLAN

Enable IGMP snooping on a VLAN to forward the multicast data to only those ports that are members of the group.

IGMP snooping is disabled by default.

Procedure steps

Step	Action

1	To enable IGMP snooping, enter the following from the VLAN Interface Configuration command mode:
	`[default] [no] ip igmp snooping`
	OR
	Enter the following from the Global Configuration command mode:
	`[default] vlan igmp <vid> [snooping {enable | disable}]`

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
default	Disables IGMP snooping on the selected VLAN.
no	Disables IGMP snooping on the selected VLAN.
enable	Enables IGMP snooping on the selected VLAN.
disable	Disables IGMP snooping on the selected VLAN.

Configuring IGMP send query on a VLAN

Use this procedure to enable IGMP send query on a snoop-enabled VLAN. When IGMP snooping send query is enabled, the IGMP snooping querier sends out periodic IGMP queries that trigger IGMP report messages from the switch or host that wants to receive IP multicast traffic. IGMP snooping listens to these IGMP reports to establish appropriate forwarding.

IGMP send query is disabled by default.

Prerequisites

- You must enable snoop on the VLAN.

Procedure steps

Step	Action
1	To enable IGMP send query, enter the following command from the VLAN Interface Configuration mode: `ip igmp send-query`

--End--

Configuring IGMP proxy on a VLAN

Use this procedure to enable IGMP proxy on a snoop-enabled VLAN. With IGMP proxy enabled, the switch consolidates incoming report messages into one proxy report for that group.

IGMP proxy is disabled by default.

Prerequisites

- You must enable snoop on the VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action

1	To enable IGMP proxy, enter the following from the VLAN Interface Configuration mode:
	`[default] [no] ip igmp proxy`
	OR
	Enter the following from the Global Configuration command mode:
	`[default] [no] vlan igmp <vid> [proxy {enable | disable}]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
default	Disables IGMP proxy on the selected VLAN.
no	Disables IGMP proxy on the selected VLAN.
<vid>	Specifies the VLAN ID.
enable	Enables IGMP proxy on the selected VLAN.
disable	Disables IGMP proxy on the selected VLAN.

Configuring the IGMP version on a VLAN

Use this procedure to configure the IGMP version running on the VLAN. You can specify the version as IGMPv1, IGMPv2, or IGMPv3 (IGMPv3 is supported for IGMP snooping only; it is not supported with PIM-SM). The default is IGMPv2.

Procedure steps

Step	Action
1	To configure the IGMP version, enter the following from the VLAN Interface Configuration mode: `[default] ip igmp version <1-3>`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
default	Restores the default IGMP protocol version (IGMPv2).
<1-3>	Specifies the IGMP version.

Configuring static mrouter ports on a VLAN

IGMP snoop considers the port on which the IGMP query is received as the active IGMP multicast router (mrouter) port. By default, the switch forwards incoming IGMP Membership Reports only to the active mrouter port.

To forward the IGMP reports to additional ports, you can configure the additional ports as static mrouter ports.

Procedure steps

Step	Action
1	To configure static mrouter ports on a VLAN (IGMPv1, IGMPv2, and IGMPv3 according to the supported version on the VLAN), enter the following from the VLAN Interface Configuration mode: **`[default] [no] ip igmp mrouter <portlist>`** **OR** To configure IGMPv1 or IGMPv2 static mrouter ports, enter the following from the Global Configuration command mode: **`[no] vlan igmp <vid> {v1-members | v2-members} [add | remove] <portlist>`**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
default	Removes all static mrouter ports.
no	Removes the specified static mrouter port.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<portlist>	Specifies the list of ports to add or remove as static mrouter ports.
{v1-members | v2-members}	Specifies whether the static mrouter ports are IGMPv1 or IGMPv2.
[add | remove]	Specifies whether to add or remove the static mrouter ports.
<portlist>	Specifies the list of ports to add or remove as static mrouter ports.

Displaying IGMP snoop, proxy, and mrouter configuration

Use this procedure to display the IGMP snoop, proxy, and mrouter configuration per VLAN.

Procedure steps

Step	Action
1	To display IGMP snoop information, enter: **`show ip igmp snooping`**

--End--

Job aid

The following table shows the field descriptions for the **`show ip igmp snooping`** command.

Field	Description
Vlan	Indicates the Vlan ID.
Snoop Enable	Indicates whether snoop is enabled (true) or disabled (false).
Proxy Snoop Enable	Indicates whether IGMP proxy is enabled (true) or disabled (false).
Static Mrouter Ports	Indicates the static mrouter ports in this VLAN that provide connectivity to an IP multicast router.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Active Mrouter Ports	Displays all dynamic (querier port) and static mrouter ports that are active on the interface.
Mrouter Expiration Time	Specifies the time remaining before the multicast router is aged out on this interface. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The Query Max Response Interval (obtained from the queries received) is used as the timer resolution.

Configuring IGMP parameters on a VLAN

Use this procedure to configure the IGMP parameters on a VLAN.

ATTENTION
The query interval, robustness, and version values must be the same as those configured on the interface (VLAN) of the multicast router (IGMP querier).

Procedure steps

Step	Action

1	To configure IGMP parameters, enter the following from the VLAN Interface Configuration mode:
	`[default] ip igmp`
	`[last-member-query-interval <last-mbr-query-int>]`
	`[query-interval <query-int>]`
	`[query-max-response <query-max-resp>]`
	`[robust-value <robust-val>]`
	`[version <1-3>]`
	OR
	enter the following from the Global Configuration command mode:
	`[default] vlan igmp <vid>`
	`[query-interval <query-int>]`
	`[robust-value <robust-val>]`

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
default	Sets the selected parameter to the default value. If no parameters are specified, snoop is disabled and all IGMP parameters are set to their defaults.
<last-mbr-query-int>	Sets the maximum response time (in 1/10 seconds) that is inserted into group-specific queries sent in response to leave group messages. This parameter is also the time between group-specific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The range is from 0–255, and the default is 10 (1 second). Nortel recommends configuring this parameter to values higher than 3. If a fast leave process is not required, Nortel recommends values above 10. (The value 3 is equal to 0.3 of a second, and 10 is equal to 1.0 second.)
<query-int>	Sets the frequency (in seconds) at which host query packets are transmitted on the VLAN. The range is 1–65535. The default value is 125 seconds.
<query-max-resp>	Specifies the maximum response time (in 1/10 seconds) advertised in IGMPv2 general queries on this interface. The range is 0–255. The default value is 100 (10 seconds).
<robust-val>	Specifies tuning for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, you must increase the robustness value. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). The range is from 2 to 255, and the default is 2. The default value of 2 means that one query for each query interval can be dropped without the querier aging out.

Configuring the router alert option on a VLAN

Use this command to enable the router alert feature. This feature instructs the router to drop control packets that do not have the router-alert flag in the IP header.

ATTENTION

To maximize your network performance, Nortel recommends that you set the router alert option according to the version of IGMP currently in use:

- IGMPv1—Disable
- IGMPv2—Enable
- IGMPv3—Enable

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To configure the router alert option on a VLAN, enter the following from the VLAN Interface Configuration mode: **`[default] [no] ip igmp router-alert`**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
default	Disables the router alert option.
no	Disables the router alert option.

Displaying IGMP interface information

Use this procedure to display IGMP interface parameters.

Procedure steps

Step	Action
1	To display the IGMP interface information, enter: **`show ip igmp interface [vlan <vid>]`** **OR** Enter: **`show vlan igmp <vid>`**

--End--

Job aids

The following table shows the field descriptions for the **`show ip igmp interface command`** command.

Field	Description
VLAN	Indicates the VLAN on which IGMP is configured.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Query Intvl	Specifies the frequency (in seconds) at which host query packets are transmitted on the interface.
Vers	Specifies the version of IGMP configured on this interface.
Oper Vers	Specifies the version of IGMP running on this interface.
Querier	Specifies the IP address of the IGMP querier on the IP subnet to which this interface is attached.
Query MaxRspT	Indicates the maximum query response time (in tenths of a second) advertised in IGMPv2 queries on this interface.
Wrong Query	Indicates the number of queries received whose IGMP version does not match the Interface version. You must configure all routers on a LAN to run the same version of IGMP. Thus, if queries are received with the wrong version, a configuration error occurs.
Joins	Indicates the number of times a group membership was added on this interface.
Robust	Specifies the robust value configured for expected packet loss on the interface.
LastMbr Query	Indicates the maximum response time (in tenths of a second) inserted into group-specific queries sent in response to leave group messages, and is also the amount of time between group-specific query messages. Use this value to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group. This does not apply if the interface is configured for IGMPv1.
Send Query	Indicates whether the ip igmp send-query feature is enabled or disabled. Values are YES of NO. Default is disabled.

The following table shows the field descriptions for the **`show vlan igmp`** command.

Field	Description
Snooping	Indicates whether snooping is enabled or disabled.
Proxy	Indicates whether proxy snoop is enabled or disabled.
Robust Value	Indicates the robust value configured for expected packet loss on the interface.
Query Time	Indicates the frequency (in seconds) at which host query packets are transmitted on the interface.
IGMPv1 Static Router Ports	Indicates the IGMPv1 static mrouter ports.
IGMPv2 Static Router Ports	Indicates the IGMPv2 static mrouter ports.
Send Query	Indicates whether the ip igmp send-query feature is enabled or disabled. Values are YES of NO. Default is disabled.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying IGMP group membership information

Display the IGMP group information to show the learned multicast groups and the attached ports.

Procedure steps

Step	Action
1	To display IGMP group information, enter: **show ip igmp group [count] [group <A.B.C.D>] [member-subnet <A.B.C.D>/<0-32>]** **OR** Enter: **show vlan multicast membership <vid>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
count	Displays the number of IGMP group entries.
group <A.B.C.D>	Displays group information for the specified group.
member-subnet <A.B.C.D>/<0-32>	Displays group information for the specified member subnet.

Job aids

The following table shows the field descriptions for the **show ip igmp group** command.

Field	Description
Group Address	Indicates the multicast group address.
VLAN	Indicates the VLAN interface on which the group exists.
Member Address	Indicates the IP address of the IGMP receiver (host or IGMP reporter). The IP address is 0.0.0.0 if the type is static.
Expiration	Indicates the time left before the group report expires. This variable is updated upon receiving a group report.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Type	Specifies the type of membership: static or dynamic.
In Port	Identifies the member port for the group. This is the port on which group traffic is forwarded and in those case where the type is dynamic, it is the port on which the IGMP join was received.

The following table shows the field descriptions for the **show vlan multicast membership** command.

Field	Description
Multicast Group Address	Indicates the multicast group address.
In Port	Indicates the physical interface or a logical interface (VLAN) that received group reports from various sources.

Configuring unknown multicast packet filtering

The default switch behavior is to flood all packets with unknown multicast addresses. Use this procedure to prevent the flooding of packets with unknown multicast addresses and enable the forwarding of these packets to static mrouter ports only.

Procedure steps

Step	Action
1	To configure unknown multicast packet flooding, enter the following from the Global Configuration mode: **[no] [default] vlan igmp <vid> unknown-mcast-no-flood {enable | disable}**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
no	Enables the flooding of multicast packets on the VLAN.
default	Enables the flooding of multicast packets on the VLAN.
enable	Prevents the flooding of multicast packets on the VLAN.
disable	Enables the flooding of multicast packets on the VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying the status of unknown multicast packet filtering

Use this procedure to display the status of unknown multicast filtering: enabled (no flooding) or disabled (flooding allowed).

Procedure steps

Step	Action
1	To display the unknown multicast flooding configuration, enter: **show vlan igmp unknown-mcast-no-flood**

--End--

Job aid

The following table shows the field descriptions for the **show vlan igmp unknown-mcast-no-flood** command.

Field	Description
Unknown Multicast No-Flood	Specifies the status of unknown multicast filtering: enabled or disabled.

Specifying a multicast MAC address to be allowed to flood all VLANs

Use this procedure to allow particular unknown multicast packets to be flooded on all switch VLANs.

To add MAC addresses starting with 01.00.5E to the allow-flood table, you must specify the corresponding multicast IP address. For instance, you cannot add MAC address 01.00.5E.01.02.03 to the allow-flood table, but instead you must specify IP address 224.1.2.3.

For all other types of MAC address, you can enter the MAC address directly to allow flooding.

Procedure steps

Step	Action
1	To allow particular unknown multicast packets to be flooded, enter the following from the Global Configuration mode: **vlan igmp unknown-mcast-allow-flood {<H.H.H> | <mcast_ip_address>}**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
<H.H.H>	Specifies the multicast MAC address to be flooded. Accepted formats are: • H.H.H • xx:xx:xx:xx:xx:xx • xx.xx.xx.xx.xx.xx • xx-xx-xx-xx-xx-xx
<mcast_ip_address>	Specifies the multicast IP address to be flooded.

Displaying the multicast MAC addresses for which flooding is allowed

Use this procedure to display the multicast MAC addresses for which flooding is allowed on all switch VLANs.

Procedure steps

Step	Action
1	To display the multicast MAC addresses for which flooding is allowed, enter: **show vlan igmp unknown-mcast-allow-flood**

--End--

Job aid

The following table shows the field descriptions for the **show vlan igmp unknown-mcast-allow-flood** command.

Field	Description
Allowed Multicast Addresses	Indicates multicast addresses that can flood.

Displaying IGMP cache information

Display the IGMP cache information to show the learned multicast groups in the cache and the IGMPv1 version timers.

Note: Using the **show ip igmp cache** command may not display the expected results in some configurations. If the expected results are not displayed, use the **show ip igmp group** command to view the information.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To display the IGMP cache information, enter: **show ip igmp cache**

--End--

Job aid

The following table shows the field descriptions for the **show ip igmp cache** command.

Field	Description
Group Address	Indicates the multicast group address.
Vlan ID	Indicates the VLAN interface on which the group exists.
Last Reporter	Indicates the last IGMP host to join the group.
Expiration	Indicates the group expiration time (in seconds).
V1 Host Timer	Indicates the time remaining until the local router assumes that no IGMP version 1 members exist on the IP subnet attached to the interface. Upon hearing an IGMPv1 membership report, this value is reset to the group membership timer. When the time remaining is nonzero, the local interface ignores IGMPv2 Leave messages that it receives for this group.
Type	Indicates whether the entry is learned dynamically or is added statically.

Flushing the router table

Use this procedure to flush the router table.

Procedure steps

Step	Action
1	To flush the router table, enter the following from the Global Configuration mode: **ip igmp flush vlan <vid> {grp-member|mrouter}**

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
{grp-member|mrouter}	Flushes the table specified by type.

Configuring IGMP selective channel block

In certain deployment scenarios it might be required not to allow multicast streaming from specific group addresses to users connected to certain ports. With the IGMP selective channel block feature this type of control can be implemented. When configured it will control the IGMP membership of ports by blocking IGMP reports received from users on that port destined for the specific group address/addresses. The filter can be configured to block a single multicast address or range of addresses.

This feature will work regardless of whether the switch is in Layer 2 IGMP snooping mode or the full IGMP mode (PIM-SM enabled). It will also be applicable for IGMPv1 and v2.

Creating an IGMP profile

Use this procedure to create an IGMP profile.

Step	Action
1	From Global Configuration mode, enter the **ip igmp profile <profile number (1-65535)>**command.
2	Enter the **deny** command.
3	Enter the **range <ip multicast address> <ip multicast address>** command.

--End--

Deleting an IGMP profile

Use this procedure to delete an IGMP profile.

Step	Action
1	To delete an IGMP profile enter the following command from Global Configuration mode: **no ip igmp profile <profile number (1-65535)>**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Applying the IGMP filter profile on interface

Use this procedure to apply the IGMP filter profile on an interface.

Step	Action
1	From Global Configuration mode enter the **interface <interface-id>** command.
2	Enter the **ip igmp filter <profile number>** command.

--End--

Removing a profile from an interface

Use this procedure to remove a profile from an interface.

Step	Action
1	From Global Configuration mode enter the **interface <interface-id>** command.
2	Enter the **no ip igmp filter <profile number>** command.

--End--

Displaying an IGMP profile

Use this procedure to display an IGMP profile.

Step	Action
1	To display an IGMP profile enter the following command from Global Configuration mode: **show ip igmp profile <cr> or <profile number>**

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

PIM-SM configuration using NNCLI

This chapter describes the procedures you can use to configure PIM-SM using the NNCLI.

Unlike dense-mode protocols, such as DVMRP that initially flood multicast traffic to all routers over an entire internetwork, PIM-SM sends multicast traffic only to routers that belong to a specific multicast group and that choose to receive the traffic. PIM-SM reduces overhead costs for processing unwanted multicast packets.

Prerequisites for PIM-SM configuration

Before you can configure PIM-SM, you must prepare the switch as follows:

Step	Action
1	Install the Advanced Routing software license. **ATTENTION** If your Ethernet Routing Switch is running an Advanced License for a release prior to Release 6.0, to enable PIM-SM you must regenerate your license file from the Nortel web site and install the new license file on the switch.
2	Enable routing globally.
3	Configure IP addresses and enable routing on the VLAN interfaces on which you want to configure PIM-SM.
4	Enable a unicast protocol, either RIP or OSPF, globally and on the interfaces on which you want to configure PIM-SM.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

ATTENTION
PIM-SM requires a unicast protocol to multicast traffic within the network when performing the Reverse Path Forwarding (RPF) check. PIM-SM also uses the information from the unicast routing table to create and maintain the shared and shortest path multicast tree. The unicast routing table must contain a route to every multicast source in the network, as well as routes to PIM entities such as the rendezvous points (RP) and bootstrap router (BSR).

--End--

PIM-SM configuration procedures

To configure PIM-SM, you must perform the following procedures:

Step	Action
1	Enable PIM-SM globally. (If desired, modify the default global PIM-SM properties.)
2	Enable PIM-SM on individual VLAN interfaces. (If desired, modify the default VLAN PIM-SM properties.)
3	Configure candidate RPs for the multicast groups in the network. (It is best to have multiple candidate-RPs in the network; however, with the Ethernet Routing Switch 5000, you can only configure one candidate-RP per switch for any number of groups.) OR Configure one (or several) static RPs for the multicast groups in the network. (To enable static RP in the PIM-SM domain, you must configure the same static RPs on every system that takes part in PIM-SM forwarding.)
4	Configure one or several candidate BSRs to propagate RP information to all switches in the network. (You can configure every PIM-enabled VLAN as a C-BSR. If Static RP is enabled, this step is not required.)

ATTENTION
Ensure that all routers in the path from the receivers to the RP and to the multicast source are PIM-enabled. Also ensure that all PIM routers have unicast routes to reach the source and RP through directly-connected PIM neighbors.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

All additional configurations listed below are optional and can be configured according to the requirements of your network.

PIM-SM configuration navigation

- "Job aid: Roadmap of PIM-SM configuration commands" (page 325)
- "Enabling and disabling PIM-SM globally" (page 327)
- "Configuring global PIM-SM properties" (page 327)
- "Enabling PIM-SM on a VLAN" (page 329)
- "Configuring the PIM-SM interface type on a VLAN" (page 330)
- "Displaying PIM-SM neighbors" (page 331)
- "Configuring PIM-SM properties on a VLAN" (page 331)
- "Displaying the PIM-SM configuration for a VLAN" (page 332)
- "Specifying the router as a candidate BSR on a VLAN" (page 333)
- "Displaying the BSR configuration" (page 334)
- "Specifying a local IP interface as a candidate RP" (page 335)
- "Displaying the candidate RP configuration" (page 336)
- "Displaying the PIM-SM RP set" (page 337)
- "Displaying the active RP per group" (page 337)
- "Enabling and disabling static RP" (page 338)
- "Configuring a static RP" (page 339)
- "Displaying the static RP configuration" (page 340)
- "Specifying a virtual neighbor on an interface" (page 340)
- "Displaying the virtual neighbor configuration" (page 341)
- "Displaying PIM multicast routes" (page 341)
- "Displaying the PIM mode" (page 342)

Job aid: Roadmap of PIM-SM configuration commands

The following table lists the commands and their parameters that you use to complete the procedures in this section.

Command	Parameter
Global Configuration mode	
`ip pim`	`bootstrap-period <bootstrap-period >`

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Command	Parameter
	rp-c-adv-timeout <rp-c-adv-time>
	disc-data-timeout <disc-data-time>
	enable
	join-prune-interval <join-prune-int>
	fwd-cache-timeout <fwd-cache-time>
	register-suppression-timeout <rgstr-suppr-time>
	unicast-route-change-timeout <unicast-rte-chge-time>
ip pim rp-candidate	group <group-addr> <group-mask> rp <rp-addr>
ip pim static-rp	enable
	<group-addr> <group-mask> <static-rp-addr>
ip pim virtual-neighbor	<if-ipaddr> <v-nbr-ipaddr>
Interface vlan mode	
ip pim	bsr-candidate priority <priority>
	enable
	interface-type <active|passive>
	join-prune-interval <join-prune-int>
	query-interval <query-int>
Privileged EXEC mode	
show ip mroute	{interface | next-hop | route}
show ip pim	
show ip pim active-rp	[group <group-addr>]
show ip pim bsr	
show ip pim interface	[vlan <vlan-id>]
show ip pim mode	
show ip pim mroute	[source <ipaddr>] [group <group>] [summary]
show ip pim neighbor	
show ip pim rp-candidate	[group <group-addr>]
show ip pim rp-hash	
show ip pim static-rp	
show ip pim virtual-neighbor	

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Enabling and disabling PIM-SM globally

To enable PIM-SM on individual interfaces, you must first enable PIM-SM globally.

By default, PIM-SM is disabled.

Procedure steps

Step	Action
1	To configure the global PIM-SM status on the switch, enter the following from the Global Configuration mode: `[no] ip pim enable`

--End--

Variable definitions

The following table describes the command variable.

Variable	Value
[no]	Disables PIM-SM globally.

Configuring global PIM-SM properties

Use this procedure to configure the global PIM-SM parameters on the switch.

Procedure steps

Step	Action

1	To configure the global PIM-SM parameters, enter the following from the Global Configuration mode:
	`ip pim`
	`bootstrap-period <bootstrap-period>`
	`disc-data-timeout <disc-data-time>`
	`fwd-cache-timeout <fwd-cache-time>`
	`join-prune-interval <join-prune-int>`
	`register-suppression-timeout <rgstr-suppr-time>`
	`rp-c-adv-timeout <rp-c-adv-time>`
	`unicast-route-change-timeout <unicast-rte-chge-time>`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
<bootstrap-period>	Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages. Range is 5–32757. The default is 60.
<disc-data-time>	After the router forwards the first source packet to the RP, this value specifies how long (in seconds) the router discards subsequent source data while waiting for a join from the RP. An IPMC discard record is created and deleted after the timer expires or after a join is received. Range is 5–65535. The default is 60.
<fwd-cache-time>	Specifies the forward cache timeout globally. This value is used in aging PIM-SM mroutes. Range is 10–86400. The default is 210.
<join-prune-int>	Specifies how long to wait (in seconds) before the PIM-SM router sends out the next join/prune message to the upstream neighbors. Range is 1–18724. The default is 60.
<rgstr-suppr-time>	Specifies the PIM-SM register suppression timeout. Range is 6–65535. The default is 60.
<rp-c-adv-time>	Specifies how often (in seconds) candidate RPs (C-RP) send C-RP advertisement messages. After this timer expires, the C-RP sends an advertisement message to the elected BSR. Range is 5–26214. The default is 60.
<unicast-rte-chge-time>	Specifies the PIM-SM unicast route change timeout. Indicates how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates to be used by PIM. Range is 2–65535. The default is 5.

Displaying global PIM-SM properties

Use this procedure to display global PIM-SM properties.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To display the global properties of PIM-SM on the switch, enter: `show ip pim`

--End--

Job aid

The following table shows the field descriptions for the **`show ip pim`** command.

Field	Description
PIM Admin Status	Indicates the status of PIM-SM.
PIM Boot Strap Period	Indicates the interval between originating bootstrap messages at the elected BSR.
PIM C-RP-Adv Message Send Interval	Indicates the candidate RPs timer (in seconds) for sending C-RP advertisement messages.
PIM Discard Data Timeout	After the router forwards the first source packet to the RP, this value indicates how long (in seconds) the router discards subsequent source data while waiting for a join from the RP. An IPMC discard record is created and deleted after the timer expires or after a join is received.
PIM Join Prune Interval	Indicates the join/prune interval in seconds.
PIM Register Suppression Timer	Indicates the register suppression timer in seconds.
PIM Uni Route Change Timeout	Indicates how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates to be used by PIM.
PIM Mode	Indicates the PIM mode. The Ethernet Routing Switch 5000 Series supports only the sparse mode (SM).
PIM Static-RP	Indicates the status of static RP.
Forward Cache Timeout	Indicates the PIM-SM forward cache expiry value in seconds. This value is used in aging PIM-SM mroutes.

Enabling PIM-SM on a VLAN

Use this procedure to enable PIM-SM on a VLAN.

By default, PIM-SM is disabled on VLANs.

Prerequisites

- You must enable PIM-SM globally.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To configure the PIM-SM status on a particular VLAN, enter the following from the VLAN Interface mode: `[no] ip pim enable`

--End--

Variable definitions

The following table describes the command variable.

Variable	Value
[no]	Disables PIM-SM on the VLAN.

Configuring the PIM-SM interface type on a VLAN

Use this procedure to change the state (active or passive) of PIM on a VLAN interface. An active interface transmits and receives PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. This feature is useful when you have a high number of PIM interfaces and these interfaces are connected to end users, not to other switches.

By default, VLANs are active interfaces.

Prerequisites

- Before you change the state of PIM on a VLAN interface, you must disable PIM on the interface to prevent instability in the PIM operations, especially when neighbors are present or when streams are received.

Procedure steps

Step	Action
1	To set the interface type on a VLAN, enter the following from the VLAN Interface mode: `ip pim interface-type <active|passive>`

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the command variables.

Variable	Value
interface-type <active|passive>	Sets the interface type on a particular VLAN: • active: allows PIM-SM control traffic to be transmitted and received. • passive: prevents PIM-SM control traffic from being transmitted or received, reducing the load on the system.

Displaying PIM-SM neighbors

Use this procedure to display PIM-SM neighbors.

Procedure steps

Step	Action
1	To display PIM neighbors, enter: **`show ip pim neighbor`**

--End--

Job aid

The following table shows the field descriptions for the **`show ip pim neighbor`** command.

Field	Description
Address	Specifies the IP address of the PIM-SM neighbor.
Vlan	Specifies the local interface.
Uptime	Specifies the elapsed time since the PIM-SM neighbor last became a neighbor of the local interface.
Expiry Time	Specifies the time remaining before this PIM-SM neighbor times out.
Total PIM Neighbors	Specifies the total number of PIM neighbors on the switch.

Configuring PIM-SM properties on a VLAN

Configure the PIM-SM properties on a VLAN to modify the join/prune interval or the query interval.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To configure the PIM-SM VLAN parameters, enter the following from the VLAN Interface mode: **ip pim** **join-prune-interval <join-prune-int>** **query-interval <query-int>**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<join-prune-int>	Specifies how long to wait (in seconds) before the PIM-SM switch sends out the next join/prune message to the upstream neighbors. Range is 1–18724, and the default is 60.
<query-int>	Sets the hello interval for the VLAN. Range is 0–18724, and the default is 30.

Displaying the PIM-SM configuration for a VLAN

Use this procedure to display information about the PIM-SM interface configuration for a VLAN.

Procedure steps

Step	Action
1	To display information about the PIM-SM interface configuration for VLANs, enter: **show ip pim interface [vlan <vid>]**

--End--

Variable definitions

The following table describes the command variables.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<vid>	Specifies the VLAN to display (1–4094).

Job aid

The following table shows the field descriptions for the **show ip pim interface vlan** command.

Field	Description
Vlan	Identifies the VLAN.
State	Indicates the state of PIM-SM on the VLAN.
Address	Specifies the VLAN IP address.
Mask	Specifies the VLAN subnet mask.
Mode	Indicates the PIM mode of this VLAN. Ethernet Routing Switch 5000 Series supports only sparse mode.
DR	Indicates the Designated Router for this interface.
Hello Interval	Indicates how long the switch waits (in seconds) between sending out hello message to neighboring switches. The default hello interval is 30 seconds.
Join Prune Interval	Indicates how long the switch waits (in seconds) between sending out join/prune message to the upstream neighbors. The default join/prune interval is 60 seconds.
CBSPR	Indicates the priority for this local interface to become a Candidate BSR. The Candidate BSR with the highest BSR priority and address is referred to as the preferred BSR. The default is –1, which indicates that the current interface is not a Candidate BSR.
Oper State	Indicates the status of PIM-SM on this interface: up or down.
Interface Type	Indicates whether the PIM-SM interface is active or passive.

Specifying the router as a candidate BSR on a VLAN

Because PIM-SM cannot run without a bootstrap router (BSR), you must specify at least one C-BSR in the domain. The C-BSR with the highest configured priority becomes the BSR for the domain. You can configure additional C-BSRs to provide backup protection in case the primary BSR fails.

If two C-BSRs have equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with the highest priority to the domain, it automatically becomes the new BSR.

With the Ethernet Routing Switch 5000 Series, you can configure every PIM-enabled interface as a C-BSR.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To enable the router as a candidate BSR on a VLAN, enter the following from the VLAN Interface mode: **`[no] ip pim bsr-candidate priority <priority>`**

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<priority>	Specifies the priority value of the candidate to become a BSR. The range is 0 to 255 and the default is -1, which indicates that the current interface is not a Candidate BSR.
[no]	Removes the candidate BSR configuration.

Displaying the BSR configuration

Use this procedure to display the current BSR configuration.

Procedure steps

Step	Action
1	To display the current BSR configuration, enter: **`show ip pim bsr`**

--End--

Job aid

The following table shows the field descriptions for the **`show ip pim bsr`** command.

Field	Description
Current BSR Address	Specifies the IP address of the current BSR for the local PIM-SM domain.
Current BSR Priority	Specifies the priority of the current BSR. The Candidate BSR (C-BSR) with the highest BSR priority and address (referred to as the preferred BSR) is elected as the BSR for the domain.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Current BSR Hash Mask	Specifies the mask used in the hash function to map a group to one of the C-RPs from the RP-Set. With the hash-mask, a small number of consecutive groups (for example, four) can always hash to the same RP.
Current BSR Fragment Tag	Specifies a randomly generated number that distinguishes fragments belonging to different bootstrap messages. Fragments belonging to the same bootstrap message carry the same Fragment Tag.
Current BSR Boot Strap Timer	Specifies the time the BSR waits between sending bootstrap messages.

Specifying a local IP interface as a candidate RP

Because PIM-SM cannot run without an RP, you must specify at least one C-RP in the domain. Use this procedure to configure a local PIM-SM interface as a candidate RP (C-RP).

With the Ethernet Routing Switch 5000 Series, you can configure only one local interface as a C-RP for any number of groups.

With the mask value, you can configure a C-RP for several groups in one configuration. For example, with a C-RP configuration with a group address of 224.0.0.0 and a group mask of 240.0.0.0, you can configure the C-RP for a multicast range from 224.0.0.0 to 239.255.255.255.

Procedure steps

Step	Action
1	To add the interface as a C-RP in the RP set, enter the following from the Global Configuration mode: **`[no] ip pim rp-candidate group <group-addr> <group-mask> rp <c-rp-addr>`**

--End--

Variable definitions

The following table describes the command variables.

Variable	value
<group-addr>	Specifies the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	value
<group-mask>	Specifies the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
<c-rp-addr>	Specifies the IP address of the C-RP. This address must be one of the local PIM-SM enabled interfaces.
[no]	Removes the configured RP candidate.

Displaying the candidate RP configuration

Use this procedure to display the candidate RP configuration.

Procedure steps

Step	Action
1	To display the candidate RP configuration, enter: **`show ip pim rp-candidate [group <group-addr>]`**

--End--

Variable definitions

The following table describes the command variables.

Variable	value
<group-addr>	Specifies the IP address of the multicast group configuration to display.

Job aid

The following table shows the field descriptions for the **`show ip pim rp-candidate`** command.

Field	Description
Group Address	Specifies the IP address of the multicast group. When combined with the group mask, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
Group Mask	Specifies the address mask of the multicast group. When combined with the group address, it identifies the prefix that the local router uses to advertise itself as a C-RP router.
RP Address	Specifies the IP address of the C-RP.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying the PIM-SM RP set

Display the RP set for troubleshooting purposes. The BSR constructs the RP set from C-RP advertisements, and then distributes it to all PIM routers in the PIM domain for the BSR.

Procedure steps

Step	Action
1	To display the PIM RP set, enter: **show ip pim rp-hash**

--End--

Job aid

The following table shows the field descriptions for the **show ip pim rp-hash** command.

Field	Description
Group Address	Specifies the IP address of the multicast group.
Group Mask	Specifies the address mask of the multicast group.
Address	Specifies the IP address of the C-RP for the specified group.
Hold Time	Indicates the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
Expiry Time	Specifies the time remaining before this C-RP times out.

Displaying the active RP per group

Use this procedure to display the active RP per group.

The active RP is displayed only when there is at least one (*,G) or (S,G) entry on the router after either joins or multicast data are received by the router.

Procedure steps

Step	Action
1	To display the active RP, enter:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
show ip pim active-rp [group <group-addr>]
```

--End--

Variable definitions

The following table describes the command variables.

Variable	value
<group-addr>	Specifies the IP address of the multicast group configuration to display.

Job aid

The following table shows the field descriptions for the **show ip pim active-rp** command.

Field	Description
Group Address	Specifies the IP address of the multicast group.
Group Mask	Specifies the address mask of the multicast group.
Active RP	Specifies the IP address of the active RP.
Priority	Specifies the RP priority.

Enabling and disabling static RP

Enable static RP to avoid the process of dynamically learning C-RPs through the BSR mechanism. With this feature, static RP-enabled Ethernet Routing Switch 5000 Series switches can communicate with switches from other vendors that do not use the BSR mechanism.

ATTENTION
When you enable static RP, all dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.

Procedure steps

Step	Action
1	To enable static RP, enter the following from the Global Configuration mode: `[no] ip pim static-rp [enable]`
2	After you enter the command, a warning message appears. To confirm the change, enter:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
Y
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
[no]	Disables static RP.
[enable]	Enables static RP.

Configuring a static RP

Use this procedure to configure a static RP entry. After you configure static RP, the switch ignores the BSR mechanism and uses only the RPs that you configure statically.

ATTENTION
You cannot configure a static RP-enabled switch as a BSR or as a C-RP.

Prerequisites

- You must enable static RP.

Procedure steps

Step	Action
1	To add a static RP to the RP set, enter: `[no] ip pim static-rp <group-addr> <group-mask> <static-rp-addr>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<group-addr>	Specifies the IP address of the multicast group. Together with the group mask, the group address identifies the range of the multicast addresses that the RP handles.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Value
<group-mask>	Specifies the address mask of the multicast group. Together with the group address, the address mask identifies the range of the multicast addresses that the RP handles.
<static-rp-addr>	Specifies the IP address of the static RP.

Displaying the static RP configuration

Use this procedure to display the static RP configuration.

Procedure steps

Step	Action
1	To display the static RP configuration, enter: **show ip pim static-rp**

--End--

Job aid

The following table shows the field descriptions for the **show ip pim static-rp** command.

Field	Description
Group Address	Indicates the IP address of the multicast group. When combined with the group mask, the group address identifies the prefix that the local router uses to advertise as a static RP.
Group Mask	Indicates the address mask of the multicast group. When combined with the group address, the group mask identifies the prefix that the local router uses to advertise as a static RP.
RP Address	Indicates the IP address of the static RP.
Status	Indicates the status of static RP.

Specifying a virtual neighbor on an interface

Configure a virtual neighbor when the next hop for a static route cannot run PIM-SM, such as a Virtual Redundancy Router Protocol address on an adjacent device. The virtual neighbor IP address appears in the Ethernet Routing Switch 5000 neighbor table.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To configure a virtual neighbor, enter the following from the Global Configuration mode: `[no] ip pim virtual-neighbor <if-ipaddr> <v-nbr-ipaddr>`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<if-ipaddr>	Specifies the IP address of the selected interface.
<v-nbr-ipaddr>	Specifies the IP address of the virtual neighbor.
[no]	Removes the configured virtual neighbor.

Displaying the virtual neighbor configuration

Use this procedure to display the virtual neighbor.

Procedure steps

Step	Action
1	To display the virtual neighbor, enter: `show ip pim virtual-neighbor`

--End--

Job aid

The following table shows the field descriptions for the **`show ip pim virtual-neighbor`** command.

Field	Description
Vlan	Indicates the VLAN interface.
Neighbor address	Indicates the IP address of the virtual neighbor.

Displaying PIM multicast routes

Use this procedure to display PIM multicast routes.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	To display the PIM multicast routes, enter: `show ip pim mroute [source <source-addr>] [group <group-addr>] [summary]`

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
<source-addr>	Specifies the IP address of the source.
<group-addr>	Specifies the IP address of the multicast group.
summary	Specifies a summary, if the multicast table is too long.

Displaying the PIM mode

Use this procedure to display the PIM mode. On the Ethernet Routing Switch 5000 Series, the only available mode is sparse mode (SM).

Procedure steps

Step	Action
1	To display the PIM mode, enter: `show ip pim mode`

--End--

Displaying multicast route information

Use this procedure to display multicast route information.

Procedure steps

Step	Action
1	To display multicast route information, enter:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
show ip mroute {interface | next-hop | route}
```

--End--

Variable definitions

The following table describes the command variables.

Variable	Value
interface	Displays multicast route information per interface.
next-hop	Displays next-hop information for the multicast routes
route	Displays multicast route information.

Job aids

The following table shows the field descriptions for the **show ip mroute interface** command.

Field	Description
Interface	Indicates the interface.
Ttl	Indicates the datagram TTL threshold for the interface. IP multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 0 means all multicast packets are forwarded out of the interface.
Protocol	Indicates the routing protocol running on this interface.

The following table shows the field descriptions for the **show ip mroute next-hop command** command.

Field	Description
Interface	Indicates the interface identity.
Group	Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.
Source	Indicates the network address, which when combined with the corresponding value of Scmask identifies the sources for which this entry specifies a next hop on an outgoing interface.
Srcmask	Indicates the network mask, which when combined with the corresponding value of Source identifies the sources for which this entry specifies a next hop on an outgoing interface.
Address	Indicates the address of the next hop specific to this entry. For most interfaces, this address is identical to Group.
State	Indicates whether the outgoing interface and next hop represented by this entry are currently forwarding IP datagrams. The value forwarding indicates the information is currently used; the value pruned indicates it is not used.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Exptime	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Closehop	Indicates the minimum number of hops between this router and members of this IP multicast group reached through this next hop on this outgoing interface. IP multicast datagrams for the group that use a TTL less than this number of hops are forwarded to the next hop
Protocol	Indicates the routing mechanism through which this next hop was learned.

The following table shows the field descriptions for the **`show ip mroute route`** command.

Field	Description
Group	Indicates the IP multicast group for which this entry specifies a next hop on an outgoing interface.
Source	Indicates the network address that, when combined with the corresponding value of Srcmask, identifies the sources for which this entry specifies a next hop on an outgoing interface.
Srcmask	Indicates the network mask that, when combined with the corresponding value of Source, identifies the sources for which this entry specifies a next hop on an outgoing interface.
Upstream_nbr	Indicates the address of the upstream neighbor from which IP datagrams from these sources to this multicast address are received, or 0.0.0.0 if the upstream neighbor is unknown.
If	Indicates the value of ifIndex for the interface on which IP datagrams sent by these sources to this multicast address are received. A value of 0 indicates that datagrams are not subject to an incoming interface check, but can be accepted on multiple interfaces (for example, in CBT).
Expir	Indicates the minimum amount of time remaining before this entry ages out. The value 0 indicates that the entry is not subject to aging.
Prot	Indicates the outgoing mechanism through which this route was learned.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

PIM-SM configuration example using NNCLI

The following figure shows a sample topology using PIM-SM.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Figure 47
PIM-SM sample topology

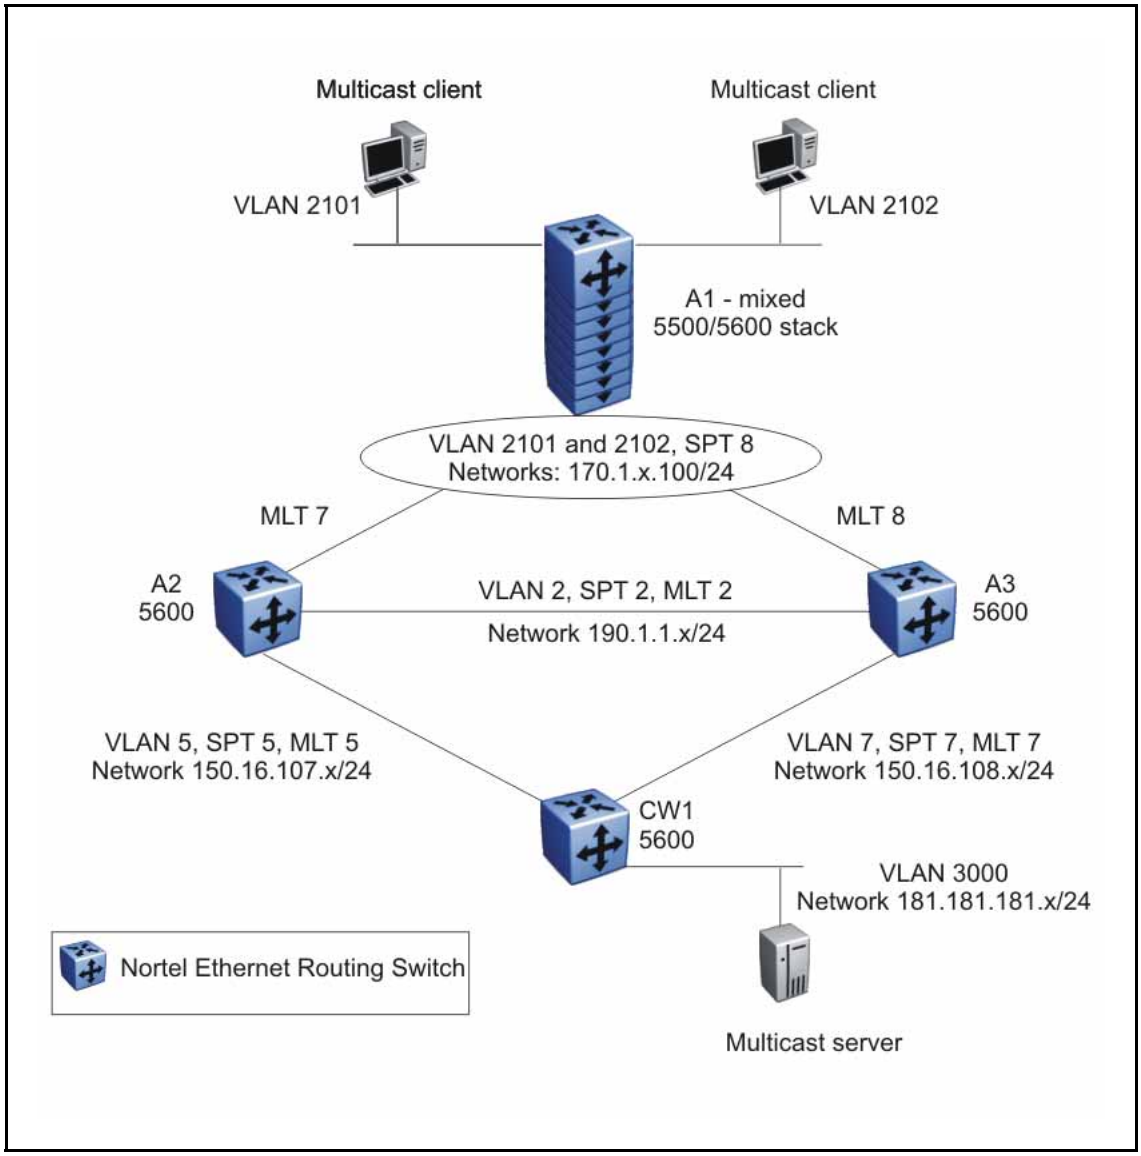

In this example, A1 is an 8-unit mixed stack of 5500 and 5600 Series switches running IGMPv2 snooping. A2, A3, and CW1 are all 5600 Series switches with PIM-SM enabled.

RIP is used as the Layer 3 routing protocol. You can also configure OSPF or static routes according to your network requirements. The PIM, MLT, VRRP, and IGMP settings provided remain unaffected by the choice of routing protocol.

The multicast group range is 224.10.10.0 255.255.255.0.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The STG, MLT, and VLAN number information are displayed in the topology diagram.

A1 description

A1 is an 8-unit mixed stack of 5500 and 5600 Series switches running IGMPv2 snooping. There are a total of two multicast clients on the access layer connected to the A1 stack, each in a different VLAN (2101 and 2102) and in a different network.

For simplicity, the configuration shows only two clients connected to the access layer stack. You can add more ports to each VLAN on the stack to have more users per VLAN.

A2 and A3 description

The distribution layer switches (A2 and A3) are configured as dynamic C-RPs or static RPs (configurations for both options are provided). You can use static RP or dynamic RP (but not both) in accordance with the requirements of your network. If you choose static RP, you must configure the same static RP on every PIM router in your network.

VRRP is enabled on A2 and A3, and all multicast clients have the VRRP virtual IP address as the default gateway for a specific VLAN.

ATTENTION
The VRRP configuration shown is an optional configuration providing a virtual IP for the host gateway. If your network does not need a virtual IP for a gateway, you do not need to configure VRRP. PIM-SM is independent of VRRP.

In this example, A3 is the DR for both PIM client VLANs (2101 and 2102), so all (S,G) entries install on A3. However, you can manage the DR election for the client VLANs by manipulating the IP address of the A2 and A3 VLAN interfaces. To load-share between A2 and A3, you can configure one of the VLAN interfaces on A2 (for example, 2101) with a higher IP address than the corresponding VLAN interface on A3. For the second VLAN, 2102, you can maintain the higher IP address on the A3 interface. In this way, A2 can become the DR for VLAN 2101, and A3 can remain the DR for VLAN 2102. This allows the (S,G) load to be split between the two switches and the system to be used to its maximum limits.

CW1 description

CW1 is configured as the BSR with priority 10 (only applicable to dynamic RP). A higher priority indicates a higher probability of being elected the BSR.

CW1 is directly connected to the multicast server. If desired, you can have a Layer 2 switch between the CW1 and the server with VLAN 3000 spanning through the switch to maintain the connection.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The CW1 connection to the multicast server is configured as a passive interface as it is on the edge and is not required to form a neighbor relationship with any other PIM router. You can configure this interface as an active interface according to the requirements of your network.

Link descriptions

The link connections (port numbers) between devices are listed below; the physical connections are in a one-to-one mapping in sequence as listed for each set of connections.

- A2 – A1:
  - 12,24,36,48,60,72,86,90 – 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2
  - MLT 7, VLAN 2101 to 2128, STG 8
- A3 – A1:
  - 2,14,26,38,50,62,74,80 -- 1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14
  - MLT 8, VLAN 2101 to 2128, STG 8
- A2 – A3:
  - 95,96 – 95,96
  - MLT 2, VLAN 2, STG 2
- A2 – CW1:
  - 91,92 –23,24
  - MLT 5, VLAN 5, STG 5
- A3 – CW1:
  - 91,92 – 21,22
  - MLT 7, VLAN 7, STG 7
- CW1 – Multicast server NIC:
  - 12 – Multicast server NIC
- A1 – Multicast client NICs:
  - VLAN 2101: 1/11 – MC1
  - VLAN 2102: 2/11 – MC2

See the following sections to configure the topology shown. In addition to the listed configurations, you can also configure the optional PIM-SM global and interface parameters, although it is advisable to leave these parameters at their default values.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

A1 configuration

The following procedure shows the configuration required for the A1 mixed stack running IGMP snooping.

Step	Action
1	Enter Global Configuration mode: `configure terminal`
2	Enable tagging on ports: `vlan port 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14 tagging enable`
3	Create the spanning tree instance: `spanning-tree stp 8 create`

4	Configure the VLANs:
	`vlan members remove 1 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14`
	
	`vlan create 2101 type port`
	`vlan members add 2101 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14,1/11`
	`spanning-tree stp 8`
	`add-vlan 2101`
	`int vlan 2101`
	`ip igmp snooping`
	`ip igmp mrouter 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14`
	
	`vlan create 2102 type port`
	`vlan members add 2102 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14,2/11`
	`spanning-tree stp 8 add-vlan 2102`
	`int vlan 2102`
	`ip igmp snooping`
	`ip igmp mrouter 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2,1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14`

5	Enable spanning tree: `spanning-tree 8 enable`
6	Configure the MLTs:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
mlt 7 member 1/2,2/14,3/14,4/38,5/12,6/14,7/2,8/2
mlt 7 enable
mlt 8 member 1/48,2/48,3/2,4/2,5/14,6/38,7/14,8/14
mlt 8 enable
```

--End--

A2 configuration

The following procedure shows the configuration required for the A2 PIM-SM-enabled distribution layer 5600 Series switch running VRRP and RIP.

Step	Action
1	Enter Global Configuration mode: `configure terminal`

2	Enable RIP and PIM:
	`ip routing`
	`router rip enable`
	`ip pim enable`

3	Enable tagging on ports: `vlan port 95-96,98,91-92,12,24,36,48,60,72,86,90 tagging enable`
4	Configure the VLANs:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
vlan members remove 1 12,24,36,48,60,72,86,90,95,96,91,
92

vlan create 2 type port
vlan members remove 1 95-96
vlan members add 2 95-96
interface vlan 2
ip address 190.1.1.2 255.255.255.0
ip pim en
ip rip en

vlan create 5 type port
vlan members remove 1 91-92
vlan members add 5 91-92
interface vlan 5
ip address 150.16.107.2 255.255.255.0
ip pim en
ip rip en

vlan create 2101 type port
vlan members add 2101 12,24,36,48,60,72,86,90
interface vlan 2101
ip address 170.1.1.1 255.255.255.0
ip pim en
ip rip en
```

5 Configure spanning tree:

```
spanning-tree stp 5 create
spanning-tree stp 5 add-vlan 5
spanning-tree stp 5 enable

spanning-tree stp 2 create
spanning-tree stp 2 add-vlan 2
spanning-tree stp 2 enable

spanning-tree stp 8 create
spanning-tree stp 8 add-vlan 2101
spanning-tree stp 8 add-vlan 2102
spanning-tree stp 8 enable
```

6 Configure the MLTs:

```
mlt 5 member 91-92
mlt 5 enable

mlt 2 member 95-96
mlt 2 enable

mlt 7 member 12,24,36,48,60,72,86,90
mlt 7 enable
```

7 Configure VRRP:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
router vrrp enable
interface vlan 2101
ip vrrp add 21 170.1.1.100
ip vrrp 21 enable

interface vlan 2102
ip vrrp add 22 170.1.2.100
ip vrrp 22 enable
```

8 For PIM-SM, configure a static RP:

```
ip pim static-rp enable
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.107.2
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.108.3
```

OR

configure a dynamic C-RP:

```
ip pim rp-candidate group 224.10.10.0 255.255.255.0 rp
150.16.107.2
```

--End--

A3 configuration

The following procedure shows the configuration required for the A3 PIM-SM-enabled distribution layer 5600 Series switch running VRRP and RIP.

Step	Action
1	Enter Global Configuration mode: `configure terminal`

2	Enable RIP and PIM:
	`ip routing`
	`router rip enable`
	`ip pim enable`

3	Enable tagging on ports: `vlan port 95-96,98,91-92,2,14,26,38,50,62,74,80 tagging ena`
4	Configure the VLANs:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
vlan members remove 1 2,14,26,38,50,62,74,80

vlan create 2 type port
vlan members remove 1 95-96
vlan members add 2 95-96
interface vlan 2
ip address 190.1.1.3 255.255.255.0
ip pim en
ip rip en

vlan create 7 type port
vlan members remove 1 91-92
vlan members add 7 91-92
interface vlan 7
ip address 150.16.108.3 255.255.255.0
ip pim en
ip rip en

vlan create 2101 type port
vlan members add 2101 2,14,26,38,50,62,74,80
interface vlan 2101
ip address 170.1.1.2 255.255.255.0
ip pim en
ip rip en

vlan create 2102 type port
vlan members add 2102 2,14,26,38,50,62,74,80,49
interface vlan 2102
ip address 170.1.2.2 255.255.255.0
ip pim en
ip rip en
```

5 Configure spanning tree:

```
spanning-tree stp 7 create
spanning-tree stp 7 add-vlan 7
spanning-tree stp 7 enable
spanning-tree stp 2 create
spanning-tree stp 2 add-vlan 2
spanning-tree stp 2 enable

spanning-tree stp 8 create
spanning-tree stp 8 add-vlan 2101
spanning-tree stp 8 add-vlan 2102
spanning-tree stp 8 enable
```

6 Configure the MLTs:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
mlt 7 member 91-92
mlt 7 enable

mlt 2 member 95-96
mlt 2 enable

mlt 8 member 2,14,26,38,50,62,74,80
mlt 8 enable
```

7 Configure VRRP:

```
router vrrp enable
interface vlan 2101
ip vrrp add 21 170.1.1.100
ip vrrp 21 enab

interface vlan 2102
ip vrrp add 22 170.1.2.100
ip vrrp 22 enab
```

8 For PIM-SM, configure a static RP:

```
ip pim static-rp enable
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.107.2
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.108.3
```

OR

configure a dynamic C-RP:

```
ip pim rp-candidate group 224.10.10.0 255.255.255.0 rp
150.16.108.3
```

--End--

CW1

The following procedure shows the configuration required for the CW1 PIM-SM-enabled 5600 Series switch running RIP. This is the source DR.

The following procedure shows the configuration required for the CW1 PIM-SM-enabled switch running RIP.

Step	Action
1	Enter Global Configuration mode: `configure terminal`
2	Enable RIP and PIM:

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

```
ip routing
router rip enable
ip pim enable
```

3 Enable tagging on ports:

```
vlan port 1-2,13-14,21-24,25,29,32 tagging ena
```

4 Configure the VLAN:

```
vlan mem remove 1 23,24,21,22,12

vlan create 5 type port
vlan members add 5 23-24
interface vlan 5
ip address 150.16.107.1 255.255.255.0
ip pim en
ip rip en

vlan create 7 type port
vlan members add 7 21-22
interface vlan 7
ip address 150.16.108.1 255.255.255.0
ip pim en ip rip en

!!  THE FOLLOWING VLAN IS A PASSIVE PIM VLAN (As it is
connected to the multicast server, and it does not need to
be part of PIM control messages.)
!!  IT CAN BE MADE ACTIVE AS PER YOUR NETWORK
vlan create 3000 type port
vlan members add 3000 12
interface vlan 3000
ip address 181.181.181.100 255.255.255.0
ip pim interface-type passive
ip pim en
ip rip en
```

5 Configure spanning tree:

```
spanning-tree stp 5 create
spanning-tree stp 5 add-vlan 5
spanning-tree stp 5 enable

spanning-tree stp 7 create
spanning-tree stp 7 add-vlan 7
spanning-tree stp 7 enable
```

6 Configure the MLTs:

```
mlt 5 member 23-24
mlt 5 enable

mlt 7 member 21-22
mlt 7 enable
```


Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

7 For PIM-SM, configure a static RP:

```
ip pim static-rp enable
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.107.2
ip pim static-rp 224.10.10.0 255.255.255.0 150.16.108.3
```

OR

for dynamic RP, configure the C-BSR:

```
interface vlan 5 ip pim bsr-candidate priority 10 exit
```

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IP routing configuration using Device Manager

This chapter describes the procedures you can use to configure routable VLANs using Device Manager.

The Nortel Ethernet Routing Switch 5000 Series are Layer 3 (L3) switches. This means that a regular Layer 2 VLAN becomes a routable Layer 3 VLAN if an IP address and MAC address are attached to the VLAN. When routing is enabled in Layer 3 mode, every Layer 3 VLAN is capable of routing as well as carrying the management traffic. You can use any Layer 3 VLAN instead of the Management VLAN to manage the switch.

IP routing configuration procedures

To configure IP routing on VLANs, perform the following steps:

Step	Action
1	Enable IP routing globally.
2	Assign an IP address to a specific VLAN.

--End--

In the above procedure, you are not required to enable IP routing as the first step. All IP routing parameters can be configured on the Nortel Ethernet Routing Switch 5000 Series before routing is actually enabled on the switch.

IP routing configuration navigation

- "Configuring global IP routing status and ARP lifetime" (page 358)
- "Configuring an IP address and enabling routing for a VLAN" (page 358)
- "Displaying configured IP Addresses" (page 359)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring global IP routing status and ARP lifetime

Use this procedure to enable and disable global routing at the switch level. By default, routing is disabled.

Step	Action
1	From the Device Manager menu, select**IP Routing > IP**. The IP dialog box appears with the Globals tab displayed.
2	To enable routing, select the **forwarding** option in the **Forwarding** text box.
3	To configure the ARP life time, modify the value in the **ARPLIfeTime** box.
4	Click **Apply**.

--End--

Variable definitions

The following table describes the Globals tab fields.

Field	Description
Forwarding	Indicates whether routing is enabled (forwarding) or disabled (nonforwarding) on the switch.
DefaultTTL	Indicates the default time-to-live (TTL) value for a routed packet. TTL is the maximum number of seconds elapsed before a packet is discarded. The value is inserted in the TTL field of the IP header of datagrams when one is not supplied by the transport layer protocol. Range is 1-255. Default value is 64 seconds.
ReasmTimeout	Indicates the maximum number of seconds that received fragments are held while they await reassembly at this entity. Default value is 60 seconds.
ARPLifeTime	Specifies the lifetime in minutes of an ARP entry within the system. Range is 5-360. Default is 360 minutes.

Configuring an IP address and enabling routing for a VLAN

Use this procedure to configure an IP address and enable routing for a VLAN.

Prerequisites

- Enable routing globally on the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	From the Device Manager menu, select **VLAN > VLANs**.
2	Select a VLAN.
3	Click **IP**. The IP, VLAN dialog box appears with the IP Address tab selected.
4	Click **Insert**. The Insert IP Address dialog box appears.
5	Type the IP address, subnet mask, and MAC address offset in the fields provided.
6	Click **Insert**.

--End--

Variable definitions

The following table describes the IP Address tab fields.

Field	Description
IpAddress	Specifies the IP address to associate with the selected VLAN.
NetMask	Specifies the subnet mask.
BcastAddrFormat	Specifies the IP broadcast address format used on this interface.
ReasmMaxSize	Specifies the size of the largest IP datagram which this entity can reassemble from fragmented incoming IP datagrams received on this interface.
VlanId	Specifies the VLAN ID. A value of -1 indicates that the VLAN ID is ignored.
MacOffset	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address. The valid range is 1-256. Specify the value 1 for the Management VLAN only. If no MAC offset is specified, the switch applies one automatically.
SecondaryIf	Indicates whether or not this entry corresponds to a secondary interface. If the value is **false**, then this is the primary IP address, if the value is **true**, then this is a secondary IP address. ***Note:*** You can assign 1 primary IP address and up to 8 secondary IP addresses to a VLAN.

Displaying configured IP Addresses

Use this procedure to display configured IP addresses on the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Step	Action
1	From the Device Manager menu, select **IP Routing > IP**.
2	Select the **Addresses** tab.

--End--

Variable definitions

The following table describes the Addresses tab fields.

Field	Description
IfIndex	Specifies the port number or VLAN ID.
IpAddress	Specifies the associated IP address.
NetMask	Specifies the subnet mask.
BcastAddrFormat	Specifies the format of the IP broadcast address.
ReasmMaxSize	Specifies the size of the largest IP datagram that this entity can reassemble from fragmented datagrams received on this interface.
VlanId	Specifies the VLAN ID number. A value of -1 indicates that the VLAN ID is ignored.
MacOffset	Specifies the value used to calculate the VLAN MAC address, which is offset from the switch MAC address.
SecondaryIf	Indicates whether or not this entry corresponds to a secondary interface. If the value is **false**, then this is the primary IP address, if the value is **true**, then this is a secondary IP address.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Static route configuration using Device Manager

Use the following procedure to configure static routes using Device Manager.

Static route configuration navigation

- "Configuring static routes" (page 361)
- "Displaying IP routes" (page 116)
- "Filtering route information" (page 363)
- "Displaying TCP information for the switch" (page 364)
- "Displaying TCP Connections" (page 365)
- "Displaying TCP Listeners" (page 365)
- "Displaying UDP endpoints" (page 366)

Configuring static routes

Use this procedure to configure static routes for the switch.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLANs to be routed.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > IP**.
2	Select the **Static Routes** tab.
3	Click **Insert**.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The Insert Static Routes dialog box appears.

4 In the fields provided, enter the information for the new static route.

5 Click **Insert**.

The new static route is displayed in the Static Routes tab.

--End--

Variable definitions

The following table describes the Static Routes tab fields.

Field	Description
Dest	Specifies the destination IP address of the route. 0.0.0.0 is considered the default route.
Mask	Specifies the destination mask of the route.
NextHop	Specifies the IP address of the next hop of this route.
Metric	Represents the cost of the static route. It is used to choose the best route (the one with the smallest cost) to a certain destination. The range is 1-65535. If this metric is not used, the value is set to -1.
IfIndex	Specifies the interface on which the static route is configured.
Enable	Specifies whether the route is administratively enabled (true) or disabled (false).
Status	Specifies the operational status of the route.

Displaying IP Routes

Use this procedure to display the different routes known to the switch.

Routes are not be displayed until at least one port in the VLAN has link.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > IP**.
2	Select the **Routes** tab.

--End--

Variable definitions

The following table describes the Routes tab fields.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Dest	Specifies the destination address of the route.
Mask	Specifies the subnet mask used by the route destination.
NextHop	Specifies the next hop in the listed route.
HopOrMetric	Specifies the RIP hop count, OSPF cost, or metric associated with the route.
Interface	Specifies the interface associated with the route.
Proto	Specifies the protocol associated with the route.
PathType	Specifies the route path type: • i: indirect • d: direct • A: alternative • B: best • E: ECMP • U: unresolved
Pref	Specifies the preference value associated with the route.

Filtering route information

Filter the Routes tab to display only the desired switch routes.

Procedure steps

Step	Action
1	With the **Routes** tab open, click **Filter**. The Filter dialog box is displayed.
2	Using the fields provided, set the filter for the tab.
3	Click **Filter**.

--End--

Variable definitions

The following table describes the Filter tab fields.

Field	Description
Condition	When using multiple filter expressions on the tab, this is the condition that is used to join them together.
Ignore Case	Denotes whether filters are case sensitive or insensitive.
Column	Denotes the type of criteria that will be applied to values used for filtering.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
All Records	Select this check box to clear any filters and display all rows.
Dest	Select this check box and enter a value to filter on the route destination value.
Mask	Select this check box and enter a value to filter on the route destination subnet mask value.
NextHop	Select this check box and enter a value to filter on the route next hop value.
HopOrMetric	Select this check box and enter a value to filter on the hop count or metric of the route.
Interface	Select this check box and enter a value to filter on the route's associated interface.
Proto	Select this check box and enter a value to filter on the route protocol.
PathType	Select this check box and enter a value to filter on the route path type.
Pref	Select this check box and enter a value to filter on the route preference value.

Displaying TCP information for the switch

Use this procedure to display Transmission Control Protocol (TCP) information for the switch.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > TCP/UDP**. The Ipv4TcpUdp dialog box appears with the TCP Globals tab displayed.

--End--

Variable definitions

The following table describes the TCP Globals tab fields.

Field	Description
RtoAlgorithm	Specifies the algorithm used to determine the timeout value used for retransmitting unacknowledged octets.
RtoMin	Specifies the minimum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
RtoMax	Specifies the maximum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds.
MaxConn	Specifies the limit on the total number of TCP connections that the entity can support. In entities where the maximum number of connections is dynamic, this object contains the value -1.

Displaying TCP Connections

Use this procedure to display information on the current TCP connections the switch maintains.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > TCP/UDP**.
2	Select the **TCP Connections** tab.

--End--

Variable definitions

The following table describes the TCP Connections tab fields.

Field	Description
LocalAddressType	Specifies the local IP address type for this TCP connection.
LocalAddress	Specifies the local IP address for this TCP connection. In the case of a connection in the listen state, which is willing to accept connections for any IP interface associated with the node, the value 0.0.0.0 is used.
LocalPort	Specifies the local port number for this TCP connection.
RemAddressType	Specifies the remote IP address type for this TCP connection.
RemAddress	Specifies the remote IP address for this TCP connection.
RemPort	Specifies the remote port number for this TCP connection.
State	Specifies the state of this TCP connection.

Displaying TCP Listeners

Use this procedure to display information on the current TCP listeners on the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > TCP/UDP**.
2	Select the **TCP Listeners** tab.

--End--

Variable definitions

The following table describes the TCP Listeners tab fields.

Field	Description
LocalAddressType	Specifies the IP address type of the local TCP listener.
LocalAddress	Specifies the local IP address of the TCP listener. The value of this field can be represented in three possible ways, depending on the characteristics of the listening application: 1. For an application willing to accept both IPv4 and IPv6 datagrams, the value of this object is a zero-length octet string, and the value of the corresponding LocalAddressType field is unknown. 2. For an application willing to accept either IPv4 or IPv6 datagrams, the value of this object must be 0.0.0.0 or ::, with the LocalAddressType identifying the supported address type. 3. For an application that is listening for data destined only to a specific IP address, the value of this object is the specific local address, with LocalAddressType identifying the supported address type.
LocalPort	Specifies the local port number for this TCP connection

Displaying UDP endpoints

Use this procedure to display information on the UDP endpoints currently maintained by the switch.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > TCP/UDP**.
2	Select the **UDP Endpoints** tab.
3	Click **Refresh** to immediately refresh the information displayed.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the UDP Listeners tab fields.

Field	Description
LocalAddressType	Specifies the local address type (IPv6 or IPv4).
LocalAddress	Specifies he local IP address for this UDP listener. In the case of a UDP listener that accepts datagrams for any IP interface associated with the node, the value 0.0.0.0 is used. 1. For an application willing to accept both IPv4 and IPv6 datagrams, the value of this object is a zero-length octet string, and the value of the corresponding LocalAddressType field is unknown. 2. For an application willing to accept either IPv4 or IPv6 datagrams, the value of this object must be 0.0.0.0 or ::, with the LocalAddressType identifying the supported address type. 3. For an application that is listening for data destined only to a specific IP address, the value of this object is the address for which this node is receiving packets, with LocalAddressType identifying the supported address type.
LocalPort	Specifies the local port number for this UDP listener.
RemoteAddressType	Displays the remote address type (IPv6 or IPv4).
RemoteAddress	Displays the remote IP address for this UDP endpoint. If datagrams from all remote systems are to be accepted, this value is a zero-length octet string. Otherwise, the address of the remote system from which datagrams are to be accepted (or to which all datagrams are to be sent) is displayed with the RemoteAddressType identifying the supported address type.
RemotePort	Displays the remote port number. If datagrams from all remote systems are to be accepted, this value is zero.
Instance	Distinguishes between multiple processes connected to the same UDP endpoint.
Process	Displays the ID for the UDP process.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Brouter port configuration using Device Manager

A brouter port is a single-port VLAN that can route IP packets as well as bridge all non-routable traffic. The following section describes the procedures necessary to configure and manage brouter ports on the Nortel Ethernet Routing Switch 5000 Series using Device Manager.

In Device Manager, brouter ports are treated as routable VLANs and are displayed on the **Basic** tab of the **VLANs** dialog box.

Use the following procedure to configure a brouter port using Device Manager.

Configuring a brouter port

Use this procedure to configure and manage brouter ports.

Procedure steps

Step	Action
1	Select a port from the Device Manager front panel.
2	From the main menu, select **Edit > Port**.
	The Port dialog box appears with the Interface tab selected.
3	Select the **IP Address** tab.
4	Click **Insert**.
	The **Insert IP Address** dialog appears.
5	Using the provided fields, create the new brouter port.
6	Click **Insert**.

--End--

Variable definitions

The following table describes the IP Address tab fields.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Table 12
IP Address tab fields

Field	Description
IpAddress	Specifies the IP address assigned to this brouter.
NetMask	Specifies the subnet mask associated with the brouter IP address.
BcastAddrFormat	Specifies the IP broadcast address format used on this interface.
ReasmMaxSize	Specifies the size of the largest IP datagram which this entity can reassemble from incoming IP fragmented datagrams received on this interface.
VlanId	Specifies the VLAN ID associated with this brouter port.
MacOffset	Specifies the MAC address offset associated with this brouter port.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

OSPF configuration using Device Manager

This chapter describes the procedures you can use to configure OSPF using Device Manager.

The 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 and the tagging of externally-derived routing information.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Assign an IP address to the VLAN that you want to enable with OSPF.

  Routing is automatically enabled on the VLAN when you assign an IP address to it.

OSPF configuration navigation

- "Configuring Global OSPF properties" (page 372)
- "Configuring an OSPF area" (page 374)
- "Configuring an area aggregate range" (page 375)
- "Configuring OSPF stub area metrics" (page 376)
- "Configuring OSPF interfaces" (page 377)
- "Configuring OSPF interface metrics" (page 379)
- "Defining MD5 keys for OSPF interfaces" (page 379)
- "Displaying OSPF neighbor information" (page 380)
- "Configuring an OSPF virtual link" (page 381)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- "Configuring an automatic virtual link" (page 383)
- "Defining MD5 keys for OSPF virtual links" (page 383)
- "Displaying virtual neighbor information" (page 384)
- "Configuring OSPF host routes" (page 385)
- "Displaying link state database information" (page 385)
- "Displaying external link state database information" (page 386)
- "Displaying OSPF statistics" (page 387)
- "Displaying VLAN OSPF statistics" (page 388)

Configuring Global OSPF properties

Use the following procedure to configure global OSPF parameters.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Using the fields provided, configure the global OSPF parameters.
3	Click **Apply**.

--End--

Variable definitions

The following table describes the General tab fields.

Field	Description
RouterId	Specifies the unique ID of the router in the Autonomous System.
AdminStat	Specifies the administrative status of OSPF on the router.
VersionNumber	Specifies the version of OSPF running on the router.
AreaBrdRtrStatus	Specifies whether this router is an Area Border Router.
ASBrdRtrStatus	Specifies whether this router is an Autonomous System Border Router.
ExternLsaCount	Specifies the number of external (link state type 5) link-state advertisements in the link state database.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
ExternLsaCksumSum	Specifies the sum of the link state checksums of the external link state advertisements contained in the link state database. This sum is used to determine if the link state database of the router changes and to compare the link state databases of two routers.
OriginateNewLsas	Specifies the number of new link state advertisements that have been originated. This number is incremented each time the router originates a new link state advertisement.
RxNewLsas	Specifies the number of link state advertisements received determined to be new instantiations. This number does not include newer instantiations of self-originated link state advertisements.
10MbpsPortDefaultMetric	Specifies the default metric of a 10 Mbps port. This is an integer value between 1 and 65535. Default value is 100.
100MbpsPortDefaultMetric	Specifies the default metric of a 100 Mbps port. This is an integer value between 1 and 65535. Default value is 10.
1000MbpsPortDefaultMetric	Specifies the default metric of a 1000 Mbps port. This is an integer value between 1 and 65535. Default value is 1.
10000MbpsPortDefault Metric	Specifies the default metric of a 10000 Mbps port. This is an integer value between 1 and 65535. Default value is 1.
TrapEnable	Specifies whether OSPF traps are enabled. The default setting is disabled.
AutoVirtLinkEnable	Specifies the status of OSPF automatic Virtual Link creation. The default setting is disabled.
SpfHoldDownTime	Specifies the SPF Hold Down Timer value, which is an integer between 3 and 60. The default value is 10. The SPF runs, at most, once per hold down timer value.
OspfAction	Specifies an immediate OSPF action to take. Select **runSpf** and click **Apply** to initiate an immediate SPF run.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Rfc1583Compatibility	Controls the preference rules used when choosing among multiple Autonomous System external link state advertisements advertising the same destination. When this is enabled, the preference rule will be the same as specified by RFC 1583. When disabled, the new preference rule, as described in RFC 2328, will be applicable. This potentially prevents the routing loops when Autonomous System external link state advertisements for the same destination have been originated from different areas.
LastSpfRun	Specifies the time the last SPF calculation was done.

Configuring an OSPF area

Use this procedure to configure an OSPF area.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Areas** tab.
3	Click **Insert**. The Insert Areas dialog box appears.
4	Using the fields provided, configure the OSPF area.
5	Click **Insert**.

--End--

Variable definitions

The following table describes the Areas tab fields.

Field	Description
AreaId	Specifies the unique identifier for the area. Area ID *0.0.0.0* is used for the OSPF backbone.
ImportAsExtern	Specifies the area type by defining its support for importing Autonomous System external link state advertisements. The options available are: • importExternal: specifies a normal area • importNoExternal: specifies a stub area • importNssa: specifies an NSSA

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
SpfRuns	Specifies the number of times that the OSPF intra-area route table has been calculated using this area link state database.
AreaBdrRtrCount	Specifies the total number of Area Border Routers reachable within this area. This is initially zero and is calculated in each SPF pass.
AsBdrRtrCount	Specifies the total number of Autonomous System Border Routers reachable within this area. This is initially zero, and is calculated in each SPF pass.
AreaLsaCount	Specifies the total number of link state advertisements in this area's link state database, excluding Autonomous System external link state advertisements.
AreaLsaCksumSum	Specifies the sum of the link state advertisements' checksums contained in this area's link state database. This sum excludes external (link state type 5) link state advertisements. The sum can be 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	Controls the import of summary link state advertisements on an ABR into a stub area. It has no effect on other areas. If the value is noAreaSummary, the ABR neither originates nor propagates summary link state advertisements into the stub area (creating a totally stubby area). If the value is sendAreaSummary, the ABR both summarizes and propagates summary link state advertisements.

Configuring an area aggregate range

Configure OSPF area aggregate ranges to reduce the number of link state advertisements that are required within the area. You can also control advertisements.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Area Aggregate** tab.
3	Click **Insert**.
4	Using the fields provided, create the new area aggregate.
5	Click **Insert**.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the Area Aggregate tab fields.

Field	Description
AreaID	Specifies the unique identifier of the Area this address aggregate is found in.
LsdbType	Specifies the type of address aggregate. This field specifies the link state database type that this address aggregate applies to. Options available are: summaryLink or nssaExternalLink.
IpAddress	Specifies the IP address of the network or subnetwork indicated by the aggregate range.
Mask	Specifies the subnet mask that pertains to the network or subnetwork.
Effect	Specifies the aggregates effect. Subnets subsumed by aggregate ranges either trigger the advertisement of the indicated aggregate (*advertiseMatching* value) or result in the subnet not being advertised at all outside the area. Select one of the following types: • AdvertiseMatching: advertises the aggregate summary LSA with the same LSID • DoNotAdvertiseMatching: suppresses all networks that fall within the entire range • AdvertiseDoNotAggregate: advertises individual networks
AdvertiseMetric	Specifies the advertisement metric associated with this aggregate. Enter an integer value between 0 and 65535 which represents the Metric cost value for the OSPF area range.

Configuring OSPF stub area metrics

Use this procedure to view the set of metrics that are advertised by a default area border router into a stub area to determine if you wish to accept the current values or configure new ones.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Stub Area Metrics** tab.
3	Using the fields provided, configure the stub area metrics.
4	Click **Apply**.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the Stub Area Metrics tab fields.

Field	Description
AreaId	Specifies the unique ID of the stub area.
TOS	Specifies the Type of Service associated with the metric.
Metric	Specifies the metric value applied to the indicated type of service. By default, this equals the least metric at the type of service among the interfaces to other areas.
Status	Displays the status of the entry; **Active** or **Not Active**. This field is read-only.

Configuring OSPF interfaces

Use this procedure to configure OSPF interfaces.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Interfaces** tab.
3	Using the fields provided, configure the OSPF interface.
4	Click **Apply**.

--End--

Variable definitions

The following table describes the Interfaces tab fields.

Field	Description
IpAddress	Specifies the IP address of the OSPF interface.
AreaId	Specifies the unique ID of the area to which the interface connects. Area ID 0.0.0.0 indicates the OSPF backbone.
AdminStat	Specifies the administrative status of the OSPF interface.
State	Specifies the DR state of the OSPF interface: up–DR, BDR, OtherDR; down–down, waiting.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
RtrPriority	In multi-access networks, specifies the priority of the interface in the designated router election algorithm. The interface with the highest priority number is the designated router. The interface with the second-highest priority becomes the backup designated router. The value 0 signifies that the router is not eligible to become the designated router on this network. This is an integer value between 0 and 255. In the event of a tie in the priority value, routers use their Router ID as a tie breaker. The default value is 1.
DesignatedRouter	Specifies the IP address of the Designated Router.
BackupDesignatedRouter	Specifies the IP address of the Backup Designated Router.
Type	Specifies the OSPF interface type. The options available are: broadcast or passive.
AuthType	Specifies the interface authentication type. The options available are: none, simplePassword, or md5.
AuthKey	Specifies the interface authentication key. This key is used when AuthType is simplePassword.
PrimaryMd5Key	Specifies the MD5 primary key if it exists. Otherwise this field displays 0. This key is used when AuthType is md5.
TransitDelay	Specifies the estimated number of seconds it takes to transmit a link state update packet over this interface. This is an integer value between 0 and 3600.
RetransInterval	Specifies the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. This value is also used when retransmitting database description and link state request packets. This is an integer value between 0 and 3600.
HelloInterval	Specifies the interval in seconds between the Hello packets sent by the router on this interface. This value must be the same for all routers attached to a common network. This is an integer value between 1 and 65535.
RtrDeadInterval	Specifies the number of seconds that a neighbor waits for a Hello packet from this interface before the router neighbors declare it down. This value must be some multiple of the Hello interval and must be the same for all routers attached to the common network. This is an integer value between 0 and 2147483647.
PollInterval	Specifies the poll interval.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
AdvertiseWhenDown	Specifies whether this interface sends advertisements even when it is non-operational.
MtuIgnore	Specifies whether the MTU value is ignored on this interface.
Events	Specifies the number of times this OSPF interface has changed its state or an error has occurred.

Configuring OSPF interface metrics

Use this procedure to configure OSPF interface metrics.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**. The OSPF dialog box appears.
2	Select the **If Metrics** tab.
3	Using the fields provided, configure the interface metrics.
4	Click **Apply**.

--End--

Variable definitions

The following table describes the If Metrics tab fields.

Table 13
If Metrics tab fields

Field	Description
IpAddress	Specifies the IP address of the interface.
TOS	Specifies the Type of Service associated with the metric.
Value	Specifies the value advertised to other areas indicating the distance from the OSPF router to any network in the range. This is an integer value between 0 and 65535.
Status	Displays the status of the entry; **Active** or **Not Active**. This field is read-only.

Defining MD5 keys for OSPF interfaces

Use this procedure to configure OSPF MD5 keys for OSPF interfaces.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Message Digest** tab.
3	Click **Insert**. The Insert Message Digest dialog box appears.
4	Using the fields provided, create the new digest entry.
5	Click **Insert**.

--End--

Variable definitions

The following table describes the Message Digest tab fields.

Field	Description
IpAddress	Specifies the IP address of the OSPF interface associated with the digest entry.
Index	Specifies an index value for the digest entry. This is an integer value between 1 and 255.
Type	Specifies the type of digest entry. Only MD5 is supported.
Key	Specifies the key value associated with the digest entry.

Displaying OSPF neighbor information

Use this procedure to display OSPF neighbors.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Neighbors** tab.
3	Click **Refresh** to update the displayed information.

--End--

Variable definitions

The following table describes the Neighbor tab fields.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
IpAddr	Specifies the IP address this neighbor is using as an IP source address. On addressless links, this will not be represented as *0.0.0.0* but as the address of another of the neighbor's interfaces.
AddressLessIndex	Specifies the corresponding value of the interface index on addressless links. This value is zero for interfaces having an IP address.
RouterId	Specifies the unique ID of the neighboring router in the Autonomous System.
Options	Specifies a value corresponding to the neighbor's Options field.
Priority	Specifies the priority of the neighbor in the designated router election algorithm. A value of 0 indicates that the neighbor is not eligible to become the designated router on this particular network. This is a value between 0 and 255.
State	Specifies the state of the relationship with this neighbor.
Events	Specifies the number of times this neighbor relationship has changed state or an error has occurred.
RetransmissionQueueLength	Specifies the current length of the retransmission queue.
NbmaNbrPermanence	Specifies the status of the entry. The values *dynamic* and *permanent* refer to how the neighbor came to be known.
HelloSuppressed	Specifies whether Hello packets are being suppressed to the neighbor.
InterfaceAddr	Specifies the neighbor's interface address.

Configuring an OSPF virtual link

Use the following procedure to create an OSPF virtual link.

Procedure steps

Step	Action
1	From the Device Manager main menu, select **IP Routing > OSPF**.
2	Select the **Virtual If** tab.
3	Click **Insert**. The OSPF, Insert Virtual If dialog box appears.
4	To configure a virtual interface, enter the information in the fields provided.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

5 Click **Insert**.

--End--

Variable definitions

The following table describes the Virtual If tab fields.

Field	Description
AreaId	Specifies the unique ID of the area connected to the interface. An area ID of 0.0.0.0 indicates the OSPF backbone.
Neighbor	Specifies the router ID of the virtual neighbor.
TransitDelay	Specifies the estimated number of seconds required to transmit a link state update packet over the virtual interface. The transit delay is expressed as an integer between 1 and 3600. The default value is 1.
RetransInterval	Specifies the number of seconds between link state advertisement retransmissions for adjacencies belonging to the virtual interface. The retransmit interval is also used to transmit database description and link state request packets. The retransmit interval is expressed as an integer between 1 and 3600. The default value is 5.
HelloInterval	Specifies the interval, in seconds, between the Hello packets sent by the router on the virtual interface. This value must be the same for all routers attached to a common network. The hello interval is expressed as an integer between 1 and 65535. The default value is 10.
RtrDeadInterval	Specifies the number of seconds that a neighbor router waits to receive transmitted hello packets from this interface before the neighbor declares it down. The retransmit dead interval is expressed as an integer between 1 and 2147483647. The retransmit dead interval must be a multiple of the hello interval and must be the same for all routers attached to a common network. The default value is 60.
AuthType	Specifies the interface authentication type. The available authentication types are: none, simplePassword, or MD5.
AuthKey	Specifies the interface authentication key used with the simplePassword authentication type.
PrimaryMd5Key	Specifies the MD5 primary key. If no MD5 primary key exists, the value in this field is 0.
State	Specifies the OSPF virtual interface state.
Events	Specifies the number of times the virtual interface has changed state or the number of times an error has occurred.
Type	Specifies whether the virtual interface is broadcast or passive.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring an automatic virtual link

Configure an automatic virtual link to provide an automatic, dynamic backup link for vital OSPF traffic. Automatic virtual links require more system resources than manually-configured virtual links.

Automatic Virtual Links are removed when the transit area is deleted or when the router is no longer an ABR.

Procedure steps

Step	Action
1	From the Device Manager main menu, select **IP Routing > OSPF**.
2	On the General tab, enter the router ID in the **RouterId** box for one of the end point ABRs.
3	Check the **AutoVirtLinkEnable** box.
4	Click **Apply**.
5	On the remote ABR to use for the virtual link, repeat the preceding steps.

--End--

Defining MD5 keys for OSPF virtual links

Use this procedure to configure OSPF MD5 keys for OSPF virtual interfaces.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Virtual If Message Digest** tab.
3	Click **Insert**. The OSPF, Insert Virtual If Message Digest window appears.
4	Enter the information in the fields provided.
5	Click **Insert**.

--End--

Variable definitions

The following table describes the Virtual If Message Digest tab fields.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
AreaId	Specifies the area ID of the area associated with the virtual interface.
Neighbor	Specifies the IP address of the neighbor router associated with the virtual interface.
Index	Specifies the index value of the virtual interface message digest entry. The value is an integer between 1 and 255.
Type	Specifies the type of digest entry. Only MD5 is supported.
Key	Specifies the key value associated with the digest entry.

Displaying virtual neighbor information

Use this procedure to view OSPF Virtual Neighbors information.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Virtual Neighbors** tab. The Virtual Neighbors tab appears.
3	Click one of the labelled buttons across the bottom of the tab to refresh the information in the view, copy settings, or print.

--End--

Variable definitions

The following table describes the Virtual Neighbors tab fields.

Field	Description
Area	Specifies the subnetwork in which the virtual neighbor resides.
RouterId	Specifies the 32-bit integer uniquely identifying the neighboring router in the autonomous system.
IpAddr	Specifies the IP address of the virtual neighboring router.
Options	Specifies a bit mask corresponding to the option field of the neighbor.
State	Specifies the state of the virtual neighbor relationship.
Events	Specifies the number of state changes or error events that have occurred between the OSPF router and the neighbor router.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
RetransQLen	Specifies the current length of the retransmission queue (the number of elapsed seconds between advertising retransmissions of the same packet to a neighbor).
HelloSuppressed	Specifies whether Hello packets to the virtual neighbor are suppressed or not.

Configuring OSPF host routes

Create OSPF hosts routes to specify which hosts are directly attached to the router and the metrics that must be advertised for them.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Hosts** tab.
3	Click **Insert**.
	The OSPF Insert Hosts dialog box appears.
4	Enter the information on the **Insert Hosts** dialog box.
5	Click **Insert**.

--End--

Variable definitions

The following table describes the Hosts tab fields.

Field	Description
IpAddress	Specifies the host IP address.
TOS	Specifies the configured route type of service. The value in this field should be 0 as TOS-based routing is not supported.
Metric	Specifies the configured cost of the host.
AreaID	Specifies the ID of the area connected to the host.

Displaying link state database information

Use this procedure to view OSPF link states.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Link State Database** tab.
3	Click **Refresh** to update the displayed information.

--End--

Variable definitions

The following table describes the Link State Database tab fields.

Field	Description
AreaId	Specifies the unique identifier of the Area from which the link state advertisement was received.
Type	Specifies the type of link state advertisement. Each link state type has a separate advertisement format.
Lsid	Specifies the Link State ID, a link state type-specific field containing either a Router ID or an IP address. This field identifies the section of the routing domain that is being described by the advertisement.
RouterId	Specifies the unique identifier of the originating router in the Autonomous System.
Sequence	This field is used to detect old or duplicate link state advertisements by assigning an incremental number to duplicate advertisements. The higher the sequence number, the more recent the advertisement.
Age	Specifies the age of the link state advertisement in seconds.
Checksum	Specifies the checksum of the complete content of the advertisement, excluding the *Age* field. This field is excluded so that the advertisement's age can be increased without updating the checksum. The checksum used is the same as that used in ISO connectionless datagrams and is commonly referred to as the *Fletcher checksum*.

Displaying external link state database information

Use this procedure to view the OSPF external link state database.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Ext. Link State Database** tab.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

3 Click **Refresh** to update the displayed information.

--End--

Variable definitions

The following table describes the Ext. Link State Database tab fields.

Field	Description
Type	Specifies the type of link state advertisement. Each link state type has a separate advertisement format.
Lsid	Specifies the Link State ID, a link state type-specific field containing either a Router ID or an IP address. This field identifies the section of the routing domain that is being described by the advertisement.
RouterId	Specifies the unique identifier of the originating router in the Autonomous System.
Sequence	This field is used to detect old or duplicate link state advertisements by assigning an incremental number to duplicate advertisements. The higher the sequence number, the more recent the advertisement.
Age	Specifies the age of the link state advertisement in seconds.
Checksum	Specifies the checksum of the complete content of the advertisement, excluding the *Age* field. This field is excluded so that the advertisement's age can be increased without updating the checksum. The checksum used is the same as that used in ISO connectionless datagrams and is commonly referred to as the *Fletcher checksum*.
Advertisement	Specifies the hexadecimal representation of the entire link state advertisement including the header.

Displaying OSPF statistics

Use this procedure to display OSPF statistics.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Stats** tab.
3	Values on the Stats tab will refresh automatically based on the value selected in the Poll Interval field. To clear the counters and start over at zero, click **Clear Counters**.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the Stats tab fields.

Field	Description
LsdbTblSize	Indicates the number of entries in the link state database.
TxPackets	Indicates the number of packets transmitted by OSPF.
RxPackets	Indicates the number of packets received by OSPF.
TxDropPackets	Indicates the number of packets dropped by OSPF before transmission.
RxDropPackets	Indicates the number of packets dropped before receipt by OSPF.
RxBadPackets	Indicates the number of bad packets received by OSPF.
SpfRuns	Indicates the total number of SPF calculations performed. This also includes the number of partial route table calculations.
BuffersAllocated	Indicates the total number of buffers allocated for OSPF.
BuffersFreed	Indicates the total number of buffers that are freed by OSPF.
BufferAllocFailures	Indicates the number of times that OSPF has failed to allocate buffers.
BufferFreeFailures	Indicates the number of times that OSPF has failed to free buffers.

Displaying VLAN OSPF statistics

Use this procedure to view VLAN OSPF statistical information on a per-interface basis.

Procedure steps

Step	Action
1	From the Device Manager menu, select **VLAN > VLANs**.
2	Select an interface from those listed on the **Basic** tab.
3	Click the **IP** button.
4	On the **IP VLAN** dialog box, select the **OSPF Stats** tab.
5	To graph the statistical information, select the desired data and click the appropriate graph button at the bottom of the screen.

--End--

Variable definitions

The following table describes the Stats tab fields.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
VersionMismatches	Specifies the number of version mismatches received by this interface.
AreaMismatches	Specifies the number of area mismatches received by this interface.
AuthTypeMismatches	Specifies the number of AuthType mismatches received by this interface.
AuthFailures	Specifies the number of authentication failures on this interface.
NetMaskMismatches	Specifies the number of net mask mismatches received by this interface.
HelloIntervalMismatches	Specifies the number of hello interval mismatches received by this interface.
DeadIntervalMismatches	Specifies the number of dead interval mismatches received by this interface.
OptionMismatches	Specifies the number of option mismatches received by this interface.
RxHellos	Specifies the number of hello packets received by this interface.
RxDBDescrs	Specifies the number of database descriptor packets received by this interface.
RxLSUpdates	Specifies the number of link state update packets received by this interface.
RxLSReqs	Specifies the number of link state request packets received by this interface.
RxLSAcks	Specifies the number of link state acknowledge packets received by this interface.
TxHellos	Specifies the number hello packets transmitted by this interface.
TxDBDescrs	Specifies the number of database descriptor packets transmitted by this interface.
TxLSUpdates	Specifies the number of link state update packets transmitted by this interface.
TxLSReqs	Specifies the number of link state request packets transmitted by this interface.
TxLSAcks	Specifies the number of link state acknowledge packets transmitted by this interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

RIP configuration using Device Manager

This section describes the Device Manager procedures used to configure and manage the Routing Information Protocol (RIP) on the Nortel Ethernet Routing Switch 5000 Series . RIP is a distance vector protocol used to dynamically discover network routes based on information passed between routers in the network. RIP is useful in network environments where using static route administration would be difficult.

Prerequisites

- Enable IP routing globally.
- Assign an IP address to the VLAN or brouter port that you want to enable with RIP.

  Routing is automatically enabled on the VLAN when you assign an IP address to it.

RIP configuration procedures

To configure RIP routing on the Ethernet Routing Switch, perform the following steps:

Procedure steps

Step	Action
1	Enable RIP globally.
2	Configure global RIP properties as required.
3	Enable RIP on the desired VLAN or brouter interfaces.
4	Configure interface RIP properties as required.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

RIP configuration navigation

- "Configuring Global RIP properties" (page 392)
- "Configuring a RIP interface" (page 393)
- "Configuring advanced RIP interface properties" (page 394)
- "Displaying RIP Statistics" (page 395)
- "Configuring RIP parameters for a VLAN" (page 396)

Configuring Global RIP properties

Use this procedure to configure global RIP parameters.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > RIP**.
2	Using the fields provided, configure the global RIP parameters.
3	Click **Apply**.

--End--

Variable definitions

The following table describes the Globals tab fields.

Field	Definition
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.
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.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Definition
TimeOutInterval	Specifies the global timeout interval parameter. If a RIP router does not receive an update from another RIP router within the configured timeout period, it moves the routes advertised by the nonupdating router to the garbage list. The timeout interval must be greater than the update timer. Range is 15–259200 seconds. Default is 180 seconds.
DefImportMetric	Sets the value of the default import metric applied to routes imported the RIP domain. For announcing OSPF internal routes into a RIP domain, if the policy does not specify a metric value, the default import metric is used. For OSPF external routes, the external cost is used.

Configuring a RIP interface

To configure a RIP interface to tailor RIP to the individual interfaces.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > RIP**.
2	Select the **Interface** tab.
3	Using the fields provided, configure the interface.
4	Click **Apply**.

--End--

Variable definitions

The following table describes the Interface tab fields.

Field	Definition
Address	The IP address of the RIP interface. This field is for organizational purposes only and cannot be edited.
Send	Sets the RIP version sent on this interface. The following values are valid: • doNotSend - No RIP updates sent on this interface. • ripVersion1 - RIP updates compliant with RFC 1058. • rip1Compatible - Broadcasts RIPv2 updates using RFC 1058 route subsumption rules. • ripVersion2 - Multicasting RIPv2 updates.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Definition
	The default is rip1Compatible.
Receive	Sets the RIP version received on this interface: rip1, rip2, or rip1OrRip2 The default is rip1OrRip2. Note that rip2 and rip1OrRip2 imply reception of multicast packets.

Configuring advanced RIP interface properties

Configure advanced RIP interface properties to fine tune and further configure a RIP interface.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > RIP**.
2	Select the **Interface Advance** tab.
3	Using the fields provided, configure the advanced RIP features.
4	Click **Apply**.

--End--

Variable definitions

The following table describes the Interface Advance tab fields.

Field	Definition
Address	The IP address of the RIP interface. This field is for organizational purposes only and cannot be edited.
Interface	The switch interface that corresponds to the listed IP address.
Enable	Enables or disables RIP on this interface.
Supply	Determines whether this interface supplies RIP advertisements.
Listen	Determines whether this interface listens for RIP advertisements.
Poison	Enables or disables poison reverse on this interface.
DefaultSupply	Determines whether this interface advertises default routes.
DefaultListen	Determines whether this interface listens for default route advertisements.
TriggeredUpdate	Enables or disables triggered updates on this interface.
AutoAggregate	Enables or disables auto aggregation on this interface.
InPolicy	Associates a previously configured switch policy with this interface for use as an in policy.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Definition
OutPolicy	Associates a previously configured switch policy with this interface for use as an out policy.
Cost	The cost associated with this interface.
HoldDownTime	Sets the holddown timer for this interface. This is an integer value in seconds between 0 and 360.
TimeoutInterval	Sets the timeout interval for this interface. This is an integer value between 15 and 259200.
ProxyAnnounceFlag	Enables or disables proxy announcements on this interface.

Displaying RIP Statistics

Use this procedure to view RIP statistics.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > RIP**.
2	Select the **Stats** tab. RIP statistics for the configured interfaces are displayed.
3	To graph these statistics, select a row from the **Stats** tab and click the **Graph** button. A new dialog box appears with the statistics for the selected interface.
4	Select a graph type by clicking the appropriate graphing button.

--End--

Variable definitions

The following table describes the RIP—Stats tab fields.

Field	Definition
Address	The RIP interface address.
RcvBadPackets	The number of RIP response packets received by the interface that have been discarded.
RcvBadRoutes	The number of RIP routes received by the interface that have been ignored.
SentUpdates	The number of triggered RIP updates actually sent on this interface. This does not include full updates sent containing new information.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring RIP parameters for a VLAN

Use this procedure to configure VLAN RIP parameters.

Procedure steps

Step	Action
1	From the Device Manager menu, select **VLAN > VLANs**.
2	On the **Basic** tab, select an interface and click the **IP** button. The IP VLAN dialog box appears with the IP Address tab selected.
3	Select the **RIP** tab.
4	Using the provided fields, configure the interface RIP parameters.
5	Click **Apply**.

--End--

Variable definitions

The following table describes the RIP tab fields.

Field	Definition
Poison	Determines whether or not poison reverse is implemented on this interface.
DefaultSupply	Determines whether or not the interface implements the default supply mechanism.
DefaultListen	Determines whether or not the interface implements the default listen mechanism.
AutoAggregateEnable	Determines whether or not auto aggregation is enabled on this interface.
AdvertiseWhenDown	Determines whether or not this interface will advertise even when non-operational.
Cost	The cost associated with this interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

ECMP configuration using Device Manager

This section describes the procedure you can use to configure ECMP using Device Manager.

The Equal Cost MultiPath (ECMP) feature allows routers to determine 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.

ATTENTION
ECMP is not supported on the Nortel Ethernet Routing Switch 5510. ECMP works in a mixed stack but cannot run on any Nortel Ethernet Routing Switch 5510 units in the stack.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Configure routing (RIP, OSPF, or static routes) on the switch.

Configuring ECMP

Use the following procedure to configure and ECMP settings for RIP, OSPF, and static routes.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > IP**.
2	Select the **ECMP** tab.
3	In the **MaxPath** field, enter the number of equal-cost paths allotted to each protocol listed. Up to 4 paths can be allotted to each. The default is 1.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

4 Click **Apply**.

--End--

Variable definitions

The following table describes the ECMP tab fields.

Field	Description
RoutingProtocol	The routing protocol to be configured.
MaxPath	The maximum number of ECMP paths assigned to the protocol.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Route policies configuration using Device Manager

This section describes the procedure you can use to configure route policies using Device Manager.

Route policies are a Nortel proprietary improvement on existing routing schemes. Using existing routing schemes, packets are forwarded based on routes that have been learned by the router through routing protocols such as RIP and OSPF or through the introduction of static routes. Route policies introduce the ability to forward packets based on rule sets created by the network administrator. These rule sets, or policies, are then applied to the learned or static routes.

Route policies configuration procedures

To configure routing policies, perform the following steps:

Step	Action
1	Create the appropriate prefix lists.
2	Assign those prefix lists to route policies.
3	Apply the route policies to the appropriate policy type.

--End--

Route policies configuration navigation

- “Creating a prefix list ” (page 400)
- “Creating a route policy” (page 401)
- “Configuring RIP in and out policies” (page 402)
- “Configuring an OSPF Accept Policy” (page 403)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- “Configuring OSPF redistribution parameters” (page 404)
- “Applying an OSPF accept or redistribution policy” (page 405)

Creating a prefix list

Use this procedure to create a new prefix list. Prefix lists are the base item in a routing policy. Prefix lists contain lists of IP addresses with their associated masks that support the comparison of ranges of masks.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > Policy**. The Policy dialog box appears with the Prefix List tab selected.
2	Click **Insert**. The Insert Prefix List dialog box appears.
3	Using the fields provided, create a new prefix list.
4	Click **Insert**.

--End--

Variable definitions

The following table describes the Prefix List tab fields.

Field	Definition
Id	Specifies the unique identifier of this prefix list.
Prefix	Specifies the IP address associated with this prefix list.
PrefixMaskLen	Specifies the subnet mask length associated with this prefix list.
Name	Specifies the name associated with this prefix list.
MaskLenFrom	Specifies the lower bound of the mask length. This value, when combined with the higher bound mask length (MaskLenUpto), specifies a subnet range covered by the prefix list. The default value is the mask length (PrefixMaskLen).
MaskLenUpto	Specifies the higher bound of the mask length. This value, when combined with the lower bound mask length (MaskLenFrom), specifies a subnet range covered by the prefix list.The default value is the mask length (PrefixMaskLen).

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Creating a route policy

Use this procedure to create a new route policy. Route policies are created and then applied to the switch as accept (in), announce (out), or redistribution policies.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > Policy**.
2	Select the **Route Policy** tab.
3	Click **Insert**.
	The Insert Route Policy dialog box appears.
4	Using the fields provided, create the new route policy.
5	Click **Insert**.

--End--

Variable definitions

The following table describes the Route Policy tab fields.

Field	Definition
Id	Specifies an index value to uniquely identify a policy.
SequenceNumber	Specifies a secondary index value that identifies individual policies inside a larger policy group.
Name	Specifies the name associated with this policy.
Enable	Specifies whether this policy sequence number is enabled or disabled. If disabled, the policy sequence number is ignored.
Mode	Specifies the action to be taken when this policy is selected for a specific route. A value of **permit** indicates that the route is allowed while **deny** indicates that the route is ignored.
MatchProtocol	If configured, matches the protocol through which the route is learned. This field is used only for RIP announce policies. Options are RIP, Static, Direct, OSPF or Any.
MatchNetwork	If configured, matches the destination network against the contents of the specified prefix list.
MatchIpRouteSource	If configured, matches the source IP address for RIP routes and advertising router IDs for OSPF routes against the contents of the specified prefix list. This option is ignored for all other route types.
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.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Definition
MatchInterface	If configured, 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.
MatchRouteType	Sets a specific route-type to be matched (applies only to OSPF routes). Externaltype1, and Externaltype2 specify the OSPF routes of the specified type only. OSPF internal refers to intra- and inter-area routes.
MatchMetric	If configured, matches the metric of the incoming advertisement or existing route against the specified value (1 to 655535). If 0, then this field is ignored. The default is 0.
NssaPbit	Sets or resets the P bit in specified type 7 LSA. By default the P bit is always set in case the user sets it to a disabled 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 enabled.
SetRoutePreference	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.
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	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.
SetInjectNetList	If configured, the switch replaces the destination network of the route that matches this policy with the contents of the specified prefix list.
SetMask	Indicates the mask to used for routes that pass the policy matching criteria.

Configuring RIP in and out policies

To configure RIP accept and announce policies, follow this procedure:

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > Policy**.
2	Select the **RIP In/Out Policy** tab.
3	Using the fields provided, configure the RIP policies.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

4 Click **Apply**.

--End--

Variable definitions

The following table describes the RIP In/Out Policy tab fields.

Field	Definition
Address	Specifies the address of the RIP interface.
Interface	Specifies the associated switch interface.
InPolicy	Specifies a previously configured policy to be used as the accept policy on this interface.
OutPolicy	Specifies a previously configured policy to be used as the announce policy on this interface.

Configuring an OSPF Accept Policy

Use this procedure to configure OSPF accept policies.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > Policy**.
2	Select the **OSPF Accept** tab.
3	Click **Insert**. The Insert OSPF Accept dialog box appears.
4	Use the fields provided to create the new accept policy.
5	Click **Insert**.
6	Using the fields provided, configure the desired accept policy.
7	Click **Apply**.

--End--

Variable definitions

The following table describes the OSPF Accept tab fields.

Field	Definition
AdvertisingRtr	Represents the IP address of the router from which advertisements are to be accepted. The value *0.0.0.0* denotes that advertisements from all routers are accepted.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Definition
Enable	Indicates whether the policy is enabled.
MetricType	Indicates the metric type associated with the policy. Available options are: type1, type2, and any.
PolicyName	Specifies a previously configured policy to be used as the OSPF accept policy.

Configuring OSPF redistribution parameters

Use this procedure to configure OSPF redistribution parameters.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > OSPF**.
2	Select the **Redistribute** tab.
3	Click **Insert**.
	The Insert Redistribute dialog box appears.
4	Using the fields provided, create the new redistribution entry.
5	Click **Insert**.

--End--

Variable definitions

The following table describes the Redistribute tab fields.

Field	Description
RouteSource	Specifies the route source protocol for redistribution (RIP, Direct or Static).
Enable	Indicates whether the redistribution entry is active.
Metric	Specifies the metric to be announced in the advertisement. This is a value between 0 and 65535.
MetricType	Specifies the metric type to associate with the route redistribution: *type1* or *type2*.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Subnets	This field indicates whether subnetworks need to be advertised individually. Options available are: allow and supress.
RoutePolicy	Specifies the name of preconfigured route policy to be used as the redistribution policy.

Applying an OSPF accept or redistribution policy

Use this procedure to configure OSPF policy application.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > Policy**.
2	Select the **Applying Policy** tab.
3	To apply a preconfigured OSPF accept policy, select the **OspfInFilterApply** check box.
4	To apply a preconfigured OSPF redistribution policy, select the **RedistributeApply** check box.
5	If you are applying OSPF redistribution policies, select the type of redistribution to apply from the available options in the **OspfApplyRedistribute** field.
6	Click **Apply**.

--End--

Variable definitions

The following table describes the Applying Policy tab fields.

Field	Definition
OspfInFilterApply	Specifies whether OSPF accept policies are enabled.
RedistributeApply	Specifies whether OSPF redistribution policies are enabled.
OspfApplyRedistribute	Specifies the type of redistribution that is applied for OSPF redistribution policies.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

DHCP relay configuration using Device Manager

The following topics describe DHCP relay configuration using Device Manager.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be set as the DHCP relay agent.
- Ensure that a route to the destination DHCP server is available on the switch.

DHCP relay configuration procedures

To configure DHCP relay using Device Manager, perform the following steps:

Step	Action
1	Configure the DHCP relay forwarding path, specifying the VLAN IP as the DHCP relay agent and the remote DHCP server as the destination, and enable DHCP relay on the VLAN.

--End--

DHCP relay configuration navigation

- “Configuring DHCP Relay ” (page 408)
- “Configuring DHCP parameters on a VLAN” (page 408)
- “Displaying and graphing DHCP counters on a VLAN” (page 409)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring DHCP Relay

Use this procedure to configure DHCP Relay.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > DHCP**. The DHCP Relay tab appears.
2	Click **Insert**. The Insert DHCP Relay dialog box appears.
3	In the **AgentAddr** field, enter the IP address of the local VLAN to serve as the DHCP relay agent.
4	In the **ServerAddr** field, enter the remote DHCP Server IP address.
5	Ensure the **Enable** check box is selected.
6	In the **Mode** box, select the desired DHCP relay mode.
7	Click **Insert**. The new DHCP entry appears in the DHCP Relay tab.

--End--

Variable definitions

The following table describes the DHCP Relay tab fields.

Field	Description
AgentAddr	The IP address of the local VLAN serving as the DHCP relay agent.
ServerAddr	The IP address of the remote DHCP server.
Enable	Enables (selected) or disables (cleared) DHCP relay.
Mode	Indicate whether the relay instance applies for BOOTP packets, DHCP packets, or both.

Configuring DHCP parameters on a VLAN

Use this procedure to configure the DHCP relay parameters on a VLAN.

Procedure steps

Step	Action
1	From the Device Manager menu, select **VLAN > VLANs**.
2	Select the VLAN for which DHCP relay is to be configured.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

3 Click **IP**.

The IP, VLAN dialog box appears.

4 Select the **DHCP** tab.

5 To configure the DHCP relay parameters, modify the values in the fields provided, as required.

6 Click **Apply**.

--End--

Variable definitions

The following table describes the DHCP tab fields.

Field	Description
Enable	Specifies whether DHCP relay is enabled or disabled.
MinSec	Specifies the min-sec value. The switch immediately forwards a BootP/DHCP packet if the 'secs' field in the BootP/DHCP packet header is greater than the configured min-sec value; otherwise, the packet is dropped.
Mode	Specifies the type of packets this VLAN interface forwards: BootP, DHCP, or both.
AlwaysBroadcast	Specifies whether DHCP Reply packets are broadcast to the DHCP clients on this VLAN interface.
ClearCounters	Specifies to clear the DHCP relay counters for the VLAN.
CounterClearTime	Specifies the last time the counter values in this entry were reset to 0.

Displaying and graphing DHCP counters on a VLAN

Use this procedure to display and graph the current DHCP counters on a VLAN.

Procedure steps

Procedure steps

Step	Action
1	From the Device Manager menu, select **VLAN > VLANs**.
2	Select the VLAN for which DHCP is configured.
3	Click **IP**. The IP, VLAN dialog box appears.
4	Select the **DHCP** tab.
5	Click **Graph**. The DHCP Stats dialog box appears.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

6 Use the buttons provided to graph selected DHCP counter information.

--End--

Variable definitions

The following table describes the DHCP Stats dialog box fields.

Field	Description
NumRequests	Indicates the number of DHCP requests.
NumReplies	Indicates the number of DHCP replies.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

UDP broadcast forwarding configuration using Device Manager

UDP broadcast forwarding is a general mechanism for selectively forwarding limited UDP broadcasts received on an IP interface to a configured IP address. This section describes the procedures used to configure and manage UDP broadcast forwarding using Device Manager.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN interface.
- Ensure that a route to the destination address is available on the switch.

UDP broadcast forwarding configuration procedures

To configure UDP broadcast forwarding using Device Manager, perform the following steps:

Procedure steps

Step	Action
1	Create UDP protocol entries that specify each UDP port and associated protocol that you want to forward.
2	Create UDP forwarding entries that specify the destination address for each UDP port that you want to forward.
3	Add UDP forwarding entries to a UDP forwarding list (you can create up to 128 UDP forwarding lists.)
4	Apply UDP forwarding lists to local VLAN interfaces.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

UDP broadcast forwarding configuration navigation

- "Configuring UDP protocol table entries" (page 412)
- "Configuring UDP forwarding entries" (page 412)
- "Configuring a UDP forwarding list" (page 413)
- "Applying a UDP forwarding list to a VLAN" (page 414)

Configuring UDP protocol table entries

Use this procedure to create UDP table entries that identify the protocols associated with specific UDP ports that you want to forward.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > UDP Forwarding**. The UDP_Forward dialog box appears with the Protocols tab selected.
2	Click **Insert**. The Insert Protocols dialog box appears.
3	In the **PortNumber** box, enter the UDP port number that you want to forward.
4	In the **Name** box, enter the protocol name associated with the UDP port number.
5	Click **Insert**.

--End--

Variable definitions

The following table describes the Protocols tab fields.

Field	Description
PortNumber	Specifies the UDP port number.
Name	Specifies the protocol name associated with the UDP port.

Configuring UDP forwarding entries

Use this procedure to configure individual UDP forwarding list entries, which associate UDP forwarding ports with destination IP addresses.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > UDP Forwarding**. The UDP_Forward dialog box appears.
2	Select the **Forwardings** tab.
3	Click **Insert**. The Insert Forwardings dialog box appears.
4	Using the provided fields, specify a destination address for a selected port.
5	Click **Insert**.

--End--

Variable definitions

The following table describes the Forwardings tab fields.

Field	Description
DestPort	Specifies the port on which the UDP forwarding originates (configured using the Protocols tab).
DestAddr	Specifies the destination IP address.
Id	Specifies an ID for the entry.
FwdListIdList	Indicates the UDP forward list with which this entry is associated (using the Forwarding Lists tab).

Configuring a UDP forwarding list

Use this procedure to add the UDP port/destination forwarding list entries (configured in the Forwardings tab) to UDP forwarding lists. Each UDP forwarding list can contain multiple port/destination entries.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > UDP Forwarding**.
2	Select the **Forwarding Lists** tab.
3	Click **Insert**. The Insert Forwarding Lists dialog box appears.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

4 In the **Id** field, assign a unique ID to the UDP forwarding list.

5 In the **Name** field, enter a unique name for the UDP forwarding list.

6 In the **FwdIdList** field, click the ellipsis (**...**), select the desired port/destination pairs from the list, and click **OK**.

7 Click **Apply**.

--End--

Variable definitions

The following table describes the Forwarding Lists tab fields.

Field	Description
Id	The unique identifier assigned to the forwarding list.
Name	The name assigned to the forwarding list.
FwdIdList	The forwarding entry IDs associated with the port/server IP pairs created using the Forwardings tab.

Applying a UDP forwarding list to a VLAN

Use this procedure to assign a UDP forwarding list to a VLAN and to configure the related UDP forwarding parameters for the VLAN.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > UDP Forwarding**.
2	Select the **Broadcast Interfaces** tab.
3	Click **Insert**. The Insert Broadcast Interface dialog box appears.
4	In the **LocalIfAddr** field, click the **Addr** button and select a VLAN IP address from the list.
5	In the **UdpPortFwdListId** field, click the **IdList** button and select the desired UDP forwarding list to apply to the VLAN from those listed.
6	To modify the maximum TTL, modify the value in the **MaxTtl** field.
7	To specify a broadcast mask, enter a mask in the **BroadCastMask** field.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

8 Click **Insert**.

--End--

Variable definitions

The following table describes the Broadcast Interface tab fields.

Field	Description
LocalIfAddr	The IP address of the local VLAN interface.
UdpPortFwdListId	The port forwarding lists associated with the interface. This ID is defined in the Forwarding Lists tab.
MaxTtl	Indicates the maximum number of hops an IP broadcast packet can take from the source device to the destination device. This is an integer value between 1 and 16.
NumRxPkts	The total number of UDP broadcast packets received by this local interface.
NumFwdPkts	The total number of UDP broadcast packets forwarded.
NumDropPktsDestUnreach	The total number of UDP broadcast packets dropped because the destination was unreachable.
NumDropPktsUnknownPort	The total number of UDP broadcast packets dropped because the destination port or protocol specified has no matching forwarding policy.
BroadCastMask	Specifies the 32 bit mask used by the selected VLAN interface to take forwarding decisions based on the destination IP address of the incoming UDP broadcast traffic. If you do not specify a broadcast mask value, the mask of the interface to which the list is attached is used.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Static ARP and Proxy ARP configuration using Device Manager

This chapter describes the procedures you can use to configure static ARP and proxy ARP using Device Manager.

Static ARP and Proxy ARP configuration navigation

- "Configuring static ARP entries" (page 417)
- "Configuring Proxy ARP" (page 418)

Configuring static ARP entries

Use this procedure to configure static ARP entries for the switch.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the target VLAN interface.

Procedure steps

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > IP**.
2	Select the **ARP** tab.
3	Click **Insert**.
	The Insert ARP dialog box appears.
4	From the **Port in VLAN** list, select the VLAN to which the static ARP entry is being added.
	A VLAN dialog box appears listing all member ports.
5	In the **VLAN** dialog box, select the port for this ARP entry.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The Interface(vlanId:Port) field updates with the appropriate VLAN and port information

6 In the **IPAddress** field, specify the IP address for the ARP entry.

7 In the **MacAddress** field, specify the MAC address for the ARP entry.

8 Click **Insert**.

--End--

Variable definitions

The following table describes the Insert ARP tab fields.

Field	Description
Interface	Specifies the VLAN and port to which the static ARP entry is being added.
MacAddress	Specifies the MAC address of the device being set as a static ARP entry.
IpAddress	Specifies the IP address of the device being set as a static ARP entry.
Type	Specifies the type of ARP entry: static, dynamic, local, or broadcast.

Configuring Proxy ARP

Use this procedure to configure proxy ARP on the switch. Proxy ARP allows the switch to respond to an ARP request from a locally attached host (or end station) for a remote destination.

Prerequisites

- Enable IP routing globally.
- Enable IP routing and configure an IP address on the VLAN to be configured as a Proxy ARP interface.

Procedure steps

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > IP** .
2	Select the **ARP Interfaces** tab.

ATTENTION
Device Manager does not display the ARP Interfaces tab if you have not enabled routing on the switch.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

3 To enable proxy ARP, in the DoProxy field, select **enable**.

4 Click **Apply**.

--End--

Variable definitions

The following table describes the ARP Interfaces tab fields.

Field	Description
IfIndex	Specifies a configured switch interface.
DoProxy	Enables or disables proxy ARP on the interface.
DoResp	Enables or disables the sending of ARP responses on the specified interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

VRRP configuration using Device Manager

This chapter describes the procedures you can use to configure VRRP using Device Manager.

The Virtual Router Redundancy Protocol (VRRP) 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.

Prerequisites

- Install the Advanced License.
- Enable IP routing globally on the switch.
- Assign an IP address to the VLAN that you want to enable with VRRP.

  Routing is automatically enabled on the VLAN when you assign an IP address to it.

VRRP configuration procedures

To enable VRRP on a VLAN, perform the following steps:

Procedure steps

Step	Action
1	Enable VRRP globally on the switch.
2	Assign a virtual router IP address to a virtual router ID.
3	Configure the priority for this router as required and enable the virtual router.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

VRRP configuration navigation

- "Configuring global VRRP status and properties" (page 422)
- "Assigning an IP address to a virtual router ID" (page 422)
- "Configuring VRRP interface properties" (page 423)
- "Graphing VRRP interface information" (page 425)
- "Viewing general VRRP statistics" (page 426)

Configuring global VRRP status and properties

Use this procedure to configure global VRRP settings.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > VRRP**. The VRRP dialog box appears with the Globals tab selected.
2	To configure the global VRRP settings modify the fields provided.
3	Click **Apply**.

--End--

Variable definitions

The following table describes the VRRP—Globals tab fields.

Field	Definition
Enabled	Indicates whether VRRP is globally enabled on the switch.
Version	Indicates the version of VRRP supported on this switch.
NotificationCntl	Indicates whether the VRRP-enabled router generates Simple Network Management Protocol (SNMP) traps based on VRRP events: • Enabled - SNMP traps are sent. • Disabled - SNMP traps are not sent.
PingVirtualAddrEnabled	Indicates whether this device responds to pings directed to a virtual router IP address.

Assigning an IP address to a virtual router ID

Use this procedure to associate an IP address with a virtual router ID on a switch interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > VRRP**.
2	Select the **Interface Address** tab.
3	Click **Insert**. The Insert Interface Address dialog box appears.
4	Using the provided fields, create the new interface.
5	Click **Insert**.

--End--

Variable definitions

The following table describes the Interface Address tab fields.

Field	Definition
IfIndex	Specifies the IP address of the interface on which to configure VRRP. Click the ellipsis (**...**) to select a previously configured interface.
VrId	Specifies the virtual router ID to associate with this interface.
IpAddr	Specifies the IP address to associate with the virtual router ID.
Status	Specifies the status of the interface; active or inactive.

Configuring VRRP interface properties

Use this procedure to configure VRRP interface properties.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > VRRP**.
2	Select the **Interfaces** tab.
3	To configure VRRP interfaces, modify the provided fields.
4	Click **Apply**.

--End--

Variable definitions

The following table describes the Interfaces tab fields.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Definition
IfIndex	Specifies the interface index of the VRRP interface.
VrId	Specifies the virtual router ID.
PrimaryIpAddr	Specifies an IP address selected from the set of real interface addresses. VRRP advertisements are always sent using the primary IP address as the source of the IP packet.
VirtualMacAddr	Specifies the virtual MAC address of the virtual router.
State	Specifies the current state of the virtual router. A virtual router can be in one of the following states: **initalize** (indicates the virtual router is waiting for a startup event), **backup** (indicates the virtual router is monitoring the availability of the master router), or **master**(indicates the router is forwarding packets for IP addresses associated with the virtual router ID).
AdminState	Indicates the administrative status of the virtual router.
Priority	Indicates the priority to be used for the virtual router master election process. This is an integer value between 1 and 255. The priority value for the VRRP router that owns the IP addresses associated with the virtual router must be 255. The default priority value for VRRP routers backing up a virtual router is 100.
MasterIpAddr	Indicates the master router's real (primary) IP address. This is the IP address listed as the source in the VRRP advertisement last received by this virtual router.
AdvertisementInterval	Indicates the time interval, in seconds, between transmission of advertisement messages. Only the master router sends VRRP advertisements. This is an integer value between 1 and 255. The default value is 1.
VirtualRouterUpTime	Indicates the amount of time this virtual router has spent out of the **initialize** state.
HolddownTimer	Specifies the amount of time (in seconds) to wait before preempting the current VRRP master. This is an integer value between 0 and 21600.
HoldDownState	Specifies the holddown state of this VRRP interface.
HoldDownTimeRemaining	Specifies the amount of time (in seconds) left before the holddown timer will expire.
Action	Triggers an action on this VRRP interface. Options available are: none (no action) or preemptHoldDownTimer.
CriticalIpAddrEnabled	Indicates whether the user-defined critical IP address is enabled. If the user-defined critical IP address is not enabled, a default critical IP address of 0.0.0.0 will be used.
CriticalIpAddr	Specifies the IP address of the interface that will cause a shutdown event.
BackupMasterEnabled	Indicates whether the backup/master functionality is enabled on this interface.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Definition
BackupMasterState	Indicates the state of the backup/master functionality.
FastAdvertisementEnabled	Indicates if the Faster Advertisement Interval should be used. The default value is false.
FastAdvertisementInterval	Specifies the fast advertisement interval, in milliseconds, between sending advertisement messages. This is an integer value between 200 and 1000. The default value is 200.

Graphing VRRP interface information

Use this procedure to display and graph VRRP interface statistical information.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > VRRP**.
2	Select the **Interfaces** tab.
3	Select a listed interface and click **Graph**. The VRRP Stats dialog box appears.
4	Using the provided fields, view and graph the VRRP statistical information.

--End--

Variable definitions

The following table describes the VRRP Stats tab fields.

Field	Definition
BecomeMaster	Specifies the total number of times that this virtual router's state has transitioned to master.
AdvertiseRcvd	Specifies the total number of VRRP advertisements received by this virtual router.
AdvertiseIntervalErrors	Specifies the total number of VRRP advertisement packets received for which the advertisement interval is different than the one configured for the local virtual router.
IpTtlErrors	Specifies the total number of VRRP packets received by the virtual router with IP TTL (Time-To-Live) not equal to 255.
PriorityZeroPktsRcvd	Specifies the total number of VRRP packets received by the virtual router with a priority of 0.
PriorityZeroPktsSent	Specifies the total number of VRRP packets sent by the virtual router with a priority of 0.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Definition
InvalidTypePktsRcvd	Specifies the number of VRRP packets received by the virtual router with an invalid value in the **type** field.
AddressListErrors	Specifies the total number of packets received for which the address list does not match the locally configured list for the virtual router.
AuthFailures	Specifies the total number of VRRP packets received that do not pass the authentication check.
InvalidAuthType	Specifies the total number of packets received with an unknown authentication type.
AuthTypeMismatch	Specifies the total number of packets received with **Auth Type** not equal to the locally configured authentication method.
PacketLengthErrors	Specifies the total number of packets received with a packet length less than the length of the VRRP header.

Viewing general VRRP statistics

Use this procedure to view and graph general VRRP statistics.

Procedure steps

Step	Action
1	From the Device Manager menu, select **IP Routing > VRRP**.
2	Select the **Stats** tab.

--End--

Variable definitions

The following table describes the VRRP—Stats tab fields.

Field	Definition
RouterChecksumErrors	Specifies the total number of VRRP packets received with an invalid VRRP checksum value.
RouterVersionErrors	Specifies the total number of VRRP packets received with an unknown or unsupported version number.
RouterVrIdErrors	Specifies the total number of VRRP packets received with an invalid VRID for this virtual router.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IGMP snooping configuration using Device Manager

This chapter describes the procedures you can use to configure IGMP snooping on a VLAN using Device Manager.

IGMP snooping configuration procedures

Step	Action
1	To configure IGMP snoop, the only required configuration is to enable snoop on the VLAN. All related configurations, listed below, are optional and can be configured to suit the requirements of your network.
	--End--

IGMP snooping configuration navigation

- "Configuring IGMP snoop, proxy, and IGMP parameters on a VLAN" (page 428)
- "Configuring IGMP snoop, proxy, and static mrouter ports on a VLAN" (page 429)
- "Configuring IGMP router alert, snoop, and proxy on a VLAN" (page 431)
- "Flushing the IGMP router tables and configuring IGMP router alert" (page 432)
- "Configuring unknown multicast filtering" (page 435)
- "Specifying a multicast MAC address to be allowed to flood all VLANs" (page 435)
- "Displaying IGMP cache information" (page 436)
- "Displaying IGMP group information" (page 437)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

- "Displaying multicast route information" (page 438)
- "Displaying multicast next hop information" (page 438)
- "Displaying multicast interface information" (page 439)
- "Configuring IGMP selective channel block using Device Manager" (page 440)

Configuring IGMP snoop, proxy, and IGMP parameters on a VLAN

Use this procedure to configure IGMP snoop, proxy, and IGMP parameters on a VLAN.

IGMP snoop and proxy are disabled by default.

ATTENTION
With IGMP snoop, the QueryInterval, Robustness, and Version values must be the same as those configured on the interface (VLAN) of the multicast router (IGMP querier).

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **VLAN > VLANs**. The VLAN dialog box appears with the Basic tab displayed.
2	Click the **Snoop** tab.
3	To enable IGMP snoop for a VLAN, select **true** from the **Enable** column.
4	To enable IGMP proxy for a VLAN, select **true** from the **ReportProxyEnable** column.
5	To configure IGMP properties, modify the values in the boxes provided.
6	Click **Apply**.

--End--

Variable definitions

The following table describes the Snoop tab fields.

Field	Description
Id	Specifies the VLAN ID.
Name	Specifies the VLAN name.
ReportProxyEnable	Specifies the IGMP proxy status: enabled (true) or disabled (false).

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Enable	Specifies the IGMP snoop status: enabled (true) or disabled (false).
Robustness	Specifies tuning for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each query interval, plus 1. If you expect a network to lose query packets, you must increase the robustness value. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). The range is from 2–255, and the default is 2. The default value of 2 means that one query for each query interval can be dropped without the querier aging out.
QueryInterval	Specifies the frequency (in seconds) at which IGMP query packets are transmitted on the VLAN. Ensure that the query interval is the same as the configured value on the multicast router (IGMP querier). The range is from 1 to 65535, and the default is 125.
MRouterPorts	Specifies the statically-configured mrouter ports in the VLAN that provide connectivity to a nonquerier IP multicast router. Multicast data and group reports are forwarded out this port to the multicast router. You do not need to configure this parameter if only one multicast router exists on the VLAN.
Ver1MRouterPorts	Displays the mrouter ports in the VLAN that use IGMP version 1.
Ver2MRouterPorts	Displays the mrouter ports in the VLAN that use IGMP version 2.
ActiveMrouterPorts	Displays all dynamic (querier port) and static mrouter ports that are active on the interface.
ActiveQuerier	Displays the IP address of the multicast querier router.
QuerierPort	Displays the mrouter port on which the multicast querier router was heard.
MrouterExpiration	Specifies the time remaining before the multicast router is aged out on this interface. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The Query Max Response Interval (obtained from the queries received) is used as the timer resolution.

Configuring IGMP snoop, proxy, and static mrouter ports on a VLAN

Use this procedure to configure IGMP snooping, proxy, and static mrouter ports on a VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

By default, IGMP snoop and proxy are disabled, and no static mrouter ports are configured.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > IGMP**. The IGMP dialog box appears with the Globals tab displayed.
2	Click the **Snoop** tab.
3	To enable IGMP snoop, select **true** in the **SnoopEnable** column.
4	To enable IGMP proxy, select **true** in the **ProxySnoopEnable** column.
5	To configure mrouter ports, select the desired ports in the **SnoopMRouterPorts** column.
6	Click **Apply**.

--End--

Variable definitions

The following table describes the Snoop tab fields.

Field	Description
IfIndex	Specifies the VLAN ID.
SnoopEnable	Specifies the IGMP snoop status: enabled (true) or disabled (false).
ProxySnoopEnable	Specifies the IGMP proxy status: enabled (true) or disabled (false).
SnoopMRouterPorts	Specifies the static mrouter ports. Such ports are directly attached to a multicast router so the multicast data and group reports are forwarded to the router.
SnoopActiveMRouter Ports	Displays all dynamic (querier port) and static mrouter ports that are active on the interface.
SnoopMRouterExpiration	Specifies the time remaining before the multicast router is aged out on this interface. If the switch does not receive queries before this time expires, it flushes out all group memberships known to the VLAN. The Query Max Response Interval (obtained from the queries received) is used as the timer resolution.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring IGMP router alert, snoop, and proxy on a VLAN

Use this procedure to configure IGMP router alert on a VLAN. The router alert feature instructs the router to drop control packets that do not have the router-alert flag in the IP header.

This procedure also allows you to configure IGMP snooping and proxy.

ATTENTION
With IGMP snoop, the QueryInterval, Robustness, and Version values must be the same as those configured on the interface (VLAN) of the multicast router (IGMP querier).

ATTENTION
To maximize your network performance, Nortel recommends that you set the router alert parameter according to the version of IGMP currently in use:

- IGMPv1—Disable
- IGMPv2—Enable
- IGMPv3—Enable

IGMP snoop, proxy, and router alert are all disabled by default.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **VLAN > VLANs**. The VLAN dialog box appears with the Basic tab displayed.
2	Select a VLAN.
3	Click **IP**. The IP, VLAN dialog box appears with the IP Address tab displayed.
4	Click the **IGMP** tab.
5	To enable the router alert option, select **enable** in the **RouterAlertEnable** box.
6	To enable IGMP snoop, select the **SnoopEnable** box.
7	To enable IGMP proxy, select the **ProxySnoopEnable** box.
8	To configure the other IGMP properties, modify the values in the boxes provided.
9	Click **Apply**.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the IGMP tab fields.

Field	Description
QueryInterval	The frequency (in seconds) at which IGMP host query packets are transmitted on the interface. Ensure that the query interval is the same as the configured value on the multicast router (IGMP querier). The range is from 1 to 65535, and the default is 125.
QueryMaxResponseTime	The maximum response time (in 1/10 seconds) advertised in IGMPv2 general queries on this interface.
Robustness	Specifies tuning for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, you must increase the robustness value. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). The range is from 2 to 255, and the default is 2. The default value of 2 means that one query for each query interval can be dropped without the querier aging out.
LastMembQueryIntvl	Specifies the maximum response time inserted into group-specific queries sent in response to leave group messages and is also the amount of time between group-specific query messages.
RouterAlertEnable	When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the router processes IGMP packets regardless of whether the router alert IP option is set or not. To maximize your network performance, Nortel recommends that you set this parameter according to the version of IGMP currently in use: IGMPv1—Disable, IGMPv2—Enable, IGMPv3—Enable.
Version	The version of IGMP (1, 2, or 3) that you want to configure on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
SnoopEnable	Specifies the IGMP snoop status: enabled (true) or disabled (false).
ProxySnoopEnable	Specifies the IGMP proxy status: enabled (true) or disabled (false).
SnoopMRouterPorts	Specifies the statically-configured mrouter ports in the VLAN that provide connectivity to a nonquerier IP multicast router. Multicast data and group reports are forwarded out this port to the multicast router. You do not need to configure this parameter if only one multicast router exists on the VLAN.

Flushing the IGMP router tables and configuring IGMP router alert

Use this procedure to flush the IGMP router tables. This procedure also allows you to configure the router alert option and query interval on a VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

ATTENTION
The QueryInterval, Robustness, and Version values must be the same as those configured on the interface (VLAN) of the multicast router (IGMP querier).

ATTENTION
To maximize your network performance, Nortel recommends that you set the router alert parameter according to the version of IGMP currently in use:

- IGMPv1 - Disable
- IGMPv2 - Enable
- IGMPv3 - Enable

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > IGMP**. The IGMP dialog box appears.
2	Click the **Interface** tab. The Interface tab appears.
3	To flush the routing tables, select the desired flush option under the **FlushAction** column.
4	To configure the IGMP properties, modify the values in the boxes provided.
5	Click **Apply**.

--End--

Variable definitions

The following table describes the Interface tab fields.

Field	Description
IfIndex	The interface on which IGMP is enabled.
QueryInterval	The frequency (in seconds) at which IGMP host query packets are transmitted on the interface. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). The range is from 1–65535, and the default is 125.
Status	Indicates whether or not the interface is active. The interface becomes active if any IGMP forwarding ports exist on the interface. If the VLAN has no port members or if all of the port members are disabled, the status is notInService.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Version	The version of IGMP (1, 2, or 3) configured on this interface. For IGMP to function correctly, all routers on a LAN must use the same version. The default is version 2.
OperVersion	The version of IGMP currently running on this interface.
Querier	The address of the IGMP querier on the IP subnet to which this interface is attached.
QueryMaxResponse Time	The maximum response time (in 1/10 seconds) advertised in IGMPv2 general queries on this interface. This field is a read-only field and is not configurable on the Ethernet Routing Switch 5000 Series .
WrongVersionQueries	The number of queries received with an IGMP version that does not match the interface. IGMP requires that all routers on a LAN be configured to run the same version of IGMP. If queries are received with the wrong version, it indicates a version mismatch.
Joins	The number of times a group membership is added on this interface; that is, the number of times an entry for this interface is added to the cache table. This number gives an indication of the amount of IGMP activity over time.
Robustness	Specifies tuning for the expected packet loss of a network. This value is equal to the number of expected query packet losses for each serial query interval, plus 1. If you expect a network to lose query packets, you must increase the robustness value. Ensure that the robustness value is the same as the configured value on the multicast router (IGMP querier). The range is from 2 to 255, and the default is 2. The default value of 2 means that one query for each query interval can be dropped without the querier aging out.
LastMembQueryIntvl	Sets the maximum response time (in tenths of a second) that is inserted into group-specific queries sent in response to leave group messages. This parameter is also the time between group-specific query messages. This value is not configurable for IGMPv1. Decreasing the value reduces the time to detect the loss of the last member of a group. The range is from 0–255, and the default is 10 tenths of seconds. Nortel recommends configuring this parameter to values higher than 3. If a fast leave process is not required, Nortel recommends values above 10. (The value 3 is equal to 0.3 of a second, and 10 is equal to 1.0 second.)

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
RouterAlertEnable	When enabled, this parameter instructs the router to ignore IGMP packets that do not contain the router alert IP option. When disabled (default setting), the router processes IGMP packets regardless of whether the router alert IP option is set or not. To maximize your network performance, Nortel recommends that you set this parameter according to the version of IGMP currently in use: IGMPv1—Disable, IGMPv2—Enable, IGMPv3—Enable.
FlushAction	Flushes the specified table type: • none • flushGrpMem: group member table • flushMrouter: mrouter table

Configuring unknown multicast filtering

The default switch behavior is to flood all packets with unknown multicast addresses. Use this procedure to prevent the flooding of packets with unknown multicast addresses and enable the forwarding of these packets to static mrouter ports only.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **VLAN > VLANs**. The VLAN dialog box appears with the Basic tab displayed.
2	Click the **Unknown Multicast Filtering** tab.
3	To enable unknown multicast flooding, select the **MulticastFloodingEnabled** box.
4	Click **Apply**.

--End--

Specifying a multicast MAC address to be allowed to flood all VLANs

Use this procedure to allow particular unknown multicast packets to be flooded on all switch VLANs.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **VLAN > VLANs**. The VLAN dialog box appears with the Basic tab displayed.
2	Click the **MAC Multicast Filter Table** tab.
3	Click **Insert**.
4	In the **AllowedAddressMacAddr** field, enter the MAC address to flood.
5	Click **Insert**.

--End--

Displaying IGMP cache information

Display IGMP cache information to show the learned multicast groups in the cache and the IGMPv1 version timers.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > IGMP**. The IGMP dialog box appears with the Cache tab displayed.

--End--

Variable definitions

The following table describes the Cache tab fields.

Field	Description
Address	Indicates the IP multicast group address.
IfIndex	Indicates the VLAN interface from which the group address is heard.
LastReporter	Indicates the last IGMP host to join the group.
ExpiryTime	Indicates the amount of time (in seconds) remaining before this entry is aged out.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
Version1Host Timer	Indicates the time remaining until the local router assumes that no IGMP version 1 members exist on the IP subnet attached to the interface. Upon hearing an IGMPv1 membership report, this value is reset to the group membership timer. When the time remaining is nonzero, the local interface ignores IGMPv2 Leave messages that it receives for this group.
Type	Indicates whether the entry is learned dynamically or is added statically.

Displaying IGMP group information

Display IGMP group information to show the learned multicast groups and the attached ports.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > IGMP**. The IGMP dialog box appears.
2	Click the **Groups** tab. The Groups tab appears.

--End--

Variable definitions

The following table describes the Groups tab fields.

Field	Description
IpAddress	Indicates the multicast group address.
IfIndex	Indicates the VLAN interface from which the multicast group address is heard.
Members	Indicates the IP address of the IGMP receiver (host or IGMP reporter).
Expiration	Indicates the time left before the group report expires on this port. This variable is updated upon receiving a group report.
InPort	Indicates the member port for the group. This is the port on which group traffic is forwarded.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Displaying multicast route information

Display multicast route information for troubleshooting purposes.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > Multicast**. The Multicast dialog box appears with the Routes tab displayed.

--End--

Variable definitions

The following table describes the Routes tab fields.

Field	Description
Group	The IP Multicast group address.
Source	The source address.
SourceMask	The source address mask.
UpstreamNeighbor	The address of the upstream neighbor that is forwarding packets for the specified source and group. 0.0.0.0 appears if the network is local.
Interface	The VLAN where datagrams for the specified source and group are received.
ExpiryTime	The amount of time remaining before this entry is aged out. The value 0 indicates that the entry is not subject to aging.
Protocol	The routing protocol through which this route was learned.

Displaying multicast next hop information

Display all multicast next hop information to find the best route to a member group.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing, Multicast**. The Multicast dialog box appears with the Routes tab displayed.
2	Click the **Next Hops** tab.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The Next Hops tab appears.

--End--

Variable definitions

The following table describes the Next Hops tab fields.

Field	Description
Group	The IP Multicast group.
Source	The source address.
SourceMask	The source address mask.
OutInterface	The VLAN ID for the outgoing interface for the next hop.
Address	The address of the next hop specific to this entry. For most interfaces, this address is identical to the next hop group.
State	An indication of whether or not the outgoing interface and next hop represented by this entry is currently being used to forward IP datagrams. A Value of *forwarding* indicates this parameter is currently being used; *pruned* indicates the parameter is not being used.
ExpiryTime	The amount of time remaining before this entry is aged out. The value 0 indicates that the entry is not subject to aging.
ClosestMemberHops	The minimum number of hops between this router and a member of this IP Multicast group reached through this next hop on this outgoing interface. IP Multicast datagrams for the group that have a TTL less than this number of hops are not forwarded to the next hop.
Protocol	The routing protocol where this next hop was learned.

Displaying multicast interface information

Use this procedure to display multicast interface information.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > Multicast**. The Multicast dialog box appears.
2	Click the **Interfaces** tab. The Interfaces tab appears.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

The following table describes the Interfaces tab fields.

Field	Description
Variable	**Value**
Interface	The VLAN ID.
Ttl	The datagram time-to-live (TTL) threshold for the interface. IP Multicast datagrams with a TTL less than this threshold are not forwarded out of the interface. The default value of 1 means that all multicast packets are forwarded out of the interface.
Protocol	The routing protocol running on this interface.

Configuring IGMP selective channel block using Device Manager

Use this procedure to configure the IGMP selective channel block feature.

Step	Action
1	From the Device Manager menu bar select **IP Routing > IGMP**. The IGMP dialog box appears.
2	Select the **Profile** tab.
3	Click **Insert**.
4	In the **Profile ID** section enter the profile ID.
5	Click **Insert**.
6	Select the newly created profile.
7	Click **Profile Range**.
8	Enter the required information in the **RangeAddressStart** and **RangeAddressEnd** fields.
9	Click **Apply**.
10	Double-click the **ProfilePortList** field to assign ports to the profile.

--End--

Variable definitions

The following table describes the IGMP Profile tab fields.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Definition
Profile ID	The Profile ID. The range is from 1 to 65535.
RangeAddressStart	Should be a valid multicast group address.
RangeAddressEnd	Should be a valid multicast group address.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

PIM-SM configuration using Device Manager

This chapter describes the procedures you can use to configure PIM-SM using Device Manager.

Unlike dense-mode protocols, such as Distance Vector Multicast Routing Protocol (DVMRP), that initially flood multicast traffic to all routers over an entire internetwork, PIM-SM sends multicast traffic only to routers that belong to a specific multicast group and that choose to receive the traffic. PIM-SM reduces overhead costs for processing unwanted multicast packets.

Prerequisites for PIM-SM configuration

Before you can configure PIM-SM, you must prepare the switch as follows:

Step	Action
1	Install the Advanced Routing software license.
	ATTENTION If your Ethernet Routing Switch is running an Advanced License for a release prior to Release 6.0, to enable PIM-SM you must regenerate your license file from the Nortel web site and install the new license file on the switch.
2	Configure IP addresses on the VLAN interfaces on which you want to configure PIM-SM.
3	Enable a unicast protocol, either RIP or OSPF, globally and on the interfaces on which you want to configure PIM-SM.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

ATTENTION
PIM-SM requires a unicast protocol to multicast traffic within the network when performing the Reverse Path Forwarding (RPF) check. PIM-SM also uses the information from the unicast routing table to create and maintain the shared and shortest multicast tree. The unicast routing table must contain a route to every multicast source in the network, as well as routes to PIM entities such as the rendezvous points (RP) and bootstrap router (BSR).

--End--

PIM-SM configuration procedures

To configure PIM-SM, you must perform the following procedures:

Step	Action
1	Enable PIM-SM globally. (If desired, modify the default global PIM-SM properties.)
2	Enable PIM-SM on individual VLAN interfaces. (If desired, modify the default VLAN PIM-SM properties.)
3	Configure candidate RPs for the multicast groups in the network. (It is best to have multiple candidate-RPs in the network; however, with the Ethernet Routing Switch 5000, you can only configure one candidate-RP per switch for any number of groups.) OR Configure one (or several) static RPs for the multicast groups in the network. (To enable static RP in the PIM-SM domain, you must configure the same static RPs on every system that takes part in PIM-SM forwarding.)
4	Configure one or several candidate BSRs to propagate RP information to all switches in the network. You can configure every PIM-enabled VLAN as a C-BSR. (If Static RP is enabled, this step is not required.)

ATTENTION
Ensure that all routers in the path from the receivers to the RP and to the multicast source are PIM-enabled. Also ensure that all PIM routers have unicast routes to reach the source and RP through directly-connected PIM neighbors.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

All additional configurations listed below are optional and can be configured according to the requirements of your network.

PIM-SM configuration navigation

- "Configuring global PIM-SM status and properties" (page 445)
- "Configuring PIM-SM status and properties for a VLAN" (page 447)
- "Configuring PIM-SM VLAN properties from the IP Routing menu" (page 448)
- "Specifying the router as a candidate BSR on a VLAN interface" (page 450)
- "Displaying the current BSR" (page 451)
- "Specifying a local IP interface as a candidate RP" (page 452)
- "Displaying the active RP" (page 453)
- "Enabling static RP" (page 454)
- "Configuring a static RP" (page 455)
- "Specifying a virtual neighbor on an interface" (page 456)
- "Displaying PIM-SM neighbor parameters" (page 456)
- "Displaying the PIM-SM RP set" (page 457)

Configuring global PIM-SM status and properties

To enable PIM-SM on individual interfaces, you must first enable PIM-SM globally.

By default, PIM-SM is disabled.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.
2	To enable PIM-SM, select the **Enable** check box.
3	To configure global PIM-SM properties, modify the values in the boxes provided.
4	Click **Apply**.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable definitions

Field	Description
Mode	Displays the PIM mode on the switch. Ethernet Routing Switch 5000 Series Software supports only sparse mode.
Enable	Enables or disables PIM globally.
JoinPruneInterval	Specifies how long (in seconds) the PIM router waits between sending join/prune messages to its upstream neighbors. The range is 1–18724, and the default is 60 seconds.
RegisterSuppTimer	Specifies how long (in seconds) the DR suppresses sending register messages to the RP after the DR receives a register-stop message from the RP. The range is 5–65535, and the default is 60 seconds.
UniRouteChgTimeOut	Specifies how often (in seconds) the switch polls the routing table manager (RTM) for unicast routing information updates to be used by PIM. The range is 2–65535, and the default is 5 seconds. **ATTENTION** Lowering this value increases how often the switch polls the RTM. This can affect the performance of the switch, especially when a high volume of traffic is flowing through the switch.
DiscardDataTimeOut	After the router forwards the first source packet to the RP, this value specifies how long (in seconds) the router discards subsequent source data while waiting for a join from the RP. An IPMC discard record is created and deleted after the timer expires or after a join is received. The range is 5–65535, and the default is 60 seconds.
CRPADVTimeOut	Specifies how often (in seconds) routers that are configured as candidate RPs send candidate rendezvous point (C-RP) advertisement messages. After this timer expires, the C-RP sends an advertisement message to the elected BSR. The range is 5–26214, and the default is 60 seconds.
BootStrapPeriod	Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages. The range is 5–32757, and the default is 60 seconds.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
StaticRP	Enables or disables the static RP feature. Static RP permits communication with routers from other vendors that do not use the BSR mechanism. By default, static RP is disabled.
FwdCacheTimeOut	Specifies the PIM forward cache expiry value in seconds. This value is used in aging PIM mroutes in seconds. The range is 10–86400, and the default is 210.

Configuring PIM-SM status and properties for a VLAN

Use this procedure to enable PIM-SM on a VLAN and configure related properties.

By default, PIM-SM is disabled on VLANs.

Prerequisites

- You must enable PIM-SM globally.
- Before you change the state (active or passive) of a PIM interface using the InterfaceType field, first disable PIM to prevent instability in the PIM operations, especially when neighbors are present or when streams are received.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **VLAN > VLANs**. The VLAN dialog box appears with the Basic tab displayed.
2	Select the VLAN ID that you want to configure with PIM.
3	Click **IP**. The IP, VLAN dialog box appears with the IP Address tab displayed.
4	Click the **PIM** tab. The PIM tab appears.
5	Select the **Enable** check box.
6	To configure PIM-SM properties for the VLAN, modify the values in the boxes provided.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

7 Click **Apply**.

8 Click **Close**.

--End--

Variable definitions

The following table describes the PIM tab fields.

Field	Description
Enable	Enables or disables PIM-SM on the VLAN.
Mode	Displays the PIM mode on the switch. Ethernet Routing Switch 5000 Series Software supports only sparse mode.
HelloInterval	Specifies the interval (in seconds) that the PIM router waits between sending out hello message to neighboring routers. The default is 30 seconds.
JoinPruneInterval	Specifies the interval (in seconds) the PIM router waits between sending out join/prune message to its upstream neighbors. The default is 60 seconds.
CBSRPreference	Specifies the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
InterfaceType	Specifies the state (active or passive) of PIM on a VLAN interface. An active interface transmits and receives PIM control traffic. A passive interface drops all PIM control traffic, thereby reducing the load on the system. Passive interfaces are useful when you have a high number of PIM interfaces and these interfaces are connected to end users, not to other switches. By default, PIM-SM interfaces are active.

Configuring PIM-SM VLAN properties from the IP Routing menu

After you have enabled PIM-SM on a VLAN, you can also view and edit PIM-SM VLAN parameters from the PIM Interfaces tab accessible under the IP Routing menu. This procedure does not provide more configuration options than are available under the VLAN menu, but it does allow you to view some additional PIM parameters (such as DR) and also to view the configuration for multiple VLANs at once.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Prerequisites

- You must enable PIM-SM globally.
- You must enable PIM-SM on a VLAN.
- Before you change the state (active or passive) of a PIM interface using the Interface Type field, first disable PIM to prevent instability in the PIM operations, especially when neighbors are present or when streams are received.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.
2	Click the **Interfaces** tab. The PIM dialog box, Interfaces tab appears.
3	To configure properties, modify the values in the boxes provided.
4	Click **Apply**.

--End--

Variable definitions

The following table describes the Interfaces tab fields.

Field	Description
IfIndex	Specifies the VLANs configured for PIM-SM.
Address	Specifies the IP address of the PIM-SM VLAN.
NetMask	Specifies the network mask for the PIM-SM VLAN.
Enable	Specifies the status of PIM-SM on the VLAN: enabled (true) or disabled (false).
Mode	Specifies the PIM mode. The Ethernet Routing Switch 5000 Series Software supports only the sparse mode.
DesignatedRouter	Specifies the router with the highest IP address on a LAN designated to perform the DR tasks.
HelloInterval	Specifies the interval (in seconds) the switch waits between sending hello message to neighboring switches. The default is 30 seconds.
JoinPruneInterval	Specifies the interval (in seconds) the switch waits between sending join/prune message to its upstream neighbors. The default is 60 seconds.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
CBSRPreference	Specifies the preference for this local interface to become a C-BSR. The C-BSR with the highest BSR-priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a C-BSR.
InterfaceType	Specifies the type of interface: active or passive.
OperState	Indicates the operating status of PIM-SM on this interface: up or down.

Specifying the router as a candidate BSR on a VLAN interface

Because PIM-SM cannot run without a bootstrap router (BSR), you must specify at least one C-BSR in the domain. Any additional C-BSRs provide backup protection in case the primary BSR fails.

The C-BSR with the highest configured priority becomes the BSR for the domain. If two C-BSRs have equal priority, the candidate with the higher IP address becomes the BSR. If you add a new C-BSR with the highest priority to the domain, it automatically becomes the new BSR.

With the Ethernet Routing Switch 5000 Series, you can configure every PIM-enabled interface as a C-BSR.

Navigation

- “Setting the C-BSR priority using the VLAN menu” (page 450)
- “Setting the C-BSR priority using the IP Routing menu” (page 451)

Setting the C-BSR priority using the VLAN menu

Use this procedure to set the C-BSR priority on a VLAN from the VLAN menu.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **VLAN > VLANs**. The VLAN dialog box appears with the Basic tab displayed.
2	Select the VLAN ID that you want to configure with PIM. Several buttons on the bottom of the dialog box become available.
3	Click **IP**. The IP, VLAN dialog box appears with the IP Address tab displayed.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

4 Click the **PIM** tab.

The PIM tab appears.

5 In the **CBSRPreference** box, type the value of the C-BSR priority.

6 Click **Apply**.

--End--

Setting the C-BSR priority using the IP Routing menu

Use this procedure to set the C-BSR priority on a VLAN from the IP Routing menu.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM** . The PIM dialog box appears with the Globals tab displayed.
2	Click the **Interfaces** tab. The PIM dialog box appears with the Interfaces tab displayed.
3	In the **CBSRPreference** box, type the value of the C-BSR priority for the associated interface. The Candidate BSR with the highest BSR-priority and address is the preferred BSR. The default is –1, which indicates that the current interface is not a Candidate BSR; the range is 0 to 255.
4	Click **Apply**.

--End--

Displaying the current BSR

Display the current BSR information to review the configuration.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.
2	Click the **Current BSR** tab.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The Current BSR tab appears.

--End--

Variable definitions

The following table describes the Current BSR tab fields.

Field	Description
Address	Specifies the IP address of the current BSR for the local PIM domain.
FragmentTag	Specifies a randomly generated number that distinguishes fragments belonging to different bootstrap messages. Fragments belonging to the same bootstrap message carry the same Fragment Tag.
HashMask	Specifies the mask used in the hash function to map a group to one of the C-RPs from the RP set. With the hash mask, a small number of consecutive groups can always hash to the same RP.
Priority	Specifies the priority of the current BSR. The candidate BSR (C-BSR) with the highest BSR priority, and address (referred to as the preferred BSR) is elected as the BSR for the domain.
BootStrapTimer	Specifies the interval (in seconds) that the elected BSR waits between originating bootstrap messages.

Specifying a local IP interface as a candidate RP

Because PIM-SM cannot run without an RP, you must specify at least one C-RP in the domain. Use this procedure to configure a local PIM-SM interface as a candidate RP (C-RP).

With the Ethernet Routing Switch 5000 Series, you can configure only one local interface as a C-RP for any number of groups.

Using the GroupMask value, you can configure a C-RP for several groups in one configuration. For example, with a C-RP configuration with a GroupAddress value of 224.0.0.0 and a GroupMask of 240.0.0.0, you can configure the C-RP for a multicast range from 224.0.0.0 to 239.255.255.255.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.
2	Click the **Candidate RP** tab. The Candidate RP tab appears.
3	Click **Insert**. The PIM, Insert Candidate dialog box appears.
4	In the **GroupAddress** box, enter the multicast group address.
5	In the **GroupMask** box, enter the multicast group mask.
6	In the **RPAddress** box, enter the address of the C-RP.
7	Click **Insert**.

--End--

Variable definitions

The following table describes the PIM, Insert Candidate RP dialog box fields.

Field	Description
GroupAddress	Specifies the IP address of the multicast group. Together with the group mask, the group address identifies the prefix that the local router uses to advertise itself as a C-RP.
GroupMask	Specifies the address mask of the multicast group. Together with the group address, the group mask identifies the prefix that the local router uses to advertise itself as a C-RP.
RPAddress	Specifies the IP address of the C-RP. This address must be one of the local PIM-SM enabled interfaces.

Displaying the active RP

Use the following procedure to display the active RP.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

2 Click the **Active RP** tab.

The Active RP tab appears.

--End--

Variable definitions

The following table describes the Active RP dialog box fields.

Field	Description
GroupAddress	Specifies the IP address of the multicast group.
GroupMask	Specifies the address mask of the multicast group.
ActiveRP	Specifies the IP address of the active RP.
Priority	Specifies the priority of the active RP.

Enabling static RP

Enable static RP to avoid the process of dynamically learning C-RPs through the BSR mechanism. With this feature, static RP-enabled Ethernet Routing Switch 5000 Series switches can communicate with switches from other vendors that do not use the BSR mechanism.

ATTENTION
When you enable static RP, all dynamically learned BSR information is lost. However, if you disable static RP, the switch loses the static RP information and regains the BSR functionality.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.
2	Verify that the **Enable** box is selected to ensure that PIM-SM is globally enabled.
3	Select the **StaticRP** check box.
4	Click **Apply**.

--End--

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Configuring a static RP

Use this procedure to configure a static RP entry. After you configure static RP, the switch ignores the BSR mechanism and uses the statically-configured RPs only.

Prerequisites

- You must enable static RP.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.
2	Click the **Static RP** tab. The Static RP tab appears.
3	Click **Insert**. The PIM, Insert Static RP dialog box appears.
4	In the **GroupAddress** box, type the multicast group address.
5	In the **GroupMask** box, type the multicast group mask.
6	In the **Address** box, enter the address of the static RP.
7	Click **Insert**.

--End--

Variable definitions

The following table describes the Static RP tab fields.

Field	Description
GroupAddress	Specifies the IP address of the multicast group. Together with the group mask, the IP address identifies the range of the multicast addresses that the RP handles.
GroupMask	Specifies the address mask of the multicast group. Together with the group address, the address mask identifies the range of the multicast addresses that the RP handles.
RPAddress	Specifies the IP address of the static RP.
Status	Shows the current status of the static RP entry. The status is valid when the switch has a unicast route to the network for the static RP and is invalid otherwise.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Specifying a virtual neighbor on an interface

Configure a virtual neighbor when the next hop for a static route cannot run PIM-SM, such as a Virtual Redundancy Router Protocol address on an adjacent device.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.
2	Click the **Virtual Neighbors** tab. The Virtual Neighbors tab appears.
3	Click **Insert**.
4	In the **NeighborIfIndex** box, click **VLAN**.
5	Select the desired VLAN, and then click **OK**.
6	In the **NeighborAddress** box, enter the IP address of the virtual neighbor.
7	Click **Insert**.

--End--

Variable definitions

The following table describes the Neighbors tab fields.

Field	Description
NeighborIfIndex	Specifies the VLAN ID of the interface used to reach this PIM virtual neighbor.
NeighborAddress	Specifies the IP address of the PIM virtual neighbor.

Displaying PIM-SM neighbor parameters

View PIM-SM neighbor parameters to troubleshoot connection problems or review the configuration.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.
2	Click the **Neighbors** tab.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The Neighbors tab appears.

--End--

Variable definitions

The following table describes the Neighbors tab fields.

Field	Description
Address	Specifies the IP address of the PIM neighbor.
IfIndex	Specifies the VLAN ID of the interface used to reach this PIM neighbor.
UpTime	Specifies the elapsed time since this PIM neighbor last became a neighbor of the local router.
ExpiryTime	Specifies the time remaining before this PIM neighbor times out.

Displaying the PIM-SM RP set

Display the RP set for troubleshooting purposes. The BSR constructs the RP set from C-RP advertisements and then distributes it to all PIM routers in the PIM domain for the BSR.

Procedure steps

Step	Action
1	From the Device Manager menu bar, select **IP Routing > PIM**. The PIM dialog box appears with the Globals tab displayed.
2	Click the **RP Set** tab. The RP Set tab appears.

--End--

Variable definitions

The following table describes the RP Set tab fields.

Field	Description
GroupAddress	Specifies the IP address of the multicast group. Together with the group mask, the group address identifies the prefix that a router uses to advertise itself as a C-RP.
GroupMask	Specifies the address mask of the multicast group. Together with the group address, the group mask identifies the prefix that a router uses to advertise itself as a C-RP.
Address	Specifies the IP address of the C-RP.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Field	Description
HoldTime(sec)	Indicates the time specified in a C-RP advertisement that the BSR uses to time out the RP. After the BSR receives an advertisement for the RP, it restarts the timer. If no advertisement arrives before the timer expires, the BSR removes that RP from the RP set.
ExpiryTime	Specifies the time remaining before this C-RP times out.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

IGMP snooping configuration using Web-based management

This chapter describes the procedures you can use to configure IGMP snooping on a VLAN using Web-based management.

IGMP snooping configuration navigation

- "Configuring IGMP snooping" (page 459)
- "Displaying multicast membership" (page 460)

Configuring IGMP snooping

Use this procedure to configure IGMP using Web-based management.

Procedure steps

Step	Action
1	From the main menu, select **Applications > IGMP > IGMP Configuration**.
2	In the **Action** column, click the icon for the VLAN to be configured. The IGMP: VLAN Configuration window appears. The configurations made in this page affect only the VLAN specified in the VLAN field.
3	To enable IGMP snooping, select **Enabled** from the **Snooping** field.
4	To enable IGMP proxy, select **Enabled** from the **Proxy** field.
5	To configure static mrouter ports, select the ports from the appropriate **Static Router Ports** box (**Version 1** for IGMPv1 or **Version 2** for IGMPv2)
6	To configure the robustness value and query interval, modify the values in the fields provided.
7	Click **Submit**.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The VLAN Configuration page refreshes and the settings are saved.

--End--

Variable definitions

The following table describes the fields in the IGMP VLAN Setting window.

Fields	Description
VLAN	Specifies the VLAN ID.
Snooping	Enables or disables the IGMP snooping feature. The default setting is Disabled.
Proxy	Enables or disables the IGMP proxy feature. This feature lets the switch consolidate IGMP Host Membership Reports received on its downstream ports and generate a consolidated proxy report for forwarding to its upstream neighbor. The default setting is Disabled.
Robust Value	Specifies the robustness value. This feature lets you set the switch to offset expected packet loss on a subnet. If packet losses on a subnet are unacceptably high, the Robust Value field can be increased to a higher value. The default setting is 2.
Query Time	Specifies the query time (in seconds). This feature lets you to control the number of IGMP messages allowed on the subnet by varying the query interval (the interval between general queries sent by the multicast router). The default setting is 125 seconds.
Static Router Ports (Version 1 and Version 2)	Specifies static mrouter ports associated with the VLAN.

Displaying multicast membership

Use the following procedure to display Multicast membership using Web-based management.

Procedure steps

Step	Action
1	From the main menu, select **Applications > IGMP > Multicast Group**.
2	From the **VLAN** list in the **Multicast Group Membership Selection (View By)** window, select the VLAN for which to display group membership.
3	Click **Submit**.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

The membership information appears in the Multicast Group Membership Table.

--End--

Variable definitions

The following table describes the Multicast Group Membership Table fields.

Field	Description
Multicast Group Address	Specifies the IP Multicast group addresses that are currently active on the associated port.
Port	Specifies the port numbers associated with the IP Multicast group addresses displayed in the IP Multicast Group Address field.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

DHCP Relay configuring using Web-based management

This chapter contains procedures for DHCP relay configuration using Web-based management.

DHCP Relay configuration navigation

- "Configuring DHCP relay status" (page 463)
- "Creating a forwarding path" (page 463)
- "Editing a forwarding path" (page 464)
- "Configuring DHCP relay VLAN settings" (page 464)

Configuring DHCP relay status

Step	Action
1	Select **Applications, DHCP Relay** from the menu.
2	Click **DHCP Relay Configuration**.
3	Select the operational status of DHCP relay from the **DHCP Relay** field in the **DHCP Relay Setting** area.
4	Click **Submit**.

--End--

Creating a forwarding path

Step	Action
1	Select **Applications, DHCP Relay** from the menu.
2	Click **DHCP Relay Configuration**.
3	In the **Agent Address** field, enter the interface address.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

4 In the **Server Address** field, enter the server address.

5 Select the DHCP relay mode from the **Mode** field.

6 Click **Submit**.

--End--

Editing a forwarding path

Step	Action
1	Select **Applications, DHCP Relay** from the menu.
2	Click **DHCP Relay Configuration**.
3	Use the fields provided in the **DHCP Forwarding Path** area to edit the existing forwarding path.
4	Click **Submit**.

--End--

Configuring DHCP relay VLAN settings

Step	Action
1	Select **Applications, DHCP Relay** from the menu.
2	Click **DHCP Relay VLAN**.
3	Using the fields provided, edit the DHCP relay settings for the desired VLAN.
4	Click **Submit**.

--End--

Variable definitions

Variable	Definition
IfIndex	Indicates the VLAN interface index.
INTERFACE Address	Indicates the interface IP address of the DHCP relay agent.
MIN_SEC	Indicates the minimum time, in seconds, to wait between receiving a DHCP packet and forwarding the DHCP packet to the destination device. A value of zero indicates forwarding is done immediately without delay.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Variable	Definition
ENABLED	Indicates whether DHCP relay is enabled on the VLAN.
MODE	Indicates the type of DHCP packets this interface supports. Options include none, BootP, DHCP, and both.
ALWAYS_BROADCAST	Indicates whether DHCP reply packets are broadcast to the DHCP client on this VLAN interface.
REQUESTS Counter	Indicates the number of DHCP requests.
REPLIES Counter	Indicates the number of DHCP replies.
Clear Counters	Use this procedure to clear the DHCP relay counters for a VLAN.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

Nortel Ethernet Routing Switch 5000 Series

Configuration — IP Routing Protocols

Release: 6.1
Publication: NN47200-503
Document revision: 05.01
Document release date: 25 May 2009

Copyright © 2005-2009 Nortel Networks. All Rights Reserved.

While the information in this document is believed to be accurate and reliable, except as otherwise expressly agreed to in writing NORTEL PROVIDES THIS DOCUMENT "AS IS" WITHOUT WARRANTY OR CONDITION OF ANY KIND, EITHER EXPRESS OR IMPLIED. The information and/or products described in this document are subject to change without notice.

Nortel, Nortel Networks, the Nortel logo, and the Globemark are trademarks of Nortel Networks.

All other trademarks are the property of their respective owners.

To provide feedback or to report a problem in this document, go to www.nortel.com/documentfeedback.

www.nortel.com