

Version 8.0

Part No. NN46110-504 02.01

315898-F Rev 01

13 October 2008

Document status: Standard

600 Technology Park Drive
Billerica, MA 01821-4130

Nortel VPN Router Configuration — Routing



Copyright © 2008 Nortel Networks. All rights reserved.

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

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

Trademarks

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

Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.

Microsoft, Windows, Windows NT, and MS-DOS are trademarks of Microsoft Corporation.

All other trademarks and registered trademarks are the property of their respective owners.

Restricted rights legend

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

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

Statement of conditions

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

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

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

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

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

Nortel Networks Inc. software license agreement

This Software License Agreement (“License Agreement”) is between you, the end-user (“Customer”) and Nortel Networks Corporation and its subsidiaries and affiliates (“Nortel Networks”). PLEASE READ THE FOLLOWING

CAREFULLY. YOU MUST ACCEPT THESE LICENSE TERMS IN ORDER TO DOWNLOAD AND/OR USE THE SOFTWARE. USE OF THE SOFTWARE CONSTITUTES YOUR ACCEPTANCE OF THIS LICENSE AGREEMENT. If you do not accept these terms and conditions, return the Software, unused and in the original shipping container, within 30 days of purchase to obtain a credit for the full purchase price.

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

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

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

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

4. General

- a. If Customer is the United States Government, the following paragraph shall apply: All Nortel Networks Software available under this License Agreement is commercial computer software and commercial computer software documentation and, in the event Software is licensed for or on behalf of the United States Government, the respective rights to the software and software documentation are governed by Nortel Networks standard commercial license in accordance with U.S. Federal Regulations at 48 C.F.R. Sections 12.212 (for non-DoD entities) and 48 C.F.R. 227.7202 (for DoD entities).

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

Contents

Preface	11
Before you begin	11
Text conventions	11
Related publications	14
Printed technical manuals	15
Finding the latest updates on the Nortel Web site	16
Getting help from the Nortel Web site	16
Getting help over the phone from a Nortel Solutions Center	16
Getting help from a specialist by using an Express Routing Code	17
Getting help through a Nortel distributor or reseller	17
New in this release	19
Features	19
IGMP proxy for client tunnels	19
Chapter 1	
Routing overview	21
Routing fundamentals	21
Integrated firewall and routing	22
Dynamic routing	23
VPN routing	23
Static routes	23
Route table	24
Routing status	24
Chapter 2	
Route table and default routes	27
Route table and default route fundamentals	27
Route table lookup	29
Route selection based on destination	30
Route selection based on precedence in route table	30
Viewing and searching the route table	31

Showing route table information	32
Configuring default routes	33
 Chapter 3	
RIP configuration	35
 RIP fundamentals	35
Protecting against routing loops	37
RIP configuration	37
Configuring RIP interfaces	37
Configuring RIP globally	39
Enabling RIP on branch office tunnels	39
Showing RIP interface information	40
Configuring RIP for branch office tunnels	42
 Chapter 4	
OSPF configuration	45
 OSPF fundamentals	45
Installing the Advanced Routing key	46
Virtual link support	47
OSPF configuration	47
Configuring OSPF interfaces	47
Configuring OSPF globally	49
Viewing global OSPF information	51
Configuring OSPF for branch offices	54
 Chapter 5	
BGP configuration	57
 BGP fundamentals	57
RFCs	58
EBGP and IBGP peers	58
BGP peering and connection processing	58
BGP update processing	59
Unfeasible route processing	59
Feasible route processing	59
Path attribute processing	59

Keep Alive processing	61
BGP policies	61
Accept and announce policies	62
Access (Prefix) lists	62
AS-Path regular expressions	63
Route maps	64
Multihop BGP	66
Route reflector	66
BGP communities	67
Health check support	69
Installing the Border Gateway key	70
BGP configuration	70
Adding a route map	71
Configuring route maps	71
Configuring BGP interfaces	72
Configuring neighbors	74
Adding a network	75
Configuring the Route Reflector	75
Configuring AS Path access lists	76
Configuring community lists	77
Chapter 6	
Static route configuration	79
Static route configuration	79
Enabling static routes	79
Configuring static routes	80
Viewing static route information	80
Configuring public default routes	81
Configuring private default routes	81
Pinging to validate public default route	82
Chapter 7	
RPS configuration	85
RPS fundamentals	85
Redistribution of routes	87

RPS configuration	88
Creating a policy list	88
Editing a policy list	89
Configuring RPS	89

Chapter 8

Client address redistribution 91

Client address redistribution fundamentals	91
Configuring client address redistribution	95
Viewing client address redistribution information	95

Chapter 9

Multicast relay configuration 97

Multicast relay fundamentals	97
Configuring multicast relay	99
Viewing multicast relay information	99

Chapter 10

IGMP configuration 101

IGMP fundamentals	101
IGMP modes	101
Router mode	102
Host mode	102
IGMP versions	103
IGMPv1	103
IGMPv2	104
IGMPv3	104
IGMP version interoperability	104
IGMP message types	104
IGMPv1 and IGMPv2 messages	105
IGMPv3 messages	107
Membership Queries	107
Membership Reports	109
Host Leave messages	112
IGMP MIB considerations	113

IGMP configuration	113
Disabling multicast relay	113
Enabling split tunneling	114
Configuring IGMP on an interface	114
Configuring IGMP globally	115
Configuring IGMP on branch offices	115
 Chapter 11	
VRRP configuration	117
 VRRP fundamentals	117
VRRP and dynamic routing for high availability	118
Interface groups and critical interface failover	122
VRRP configuration	123
Configuring VRRP for LAN and VLAN	123
Configuration examples of IP addresses for backups	125
Configuring interface groups	128
 Chapter 12	
ECMP configuration	129
 ECMP fundamentals	129
Configuring ECMP	129
Index	131

Preface

This guide describes the Nortel VPN Router routing. It also provides information to help you configure routing.

Before you begin

This guide is for network managers who set up and configure the Nortel VPN Router. This guide is based on the assumption that you have experience with windowing systems or graphical user interfaces (GUI) and that you are familiar with network management.

Text conventions

This guide uses the following text conventions:

- | | |
|--------------------------|--|
| angle brackets (<>) | Indicates that you choose the text to enter based on the description inside the brackets. Do not enter the brackets when you enter the command.
Example: If the command syntax is ping <ip_address> , you enter ping 192.32.10.12 |
| bold Courier text | Indicates command names and options and text that you need to enter.
Example: Use the show health command.
Example: Enter terminal paging {off on} . |

braces ({ })	<p>Indicates required elements in syntax descriptions if more than one option exists. You must choose only one option. Do not enter the braces when you enter the command.</p> <p>Example: If the command syntax is ldap-server source {external internal}, you must enter either ldap-server source external or ldap-server source internal, but not both.</p>
brackets ([])	<p>Indicates optional elements in syntax descriptions. Do not enter the brackets when you enter the command.</p> <p>Example: If the command syntax is show ntp [associations], you can enter either show ntp or show ntp associations.</p> <p>Example: If the command syntax is default rsvp [token-bucket {depth rate}], you can enter default rsvp, default rsvp token-bucket depth, or default rsvp token-bucket rate.</p>
ellipsis points (. . .)	<p>Indicates that you repeat the last element of the command as needed.</p> <p>Example: If the command syntax is more diskn:<directory>/...<file_name>, you enter more and the fully qualified filename.</p>
<i>italic text</i>	<p>Indicates new terms, book titles, and variables in command syntax descriptions. If a variable is two or more words, if an underscore connects the words.</p> <p>Example: If the command syntax is ping <ip_address>, <i>ip_address</i> is one variable and you substitute one value.</p>
plain Courier text	<p>Indicates system output, for example, prompts and system messages.</p> <p>Example: File not found.</p>

separator (,)	Shows menu paths. Example: Choose Status, Health Check.
vertical line ()	Separates choices for command keywords and arguments. Enter only one choice. Do not enter the vertical line when you enter the command. Example: If the command syntax is terminal paging {off on} , you enter either terminal paging off or terminal paging on , but not both.

Related publications

For more information about the Nortel VPN Router, see the following publications:

- Release notes provide the most recent information, including brief descriptions of the new features, problems fixed in this release, and known problems and workarounds.
- *Nortel VPN Router Configuration—Client* (NN46110-306) provides information to install and configure client software for the Nortel VPN Router.
- *Nortel VPN Router Configuration—TunnelGuard* (NN46110-307) provides information to configure and use the TunnelGuard feature.
- *Nortel VPN Router Upgrades—Server Software Release 8.0* (NN46110-407) provides information to upgrade the server software to the most recent release.
- *Nortel VPN Router Installation and Upgrade—Client Software Release 8.01* (NN46110-409) provides information to upgrade the Nortel VPN Client to the most recent release.
- *Nortel VPN Router Configuration—Basic Features* (NN46110-500) introduces the product and provides information about initial setup and configuration.
- *Nortel VPN Router Configuration—SSL VPN Services* (NN46110-501) provides instructions to configure services on the SSL VPN Module 1000, including authentication, networks, user groups, and portal links.
- *Nortel VPN Router Configuration—Advanced Features* (NN46110-502) provides configuration information for advanced features such as the Point-to-Point Protocol (PPP), Frame Relay, and interoperability with other vendors.
- *Nortel VPN Router Configuration—Tunneling Protocols* (NN46110-503) provides configuration information for the tunneling protocols IPsec, Layer 2 Tunneling Protocol (L2TP), Point-to-Point Tunneling Protocol (PPTP), and Layer 2 Forwarding (L2F).
- *Nortel VPN Router Using the Command Line Interface* (NN46110-507) provides syntax, descriptions, and examples for the commands that you can use from the command line interface (CLI).
- *Nortel VPN Router Configuration—Firewalls, Filters, NAT, and QoS* (NN46110-508) provides instructions to configure the Stateful Firewall and Nortel VPN Router interface and tunnel filters.

- *Nortel VPN Router Security—Servers, Authentication, and Certificates* (NN46110-600) provides instructions to configure authentication services and digital certificates.
- *Nortel VPN Router Troubleshooting—Server* (NN46110-602) provides information about system administrator tasks such as recovery and instructions to monitor Nortel VPN Router status and performance. This document provides troubleshooting information and event log messages.
- *Nortel VPN Router Administration* (NN46110-603) provides information about system administrator tasks such as backups, file management, serial connections, initial passwords, and general network management functions.
- *Nortel VPN Router Troubleshooting—Client* (NN46110-700) provides information to troubleshoot installation and connectivity problems with the Nortel VPN Client.

Printed technical manuals

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

How to get Help

This section explains how to get help for Nortel products and services.

Finding the latest updates on the Nortel Web site

The content of this documentation was current at the time the product was released. To check for updates to the latest documentation and software for **Nortel VPN Router**, click one of the following links:

Link	Website
Most recent software	Nortel page for Nortel VPN Router software located at support.nortel.com/go/main.jsp?cscat=SOFTWARE&poid=12325
Most recent documentation	Nortel page for Nortel VPN Router documentation located at support.nortel.com/go/main.jsp?cscat=DOCUMENTATION&poid=12325

Getting help from the Nortel Web site

The best way to get technical support for Nortel products is from the Nortel Technical Support Web site:

www.nortel.com/support

This site provides quick access to software, documentation, bulletins, and tools to address issues with Nortel products. From this site, you can

- download software, documentation, and product bulletins
- search the Technical Support Web site and the Nortel Knowledge Base for answers to technical issues
- sign up for automatic notification of new software and documentation for Nortel equipment
- open and manage technical support cases

Getting help over the phone from a Nortel Solutions Center

If you do not find the information you require on the Nortel Technical Support Web site, and you have a Nortel support contract, you can also get help over the phone from a Nortel Solutions Center.

In North America, call 1-800-4NORTEL (1-800-466-7835).

Outside North America, go to the following web site to obtain the phone number for your region:

www.nortel.com/callus

Getting help from a specialist by using an Express Routing Code

To access some Nortel Technical Solutions Centers, you can use an Express Routing Code (ERC) to quickly route your call to a specialist in your Nortel product or service. To locate the ERC for your product or service, go to:

www.nortel.com/erc

Getting help through a Nortel distributor or reseller

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

New in this release

The following sections detail what is new in *Nortel VPN Router Configuration — Routing* for Release 7.0.

Features

See the following sections for information about feature changes.

IGMP proxy for client tunnels

The Internet Group Management Protocol (IGMP) is the communications protocol used to manage the membership of Internet Protocol (IP) multicast groups. Multicast hosts use IGMP to signal requests to the Nortel VPN Router to join specific multicast groups and to begin receiving group traffic. Using the Query-Response Model, the multicast router can determine host membership for various multicast groups.

For more information about IGMP Proxy for client tunnels, see [“IGMP fundamentals” on page 101](#).

Chapter 1

Routing overview

This chapter contains an overview of routing for the Nortel VPN Router, including the following topics:

- [“Routing fundamentals” on page 21](#)
- [“Integrated firewall and routing” on page 22](#)
- [“Dynamic routing” on page 23](#)
- [“VPN routing” on page 23](#)
- [“Static routes” on page 23](#)
- [“Route table” on page 24](#)
- [“Routing status” on page 24](#)

Routing fundamentals

The Nortel VPN Router uses Secure Route Technology (SRT) to forward network traffic. SRT operates on the premise that trusted and untrusted portions exist within the network. Trusted interfaces are placed on secure network segments (such as the private LAN) and behave like traditional routed interfaces. Untrusted interfaces are placed on unsecure network segments (such as the Internet) where all insecure services are disabled. Only services considered secure can run on, or are accessible through, untrusted interfaces.

To provide this protection, you use features such as packet filtering and antispoofing to enable either the integrated Nortel VPN Router Stateful Firewall or the Nortel VPN Router tunnel filter.

[“Forwarding capabilities” on page 22](#) is a matrix of Nortel VPN Router forwarding capabilities between the source interface and destination interfaces.

Table 1 Forwarding capabilities

	Private	Public	Client tunnel	Branch office tunnel	System management
private	yes (1)	yes (1)	yes	yes	yes
public	yes (1)	yes (1)	yes (1)	yes (1)	yes (3)
client tunnel	yes	yes (1)	yes (2)	yes (2)	yes
branch office tunnel	yes	yes (1)	yes (2)	yes (2)	yes
system management	yes	yes (3)	yes	yes	not applicable

- 1.Nortel VPN Router Stateful Firewall must be enabled.
- 2.Must be enabled under SystemForwarding (disabled by default).
- 3.Only RADIUS, CMP, and CRL retrieval permitted.

Integrated firewall and routing

The Nortel VPN Router is a security device. Therefore, the routing configuration takes effect as it relates to the integrated firewall configuration of the Nortel VPN Router. In the following sections, references to “integrated firewall” mean the Nortel VPN Router Firewall option on the Services, Firewall window. Use this option by selecting either Nortel VPN Router Stateful Firewall or Nortel VPN Router interface filter. However, if you use the Nortel VPN Router interface filter option, you do not need a firewall license.

Dynamic routing

Dynamic routing protocols are available for private physical interfaces or branch office tunnel interfaces. Public interfaces are not trusted and therefore you cannot configure them to run a dynamic routing protocol. The exception is Border Gateway Protocol (BGP), which you can enable on public interfaces on request. You can configure physical LAN and WAN interfaces as either a private or public interface except slot 0 interface 1, which is always a LAN and private.



Note: The Advanced Routing License Key is required to enable features such as Open Shortest Path First (OSPF), Equal Cost Multiple Paths (ECMP), and Inter Group Multicast Protocol (IGMP). Static routes, Routing Information Protocol (RIP), and route redistribution do not need this license. The Border Gateway Protocol License Key is required to enable BGP. Another option is to purchase the Premium Routing License to enable OSPF, ECMP, IGMP, and BGP.

VPN routing

VPN routing forwards traffic between tunnels or between tunnels and private interfaces. With VPN routing, traffic enters or exits the Nortel VPN Router through a tunnel.

Enhanced routing provides additional traffic patterns beyond traditional VPN routing. You must enable either the Nortel VPN Router Stateful Firewall or Nortel VPN Router filter to support the enhanced routing feature.

Static routes

You can configure static routes between Nortel VPN Routers if you do not have a dynamic routing protocol, such as OSPF, RIP, or BGP. Even if you have dynamic routing protocols, you can use static routes because they provide strong security. The Nortel VPN Router supports multiple default and static routes.

Route table

The route table contains the routes submitted by the routing protocols and the static route application and dynamic protocols, such as OSPF, RIP, and BGP. The route table manager (RTM) chooses the best routes from the route table to populate the IP forward table. The Nortel VPN Router uses the IP forward table to determine how forwarding occurs; it selects the best routes based on the following order of protocol preference:

- direct route
- static route
- BGP route
- OSPF route
- RIP route
- default route

The route preference and the weight and cost of the route factor into the RTM route selection.

Routing status

The Routing, Status window provides access to information about each routing protocol. It also provides access to the route table and route table manager (RTM) statistics. [“Routing status window options” on page 24](#) shows routing status window options.

Table 2 Routing status window options

Column	Description
BGP Summary	Displays the overall summary of BGP running on the Nortel VPN Router, including the router ID, Local AS, Admin state (enabled or disabled), Hold Interval, Keep Alive Interval, Local Preference, Default Metric, Route Reflector, Client Reflection, Cluster ID, Always Compare MED, Auto summary, Redistribute Internal, Synchronization, Max paths, and Number of Peers.
BGP Routes	Displays Search Type, IP Address, Mask, and Mask Type.
BGP Redistributed Routes	Displays IP Address, IP Mask, and Origin Type.

Table 2 Routing status window options

Column	Description
BGP Neighbors Routes	Displays Routes Type and Neighbor.
BGP Neighbors Summary	Displays overall summary of Foreign Host, Remote AS, External Link, Remote Router ID, BGP state, Up For, Hold Time, KeepAlive Interval, Advertisement Runs, Received, Received Notifications, Sent, Community Attribute, Accepted Prefixes, Prefix Advertised, Local Host, Local Port, Foreign Host, Foreign Port, Connections Established, Elapsed Time Between Updated Msg, MinASOriginationInterval Timer.
IGMP Group Summary	Displays Group Address, In Interface, Time Left, Up, Last Reporter, Static.
IGMP Interfaces Summary	Includes Interface IP address, Upstream status, number of groups, DR IP address, and counts for queries and reports.
IGMP Groups per Interfaces Summary	Displays information for each interface, including Group Address, In Interface, Time Left, Up, Last Reporter, Static.
OSPF LSDB	Displays information about link state databases in all areas that are known to OSPF, including link state type, ID, advertising router address, metric, ASE, forward address, age, and sequence number for each area.
OSPF Neighbor	Displays information about neighbors on all interfaces that run OSPF, including the IP interface address, router ID, neighbor IP address, state, and dead time priority.
OSPF Interfaces	Displays information about interfaces configured for OSPF, including the IP address of the interface, the area to which the interface belongs, the type of interface, the state, cost and the designated router in the area to which the interface belongs.
OSPF Summary	Displays overall summary of OSPF running on the Nortel VPN Router, including the router ID, global state (up or down), whether an area border router or autonomous system border router.
OSPF Statistics	Displays information about System-wide OSPF statistics.
RIP Database	Displays all routes that RIP can distribute (based on routing priorities).
RIP Interfaces	Displays interfaces that you configured for RIP.
RIP Statistics	Displays system-wide RIP statistics.
VRRP Config	Displays VRRP configuration information.
VRRP Errors	Displays system-wide VRRP errors that occurred.
VRRP Statistics	Displays system-wide VRRP statistics.
Route Table	Displays full routing for all routes, including next hops and best routes.

Table 2 Routing status window options

Column	Description
Next Hop Table	Displays next hop address for each route.
Best Route Table	Displays forwarding table used to determine the best route.
Route Table Stats	Displays statistics about route table management that provides information about Nortel VPN Router traffic.
IP Forward Table	Displays information about the IP routes used to forward traffic.

Chapter 2

Route table and default routes

This chapter contains information about the route table and default routes and the procedures to view, search, and configure the tables for the Nortel VPN Router.

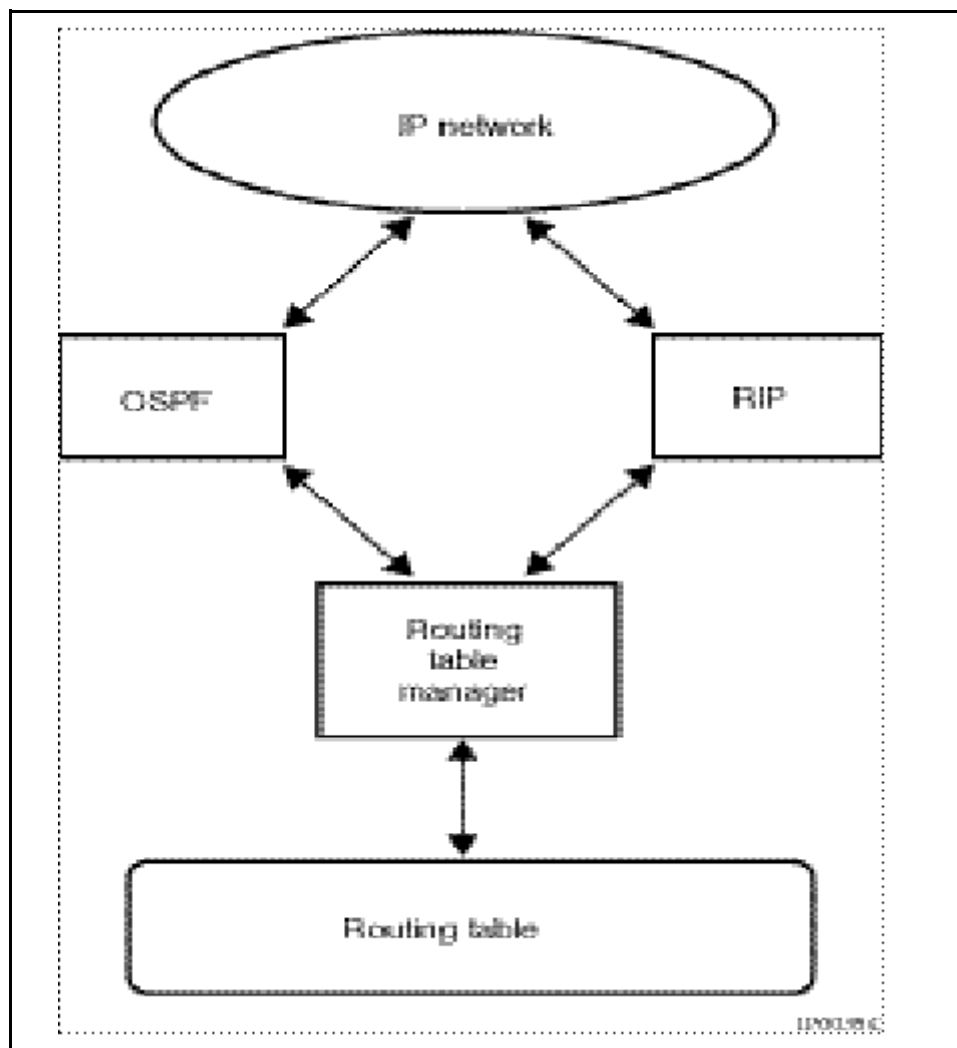
This chapter includes the following topics:

- [“Route table and default route fundamentals” on page 27](#)
- [“Route table lookup” on page 29](#)
- [“Viewing and searching the route table” on page 31](#)
- [“Configuring default routes” on page 33](#)

Route table and default route fundamentals

The route table defines where traffic is forwarded to reach its destination. The route table contains both static and dynamic routes. Static routes are manually configured routes that do not change. Dynamic routes are learned from Routing Information Protocol (RIP), Open Shortest Path First (OSPF) routing protocols, or Border Gateway Protocol (BGP) routing protocols.

[“Interaction of OSPF, BGP, and RIP with the routing table” on page 28](#) shows how various routing protocols interact with the route table manager.

Figure 1 Interaction of OSPF, BGP, and RIP with the routing table

The route table entries are divided into two groups: public and private. Because private interfaces are trusted and public interfaces are untrusted, dynamic routing protocols RIP and OSPF are permitted only on private interfaces and branch office tunnel interfaces. BGP is permitted on a public interface.

Public traffic has the following public routes:

- static routes to public interfaces

- dynamic (BGP only) routes to public interfaces
- default route to public interface

Private traffic has the following private routes:

- static routes to private interfaces
- dynamic routes to private interfaces
- static routes to branch office tunnel interfaces
- dynamic routes to branch office tunnel interfaces
- default route to private interface
- routes used for tunnels

After a packet arrives, the Nortel VPN Router performs a full lookup in the IP forwarding table to determine which route to use:

- If firewall support is enabled, all public and private routes in the IP forwarding table are available to the traffic.
- If firewall support is not enabled, only the private portion of the IP forwarding table is available.
- If the destination route is not found in the table, the public or private default route is invoked as described in [“Route table lookup” on page 29](#).

Route table lookup

The route table has two parts. One part contains the routes for traffic that uses the Nortel VPN Router public interfaces (untrusted network), and a second part has the routes for traffic that uses the private interfaces (trusted network). Tunnels are virtual interfaces and are treated as private interfaces.

The following list shows the types of routes in the Nortel VPN Router route table:

- Static routes
 - to public interfaces
 - to private interfaces
 - to branch office tunnel interfaces
- Dynamic routes

- To private interfaces
- To branch office tunnel interfaces
- To public interfaces (BGP only)
- Default routes
 - To public interfaces
 - To private interfaces
- Host routes
 - Routes added for VPN users (for example, Nortel VPN Router Clients or PPTP clients)
- Utunnel routes
 - Host or network routes for clients that log on using the client address redistribution feature

Route selection based on destination

The route to a specific destination is based on the most specific match. For example, if you have a route to the network 10.1.0.0/16 through next-hop router A and another route to 10.1.2.0/24 through next-hop router B, traffic destined to 10.1.2.1 transmits through router B, even though the address matches both 10.1.0.0/16 and 10.1.2.0/24. If B is not available, the traffic is forwarded to A.

Route selection based on precedence in route table

The route table selects the best routes submitted by the routing protocols and submits them to the forwarding table. The selection of best routes is based on the following order of precedence:

- 1** Direct routes
- 2** Static routes
- 3** BGP routes
- 4** OSPF routes
- 5** RIP routes
- 6** Default—static routes (locally defined default routes)
- 7** Default—BGP routes (learned from other routers through BGP redistribution)

- 8 Default—OSPF routes (learned from other routers through OSPF redistribution)
- 9 Default—RIP routes (learned from other routers through RIP redistribution)

You can use ECMP to load balance traffic across multiple paths for static routes, BGP routes, OSPF routes, or RIP routes of the same cost.

Viewing and searching the route table

You can view or search the route table, save it to a file, or view the IP forward table.

- 1 To view the route table, choose **Routing, Route Table**.

To search the route table, perform the following steps:

- 2 To search the route table, select **All**, **Host**, or **Network** from the **Destination** list.

If you select **Host** or **Network**, perform the following steps:

- a From the **Interface** list, select **All** or the address.
- b From the **Protocol** list, select **All** or the protocol (**BGP**, **OSPF**, **RIP**, **Static**, **Direct**, **Utnnel**, **NAT**, or **CLIP**). You must enter the IP address in the box.
- c Enter the IP address.

If you select **Network**, perform the following steps:

- a Enter the network mask.
- b From the **Search Type** list, choose **Exact** or **Best Match**.

- 3 Click **Search**.

To save the route table, perform the following steps:

- 4 Enter the file name in the **Filename** box.

You can save the route table as a text file in the directory `ide0/system/xxx`, where `xxx` is the name of the file that you specify.

- 5 Under **Route Filter**, select **Best Routes** to view all routes to a single or All Routes to view all destinations. The default is **Best Routes**.
- 6 Click **Save**.

Showing route table information

You can view route table information using the status buttons at the bottom of the Routing, Route Table window.

- 1 To check the route table status, click **IP Forward Table** to display the **IP Route Network Table**, the **IP Route Host Table**, and the **IP Public Address Table**.

[“IP Forward Table window” on page 32](#) describes the fields in the IP Forward Table window.

Table 3 IP Forward Table window

Column	Description
Destination/Mast	Displays the network address and mask
Gateway	Displays the IP address of next-hop Nortel VPN Router
Flags	Displays internal use flags
Ref	Displays the reference count
Use	Displays the number of times used
Intf	Displays the interface identifier
MTU	Displays the size of packet
OuterCtxt	(For internal use only)
CircMap	(For internal use only)
RtEntryP	(For internal use only)

- 2 Click **Route Table** to display the full internal route table.

“IP Route Table window” on page 33 describes the fields in the IP Route Table window.

Table 4 IP Route Table window

Column	Description
Seq	Displays the sequence number that shows the best route
P	Displays the protocol
IP Address/Netmask	Displays the IP address and network mask
Weight Cost	Displays combination of cost and priority for the best route
NextHop	Displays the IP address of the next hop
NextHopInterface	Displays the IP address of the next-hop interface
CId	Displays the Circuit ID

Configuring default routes

When the Nortel VPN Router receives traffic for which no matching route exists in the route table, it can use a default route. The use of default routes depends on several factors, such as whether integrated firewall support is enabled and where the traffic originated (for example, from the public or private interface).

A default Nortel VPN Router is the address of the next-hop router. Packets are routed through the default Nortel VPN Router onto the private or public network when the route table has no specific route to the destination.

You can configure both private and tunnel interfaces to be Public or Private; therefore, the following procedure applies to both interfaces.

- 1 Choose **Routing, Configuration**.
- 2 Select **Public** or **Private** as the **Outbound Routing Preference**:

- When **Public** is enabled, all packets that do not transmit across a tunnel to defined remote networks continue to transmit from the public interface using the public default Nortel VPN Router (0.0.0.0/32 in the forwarding table). Packets that go to defined remote networks transmit across the branch office tunnel and cannot have a remote network equal to 0.0.0.0/0.0.0.0 (default route). For example, if you want to reach the DNS server on the public network, select private-to-public for the routing decision.
- When **Private** is enabled, all packets transmit over your branch office tunnel and not out the public interface because the branch office tunnel has a 0.0.0.0/0.0.0.0 remote network (statically defined or received by RIP). For example, if you want to reach the DNS servers on the corporate side of the branch office tunnel, select private-to-private for the routing decision.

3 Click **OK**.

Chapter 3

RIP configuration

This chapter contains the information about Routing Information Protocol (RIP) fundamentals and the procedures to configure RIP on the Nortel VPN Router.

This chapter includes the following topics:

- [“RIP fundamentals” on page 35](#)
- [“Protecting against routing loops” on page 37](#)
- [“RIP configuration” on page 37](#)

RIP fundamentals

RIP is a distance-vector routing protocol by which routers exchange routing information by means of periodic RIP updates. Routers transmit RIP updates to neighboring subnets, and listen for RIP updates from the routers on those neighboring subnets. Routers use the information in the RIP updates to keep the internal routes current.

RIP computes distance as the number of hops (or routers) from the source subnet to the target subnet. RIP has a maximum hop count of 15 hops. Networks beyond 15 hops are considered unreachable.

RIP is one of the most common interior Nortel VPN Router protocols used. RIP Version 2 is backward compatible with RIP Version 1 and corrects many RIP Version 1 limitations, such as subnet routing, authentication, and multicast support for route messages.

The Nortel VPN Router supports RIP for routing traffic within the private network and between branch office connections. The Nortel VPN Router sends RIP broadcast or multicast messages at regular intervals. These messages contain information about routes that the Nortel VPN Router can reach. Other routers on the network listen for these messages, update the route tables, and then send route messages to the peer routers. Using Nortel VPN Router RIP, you can enable or disable propagation of RIP messages from the Nortel VPN Router private and branch office tunnel interfaces.



Note: The interface filters setting affects the behavior of routing protocols. For example, RIP uses User Datagram Protocol (UDP) as the transport mechanism, so if the interface filters are configured to deny UDP, RIP advertisements are dropped.

The Nortel VPN Router supports RIP Version 1 and Version 2. For more information about RIP, see the RFCs on the Internet Engineering Task Force (IETF) Web site at www.ietf.org.

- RFC 1058—Routing Information Protocol: Describes the Routing Information Protocol (RIP), which is loosely based on the program Routed, distributed with the 4.3 Berkeley Software Distribution. The specifications in this RFC represent a combination of features from various implementations of this program.
- RFC 1721—RIP Version 2 Protocol Analysis: Describes the key features of the RIP Version 2 protocol and the current implementation experience.
- RFC 1722—RIP Version 2 Protocol Applicability Statement: Describes how RIP Version 2, which is an extension to RIP Version 1, can be useful on the Internet.
- RFC 1723—RIP Version 2 Carrying Additional Information: Specifies an extension of the Routing Information Protocol (RIP) that expands the amount of useful information carried in RIP messages and that adds a measure of security.

Protecting against routing loops

A routing loop occurs when two or more routers continuously forward the same packet to each other until the time-to-live counter expires, or the network goes down. Loops typically occur after a new router is added to the network or after a router in an existing network is removed and the remaining routers must recalculate routes. A loop detection protocol helps prevent a routing loop and speeds up convergence while the situation corrects itself.

The Nortel VPN Router supports the following methods used by RIP minimize loops and to speed up the convergence that is caused by the normal correction of a loop:

- split horizon, where the Nortel VPN Router does not send routes that it learns from a neighboring router back to that same neighbor
- split horizon with poison reverse, where the Nortel VPN Router sends back the routes that it learns from a neighboring router and configures the metric for that route to infinity
- triggered updates, where the Nortel VPN Router sends an update immediately after a routing change on the Nortel VPN Router. By default, RIP updates the routes at regular intervals

RIP configuration

To configure RIP, perform the following procedures:

- 1 [“Configuring RIP interfaces” on page 37](#)
- 2 [“Configuring RIP globally” on page 39](#)
- 3 [“Enabling RIP on branch office tunnels” on page 39](#)
- 4 [“Configuring RIP for branch office tunnels” on page 42](#)

Configuring RIP interfaces

To configuring RIP interfaces, perform the following steps:

- 1 Choose **Routing, Interfaces**.

- 1 Click **Configure** for the RIP protocol.

The Routing Interfaces > Configure RIP window appears. The enabled check box indicates that you globally enabled RIP.

- 2 Select **V2**, **V1**, or **Off** as the transmit mode.

Using Transmit mode, you can specify which version of RIP to use to route traffic from this Nortel VPN Router. The default is V2. Selecting OFF specifies that RIP is not used.

- 3 Select **V2**, **V1**, **Both**, or **Off** as the receive mode.

Using Receive mode, you can specify which version of RIP accepts incoming traffic. The default is V2. Select OFF to specify that RIP is not used. Select BOTH to specify that incoming transmissions using either version of RIP are accepted.

- 4 Select **None**, **Simple**, or **MD5** as the authentication type that is used as part of the RIP transmission.

This authentication is specific to RIP and has no bearing on the authentication type as part of the connection to the Nortel VPN Router. The default is None, which specifies that no authentication is required. Simple indicates that authentication uses a simple password. MD5 specifies that authentication uses an MD5 secret. If you select either Simple or MD5, password and password confirmation boxes appear.

- 5 Enter a metric value for the **Cost**.

The Cost value is the cost of sending a packet on the interface expressed in the link state metric.

- 6 Select **Enabled** or **Disabled** for poison reverse.

Poison reverse updates routing loops in large networks.

- 7 If no default route is configured, you can select **Enabled** for Import Default Route to use the default route learned during RIP updates.

Typically, you specify a default route in the route table on the **Routing , Static Routes** window. The default is disabled.

- 8 Select **Enabled** to specify that the default route is exported during RIP updates or enter a metric value (1 through 15) to the default route.

- 9 Select **Enabled** to specify that static routes are exported during RIP updates or enter a metric value (1 through 15) to the default route.

- 10** Select **Enabled** to specify that OSPF routes are exported during RIP updates or enter a metric value (1 through 15) to the default route.
- 11** Select **Enabled** to specify that BGP routes are exported during RIP updates or enter a metric value (1 through 15) to the default route.
- 12** Select a metric value (1 through 15) to export the static routes metric if you have a branch office connection.

This metric informs the remote branch office connection of the routes that are used and provides the assigned metric value. The default is 1 and the maximum value is 15.

Configuring RIP globally

To configure RIP globally, perform the following steps:

- 1** Choose **Routing, RIP**.
- 2** Click **Enable**.
- 3** Enter the amount of time in seconds that you want RIP to update the routes.
The default is 30 seconds and the range of values is 5 to 65535 seconds. The hold-down timer is six times the update timer.
- 4** Select a metric value (1 through 4) Equal Cost MultiPath for the maximum number of RIP paths.

Enabling RIP on branch office tunnels

To enable RIP interfaces, perform the following steps:

- 1** Choose **Routing, RIP**.
- 2** Check **Enabled** to enable RIP on the interface. By default RIP is disabled.
- 3** Enter the interval of time in seconds for RIP to update the routes.
The supported range is 5 to 65535 seconds, with the default setting at 30 seconds. The RIP hold-down timer is automatically six times the update timer.

Showing RIP interface information

The **Configured Physical Interfaces** section lists the IP address and RIP configuration state (enabled or disabled) of each physical interface.

To view RIP interface information, perform the following steps:

- 1 Click **Statistics** to display statistics about RIP in the Nortel VPN Router.

[“RIP Statistics window” on page 40](#) describes the fields in the RIP Statistics window.

Table 5 RIP Statistics window

Column	Description
Global RIP Status	Displays the status: enabled or disabled
Update interval	Displays the interval in seconds
Route Change	Displays the number of routes changed
Query	Displays the number of queries sent
Interface	Displays the interface IP address
Cid	Displays the circuit ID
RxUpdates	Displays the number of receipt updates
TxUpdates	Displays the number of transmit updates
TxTrigUpd	Displays the number of transmit trigger updates
RxBadPkts	Displays the number of bad packets received
RxBadRoutes	Displays number of bad routes received

- 2 Click **Database** to display information for all RIP interfaces.

“[RIP Database window](#)” on page 41 describes the fields in the RIP Database window.

Table 6 RIP Database window

Column	Description
Circuit	Displays the circuit ID
Address	Displays the IP address
Mask	Displays the network mask of IP address
Owner	Displays the protocol
Cost	Displays the import cost of RIP routes
Metric	Displays the export metric of RIP routes
Gw	Displays the Nortel VPN Router IP address

- 3 Click **Interfaces** to display information for all RIP interfaces, including tunnels that run RIP.

“[RIP Interfaces window](#)” on page 41 describes the fields in the RIP Interfaces window.

Table 7 RIP Interfaces window

Column	Description
Ip	Displays the RIP interface IP address
IntfState	Displays whether the interface is up or down
Cid	Displays the circuit ID
PoisonRev	Displays whether enabled or disabled
ExpSMetric	Displays the metric (1–15) for export static route
ExpBgpMetric	Displays the metric (1–15) for export BGP route
Subnet	Displays the network mask of IP address
Auth	Displays the authentication type
RxMode	Displays whether RIP receive version supported
ImpDRoute	Displays whether the default import route is enabled or disabled
ExpDMetric	Displays the metric (1–15) for export default route
RipCost	Displays the metric for the RIP cost
RipEnabled	Displays whether RIP is enabled or disabled

Table 7 RIP Interfaces window

Column	Description
Type	Displays the interface type
TxMode	Displays whether RIP transmit version supported
ExpBOMetric	Displays the metric (1–15) export tunnel static route
ExpOspfMetric	Displays the metric (1–15) export OSPF route

Configuring RIP for branch office tunnels

To configure RIP for branch office tunnels, perform the following steps:

- 1 Choose **Profiles, Branch Office**.

- 2 Select a group.

- 3 Click **Configure**.

The Branch Office > Edit Group window appears.

- 4 Click **Configure** in the RIP section.

The Branch Office > Edit > RIP window appears.

- 5 Select **V2**, **V1**, or **Off** as the transmit mode.

Select Transmit mode to specify which version of RIP to use to route traffic from this Nortel VPN Router. The default is V2. Select OFF to specify that RIP is not used.

- 6 Select **V2**, **V1**, **Both**, or **Off** as the receive mode.

Select Receive mode to specify which version of RIP accepts incoming traffic. The default is V2. Select OFF to specify that RIP is not used. Select BOTH to specify that incoming transmissions using either version of RIP are accepted.

- 7 From the **Import Default Route** list, select **Enabled** to specify to use the default route learned during RIP updates.

Typically, you specify a default route in the route table in the **Routing, Static Routes** window. The default is Disabled.

- 8 From the **Export Default Route Metric** list, select **Enabled** to specify to export the default route during RIP updates or select a metric value (1 to 15) to the default route.

- 9 From the **Export Static Routes Metric** list, select **Enabled** to specify to export static routes during RIP updates or select a metric value (1 to 15) to the default route.
- 10 From the **Export Branch Office Static Routes Metric** list, select a metric value (1 to 15) to export the static routes metric if you have a branch office connection.

This metric informs the remote branch office connection of the routes that are used and provides the assigned metric value. The default is 1 and the map value is 15.
- 11 From the **Export OSPF Routes Metric** list, select **Enabled** to specify to export OSPF routes during RIP updates or select a metric value (1 to 15) to the default route.
- 12 From the **Export BGP Routes Metric** list, select **Enabled** to specify to export BGP routes during RIP updates or select a metric value (1 to 15) to the default route.
- 13 In the **Authentication Type** section, select **None**, **Simple**, or **MD5**.

This authentication is specific to RIP and has no bearing on the authentication performed as part of the connection to the Nortel VPN Router. The default is **None**, which specifies that no authentication is required. **Simple** indicates that authentication uses a simple password. **MD5** specifies that authentication uses an MD5 secret. If you select either **Simple** or **MD5**, password and password confirmation boxes appear.
- 14 In the **Password** box, enter a password.
- 15 Confirm the password by reentering the password in the **Confirm Password** box.
- 16 Click **OK**.

Chapter 4

OSPF configuration

This chapter contains information about Open Shortest Path First (OSPF) fundamentals and the procedures to configure OSPF on the Nortel VPN Router.

This chapter includes the following topics:

- [“OSPF fundamentals” on page 45](#)
- [“Installing the Advanced Routing key” on page 46](#)
- [“Virtual link support” on page 47](#)
- [“OSPF configuration” on page 47](#)

OSPF fundamentals

OSPF is a link-state routing protocol. With the link state information, a device that runs OSPF builds a shortest-path tree with itself as the root of the tree. The device can then identify the shortest path from itself to each destination and build the route table. Some of the benefits of OSPF are

- fast convergence with minimal routing protocol-related traffic after convergence
- Variable-length Subnet Masks (VLSM)
- hierarchical segmentation
- area routing to provide additional routing protection and a reduction in routing protocol traffic
- authentication
- virtual link support
- Equal Cost Multipath (ECMP) support

- multicast- or unicast-based route advertisement messages instead of broadcast-based advertisements

Using the Nortel VPN Router OSPF support, you can enable or disable OSPF on the Nortel VPN Router private and tunneled interfaces. OSPF supports broadcast and point-to-point network types and can act as autonomous boundary router (ABR), information retrieval (IR), autonomous system boundary router (ASBR), designated router (DR), and system designated router (SDR) router types. The Nortel VPN Router OSPF implementation conforms to OSPF 2 (RFC 2178).

The interface filters setting affects the behavior of routing protocols. For example, OSPF uses IP as the transport mechanism; therefore, if the interface filters are configured to deny IP, OSPF advertisements are not sent or received.

Installing the Advanced Routing key

You must install the Advanced Routing License key to enable OSPF on the Nortel VPN Router. The Firewall License Key is required only when the redistribution capabilities of RIP and OSPF are necessary.

To install a software license key, perform the following steps:

- 1 Choose **Admin, License Keys**.
- 2 Enter the key that you obtained from Nortel Customer Support in the box to the right of **Advanced Routing**.

If you obtained a Premium Routing Key, enter the key in the box to the right of **Premium Routing**.

- 3 Click **Install**.

After you install the key, the label **Installed** appears. You need to install a key only once on each Nortel VPN Router. Click **Remove** to remove the key. A confirmation message appears and, if you click **Yes**, the key is removed.



Note: The presence of the Advanced Routing License key is checked only when OSPF is globally enabled. If you enter the Advanced Routing Key, globally enable OSPF, and then delete the Advanced Routing Key, OSPF continues to run. However, if you then disable and reenabling OSPF, it no longer runs.

Virtual link support

OSPF requires that all nonbackbone areas have at least one connection to the backbone area (area 0). If an area has no physical connection to the backbone, a virtual link can be used to traverse an intermediate area to connect to the backbone area. The Nortel VPN Router must be an area border router for the automatic configuration of virtual links to operate properly.

OSPF configuration

To configure OSPF, perform the following procedures:

- 1 [“Configuring OSPF interfaces” on page 47](#)
- 2 [“Configuring OSPF globally” on page 49](#)
- 3 [“Configuring OSPF for branch offices” on page 54](#)

Configuring OSPF interfaces

To configure OSPF interfaces, perform the following steps:

- 1 Choose **Routing, Interfaces**.
- 2 Click **Configure** next to OSPF.

The Interfaces > Routing Interfaces > Configure OSPF window appears. Interface indicates the type of interface. IP address indicates the IP address of the interface.

OSPF is enabled by default.

- 3** If **Enabled** is not already selected, select **Enabled** to enable the OSPF State.

- 4** Select an **Area ID** for the OSPF area to which the attached network belongs.

If the Area ID you require is not present in the list, click the **Add an Area** link to add the area.

- 5** From the **Type** list, select **Broadcast** or **Point to Point** for the OSPF network type. The default is Broadcast.

- 6** From the **Authentication Type** list, select **None**, **Simple**, or **MD5** to indicate the authentication type to use as part of the OSPF transmission.

Simple indicates that authentication uses a simple password. MD5 specifies that authentication uses an MD5 secret. If you select either Simple or MD5, password and password confirmation boxes display.

- 7** In the **Cost** box, enter a metric value.

The Cost value is the cost of sending a packet on the interface expressed in the link state metric. The value must always be greater than 0 and the default is 10.

- 8** In the **Priority** box, enter the Priority level of the routers on this interface.

The router with the highest priority takes precedence and is the designated router (DR). If a tie occurs, the router with the highest Router ID takes precedence. A priority setting of 0 is ineligible to become a designated router on the attached network. Router priority applies only to broadcast networks. The default is 1.

- 9** In the **Hello Interval** box, enter the length of time in seconds between the Hello packets that the router sends on the interface.

The Hello Interval value must be the same for all routers attached to a common network. The default is 10.

- 10** In the **Dead Interval** box, enter the number of seconds after a router ceases to hear Hello packets before it declares that the router is down.

The number must be the same for all routers attached to a common network. The default is 40.

- 11** In the **Poll Interval** box, enter the number of seconds when, if a neighboring router becomes inactive, the router sends packets at a reduced rate in seconds. The default is 120.
- 12** In the **Retransmission Interval** box, enter the number of seconds between link state advertisement (LSA) retransmission for adjacencies that belongs to this interface.

The Retransmission Interval is also used to retransmit Database Description and Link State Request packets. Nortel recommends that you configure this setting to be considerably longer the expected round trip delay between two routers on the attached network. The default is 5.
- 13** In the **Transmission Delay** box, enter the number of seconds to transmit a Link State Update Packet over this interface. The default is 1.
- 14** Click **OK**.

Configuring OSPF globally

To configure OSPF globally, perform the following steps:

- 1** Choose **Routing, OSPF**.

Enabled indicates that OSPF is enabled on this window. The default setting is Disabled.
- 2** In the **Router ID** box, enter the IP address used to uniquely identify the OSPF router in the OSPF network.

The default address is the lowest IP address of the management or physical interfaces defined on the Nortel VPN Router. You can change this address if it is unique within the area.
- 3** From the **AS-Boundary-Router** list, select **True** or **False**.

If this Nortel VPN Router is an autonomous system (AS) boundary router, select **True** from the **AS-Boundary-Router** list. This parameter must be True to enable the redistribution of nonOSPF routes into OSPF. An AS boundary router is a router that exchanges routing information with routers that belong to other autonomous systems and advertises AS external routing information throughout the AS. The default is False.
- 4** From the **Auto Virtual Link** list, select **True** or **False**.

To automatically create virtual links to the backbone network, select **True**. The default is False.

- 5** From the **External Metric Type** list, select metric **Type 1** or **Type 2**.

Type 1 is the default. Type 1 external metrics are expressed in the same units as OSPF interface cost (in terms of the link state metric).

Type 2 external metrics are an order of magnitude larger; Type 2 metrics are considered greater than the cost of any path internal to the AS boundary router. Use of Type 2 external metrics is based on the assumption that routing between AS boundary routers is the major cost of routing a packet and eliminates the need to convert external costs to internal link state metrics.

- 6** From the **OSPF Maximum Paths** list, select the maximum number of ECMP paths (1 to 4).

Equal Cost Multipath provides load balancing of packets to a destination that is reachable over more than one physical interface.

- 7** To add an **Known OSPF Area**, click **Add**.



Note: The Known OSPF Areas section displays all OSPF areas defined locally to the Nortel VPN Router. The area information is not shared among Nortel VPN Routers. If you want two Nortel VPN Routers to have one interface in a common area, you must configure both Nortel VPN Routers to define the area information. Area IDs are used as representations of parts of the OSPF network. They help to manage large numbers of networks so that they can exchange information within an area. Each Area ID must be unique for OSPF. By default, all Nortel VPN Routers have an area named 0.0.0.0.

The Add > OSPF Area window appears.

- 8** Enter the **IP address** in the **Area ID** box.
- 9** For **Stub**, select **True** or **False** from the list. The default is False.
- 10** For **Default Cost**, enter the number of the default cost. The default is 1.
- 11** For **Import Summaries**, select **True** or **False** from the list. The default is False.
- 12** Click **OK**.

13 In the **Save LSDB Table** section, enter the name of the LSDB table that you want to save as a text file in the ide0/system/routing directory.

Viewing global OSPF information

In the **Status** section, you can display LSDB (link state database), Neighbor, Interfaces, Summary, or Statistics.

To view global OSPF information, perform the following steps:

- 1 Click **LSDB** to display the link state databases in all areas configured for the Nortel VPN Router.

[“LSDB window” on page 51](#) describes the OSPF LSDB window parameters.

Table 8 LSDB window

Parameter	Description
Link State ID	Displays the link state address
Adv Router	Displays the advertising router address
Age	Displays the age in seconds
Seq Nbr	Displays the sequence number
Checksum	Displays the checksum
Links	Displays the number of links

- 2 Click **Neighbor** to display a list of neighbors for all the interfaces that run OSPF.

[“OSPF Dynamic Neighbors window” on page 51](#) describes the OSPF Neighbors window parameters.

Table 9 OSPF Dynamic Neighbors window

Parameter	Description
Router ID	Displays the OSPF ID of the neighbor
P	Displays the priority number

Table 9 OSPF Dynamic Neighbors window

Parameter	Description
State	Displays the state of neighbor connection
Dead Time	Displays the time until neighbor is declared dead
Address	Displays the neighbor IP address
Interface	Displays the local IP interface address

3 Click **Interfaces** to display the list of interfaces that you configured for OSPF.

“[OSPF Interfaces window](#)” on [page 52](#) describes the fields in the OSPF Interfaces window.

Table 10 OSPF Interfaces window

Column	Description
IP Address-CId-State	Displays the IP address of the OSPF interface plus the circuit ID. If an asterisk (*) appears next to the interface, it designates that OSPF is configured, but it is administratively disabled.
Area ID	Displays the OSPF area for the interface.
Interface Type	Displays the Broadcast (BCAST) or Point to Point (PTPT).
Interface State	Displays the state of interface: Designated Router (DR), Backup Designated Router (BDR), or DR Other.
Metric Cost	Displays the cost associated with the interface.
Priority	Displays the priority used to negotiate DR/BDR state.
Designated Router	Displays the designated router IP address (0.0.0.0 for PTPT).

4 Click **Summary** to display the overall summary of OSPF that runs on the Nortel VPN Router.

“OSPF Summary window” on page 53 describes fields in the OSPF Summary window.

Table 11 OSPF Summary window

Column	Description
Router ID	Displays the unique OSPF ID of router
Router State	Displays the OSPF global configured state (up or down)
Supports TOS	Displays the type of service support
SPF schedule delay	Displays the delay time before calculating changes to SPF
Hold time between two SPF	Displays the time between shortest path first calls
Minimum LSA interval	Displays the link state advertisement interval
Minimum LSA arrival	Displays the link state advertisement arrival minimum
Number of external LSA	Displays the number of link state advertisements
Link State Update Interval	Displays the time between link state updates
Link State Age Interval	Displays the time between link state aging intervals
Number of Areas in this router	Displays the number of areas
RTM Stats	Displays the route table manager changes for route table changes
Area	Displays the area ID
Number of interfaces in this area	Displays the number of interfaces in this area
SPF algorithm has executed	Displays the number of times shortest path algorithm runs

5 Click **Statistics** to display statistical information about OSPF.

“OSPF Statistics window” on page 54 describes the fields in the OSPF Statistics window.

Table 12 OSPF Statistics window

Column	Description
Interface-CID	Displays the IP address for OSPF interface and circuit ID
Hellos	Displays the number of Hello packets received (RX) and transmitted (TX)
DBs	Displays the number of DB (Database Exchange) packets received (RX) and transmitted (TX)
LS Req	Displays the link state requests received (RX) and transmitted (TX)
LS Upd	Displays the link state updates received (RX) and transmitted (TX)
LS Ack	Displays the link state acknowledgements received (RX) and transmitted (TX)

Configuring OSPF for branch offices

To configure OSPF for branch offices, perform the following steps:

- 1 Choose **Profiles, Branch Office**.

The Branch Office window appears.

- 2 Select a group.

- 3 Click **Configure** next to the **Group** list.

The Branch Office > Edit Group window appears.

- 4 Click **Configure** in the **OSPF** section.

- 5 In the **Priority** box, enter the priority level of the routers on this interface.

The router with the highest priority takes precedence and is the designated router (DR). If a tie occurs, the router with the highest Router ID takes precedence. A priority setting of 0 is ineligible to become a designated router on the attached network. Router priority applies only to broadcast networks. The default is 1.

- 6 In the **Dead Interval** box, enter the time in seconds until a neighbor is declared dead. The default is 40.

- 7** In the **Hello Interval** box, enter the length of time in seconds between the Hello packets that the router sends on the interface. The default is 10.

The Hello Interval value must be the same for all routers attached to a common network.

- 8** In the **Retransmission Interval** box, enter the number of seconds between LSA retransmission for adjacencies that belong to this interface.

The Retransmission Interval value is also used for retransmitting Database Description and Link State Request packets. Nortel recommends that you configure this setting to be considerably longer than the expected round trip delay between two routers on the attached network and must be conservative. The default is 5.

- 9** In the **Transmission Delay** box, enter the number of seconds for the transmission delay. The default is 1.

- 10** In the **Authentication Type** section, select **None**, **Simple**, or **MD5** as the authentication type that is used as part of the OSPF transmission.

Simple indicates that authentication uses a simple password. MD5 specifies that authentication uses an MD5 secret.

Chapter 5

BGP configuration

This chapter contains information about Border Gateway Protocol (BGP) fundamentals and the procedures to configure BGP on the Nortel VPN Router.

This chapter includes the following topics:

- [“BGP fundamentals” on page 57](#)
- [“Installing the Border Gateway key” on page 70](#)
- [“BGP configuration” on page 70](#)

BGP fundamentals

BGP is a path vector protocol used to carry routing information between Autonomous Systems (AS). BGP imposes no restrictions on the underlying network topology. It is based on the assumption that routing within an AS is done through an intra-AS routing protocol. BGP considers the entire Internet to be a graph of ASs, with each AS identified by a unique autonomous number. Connections between ASs form a path, and the collection of path information forms a route to reach a specific destination. BGP uses the path information associated with a destination to ensure loop-free inter-domain routing.

BGP runs over a reliable transport protocol. The Nortel VPN Router can use a transport protocol authentication scheme in addition to the BGP authentication mechanisms. The BGP error notification mechanism is based on the assumption that the transport protocol supports a graceful close in which all outstanding data is delivered before the connection closes.

BGP-4 provides a new set of mechanisms to support classless inter-domain routing. These mechanisms include support to advertise an IP prefix and eliminate the concept of network class within BGP. BGP-4 also introduces mechanisms that allow aggregation of routes, including aggregation of AS paths.

BGP is supported over IPSEC, L2TP, L2TP/IPSEC, and PPTP tunnels.

RFCs

[“RFCs” on page 58](#) shows the BGP RFCs supported on Nortel VPN Router.

Table 13 RFCs

RFC	Description
RFC 1771 BGP4	RFC 1771 renders RFC 1654 obsolete. All implementations of the BGP protocol must conform to this RFC to ensure complete interoperability.
RFC 1966 Route Reflection	RFC 1966 describes the use and design of Route Reflection to alleviate the need for full mesh Internal BGP (IBGP).
RFC 1997 Community Attributes	RFC 1997 describes an extension to BGP that can be used to pass additional information to both neighboring and remote BGP peers.
RFC 1657 MIB	RFC 1657 describes managed objects used for managing the Border Gateway Protocol Version 4 or lower.

EBGP and IBGP peers

BGP supports two types: External BGP (EBGP) and Internal BGP (IBGP). EBGP is BGP between two ASs. If the TCP connection has hops between endpoints, EBGP must be enabled. IBGP is BGP within the same AS. With IBGP, all BGP speakers must have a peer relationship with each other.

BGP peering and connection processing

To begin a BGP peering session, one or both BGP speakers initiate a TCP connection. Both speakers can simultaneously initiate a connection, which results in two active TCP sessions between peers. The BGP protocol provides negotiation rules to determine which connection remains and which is deleted. After the TCP

connection is established, the BGP protocol negotiates with the peer by using the OPEN message to move into BGP Established State. At this time, each BGP speaker sends BGP update messages to distribute routing information between the speakers.

BGP update processing

Update messages advertise routes between a pair of BGP speakers. The destination is the systems whose IP addresses are reported in the Network Layer Reachability Information (NLRI) field, and the path is the information reported in the path attributes fields of the same Update message.

Update messages contain single reachable route updates and multiple unfeasible routes that must be withdrawn. Update messages are processed only in the BGP Established State.

Unfeasible route processing

Unfeasible routes are routes are unreachable and must be withdrawn. The first field in the BGP Update message is the Unfeasible Routes length field. If this field is zero, no unfeasible routes are present. Otherwise, this field contains the total length (in octets) of Withdrawn Routes elements. Withdrawn Routes elements consist of <prefix length, prefix> tuples as defined in RFC 1771.

Feasible route processing

A single feasible route is a set of path attributes that are associated with a number of destinations or networks. By sending an UpdateUpdate message, the peer specifies that a certain path is available and that from this path, you can transmit to certain destinations.

Path attribute processing

Path attributes fall into four separate categories:

- well-known mandatory
- well-known discretionary
- optional transitive
- optional nontransitive

All BGP implementations recognize well-known attributes. Some of these attributes are mandatory and must be included in every Update message. Others are discretionary and can or cannot be sent in a particular Update message. You can modify attribute values using route filters, thus influencing the best path selection. The path attribute information applies to all prefix destinations listed in the NLRI.

Path attribute types are listed in [“Path attribute types” on page 60](#).

Table 14 Path attribute types

Path attribute type	Code	Description
ORIGIN	1	Well-known mandatory Defines the origin of the path. 0—IGP NLRI information is interior to originating AS. 1—EGP – NLRI information is learned using EGP. 2—Incomplete – NLRI I learned by other means.
AS_PATH	2	Well-known mandatory sequence of AS Path Segments (tuple) <type, len, value> type = AS_SE —unordered set of ASs traversed by the update message in the path. AS_SEQUENCE—ordered set of ASs traversed by the update message on the path.
NEXT_HOP	3	Well-known mandatory IP address of the border router to be used as the nexthop to the destinations listed in the NLRI of the update message.
MULTI_EXIT_DISC	4	Optional nontransitive Value used by BGP speaker to discriminate among multiple exit points if more than one path exists to a neighboring AS.
LOCAL_PREF	5	Well-known discretionary Number used by BGP speaker to inform other speakers in its own AS of the originating speaker degree of preference for an advertised route.
ATOMIC_AGGREGATE	6	Well-known discretionary Informs other BGP speakers that the local system chose a less specific route, even though it had a more specific route available.

Table 14 Path attribute types

Path attribute type	Code	Description
AGGREGATOR	7	Transitive—optional Contains AS number and IP address of the BGP speaker that formed the aggregate route.
BGP Community	8	Identifies the community to which the route belongs.
Originator ID	9	Identifies the originator of the route into a route reflector cluster.
Cluster List	10	Lists the members of a route reflector cluster.

Keep Alive processing

BGP speakers use a Keepalive message to determine if the peers are reachable. You can disable the Keepalive message, but if you use it, you must configure it so that the Keepalive message is not sent more frequently than once per second.

Each BGP connection requires a Hold Time Interval. If the BGP speakers do not receive a Keepalive message or an Update message within the hold time period, a connection is considered unreachable. The BGP peer Hold Time Interval is configurable.

A Keepalive message must be sent between BGP peers at an interval frequent enough so that the Hold Timer intervals do not expire. RFC 1771 recommends a maximum time between Keepalive messages to be one-third of the Hold Time Interval.

Hold Time Interval between BGP peers is negotiable. If the two peers negotiate a Hold Time Interval of zero, Keepalive messages must not be sent. Configure the Hold Time Interval in the BGP, Configure window.

BGP policies

You can apply policy rules to either permit or deny a route. Policies provide a way to filter information based on IP prefixes, AS path information, BGP attributes, or source and destination addresses.

The two types of policies are as follows:

- interface-based policy—An inbound interface-based policy specifies that if a packet comes in on interface IX, apply policy PY to that packet.
- peer-based policy—An inbound peer-based policy (neighbor policy) specifies that if a packet comes in from peer PH, apply policy PZ to that packet.

Outbound policies are the reverse.

Accept and announce policies

In the Nortel VPN Router policy filtering model, both accept and announce policies apply only to peer-based filtering. Accept policies are rules that apply to incoming packets, and announce policies are rules that apply to outgoing packets.

You apply accept policies to incoming routes before you add routes to the BGP RIB IN table. You apply peer-based accept policies to packets received from a particular peer.

You apply announce policies to the Local RIB table before the Nortel VPN Router can advertise routes to the BGP peers. You apply peer-based announce policies to BGP updates destined for a particular peer. Outgoing routes matching the announce policy rule are either permitted or denied, depending on the rule.

Access (Prefix) lists

BGP uses access lists, another policy-filtering mechanism, to permit or deny routes. You define access lists as an address–mask pair. You specify whether you want an address–mask pair to be an exact match or a range match. If you specify a range match, addresses within the subnet range matches the rule. If you specify an exact match, only an address that exactly matches the address–mask pair satisfies the rule. You create access lists from Routing, Access List window.

- Access list example 1:

```
CES(config)# ip access-list 3 permit 55.1.0.0 255.255.0.0 range
```

This rule specifies that route updates that are in the range of 55.1.0.0 to 55.1.255.255 match the rule.

- Access list example 2:

```
CES(config)# ip access list 4 permit 55.1.0.0 255.255.0.0 exact
```

This rule says that only route updates containing the route 55.1.0.0 match the rule.

- Example using neighbor—using route maps (peer based)

```
CES(config-bgp)# neighbor 55.1.1.1 route-map EXAMPLE_MAP in
CES(config)# route-map EXAMPLE_MAP permit 10
CES(config-route-map)# match ip address 3
CES(config-route-map)# set metric 15
CES(config)# ip access-list 3 permit 44.1.0.0 255.255.0.0 range
```

In this example, IP access list 3 identifies all routes in the range 44.1.0.0 to 44.1.255.255. Any route in this range matches the access list and is propagated with a new metric of 15.

- AS path regular expression example:

A particular AS (AS = 5) consistently advertises bad routes, so you do not want to accept routes advertised by that AS. You set up a route map deny filter for routes containing AS path sequences that end in AS 5. You use a regular expression pattern-matching filter as follows:

```
ip as-path access-list 2 deny "*" 5$" (* is wildcard; $ symbolizes
ends with)
```

Any route advertisements with an AS path sequence ending in 5 are discarded.

AS-Path regular expressions

A BGP path is a sequence of alphabetic characters and consists of a set of AS numbers plus the following punctuation characters:

- ^ — the start of a path
- \$ — the end of a path
- { — the start of an AS_SET
- } — the end of an AS-SET



Note: An AS number such as 1234 is a single character in the alphabet. Although white space is used to make characters unambiguous, white space is not considered part of the alphabet. For example, to specify an AS number of 23 followed by 45, use the string “23 45”.

To match a single character in a path, the following forms can be used:

- The character itself
- `.` — matches any character
- `.*` — matches 0 or more characters
- `.+` — matches 1 or more characters
- `_` — matches 0 or 1 instance of any punctuation character (^, \$, {, })
- `[]` — specifies a set of characters. For example, `[1234 45 6789]` or `[{$}]`. All members of a set must be the same type, either AS numbers or punctuation.
- `-` — is used within brackets to specify a range of AS numbers. For example, `"[23-45]"` matches any number from 23 to 45.
- `^` — when used as the first item within brackets, specifies any AS number except the set specified. For example, to specify any AS number except 11 or 13, use `[^ 11 13]`. The caret (^) can also be used in conjunction with the hyphen (-) to specify any AS number except the specified range. For example, `[^100-200]` matches any AS number except those from 100 to 200.

You can create, delete, and modify AS path access lists. You can also apply access lists directly to neighbors for filtering. To configure AS path access lists, go to [“Configuring AS Path access lists” on page 76](#).

Route maps

You use route maps to filter routes and manage attribute. Route maps specify a certain set of criteria that need to be matched. If a match is found, an associated set of actions need to be applied to the matching route update. These filters are called Match-Set rules.

You can apply a route map to either inbound or outbound updates. Only the routes that pass the route map are sent or accepted in updates. You can create, delete, or modify route maps.

A route map can have several parts. Any route that does not match at least one match clause related to a route map command is ignored. The route is not advertised for outbound route maps and is not accepted for inbound route maps.

The route maps can be matched on

- as-path
- community-list
- ip address

The route maps can set

- as-path
- community
- local-preference
- metric
- next-hop
- origin
- weight

The following example illustrates how route maps are used:

- Route map example:

```
Format: route-map map-tag [permit | deny] [sequence number]
CES(config-bgp)# Neighbor 55.1.1.1 route-map EXAMPLE_MAP in
CES(config)# route-map EXAMPLE_MAP permit 10
CES(config-route-map)# match ip address 1
CES(config-route-map)# set metric 8
CES(config)# route-map EXAMPLE_MAP permit 20
CES(config-route-map)# match ip address 2
CES(config-route-map)# set metric 12
CES(config)# ip access-list 1 permit 33.1.0.0 255.255.0.0 exact
CES(config)# ip access-list 2 permit 44.1.0.0 255.255.0.0 exact
```

Route updates received from neighbor 55.1.1.1 are checked against this route map. First, the sequence number 10 rule states that routes matching ip access list 1 configures the metric to 8. If that check fails to match, the sequence number 20 rule is checked. This states that routes matching ip access list 2, configure the metric to 12.

If a route update comes in with network 33.1.0.0, the route is assigned metric 8. Similarly, if a route update comes in with network 44.1.0.0, it is assigned metric 12.

Multihop BGP

To configure a remote BGP peer that does not reside on a directly connected subnet, the EBGP peer must be accessible from the VPN Router and must reside on a network or subnet that exists in the IP routing table.

For IBGP peers, no restriction is specified in the protocol regarding multihop peering. Therefore, internal connection requests from neighbors not directly connected are accepted.

Multihop is configured on the BGP, Neighbor, Configuration window. By default, multihop BGP is disabled.

Route reflector

A route reflector organizes BGP peers into clusters and assigns each cluster an ID. Each member of the cluster advertises the routes only to the route reflector. The route reflector, in turn, collects all routes from all cluster members and advertises them to each IBGP peer in the cluster and to other route reflectors within the AS. Routes learned by the route reflector from other route reflectors are also forwarded to each cluster member.

All route reflectors must be fully meshed. By default, the clients of a route reflector are not required to be fully meshed. The routes from a client are reflected to other clients, and client-to-client reflection is enabled.

To increase redundancy and to avoid a single point of failure, a cluster can have more than one route reflector. In that case, configure all route reflectors in the cluster with the 4-byte cluster ID so that a route reflector recognizes updates from route reflectors in the same cluster. You can also configure the route reflector client list from a neighbor list. The clients of a route reflector cannot be members of a peer group.

Route reflector is disabled by default. To configure the route reflector, see [“Configuring the Route Reflector” on page 75](#).

BGP communities

A community is a group of destinations that share a common property. A BGP route can be a member of more than one community. Each AS administrator defines to which communities a destination belongs. Community lists are associated only with route maps. By default, all destinations belong to the general Internet community.

BGP communities simplify the route distribution based on membership to the community. You assign a community identifier to a set of destination addresses and establish a policy for the community instead of a separate policy for each prefix. All route updates received for members of a community have the same route redistribution characteristics. Control over the distribution of routing information is based on

- IP address prefixes
- value of the AS_PATH attribute (or part of it)
- identity of a group

You can create, delete, and modify community lists. The well-known communities are

- internet—the Internet community.
- no-export—routes with this community are sent to peers in other subautonomous systems within a confederation. Do not advertise this route to an EBGp peer.
- local-as—do not advertise this route to an external system.
- no-advertise—do not advertise this route to peers (internal or external).

A route is a member of a community if the Update message for the route contains a community attribute that includes that value. A BGP speaker uses this attribute to control which routing information it accepts, prefers, or distributes to other neighbors.

A BGP speaker receiving a route that does not have the COMMUNITIES path attribute can append this attribute to the route while propagating the attribute to the peers. A BGP speaker receiving a route with the COMMUNITIES path attribute can modify this attribute according to the local policy.

[“BGP communities” on page 69](#) illustrates the following example.

You do not want ISP 1 to announce ISP 2 routes to ISP 3. Likewise, you do not want ISP 3 to announce ISP 2 routes to ISP 1.

ISP 2 (AS 20) and ISP 3 (AS 30) belong to community 444.

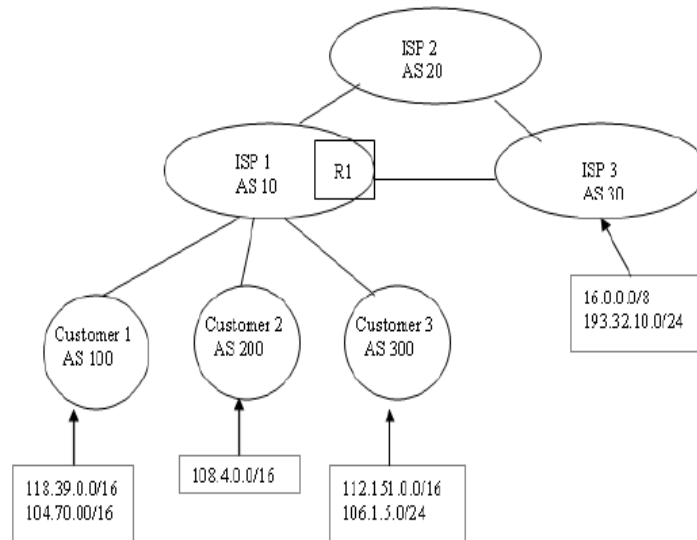
ISP 1 (AS 10) belongs to community 888.

AS 10 wants to offer transit service to customers in AS 100, AS 200 and AS 300 but nontransit service to customers in AS 20 and AS 30.

Assume that AS 100, AS 200 and AS 300 do not belong to a community. AS 10 labels all routes learned from AS 100, AS 200, and AS 300 as 10:888. Community 10:888 identifies routes that receive transit service.

AS 10 labels all routes learned from AS20 and AS30 as 10:444. This community label represents routes that receives nontransit service.

AS 10 can now have a policy that only announces routes that belong to community 10:888 and do not announce routes belonging to community 10:444.

Figure 2 BGP communities

To configure a BGP community list, see [“Configuring community lists” on page 77](#).

Health check support

The BGP-4 protocol provides a basic health check support by using the following mechanisms:

- **BGP initialization:** This returns a value of Success if BGP initialized properly. If the return value is a failure, a “Warning” is displayed on the window.
- **BGP global enable:** The RIP global enable value is checked. If the BGP protocol is disabled globally, the message “Disabled” appears in the window.

Installing the Border Gateway key

You must install the BGP-4 (Border Gateway Protocol Version 4) License key to enable BGP on the Nortel VPN Router.



Note: The Premium Routing License enables BGP-4 and the features included in the Advanced Routing License and DLSw License.

To install a software license key, perform the following steps:

- 1 Choose **Admin, License Keys**.
- 2 In the box to the right of **Advanced Routing**, enter the key that you obtained from Nortel Customer Support.
- 3 Click **Install**.

After you install the key, the label Key Installed appears. You need to install a key only once on each Nortel VPN Router.

To delete a software license key, perform the following steps:

- 1 Click **Delete** to remove the key.
A confirmation message appears.
- 2 Click **Yes**.



Note: The presence of the Border Gateway License key is checked only while BGP is globally enabled. If you enter the Border Gateway key, globally enable BGP, and then delete the Border Gateway key, BGP continues to run. However, if you then disable and re-enable BGP, it no longer runs.

BGP configuration

To configure BGP, perform the following procedures:

- 1 [“Adding a route map” on page 71](#)

- 2 [“Configuring route maps” on page 71](#)
- 3 [“Configuring BGP interfaces” on page 72](#)
- 4 [“Configuring neighbors” on page 74](#)
- 5 [“Adding a network” on page 75](#)
- 6 [“Configuring the Route Reflector” on page 75](#)
- 7 [“Configuring AS Path access lists” on page 76](#)
- 8 [“Configuring community lists” on page 77](#)

Adding a route map

If no route maps exist on the Nortel VPN Router, you can add one.

To add a route map, perform the following steps:

- 1 Select **Routing, Route Map**.
- 2 Click **Add**.
- 3 Enter a name in the **Name** text box.
- 4 Click **OK**.

The Route Maps window appears. The name you entered appears in the Route Map menu.

Configuring route maps

To configure route maps, perform the following steps:

- 1 Select **Routing, Route Map**.
- 2 Select a route map.
If no route maps are listed, you can add one. See [“Adding a route map” on page 71](#).
- 3 Click **Add** beside **Rule Number** to add a rule number.
The Route Map > Rule Add window appears.
- 4 Enter a number in the **Number** text box.

5 Click OK.

The Route Maps window reappears. The number you entered appears in the Number menu.

6 Select a type from the **Type menu.**

7 Click **Add in the **Match** section to add a match.**

The Rule > Match Add window appears.

8 Select an attribute from the **Attribute menu.**

9 Select a value from the **Value menu.**

10 Click OK.

The Route Maps window reappears with the information you selected showing under Match.

11 Click **Add below **Set** to add a set.**

The Rule > Set Add window appears.

12 Select an attribute from the **Attribute menu.**

13 Enter a value in the **Value text box.**

14 Click OK.

The Route Maps window reappears with the information you selected showing under Set.

15 Click OK.

Configuring BGP interfaces

To enable BGP interfaces, perform the following steps:

1 Select **Routing, BGP.**

The BGP window appears.

2 Select **Enabled or **Disabled** for State.**

If enabled, BGP enables all neighbors that are in enabled state. If disabled, BGP disables all neighbors that are in enabled state.

3 Enter the **Router ID.**

4 Enter a value in **Local AS.**

- 5 Enter the **Hold Timer** value. The default value is 90 seconds.
- 6 Enter the **Keep Alive Timer** value. The default value is 30 seconds.
- 7 Enable **Synchronization** if routers exist in the AS not speaking BGP.

Synchronization allows routers within an AS to access a route before BGP makes it available to other ASs.

- 8 Enter the **Local Preference** value. The default is 100.
- 9 Enter the **Default Metric** value.

Default metric value specifies the appropriate metric for the specified routing protocol. The default metric command is used with the redistribute router configuration command to cause the current routing protocol to use the same metric value for all redistributed routes. This configures the Multi Exit Discriminator (MED) metric as a hint to external neighbors about preferred paths. The MED value can also be configured using a route map. By default, during the best-path selection process, MED comparison occurs among paths from the same AS.

- 10 Check the **Always Compare MED** option if you want to allow the comparison of the MED for paths from neighbors in different ASs.
- 11 Enter the **Maximum Paths** value.

This configuration controls the number of paths allowed. By default, only one path is installed in the IP routing table. If BGP multipath support is enabled and the EBGP paths are learned from the same neighboring AS, instead of picking one best path, multiple paths are installed in the IP routing table. Support exists for a maximum of six paths and load balancing occurs among multiple paths.

- 12 To enable the Auto Summary feature, select **Auto Summary**.
- 13 To enable the Redistribute Internal feature, select **Redistribute Internal**.

You configure Neighbors, Networks, Route Reflector, AS-Path Access Lists, or Community Lists from the BGP window. You can also see a Summary window, the BGP Routes, Redistributed Routes, and Neighbors Routes from this window.

Neighbors, Networks, Route Reflector, AS-Path Access Lists, and Community Lists are described in the following sections.

Configuring neighbors

You can create, delete, or modify neighbors. The maximum number of neighbors you can create is a configurable parameter, depending on the hardware.

To configure neighbors, perform the following steps:

- 1 Choose **Routing, BGP**.
- 2 Click **Neighbors**.
- 3 Click **Add** beside **Neighbor** at the top of the window.
- 4 Enter the IP address of the neighbor.
Use the **Delete** button to remove an existing route.
- 5 Click **OK**.
- 6 Select **Enabled** or **Disabled** for **State**.
- 7 Enter your password and confirm your password.
- 8 Enter a value in **Remote AS**. At a minimum, you must configure remote-AS for neighbors to be enabled.
- 9 Enter the **Hold Timer** value. The default value is 90 seconds.
- 10 Enter the **Keep Alive Timer** value. The default value is 30 seconds.
- 11 Enter the **Advertisement Interval** value. The minimum advertisement interval is 30 seconds.
- 12 Enter the **Retry Interval** value. The default is 30 seconds.
- 13 Enter the **Source IP Address**.



Note: The source IP address typically comes from the route table, but you can enter it in the Source IP Address text box.

- 14 Enter the **Weight** value.

The administrative weight is local to the router. Any path that a VPN router originates has a default weight of 32768 and other paths have a weight of 0. You can also assign the weight through filter-lists and route maps.

- 15 Enable **NH Self** to allow BGP neighbors to have access to all neighbors on the same IP subnet.

You can also specify the next-hop address to be used by route maps.

- 16 Enable **EBGP** to allow BGP sessions, even if the neighbor is not on a directly connected segment.
- 17 Enable **Send Community** if you want to include the community parameters in the message when the BGP route is announced to a neighbor.

To see a display of the Summary of the Neighbors, go to the Routing, BGP, Neighbors, Summary window.

Adding a network

To add a network, perform the following steps:

- 1 Choose **Routing, BGP**.
- 2 Click **Networks**.
The BGP > Networks window appears.
- 3 Click **Add**.
The BGP > Networks Add window appears.
- 4 Enter an IP address in the **IP Address** box.
- 5 Enter a Mask in the **Mask** box.
- 6 Click **OK**.

Configuring the Route Reflector

To configure the Route Reflector, perform the following steps:

- 1 Choose **Routing, BGP**.
- 2 Click **Route Reflector**.
The Route Reflector window appears.
- 3 Select the **Status** of the route reflector.
The status globally enables or disables the feature.

4 Enter the Cluster ID.

The router ID of the route reflector identifies the cluster.

5 Select the Client to Client Route Reflector value.

The default is Enabled. However, if the clients are fully meshed, route reflection is not required and you must disable the route reflector.

To add or remove members from Route Reflector Client lists, perform the following steps:

1 Under Clients, select a non-member from the Non Member RR Client List.**2 Click Make RR Client.**

The non-member becomes a member of the Member RR Client List.

3 Select a member from the Member RR Client List. Click Remove RR Client.

The member is removed from the list.

Configuring AS Path access lists

To configure the AS-Path access list, perform the following steps:

1 Choose Routing, BGP.

The Routing > BGP window appears.

2 Click AS-Path Access List.

The AS-Path Access List window appears.

3 To add an Access List number, click Add beside Access List Number.

The AS-Path Access List > Add window appears.

4 Enter a number in the Number text box.

The number uniquely identifies the AS-Path access list. The acceptable range is from 1 to 99.

5 Click OK.

The BGP AS-Path Access List Number window reappears with the number you typed in the Number box showing in the Access List.

- 6 To create an **Access List** entry, click **Add** below Access List.

The BGP > AS-Path Access List > Add Entry window appears.

- 7 Select **Permit** or **Deny** from the **Type** list.

- 8 Enter an entry in **AS-Path Regular Expression**.

For more information about AS-Path regular expressions, see [“AS-Path regular expressions” on page 63](#).

- 9 Click **OK**.

The BGP > AS-Path Access List window reappears with your information. At the top of the window is the statement “Add operation completed successfully”.

- 10 To delete an **Access List**, select the list to delete and click **Delete**.

A new window appears prompting you to confirm that you want to delete the as-path access list number.

- 11 Click **OK**.

The BGP > AS-Path Access List window reappears with the number you deleted removed from the list.

At the top of the window is a note stating “Delete operation completed successfully”.

- 12 To delete an **Access List Entry**, select the entry you want to delete.

A new window appears prompting you to confirm that you want to delete the as-path access list entry.

- 13 Click **OK**.

The BGP > AS-Path Access List window reappears with the entry you deleted removed. At the top of the window is a note stating “Delete operation completed successfully”.

Configuring community lists

To configure a community list, perform the following steps:

- 1 Choose **Routing, BGP** window,
- 2 Click **Community List**.

The Community List window appears.

- 3 To add a community list number, click **Add**.

The Community List > Add Web window appears.

- 4 Enter a number in the **Number** text box.

The number uniquely identifies the community list. The acceptable range is from 1 to 99.

- 5 Click **OK**.

The Community List window reappears.

- 6 To add a community entry, click **Add**.

The Community List > Add Entry window appears.

- 7 Select either **Permit** or **Deny** from the **Type** menu.

- 8 Enter a name in the **Name** text box.

The valid range for the community entry number is 1 to 4294967200. You can also configure the name to no-export, no-advertise, or local-as.

- 9 Click **OK**.

The Community List window reappears with the information you entered.

- 10 To delete a community list number, select a number from the **Community List** list.

- 11 Click **Delete**.

A new window appears with a warning prompting you to confirm that you want to delete the community list number.

- 12 Click **OK**.

The Community List window reappears with the community list number deleted.

Chapter 6

Static route configuration

You can statically define available routes instead of dynamic routing protocols, such as Open Shortest Path First (OSPF) or Routing Information Protocol (RIP) learning available routes. Even if you use dynamic routing protocols, you can use static routes in certain situations where strong security is required. The Nortel VPN Router supports multiple default and static routes.

This chapter includes the following topics:

- [“Static route configuration” on page 79](#)
- [“Pinging to validate public default route” on page 82](#)

Static route configuration

To configure static routes, perform the following procedures:

- 1 [“Enabling static routes” on page 79](#)
- 2 [“Configuring static routes” on page 80](#)
- 3 [“Viewing static route information” on page 80](#)
- 4 [“Configuring public default routes” on page 81](#)
- 5 [“Configuring private default routes” on page 81](#)

Enabling static routes

To enable static routes:

- 1 Choose **Routing, Static Routes**.
- 2 Check the **Enabled** box.

If static routes are disabled, all static routes and default routes are disabled globally. Even if a static route is enabled, the route is not used. If static routes are enabled, traffic flow depends on other configuration settings.

Configuring static routes

To add static routes to the route table, perform the following steps:

1 Choose **Routing, Static Routes**.

2 Click **Add** under **Static Routes through Physical Interfaces**.

The Static Routes > Add Static Route window appears. If you add a static route, the Nortel VPN Router checks whether the next-hop interface address belongs to an attached network. If it does not, the Nortel VPN Router does not allow the static route.

3 Select **Enabled** or **Disabled** for the **Admin state**. The default is Enabled.

4 Select the relative cost for the Nortel VPN Router.

Use a lower cost number, such as 1, for the least expensive route. If there are multiple paths, the Nortel VPN Router chooses the route with the least cost as the preferred route. The default is 10.

5 In the **Network Address** box, enter the network address for the static route to the destination network.

6 In the **Subnet Mask** box, enter the subnet mask for the static route to the destination network.

7 In the **Gateway Address** box, enter the Nortel VPN Router address to the next-hop router to reach the destination network.

8 Click **OK**.

Viewing static route information

To view static route information, perform the following steps:

1 Choose **Routing, Static Routes**.

2 Click **Show Branch Office Routes** to display the configured branch office tunnels that are set up as static routes.

By default, a tunnel is configured as a static route between the tunnel endpoints.

- 3 Click **Adjacent Hosts** to display adjacent host routes.

Configuring public default routes

To add a public default route, perform the following steps:

- 1 Choose **Routing, Static Routes**.

- 2 Click **Add Public Route**.

The Static Route > Add Public Default Route window appears.

- 3 Click **Enabled** or **Disabled** to select the **Admin State**.

- 4 Enter the relative cost in the **Cost** box.

Use a lower cost number, such as 1, for the least expensive route. If multiple default paths exist, the Nortel VPN Router chooses the route with the least cost as the preferred route. The default cost is 10.

- 5 Enter the **IP address** for the next-hop default router in the **Gateway Address** box.

- 6 Select the **Validate at Ping Interval** box.

- 7 Enter an IP address in the **Ping Address** box to identify which IP address to base ping interval validation.

- 8 Click **OK**.

Configuring private default routes

To add a private default route, perform the following steps:

- 1 Choose **Routing, Static Routes**.

- 2 Click **Add Private Route**.

The Static Route > Add Private Default Route window appears.

- 3 Click **Enabled** or **Disabled** to select the **Admin State**.

- 4 Enter the relative cost in the **Cost** box.

Use a lower cost number, such as 1, for the least expensive route. If multiple default paths exist, the Nortel VPN Router chooses the route with the least cost as the preferred route. The default cost is 10.

- 5 Enter the **IP address** for the next-hop default router in the **Gateway Address** box.
- 6 Click **OK**.

Pinging to validate public default route

You can configure the Nortel VPN Router to use the ping utility to verify the status of a link from a public interface through an Asymmetric Digital Subscriber Line (ADSL) modem to a remote endpoint. Use the ping utility to detect a link failure at a point beyond the modem. The ping utility detects whether a broadband remote access server (BRAS) is available and if so, forwards traffic through it. The Nortel VPN Router has a public default route out of the modem interface.

The ADSL modem operates in either bridge mode or router mode. In bridge mode, the VPN Router is the BRAS interface to the digital subscriber line access multiplexer (DSLAM) and traffic is bridged from the ADSL. In router mode, the VPN Router is the ADSL and traffic is routed from the ADSL to the BRAS on a different network.

If validation is globally enabled, the following situations occur:

- At the expiration of each ping interval, the utility pings the ping address of each public default route for which individual route validation is enabled.
- If the ping address is not the Nortel VPN Router address for the route, the utility configures and enables a static route.
- If a static route with that address exists, the utility uses the route is for validation and save the state.

You can not edit or delete static routes used to validate public default routes. Static routes are deleted or returned to their original state if one of the following conditions occurs:

- Validation is globally disabled.
- The public default route is disabled or deleted.

- Validation is disabled for the public default route.
- The address to ping for validation of the public default route changes.

If validation is globally disabled, public default routes that were disabled because of validation are enabled.

To configure ping to validate a public default route, perform the following steps:

- 1** Choose **Routing, Static Routes**.
- 2** Select **Validate Public Default Routes**.
- 3** Enter a ping interval in the **Ping Interval** box.

The minimum (and default) is 30 seconds and the maximum is 5 minutes.

- 4** Click **OK**.

Chapter 7

RPS configuration

This chapter contains information about route policy services (RPS) and the procedures to configure them.

This chapter includes the following topics:

- [“RPS fundamentals” on page 85](#)
- [“RPS configuration” on page 88](#)

RPS fundamentals


Using the route policy service (RPS), you can control the flow of routing data to and from the route tables. The route policy service provides IP accept and announce policies that you enable or disable as needed.

Accept policies govern the addition of new RIP- or OSPF-derived routes to the route tables. After the routing protocol receives a new routing update, it consults the accept policies to validate the information before entering the update into the route tables. Accept policies contain search information to match fields in incoming routing updates and action information to specify the action to take with matching routes.

Announce policies govern the propagation of RIP or OSPF routing information. When OSPF prepares a routing advertisement, it consults the area boundary router to determine whether the routes to specific networks are advertised and how they are propagated. Announce policies contain network numbers (to associate a policy with a specific network) and action information (to specify a route propagation procedure). For OSPF, announce policies apply only to external routes. For RIP, announce policies apply to all routes, including external routes that are redistributed into RIP and RIP-generated routes.

[“Accept and announce policies” on page 86](#) shows the interaction between the route table manager and the accept and announce policies.

Figure 3 Accept and announce policies



The route table manager forwards a route for advertisement to the protocol. The protocol consults an announce policies to determine whether to advertise the route to the network.

OSPF link state advertisements (LSA) are received and placed in the link state database (LSDB) of the router. The information in the LSDB also propagated to other routers in the OSPF routing domain. According to the OSPF standard, all routers in an area must maintain a similar database. To maintain database integrity across the network, a router must not manipulate received LSAs before propagating them to other routers.

To accomplish this goal, OSPF accept and announce policies act in the following manner:

- The accept policies control only the information that the local router uses; they do not affect the propagation of OSPF internal and OSPF nonself-originated external information to other routers.
- OSPF announce policies control which self-originated external routing updates are placed into the LSDB for distribution according to the OSPF standard. OSPF announce policies affect what other routers learn, but only with regard to the local router self-originated information.

Redistribution of routes

The Nortel VPN Router can redistribute static, direct, BGP, and RIP routes into OSPF. It can redistribute static, direct, BGP, and OSPF routes into RIP. It can also redistribute static, direct, OSPF, and RIP routes into BGP. Access lists control the redistribution of routes from BGP to OSPF. Such a redistribution can be further controlled for each interface in RIP. Route redistribution is also based on security configurations.

[“Redistribution rules” on page 87](#) describes the rules of redistribution for RIP, OSPF, and BGP with the firewall enabled or disabled.

Table 15 Redistribution rules

Redistributed route	Firewall on	Firewall off
Public direct route	Yes	No
Public default route	Yes	No
Public static route	Yes	No
Private direct route	Yes	Out physical—No; out tunnel—Yes
Private default route	Yes	Out physical—No; out tunnel—Yes
Private static route	Yes	Out physical—No; out tunnel—Yes
Tunnel static route	Yes	OSPF—Always Yes RIP—In general, Yes, but can be controlled for each interface
Tunnel dynamic route	Yes	Yes
U tunnel routes	Yes	Yes

When a dynamic routing protocol redistributes default routes (public or private), the receiving router treats these routes as protocol-specific default routes. Therefore, locally defined default routes have a higher precedence over routes learned by redistribution.

Even though a public default route is represented by 0.0.0.0/32 when redistributed, it is represented as 0.0.0.0/0 to conform with the routing protocols. If static routes are redistributed by a routing protocol, default routes are also redistributed. However, if you have both private and public default routes, only one of them is redistributed, thus reducing the number of redundant routes to the same destination through the same next-hop interface.

RPS configuration

To configure RPS, perform the following procedures:

- 1 [“Creating a policy list” on page 88](#)
- 2 [“Editing a policy list” on page 89](#)
- 3 [“Configuring RPS” on page 89](#)

Creating a policy list

To create a policy list, perform the following steps:

- 1 Choose **Routing, Access List**.
- 2 Enter a new access list name.
Use a name or number that you choose to a maximum length of 64 characters.
- 3 Click **Create**.
The Access List, Policy window appears.
- 4 Select an option from the **Action** list.
Permit or Deny is the action applied to a route update if the subnet and mask matches the route update. If you choose Permit or Deny, enter the subnet mask, mask and mask type (Exact or Range).
- 5 Click **Add**.

- 6 Click **Close** to apply the new rule.

Editing a policy list

To edit a policy list, perform the following steps:

- 1 Choose **Routing, Access List**.
- 2 Select an access list.
- 3 Click **Edit** to change an existing rule for the selected policy.

The current information appears for each policy. You can use either an exact network address or a range of network addresses.

- 4 Enter a number in the **Move selected rule to position** box to move the position of an existing rule.

For example, if you select the third rule and enter 2, the third rule moves to the second position. The order of the rules is important because the first match causes the action to occur. If no matches occur, all traffic is denied. Therefore, build your filter rules by first permitting the services that you want to allow. You can also add a Deny rule early in the rules sequence so that an unwanted packet is dropped before all of the rules are processed.

- 5 Click **Close**.

Configuring RPS

To configure route policy services, perform the following steps:

- 1 Choose **Routing, Policy**.
- 2 Check the **Enabled** box. The default setting is Disabled.
- 3 Under **Redistribution Table**, select the source of the route for each protocol.
The route source can be Static, Direct Nets, Direct Hosts, RIP, BGP, Utunnel, CLIP, MGMT, or NAT. For correct operation, do not select more than one route source for each protocol.
- 4 Under **Policy List**, click **Add** to add a policy.

5 Select an **Access Name/Number**.



Note: You must create an access list before you can create policy entries. To create the access list, click **New Access List** to display the Access Lists window. You can edit or delete a selected list name or create a new one by typing the name in the edit box.

c Select either **OSPF, RIP, or BGP**.

d Enter or select the **Interface IP address**, which is the IP address of the physical interface where you want to apply the policy.

Select **Global** if you want to apply the policy to all interfaces. If the interface is a branch office, select the group name and enter the connection name.

e Select the **policy type**, either the **accept** or **announce**.

You can configure only one accept or announce policy for each protocol for each interface.

6 Click **OK**.

Chapter 8

Client address redistribution

This chapter contains information about client address redistribution and the procedure to configure it.

This chapter includes the following topics:

- [“Client address redistribution fundamentals” on page 91](#)
- [“Configuring client address redistribution” on page 95](#)

Client address redistribution fundamentals

After a client initiates a user tunnel, the Nortel VPN Router assigns an inner address to the client. Sources for these addresses can be

- a predefined address pool in the Nortel VPN Router with an address range that belongs to a locally attached private network
- a predefined address pool in the Nortel VPN Router with an address range that does not belong to a locally attached private network
- a static address configured in the Nortel VPN Router
- a Remote Authentication Dial-in User Service (RADIUS) or Dynamic Host Configuration Protocol (DHCP) address
- a client-supplied private address

If the client address does not belong to a locally attached Nortel VPN Router network, you must enable client address redistribution to ensure that these addresses are advertised in the dynamic route updates sent by the Nortel VPN Router. Client address redistribution uses a route type called a Utunnel. Utunnel routes can be either host or network routes.


If client address redistribution is in host mode, the Nortel VPN Router creates and advertises a user tunnel host route whenever a client tunnel is created, using an inner address that does not belong to a locally attached network. If the tunnel is taken down, the corresponding host route is deleted.

If inner addresses are allocated from an address pool with a range that does not belong to a locally attached network, use the aggregation option to reduce the number of entries in the route table and the route redistribution overhead. Aggregation creates and advertises a single Utunnel network route covering the address pool range if you create a client tunnel by using an inner address from this address pool. In Dynamic Aggregation mode, the network route remains in the route table until the last tunnel that uses an inner address from this address pool is taken down. In Static Aggregation mode, the network route remains in the route table until the user address pool is deleted.



Note: The maximum number of Utunnel routes cannot exceed the maximum number of client tunnels supported by the corresponding hardware platform. The default value is 200.

“[Client address redistribution](#)” on [page 93](#) shows an example of client address redistribution where the client has an inner address that is not within the local subnet of the private network. The Nortel VPN Router creates a Utunnel route that then propagates over the network. The Utunnel route allows the router on the private network to recognize the 200.168.1.100 address and correctly route responses back to it.

Figure 4 Client address redistribution

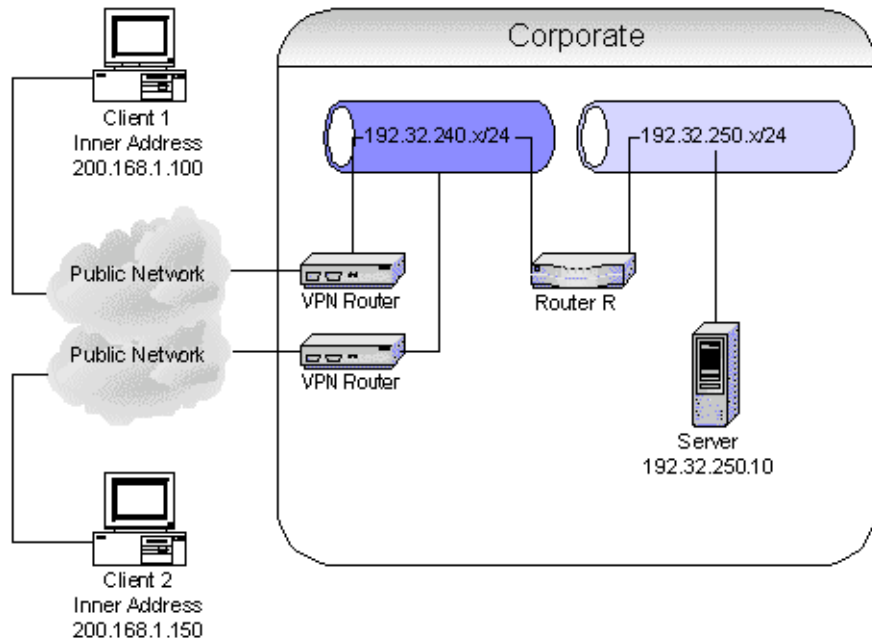
If you enable aggregation, the Nortel VPN Router identifies the subnet from the address pools where this address belongs and inserts a user tunnel network route for this subnet into the route table manager.

Aggregation is useful for large networks where route summary optimization reduces the number of Utunnel host entries in the RTM. However, if you enable aggregation, you can potentially have routing problems if the subnets of the address ranges span multiple Nortel VPN Routers. If you have two Nortel VPN Routers assigning addresses that belong to the same IP subnet, do not use the aggregation option.

For example, in [“Aggregation for client address redistribution” on page 94](#), Nortel VPN Router A has an address range of 200.168.1.100 to 200.168.1.120 and Nortel VPN Router B has an address range of 200.168.1.150 to 200.168.1.170. Both ranges are part of Class C subnet 200.168.1.x/24. Client 1 logs on to Nortel VPN Router A and Client 2 logs on to Nortel VPN Router B. Both clients have inner addresses that are not within the local subnet of the private network, but are in the same IP subnet. Nortel VPN Router A and Nortel VPN Router B running client

address redistribution create Utunnel host routes. These routes propagate over the network. The router on the private network recognizes addresses 200.168.1.100 and 200.168.1.150 and routes responses back to them through the designated VPN Router.

Figure 5 Aggregation for client address redistribution



If you enable aggregation on both Nortel VPN Routers, both routers advertise routes to 200.168.1.x/24. Router R uses one of these routes, causing either Client 1 or Client 2 to have communication problems.

The route table manager handles Utunnel routes similarly to other route types (RIP or OSPF). You can view Utunnel routes using the Routing > Route Table Manager window. The route policy service handles redistribution (advertisement) of Utunnel routes similarly to redistribution of other route types.

Configuring client address redistribution

To configure client address redistribution, perform the following steps:

- 1 Choose **Routing, Client-Addr-Dis**.
- 2 Select one of the following **CAR Options**:
 - **Disable**—Disables CAR and redistribution of client routes does not occur.
 - **Host Only**—Enables CAR and the redistribution of client routes is limited only to host routes. Host routes are added to both the forwarding table and the routing table. RIP and OSPF advertise the host routes of the VPN clients to their peers.
 - **Dynamic Aggregation**—Enables CAR and the client host addresses are added only to the forwarding table. The subnet of the user address pool from which the client address was assigned is added to the routing table. RIP and OSPF advertise only the subnet of the address pool and not the client host addresses. After the last client using this user address pool disconnects, the subnet route is removed from the routing table. RIP and OSPF propagate the route deletion to the surrounding networks.
 - **Static Aggregation**—Enables CAR and the client host addresses are added only to the forwarding table. The subnet of the user address pool from which the client address was assigned is added to the routing table. RIP and OSPF advertise only the subnet of the address pool and not the client host addresses. After the last client using this user address pool disconnects, the subnet route remains in the routing table. The subnet of the user address pool remains in the routing table as long as the user address pool remains valid. If you delete the user address pool, the subnet for the pool is then deleted from the routing table.
- 3 Enter a **Maximum Number of UTunnel Host Routes**. The default value is 200.

Viewing client address redistribution information

The Current Number of UTunnel Host Routes field displays the current number of user tunnel hosts logged on to the system.

- 1 Click **Show User Tunnel Routes** to display the user tunnel routes. “[Show user tunnel routes](#)” on page 96 describes the fields.

Table 16 Show user tunnel routes

Column	Description
IP address	Displays the IP address
Mask	Displays the IP network mask
Next Hop	Displays the next hop address
Interface	IP interface address
Cost	Displays the relative cost for the Nortel VPN Router

- 2 Click **Statistics** to display the configuration of client address redistribution, including mode, the UTunnel limit, and current UTunnel count.
- 3 Click **Refresh** to view changes.

Chapter 9

Multicast relay configuration

This chapter contains information about multicast relay and the procedure to enable it.

This chapter includes the following topics:

- [“Multicast relay fundamentals” on page 97](#)
- [“Configuring multicast relay” on page 99](#)
- [“Viewing multicast relay information” on page 99](#)

Multicast relay fundamentals

IP multicast is an extension to the standard IP network-level protocol. It provides efficient delivery of information from a single source to multiple destinations. IP multicast is useful for applications such as video conferences, shared white boards, and news feeds. IP multicast uses Class D addresses, ranging from 224.0.0.0 to 239.255.255.255. Multicast routing protocols establish the distribution tree for a multicast group.

A multicast relay listens to incoming multicast traffic and forwards it to one or more interfaces in the absence of multicast routing. Support is unavailable for multicast relay on public interfaces.

By default, multicast relay is globally disabled. If multicast relay is disabled, the Nortel VPN Router processes multicast requests in the range of 224.0.1.0 to 239.255.255.255. If enabled, multicast traffic is filtered according to interface filter lists and access lists.



Note: Multicast relay and Inter Group Multicast Protocol (IGMP) cannot run at the same time. If you want to enable IGMP and multicast relay is enabled, you must disable multicast relay. Conversely, if you want to enable multicast relay and IGMP is enabled, you must disable IGMP. For more information about IGMP, see [“IGMP configuration” on page 101](#).

The congestion threshold is configured relative to the amount of network processing memory buffers available to process the multicast traffic. The allowable range is 1 to 3000, where 3000 is the default value. If forwarding performance for unicast traffic decreases due to the multicast traffic burden, Nortel recommends that you reduce the threshold. To view network processing buffer statistics, choose the Status, Statistic, snpbufStats command.



Note: To receive multicast packets over a static tunnel, enter the multicast range of addresses as part of the list of local networks on the receiving side.

Forward multicast packets over a tunnel using the default filter (permit all). For example, to allow multicast packets received over the interface to be relayed over tunnel B01 and not over tunnel B02, define the interface-specific rules as shown in [“Multicast interface-specific rules example” on page 98](#).

Table 17 Multicast interface-specific rules example

	type	src intf	dst intf	source	dst	service	action
receiving	SRC	LAN	ANY	S	231.0.01	voice	allow
relay	DST	ANY	BO1	S	231.0.0.1	voice	allow
relay	DST	ANY	BO2	S	231.0.0.1	voice	drop

Configuring multicast relay

To configure multicast relay, perform the following steps:

- 1 Choose **Routing, Multicast**.
- 2 Check the **Enabled**.

If you enable multicast relay, received traffic is filtered according to filter lists and access lists.



Note: Multicast requires use of the Permit all interface filter.

- 3 Enter the **Congestion Threshold** value. The default value is 3000.
- 4 To add an interface to the **multicast boundary list**, click **Add**.
- 5 The Multicast, Add window appears.
- 6 Enter the **Access Name/Number** in the box.
You can click the **New Access List** link to view or add access lists.
- 7 Select the **IP address** for the interface.
- 8 Select **Enabled** for the **State**.
- 9 Click **OK**.

Viewing multicast relay information

To view multicast relay information, perform the following steps:

- 1 Click **Statistics** to display the global multicast relay status and the statistics of the configured multicast interfaces, including branch office interfaces.

“[Multicast Statistics window](#)” on [page 100](#) describes fields on the Multicast Statistics window.

Table 18 Multicast Statistics window

Column	Description
Interface	Displays the IP address of the interface
CID	Displays the circuit ID
PktsRcvd	Displays the number of packets received
PktsSent	Displays the number of packets sent
PktsDropped	Displays the number of packets dropped

- 2 Click **Interfaces** to display all configured information about enabled interfaces, including private physical and branch office tunnel interfaces.

“[Multicast Interfaces window](#)” on [page 100](#) describes fields on the Multicast Interfaces window.

Table 19 Multicast Interfaces window

Column	Description
Interface	Displays the IP address of interface
Access-list	Displays the name of the access list

Chapter 10

IGMP configuration

This section contains conceptual information about the Internet Group Management Protocol (IGMP) for the Nortel VPN Router. Use this information to configure IGMP on the Nortel VPN Router.

This chapter includes the following topics:

- [“IGMP fundamentals” on page 101](#)
- [“IGMP configuration” on page 113](#)

IGMP fundamentals

IGMP is a protocol used by the Nortel VPN Router to route multicast traffic on the private side of the Nortel VPN router to serve multiple VPN client users that are registered for membership in a multicast group. This capability is known as multicast proxy, where the VPN Router forwards multicast joins to clients who request access to a multicast stream. This reduces network traffic on the private side of the Nortel VPN Router. Instead of multiple unicast streams going to the Nortel VPN Router and out to the clients, a single multicast stream is received on the private side and then duplicated by the Nortel VPN router for each client subscriber.

IGMP requires the Advanced Routing Option license or Premium license.

IGMP modes

IGMP operates in two modes:

- router mode
- host mode

Router mode

In router mode, IGMP uses a query and response mechanism to solicit membership status from hosts. Routers (queriers) periodically send a query message and hosts respond with a membership report for each group base (IGMPv1 and IGMPv2) or for several groups (IGMPv3). After a router receives a report from a host, it processes the report message and adds the reported group addresses and sources onto the group database.

Host mode

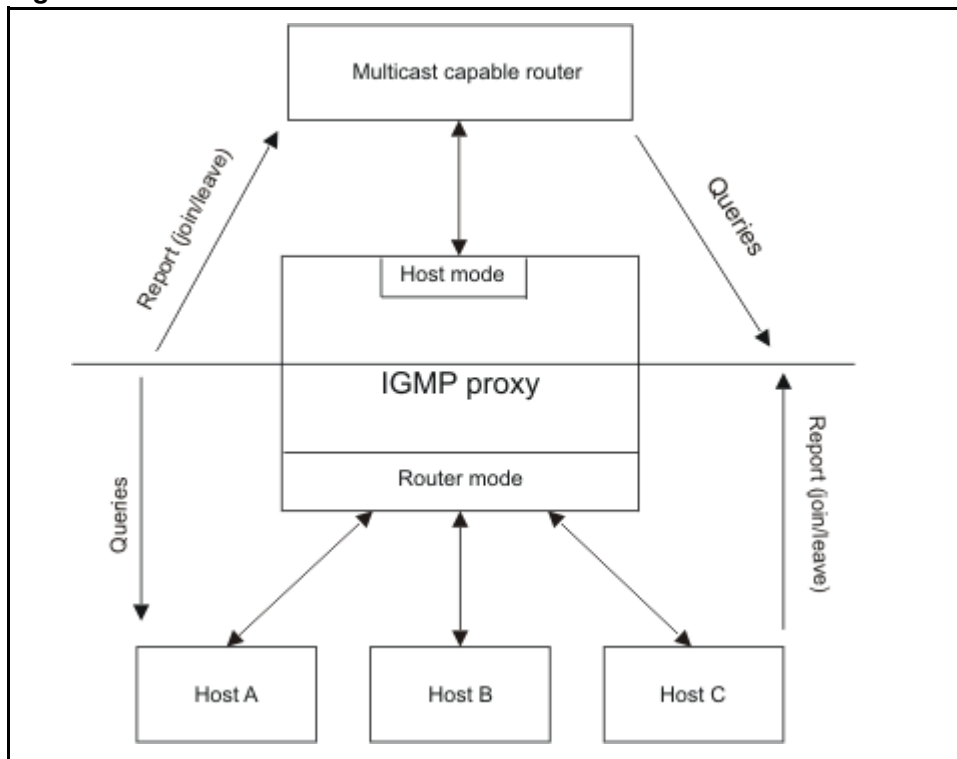
In host mode, the Nortel VPN Router implements an IGMP proxy on one interface. The IGMP proxy serves as a multicast group receiver to receive multicast traffic. Therefore, no multicast protocol implementation, such as PIM-SM, is required.

Both host and router interfaces can be the following types:

- physical interfaces
- user tunnels
- branch office tunnels
- multinets
- virtual interfaces

The host interface cannot be the end of an user tunnel.

[“IGMP modes” on page 103](#) demonstrates how the IGMP modes operate.

Figure 6 IGMP modes

IGMP versions

The VPN Router supports IGMPv3 and is backward compatible with IGMPv1 and IGMPv2. A Nortel VPN router configured for IGMPv3 can communicate with IGMPv1 or IGMPv2 hosts or routers. In addition, you can configure the IGMP feature to work in pure IGMPv1 and IGMPv2 versions. The following sections describe the main features of each version.

IGMPv1

IGMPv1 provides support for IP multicast routing. IGMPv1 specifies the mechanism for communicating IP multicast group membership requests from a host to the locally attached IGMP routers. Hosts use IGMPv1 to join multicast groups. Because Leave messages are not used in IGMPv1, a time-out mechanism is used to end group membership. The VPN Router does not use IGMPv1. For more information about IGMPv1, see RFC 1112.

IGMPv2

IGMPv2 extends the features in IGMPv1 by quickly reporting group membership termination by using Leave messages. This feature is important for multicast groups with highly volatile group membership and for high bandwidth multicast groups.

IGMPv3

IGMPv3 supports the PIM Source-Specific Multicast (SSM) protocol. IGMPv3 provides the ability for a host to selectively request or filter traffic from individual sources within a multicast group. Unwanted traffic is blocked.

IGMPv3 filters multicast packets based on the Source IP address–Destination Group Address pair rather than by the destination Group Address only.

IGMPv3 uses two datagram types: the Membership Query message (Type=0x11) and the Membership Report message (Type=0x22).

IGMP version interoperability

IGMPv3 routers process Query messages to determine the IGMP version used by other routers:

- If the Query message has a length of 8 bytes and the Maximum Response Code field is zero, the version is IGMPv1.
- If the Query message has a length of 8 bytes and the Maximum Response Code field is nonzero, the version is IGMPv2.
- If the Query message has a length of twelve bytes or greater, the version is IGMPv3.

If multiple versions of IGMP are used on the network, all routers must be configured to use the lowest common version of IGMP. IGMPv3-capable routers support IGMPv1 Compatible Mode and IGMPv2 Compatible Mode.

IGMP message types

For group addresses, IGMP uses class D IP addresses, ranging from 224.0.0.0 to 239.255.255.255. The three message types are as follows:

- Membership Queries
 - general query
 - group-specific query
 - group-and-source-specific query
- Membership Report
- Leave Group

IGMPv1 and IGMPv2 messages

IGMP encapsulates messages IP datagrams with protocol number 2 and sends the messages with an IP TTL of 1. The total message length is 8 bytes for IGMPv1 and IGMPv2. IGMPv3 supports a variable message length.

[“IGMPv1 and IGMPv2 datagram” on page 106](#) shows the format of IGMPv1 and IGMPv2 datagrams.

Figure 7 IGMPv1 and IGMPv2 datagram

“[IGMPv1 and IGMPv2 datagram fields](#)” on page 106 describes the IGMPv1 and IGMPv2 datagram fields.

Table 20 IGMPv1 and IGMPv2 datagram fields

Field	Description
Type	<p>Specifies the type of IGMP message:</p> <ul style="list-style-type: none">• 0x11—the datagram is a Membership Query, which is a General Query, is used to learn which groups have members. A Group-Specific Query is used to learn if a particular group has members.• 0x12—the datagram is a Version 1 Membership Report. This provides backward compatibility with IGMPv1.• 0x16—the datagram is a Membership Report.• 0x17—the datagram is a Leave Group message. <p>IGMPv2 ignores all other datagram types.</p>
Max Resp Time	<p>Specifies the maximum amount of time that can elapse before a response is returned after receipt of a Membership Query. This parameter is specified in tenths of a second and is a 1-byte field. The Maximum Response Time field is used only for Membership Query messages; the field is configured to zero for all other message types.</p>
Checksum	<p>Ensures that the IGMP packet is not corrupt. Checksum is a 2-byte field and is the ones complement sum of the whole IGMP packet. The checksum is computed and placed in the Checksum field prior to packet transmission. Before the packet is processed by the receiver, the checksum is verified.</p>
Group Address	<p>Communicate the group IP address in Membership Reports, Group-Specific Queries, and Leave Group datagrams. The group address is a 4-byte field. This field is configured to zero for General Queries.</p>

For more information about IGMPv2, see RFC 2236.

IGMPv3 messages

IGMPv3 supports two message types that previous versions do not support:

- group-and-source-specific query, which the Nortel VPN Router uses to learn if neighboring interfaces requests to join a group based on a specified list of sources
- membership report, which include Current-State, Filter-Mode-Change, and Source-List Change records.

Membership Queries

Membership Queries are sent to query the multicast state of neighboring interfaces. [“Membership Query message” on page 108](#) shows a Membership Query message.

Figure 8 Membership Query message

Type = 0x22	Max Resp Code	Checksum
Group Address		
Resv [S] QRV	QQIC	Number of Sources
Source Address [1]		
Source Address [2]		
...		
Source Address [n]		

“[IGMPv3 Membership Query fields](#)” on page 108 describes the IGMPv3 Membership Query fields.

Table 21 IGMPv3 Membership Query fields

Field	Description
Type	<p>Specifies the type of IGMP message:</p> <ul style="list-style-type: none"> • 0x11—the datagram is a Membership Query, which is a General Query, is used to learn which groups have members. A Group-Specific Query is used to learn if a particular group has members. • 0x12—the datagram is a Version 1 Membership Report. This provides backward compatibility with IGMPv1. • 0x16—the datagram is a Membership Report. • 0x17—the datagram is a Leave Group message. <p>IGMPv2 ignores all other datagram types.</p>
Checksum	<p>Ensures that the IGMP packet has not been corrupted. Checksum is a 2-byte field and is the ones complement sum of the whole IGMP packet. The checksum is computed and placed in the Checksum field prior to packet transmission. Before the packet is processed by the receiver, the checksum is verified.</p>
Maximum Response Code	<p>Calculates the Maximum Response Time. With the VPN Router, use the Maximum Response Time to control the burstiness of Query messages on the network. The larger the Maximum Response Time, the less bursty the traffic because the host replies spread over a larger time interval.</p>

Table 21 IGMPv3 Membership Query fields

Field	Description
Group Address	Identifies the IP multicast address queried for Group-Specific Queries and Group- and- Source-Specific Queries. The Group Address is configured to zero for General Queries.
Resv [S] QRV	<p>Contains the Robustness variable that the Querier uses. In RFC 3376, Internet Group Management Protocol, Version 3, the QRV is described as follows:</p> <p>If nonzero, the QRV field contains the Robustness Variable used by the querier, that is, the sender of the Query. If the querier Robustness Variable exceeds 7 (the maximum value of the QRV), the QRV is configured to zero. Routers adopt the QRV from the most recently received Query as their own Robustness Variable value, unless that most recently received QRV is zero, in which case the receivers use the default Robustness Variable or a statically configured value.</p> <p>Use the QRV to compensate for packet loss on the network. If the network is expected to be lossy, increase the QRV.</p>
QQIC	Specifies the query interval in seconds used by the querier. Use this value to control how many IGMP messages are sent. The larger the querier Query Interval Code (QQIC) value, the fewer the Query messages that traverse the network.
Number of Sources (N)	Identifies the number of source addresses
Source Address [N]	Identifies the source IP address.

Membership Reports

“[Membership Query report](#)” on page 110 shows a Membership Query report.

Figure 9 Membership Query report

Type = 0x22	Reserved	Checksum
Reserved		Number of Records (M) Group
Group Record [1]		
Group Record [2]		
...		
Group Record [M]		

[“IGMPv3 Membership Report fields” on page 110](#) describes the IGMPv3 Membership Report fields.

Table 22 IGMPv3 Membership Report fields

Field	Description
Type	<p>Specifies the type of IGMP message:</p> <ul style="list-style-type: none">• 0x11—the datagram is a Membership Query, which is a General Query, is used to learn which groups have members. A Group-Specific Query is used to learn if a particular group has members.• 0x12—the datagram is a Version 1 Membership Report. This provides backward compatibility with IGMPv1.• 0x16—the datagram is a Membership Report.• 0x17—the datagram is a Leave Group message. <p>IGMPv2 ignores all other datagram types.</p>
Checksum	<p>Ensures that the IGMP packet is not corrupt. Checksum is a 2-byte field and is the ones complement sum of the whole IGMP packet. The checksum is computed and placed in the Checksum field prior to packet transmission. Before the packet is processed by the receiver, the checksum is verified.</p>

Table 22 IGMPv3 Membership Report fields

Field	Description
Number of Records (M) Group	Identifies the number of records.
Group Record [M]	Identifies the group record. The Group Record contains information about the membership of the sender in a single multicast group on the interface from which the report is sent. Different types of Group Records can be included in the Report message.

Each Group Record has the following internal format:.

Record Type	Aux Data Len	Number of Sources (N)
Multicast Address		
Source Address [1]		
Source Address [2]		
...		
Source Address [N]		
Auxiliary Data		

“IGMPv3 Group Record fields” on page 112 describes the IGMPv3 Group Record fields.

Table 23 IGMPv3 Group Record fields

Field	Description
Record Type	<p>Specifies the Group Record type. The types include:</p> <ul style="list-style-type: none"> Current-State Record, which a host system sends in response to a query received on the interface. The two possible values are <code>MODE_IS_INCLUDE</code> and <code>MODE_IS_EXCLUDE</code>. The Nortel VPN Router uses the filter mode to decide whether a packet is dropped. Filter-Mode-Change Record, which the host system sends after a local invocation of <code>IPMulticastListen</code> causes a filter mode change. The two possible values are <code>CHANGE_TO_INCLUDE_MODE</code> and <code>CHANGE_TO_EXCLUDE_MODE</code>. Source-List-Change record, which the local host sends after a local invocation of <code>IPMulticastListen</code> causes a source list change that does not coincide with a filter mode change of the interface-level state entry for a multicast address. The two possible values are <code>ALLOW_NEW_SOURCES</code> and <code>BLOCK_OLD_SOURCES</code>. <p>The VPN Router ignores Membership Report messages containing unrecognized record type values.</p>
Auxiliary Data Length	Specifies the length of Auxiliary data. The VPN Router configures this field to zero.
Source Address [N]	Specifies the source address.
Auxiliary Data	Specifies auxiliary data.

Host Leave messages

If the IGMPv2 host that issued the most recent report leaves a group, it issues a Leave group message. The multicast router on the network issues a group-specific query to determine whether other group members are present on the network. If no host responds to the query, the IGMP router assumes that no members of that group exist on that interface.

IGMP MIB considerations

SNMP implementation on the Nortel VPN Router supports IGMP MIB Version 1 and Version 2 according to RFC 2933.

The MIB module contains two tables

- the IGMP Interface Table, which contains one row for each interface on which IGMP is enabled
- the IGMP Cache Table, which contains one row for each IP multicast group for which members are on a particular interface

Both tables are intended to be implemented by hosts and routers, but some columnar objects in each table apply only to routers.

IGMP configuration

To configure IGMP, perform the following procedures:

- 1 [“Disabling multicast relay” on page 113](#)
- 2 [“Enabling split tunneling” on page 114](#)
- 3 [“Configuring IGMP on an interface” on page 114](#)
- 4 [“Configuring IGMP globally” on page 115](#)
- 5 [“Configuring IGMP on branch offices” on page 115](#)

Disabling multicast relay

Before you enable IGMP, you must disable multicast relay because IGMP and multicast relay cannot function together.

To disable multicast relay, perform the following steps:

- 1 Go to the **Routing, Multicast** window.
- 2 Clear the **Enabled** check box for **Multicast Relay**.
- 3 Click **OK**.

Enabling split tunneling

For IGMP multicast traffic to pass on the client tunnels, you must enable split tunneling on the Nortel VPN Router.

To enable split tunnelling, perform the following steps:

- 1 Go to the **Profiles, Groups** window.
- 2 Click **Edit** next to the group for which you want to enable split tunneling.
The Groups > Edit window appears.
- 3 Click **Configure** for the IP Security feature.
The Groups > Edit > IPSec window appears.
- 4 Click **Configure** for **Split Tunneling**.
- 5 Select **Enabled** from the **Split Tunneling** list.
- 6 Click **OK**.

Configuring IGMP on an interface

To configure IGMP on an interface, perform the following steps:

- 1 Go to the **Routing, Interfaces**.
- 2 Select an interface.
- 3 Click **Configure** next to IGMP.
The Routing Interfaces > Configure IGMP window appears.
- 4 Enter a value in the **Robustness Value** box. The default is 2.
- 5 In the **Proxy Type** list, select **Upstream**.
- 6 Enter a value in the **Unsolicited Report Interval** box. The default is 1.
- 7 Select the **Enabled** check box.
- 8 Click **OK**.

Configuring IGMP globally

To enable IGMP globally, perform the following steps:

- 1 Go to the **Routing, IGMP** window.
- 2 Select the **Enabled** check box. IGMP is disabled by default.
- 3 Select an option from the **Software Version** list. The default software version is V2.
- 4 Enter a value in the **Membership Group Limit** box. The default limit is 80.
- 5 Click **OK**.

Configuring IGMP on branch offices

To configure IGMP on branch offices, perform the following steps:

- 1 Go to the **Profiles, Branch Office** window.
- 2 Select a group.
- 3 Click **Configure**.
The Branch Office > Edit Group window appears.
- 4 Click **Configure** for IGMP.
The Branch Office > Edit > IGMP window appears.
- 5 Enter a value in the **Query Interval** box.
- 6 Enter the maximum query response time in the **Query Max Response** box.
- 7 Enter the last member query interface value in the **Last Member Query Interval** box.
- 8 Enter the designated router life interval value in the **Designated Router Life Interval** box.
- 9 Enter the robustness value in the **Robustness Value** box.
- 10 Enter the unsolicited response value in the **Unsolicited Report Interval** box.
- 11 Select an access filter from the **Access Filter** list.

If no access filter exists, click the **New Access List** link to add one. when access list is Exact, the default group mask is 32-bit.

- 12** Click **OK**.
- 13** Click the **Return to Branch Office** link.
- 14** Select a connection.
- 15** Under the **Connection** section, click **Configure**.
- 16** In the **Multicast Configuration** section, select **Enabled** from the **IGMP State** list.
- 17** From the **IGMP Proxy Type** list, select the **Upstream** or **Downstream**.
- 18** Click **OK**.

Chapter 11

VRRP configuration

This chapter contains information about Virtual Router Redundancy Protocol (VRRP) and the procedures to configure it.

This chapter includes the following topics:

- [“VRRP fundamentals” on page 117](#)
- [“VRRP configuration” on page 123](#)

VRRP fundamentals

VRRP is a standard protocol used by the Nortel VPN Router to handle private interface failures. VRRP is one method to help maintain a state of high availability. Hosts configured with static or default Nortel VPN Routers obtain a resilient next-hop address. VRRP provides Nortel VPN Router-level failover if a private physical interface fails. Use VRRP and dynamic routing to provide a high degree of redundancy.

A virtual router ID is a software-defined object that corresponds to an IP address on a LAN or VLAN segment. You define the state and rate of each Nortel VPN Router within the virtual router group. The rate determines how quickly failover occurs. VRRP also handles information that determines the rate and state of each Nortel VPN Router within the virtual router group. This information relates to an interface and the role that the interface plays in VRRP (master or backup). This information is kept in the normal configuration file stored in the Nortel VPN Router configuration file.

For a LAN, VRRP associates one IP address with two physical routes. This association is a virtual router. On a LAN segment, a virtual router has the following properties:

- Virtual router ID
- Rate or frequency of messages between VRRP and spokes on the LAN

For a VLAN, VRRP associates one IP address with two virtual routes. This association is a virtual router. On a VLAN segment, a virtual router has the following properties:

- Virtual router ID
- Rate or frequency of messages between VRRP and the VLAN

An external Lightweight Directory Access Protocol (LDAP) server is not a requirement, but can make VRRP easy to use. The LDAP server provides a common location to maintain information for each Nortel VPN Router. Use of an LDAP server enables each Nortel VPN Router to determine the virtual router settings of other Nortel VPN Routers in the system.

To configure VRRP, the virtual router ID (VRID) for the virtual router group must be identical to all Nortel VPN Routers. If you use the internal LDAP server, you must configure the virtual router parameters the same as on the Nortel VPN Routers.

Nortel recommends that you do not use a four-port switch (Lan0) in a VRRP configuration for Nortel VPN Router 1100 platforms.

Support is unavailable for VRRP on Nortel VPN Router 1050 platforms.

VRRP and dynamic routing for high availability

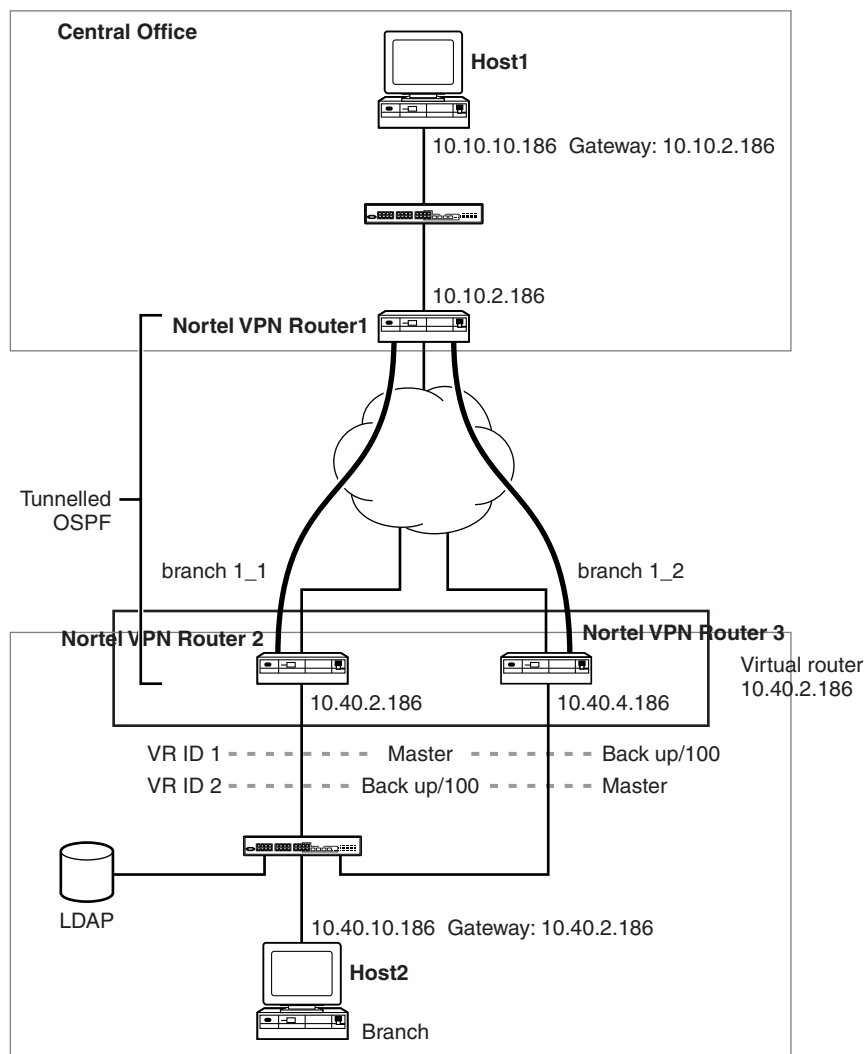
High availability (HA) depends on the core routing features, VRRP, a dynamic routing protocol (RIP, BGP, or OSPF), and consistent configuration.

[“Sample high-availability environment” on page 120](#) shows a deployment in which the central office is configured with one Nortel VPN Router, VPN1, and Host1 with the default Nortel VPN Router pointing to VPN1. The branch is configured with two Nortel VPN Routers: VPN2 and VPN3. VRRP is configured on the private side of VPN2 and VPN3 backing up the physical interface of the other router. Host2 in the branch has a default VPN Router pointing to VPN2. Two branch office tunnels, as indicated, connect the VPN Routers.

Consider the traffic flow between Host2 and Host1. If VPN2 fails or the private interface of VPN2 fails, VPN3 becomes the master of the private interface of VPN2. The VPN3 IP address changes to 10.40.2.186 and takes the MAC address of VPN2 interface. All IP traffic from Host2 to Host1 now flows through VPN3. VPN3 forwards all VPN2 routed packets, but drops packets destined to VPN2. For example, a data packet from Host2 to Host1 is forwarded, but a ping request to 10.2.40.186 is dropped. Host1 is not aware of the change.

Routing configuration plays a vital role in this failover operation. VPN2 and VPN3 must identify that the path to Host1 is through VPN1. VPN1 must identify the two paths to Host2: one through VPN2 and another through VPN3. You can manually populate the routing information about the each Nortel VPN Router causing static routes; however, dynamic routing protocols such as RIP, BGP, or OSPF provide reliable route information in networks that are considered dynamic or volatile (route information changes often). In this case, OSPF, BGP, or RIP updates VPN2 so that VPN1 no longer has a route to Host2.

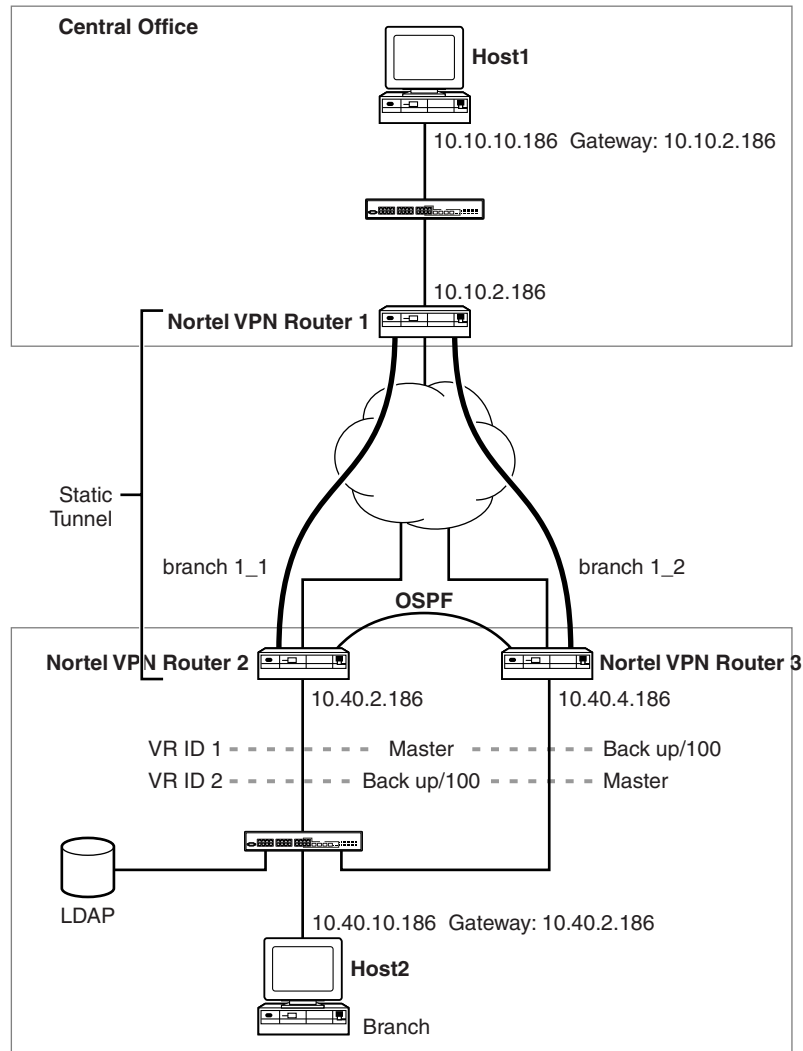
The VRRP failover occurs within 3 seconds based on the default configuration. Use of OSPF on the tunnels guarantees a maximum failover time of 40 seconds based on the default configuration. However, you can reduce the failover time by configuring the appropriate value for the OSPF hello interval. Use of RIP takes a maximum of 2.5 minutes based on the default configuration. You can modify the RIP parameters to reduce this time.

Figure 10 Sample high-availability environment

If the branch office tunnels are static routes and the Host2 default gateway VPN2 encounters a public interface failure (private interface 10.40.2.186 remained active), VRRP does not failover. If VPN2 is unaware that another route to Host1 exists through VPN3, it drops traffic from Host2 destined to Host1. To correct this, you must add another route to Host1 through VPN3 to the VPN2 route table. One way to add this route is shown in [“VRRP and static tunnels” on page 121](#).

In “[VRRP and static tunnels](#)” on page 121, an OSPF branch office tunnel between VPN2 and VPN3 provides both routers with a secondary route to Host1. Because static routes are preferred over OSPF, both VPN2 and VPN3 always use the static route to Host1 through VPN1 if it is available. This inter-VPN Router branch office tunnel does not have to use OSPF. RIP or a static route of higher cost works equally well.

Figure 11 VRRP and static tunnels



Interface groups and critical interface failover

Interface groups support the backup interface services (BIS), which is an automated mechanism to back up an interface after a designated primary connection fails, and VRRP critical interface failover (CIF). An interface group is a logical group of interfaces (physical or tunnel) defined in a Nortel VPN Router. The group can consist of the following elements

- single physical interface
- IP address
- list of physical interfaces
- tunnel
- list of tunnels
- combination of physical interfaces and tunnels

Status of a critical interface is defined to be up if at least one member is up and is down if all members are down.



Note: Branch office tunnels in an interface group for a critical interface for a VRRP must be nailed-up rather than on-demand.

For VRRP, a physical interface on which a VRRP is configured to run as master is called a VRRP master interface. With each VRRP master interfaces, you can associate a maximum of three interface groups. If one of these three interface groups goes down, the Nortel VPN Router behaves as if the VRRP master interface is down and forcing a VRRP failover. The VRRP master interface stays in this down state until all associate interface groups come up, and then claims the mastership.

For Demand Services, if all interfaces in the critical interfaces group fail to operate properly, an event is triggered and the backup interface services associated with that critical interface group are enabled.



Note: The interface IP address and the management IP address share the same interface. If the interface is down, all IP addresses on that interface are also down.

The Configured Interface Groups section of the Routing, Interface Group window lists the names of configured interface groups, the number of IP interfaces included in the group, and the current administrative and operational states of the group.

If you delete an active interface group, you must then go to the Routing, Interface Group window and click OK.



Note: When you configure a critical interface or interface groups for critical interface failover (CIF), you cannot have VRRP configured on the interfaces. If you include an interface that is running VRRP as a critical interface or part of an interface group (for CIF), it is an unsupported configuration. Where VRRP is configured on the interface, there is already a failover/availability solution provided in case of loss of that interface.

VRRP configuration

To configure VRRP, perform the following procedures:

- [“Configuring VRRP for LAN and VLAN” on page 123](#)
- [“Configuring interface groups” on page 128](#)

Configuring VRRP for LAN and VLAN

To configure VRRP for LAN and VLAN, perform the following steps:

- 1 Choose **Routing, VRRP** window.
- 2 Check **Enable**.
- 3 Enable the **Respond ICMP Packet** to make the Nortel VPN router respond to ICMP packets (PING) whenever VRRP becomes master for a backed-up IP address.
- 4 Enter the IP address for the virtual router.
- 5 Click **Create**.

The VRRP > Addresses Configured for VRRP window appears.

- 6** Enter an ID in the **VRID** box.

The value must be a decimal value from 1 to 255 for the VRID. This number must be unique to the LAN or VLAN segment that runs VRRP and common to all Nortel VPN Routers that participate within this virtual router group.

- 7** In the **Advertise Interval** box, enter the rate the virtual router advertises the hello messages.

The range is 1 to 255 seconds and the default is 1 second.

- 8** From the **Preempt Mode** list, select **True** to enable a high priority backup router to preempt a low priority master. The default is False.

If Preempt is enabled and a VRRP router comes online, it compares the parameters with the current master advertisement. If the new candidate has a higher priority, it becomes the new master.

- 9** Select an **Authentication Type** for this virtual router.

None means that VRRP protocol exchanges are not authenticated; Simple means they are authenticated by a simple text password.

If you choose Simple authentication, enter up to eight characters of text for the authentication string and confirm it.

- 10** Select a **Master Delay Mode**.

The Master Delay Mode controls when a Nortel VPN Router takes mastership of an address it owns. Normally, this occurs as soon as the interface is enabled. Master Delay mode makes it is possible to delay the master assertion. Master Delay mode operates in one of two possible ways: Delay or Time of Day. The default for a VR in Master Delay mode is disabled (None).



Note: If enable safe mode is enabled, a boot after an unclean failure starts the Safe mode image, instead of the normal boot image. If the Safe mode image is configured with VRRP, Master Delay mode works. However, Safe mode automatically boots the normal image after a configured delay. This boot appears as clean shutdown, and Master Delay mode is not invoked.

- 11** Click **OK**.

- 12** Choose **Routing, Interfaces**.

- 13** Click **Configure** next to VRRP for the appropriate interface.

The LAN (with corresponding physical address on the box) and VLAN interfaces automatically appear in the Master Status section and all others appear in the Current Backed up Addresses section.

- 14 Check **Enable** to enable VRRP for this interface.
- 15 In the **Master Status** section, enable all interfaces that you want to be masters.
- 16 Click **OK**.

The Current Backed up Addresses section displays information about the currently configured backups, including IP addresses that this subinterface backs up, VRID, configured state, current operational state, and priority.

- 17 In the **New Backed up Address** section, back up an IP address by selecting an IP address from the list.
- 18 Enter a priority number in the **Priority** box.
- 19 Click **Add**.
- 20 Click **OK**.

Configuration examples of IP addresses for backups

In the Routing, VRRP window, you can configure the IP addresses for the virtual router and the remote addresses that you need to back up. The IP address of the virtual router must be one of the Nortel VPN Router interfaces, but it does not have to be the master. The addresses you backed up are not on the local Nortel VPN Router.

For example, for VPN2 to be the master of VRID 1 and VPN3 to be the backup, configure the following:

- 1 On VPN2, choose **Routing, VRRP**.
- 2 Enter the IP address **10.40.2.186**.
- 3 Click **Create**.
- 4 Enter **1** in the **VRID** box.
- 5 Click **OK**.
- 6 Choose **Routing, Interfaces**.

- 7 Select 10.40.2.186.
- 8 Click **Configure**.
- 9 Check the **Master** box.
- 10 On VPN3, choose **Routing, VRRP**.
- 11 Enter the IP Address **10.40.4.186**.
- 12 Click **Create**.
- 13 Enter **2** in the **VRID** box.
- 14 Choose **Routing, Interfaces**.
- 15 Select 10.40.4.186.
- 16 Click **Configure**.
- 17 From the New Backed up Address, select **10.40.2.186, VRID 1**, and click **ADD**.

To configure VPN3 to be the master of VRID 2, configure the following:

- 1 On VPN2, choose **Routing, VRRP**.
- 2 Enter the IP address **10.40.2.186**.
- 3 Click **Create**.
- 4 Enter **1** in the **VRID** box.
- 5 Click **OK**.
- 6 Enter the IP address **10.40.4.186**.
- 7 Click **Create**.
- 8 Enter **2** in the VRID box.
- 9 Choose **Routing, Interfaces**.
- 10 Select **10.40.2.186**.
- 11 Select the **Master** for 10.40.2.186.
The Backed Up list contains 10.40.4.186 VRID 2.
- 12 On VPN3, choose **Routing, VRRP**.
- 13 Enter the IP address **10.40.4.186**.

- 14** Enter **2** in the **VRID** box.
- 15** Choose **Routing, VRRP**.
- 16** Enter the IP address **10.40.2.186**.
- 17** Enter **1** in the **VRID** box.
- 18** Choose **Routing, Interfaces**.
- 19** Select the **Master** box for 10.40.4.186.
The Backed Up list contains 10.40.2.186 VRID 1.

For a VLAN to be the master of VRID 1 and VPN3 to be the backup, configure the following:

- 1** On VLAN, choose **Routing, VRRP** window
- 2** Enter the IP address **1.1.1.1**.
- 3** Click **Create**.
- 4** Enter **1** in the **VRID** box.
- 5** Choose **Routing, Interfaces**.
- 6** Select **Configure** for 1.1.1.1.
- 7** Check the **Master** box.
- 8** On VPN3, choose **Routing, VRRP**.
- 9** Enter the IP Address **10.40.4.186**.
- 10** Click **Create**.
- 11** Enter **2** in the **VRID** box.
- 12** Enter the IP Address **1.1.1.1**.
- 13** Click **Create**.
- 14** Enter **1** in the **VRID** box.
- 15** Choose **Routing, Interface**.
- 16** Select **Configure** for 10.40.4.186.
- 17** From the **New Backed up Address** list, select **1.1.1.1, VRID 1**.
- 18** Click **ADD**.

Configuring interface groups

To configure interface groups, perform the following steps:

- 1** Choose **Routing, Interface Grp.**
- 2** Click **Add.**
- 3** Enter a name for the group in the **Name** box.
- 4** Select an interface from the **Available Interfaces** list.
- 5** Move the available interface into the **Interfaces in Group** list.
You can add more than one interface to an interface group.
- 6** After you add interfaces to the interface group, click **OK.**
- 7** To find interface groups with an interface, enter the IP address, and click **Search.**
- 8** Click **OK** or **Close.**
- 9** Choose **Routing, Interfaces, Configure VRRP** window for the VRRP interface you want to associate with the critical interface group.
- 10** Under **Master Status**, select the interface group from the list.
- 11** Click **Enabled.**
- 12** Click **OK** to enable the VRRP critical interface.

Chapter 12

ECMP configuration

This chapter contains information about Equal-cost Multipath (ECMP) fundamentals and the procedure to configure it.

This chapter includes the following topics:

- [“ECMP fundamentals” on page 129](#)
- [“Configuring ECMP” on page 129](#)

ECMP fundamentals

ECMP provides load balancing of packets to a destination that is reachable over more than one network path. ECMP increases the forwarding capacity of a Nortel VPN Router that is media bound and balances loads on individual packets or a packet-stream. ECMP balances traffic across tunnels whether packets are going out single or multiple physical interfaces. ECMP is supported for routes originating from the static, BGP, RIP, or OSPF routing applications.

ECMP allows the static, OSPF, BGP, and RIP routing applications to submit multiple routes to a single destination of the same cost. The route table manager passes the set of equal-cost best paths to the forwarding table. The Nortel VPN Router supports up to four equal cost paths for OSPF, BGP, and RIP and up to eight paths for static.

You must have the Advanced Routing license key installed to use ECMP.

Configuring ECMP

To configure ECMP, perform the following steps:

- 1 Choose **Routing, Configuration**.
- 2 Select an option in the **Maximum Paths** list to configure the maximum equal-cost paths allowed globally by the Nortel VPN Router.
- 3 If you use **OSPF**, select an option in the **OSPF Maximum Paths** to configure the maximum equal-cost paths for OSPF.
- 4 If you use **BGP**, select an option in the **BGP Maximum Paths** to configure the maximum equal-cost paths for BGP.
- 5 If you use **RIP**, select an option in the **RIP Maximum Paths** to configure the maximum equal-cost paths for RIP.
- 6 Select an option in the **Forwarding Algorithm** list.

You can select a forwarding algorithm without affecting route or forwarding tables. The load balancing and resource sharing is controlled by the following forwarding algorithms:

- **Per-packet**—packets are forwarded in a round-robin fashion. If the Nortel VPN Router Stateful Firewall is enabled, this policy can cause some overhead in switching the firewall context.
 - **Per-destination**—packets are forwarded based on source and destination address pair.
 - **Per-source**—packets are forwarded based on source address.
- 7 Click **OK**.

Index

A

Accept policies 85
advanced routing key 46, 70
Announce policies 85

B

BGP 58

C

client address redistribution 91
 configuring 95
 sample 92
 summarization 93
Configuring IGMP 113
Configuring IGMP globally 115
Configuring IGMP on an interface 114
Configuring IGMP on branch offices 115

D

default route 33
Disabling multicast relay 113

E

Enabling split tunneling 114
equal cost multipath (ECMP) 129

H

Host Leave messages 112

I

IGMP configuration 101
IGMP fundamentals 101
IGMP Message types 104
IGMP MIB considerations 113
IGMP modes 101
IGMP version interoperability 104
IGMP versions 103
IGMPv1 103
IGMPv1 and IGMPv2 messages 105
IGMPv2 104
IGMPv3 104
IGMPv3 messages 107
interface filter
 Permit All 99

L

load balancing 130

M

Membership Reports 109
multicast 99
multicast relay 97

O

OSPF
 configuration 49
 overview 45

P

- Permit All 99
- ping
 - validating public default route 82
- poison reverse 37
- publications
 - hard copy 15

R

- RIP 35
 - using 36, 46
- route redistribution 87
- route selection 30
- route table 24, 27
 - lookup 29
- routes
 - default 33
 - dynamic 27
 - static 27
- routing
 - dynamic 23
 - enhanced 23
 - integrated firewall 22
 - loops 37
 - overview 21
 - policy 89
 - policy service 85
 - route lookup 29
 - route table types 29
 - rules of redistribution 87
 - table 27

S

- split horizon 37
- static routes 23, 79
- status 24

T

- technical publications 15
- triggered updates 37

U

- Utnunnel 91

V

- virtual links 47
- VRRP
 - configuring 123
 - failover 119
 - high availability 118
 - master interface 120
- VRRP overview 117