

Version 8.00

Part No. NN46110-508 01.01

324659-A Rev 01

13 October 2008

Document status: Standard

600 Technology Park Drive

Billerica, MA 01821-4130

Nortel VPN Router Configuration — Firewalls, Filters, NAT, and QoS

The Nortel logo consists of the word "NORTEL" in a bold, sans-serif font. The letter "O" is stylized with a circular graphic element that resembles a globe or a network node, partially overlapping the letters "N" and "R".

NORTEL

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.

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, and the Globemark are trademarks of Nortel Networks.

Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.

Cisco and Cisco Systems are trademarks of Cisco Systems, Inc.

Java and Solaris are trademarks of Sun Microsystems.

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

Netscape, Netscape Communicator, Netscape Navigator, and Netscape Directory Server are trademarks of Netscape Communications Corporation.

SPARC is a trademark of Sparc International, Inc.

All other 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
How to get help	15
Finding the most recent updates on the Nortel Web site	16
Getting help from the Nortel Web site	17
Getting help over the phone from a Nortel Solutions Center	17
Getting help from a specialist by using an Express Routing Code	17
Getting help through a Nortel distributor or reseller	18
New in this release	19
Features	19
Interface filters	19
Branch office NAT Traversal	19
QoS information	20
Other changes	20
Document changes	20
Title change	20
Chapter 1	
Overview of firewalls, filters, and NAT	21
VPN Router Stateful Firewall concepts	22
Stateful inspection	23
Interfaces	24
Filter rules	24
Antispoofing	25
Attack detection rules	25

Filters for access control	26
Network Address Translation	27

Chapter 2

Stateful Firewall configuration 29

Configuration prerequisites	30
Java 2 software installation	31
Using Internet Explorer	31
Using Firefox	32
Enabling firewall options	33
Rule enforcement	36
Log options	36
Application-specific logging	37
Configuring remote system logging	37
Configuring antispoofing	38
Configuring malicious scan detection	39
Policy configuration	39
Firewall policy creation and modification	41
Policy creation	41
Adding a policy	41
Deleting a policy	42
Copying a policy	42
Renaming a policy	43
Navigating rules	43
Implied rules	43
Override rules	46
Interface-specific rules	47
Default rules	48
Rule creation	49
Header row menu	49
Row menu	49
Cell menus	49
Rule columns	49
Creating a new policy	54
Verifying the configuration	55

Configuring a sample security policy	55
Firewall deployment examples	57
Residential firewall example	58
Business firewall example	58
Chapter 3	
Filter configuration	61
Adding and editing filters	61
Management access restrictions	63
Configuring next-hop traffic filters	65
Chapter 4	
NAT configuration	67
Address translations	68
Dynamic many-to-one—port translation	68
Dynamic many-to-many—pooled translation	69
Static one-to-one translation	70
Port forwarding	71
Double NAT	72
IPsec-aware NAT	73
NAT modes	74
Full Cone NAT	74
Restricted Cone NAT	75
Port Restricted Cone NAT	76
Symmetric NAT	77
NAT Traversal	78
NAT and VoIP	81
Address and port discovery	82
Network address port translation (NAPT)	83
Configuring Cone NAT	84
NAT usage	85
Branch office tunnel NAT	85
Interface NAT	87
Dynamic routing protocols	88
NAT policy configuration	89

NAT policy sets	90
Rule creation	90
Creating a new policy	92
Adding a policy	93
Deleting a policy	94
Copying a policy	94
Renaming a policy	94
Sample NAT procedures	95
Configuring interface NAT with RIP	95
Configuring interface NAT with OSPF	95
Configuring branch office NAT with RIP	96
Configuring branch office NAT with OSPF	97
Configuring branch office NAT	97
Configuring NAT with the VPN Router Stateful Firewall	98
NAT ALG for SIP	99
Application level gateways	100
Configuring NAT ALG for SIP	101
Firewall SIP ALG	101
Configuring Firewall Virtual ALG	102
Hairpinning	104
Hairpinning with SIP	104
Hairpinning with a UNISTim call server	105
Hairpinning with a STUN server	108
Hairpinning requirements	108
Enabling hairpinning	109
Timeouts	109
NAT statistics	110
Proxy ARP	110
Chapter 5	
Firewall user authentication configuration	113
Chapter 6	
QoS configuration	121
Admission control	122
Globally enabling Admission Control	122

Over-subscription example	124
Bandwidth Management	124
Configuring Bandwidth Management	124
Call Admission Priority	125
Forwarding Priority	127
NNSC queues	128
Critical and Network service classes	128
Premium service class	129
Metal service classes	129
Standard service class	130
Queuing mechanisms	131
Weighted fair queuing	132
Strict priority	132
Congestion avoidance	132
Differentiated Services	133
Assured Forwarding PHB group	135
Expedited Forwarding PHB group	136
Classifier configuration	137
Configuring an MF classifier	140
Using a BA classifier and the current DSCP	140
Configuring DiffServ	141
DSCP to 802.1p mapping	142
Configuring DSCP to 802.1p mapping	145
Router-generated packets	145
Traffic conditioning	146
EF outbound traffic conditioning	148
Configuring traffic conditioning	148
Configuring interface shaping	149
RSVP	150
.....	150
Index	151

Preface

This guide provides overview and configuration information for the Nortel VPN Router Stateful Firewall and VPN Router filters.

Before you begin

This guide is for network managers who set up and configure the VPN Router. This guide assumes that you have experience with windows-based 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 (<>) | Indicate that you choose the text to enter based on the description inside the brackets. Do not type the brackets when entering the command.
Example: If the command syntax is 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>Indicate required elements in syntax descriptions where there is more than one option. You must choose only one of the options. Do not type the braces when entering the command.</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>Indicate optional elements in syntax descriptions. Do not type the brackets when entering 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>Indicate 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 name of the file.</p>
<i>italic text</i>	<p>Indicates new terms, book titles, and variables in command syntax descriptions. Where a variable is two or more words, the words are connected by an underscore.</p> <p>Example: If the command syntax is ping <ip_address>, <i>ip_address</i> is one variable and you substitute one value for it.</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 type the vertical line when entering 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 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 Configuration — Routing* (NN46110-504) provides instructions to configure the Border Gateway Protocol (BGP), Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Virtual Router Redundancy Protocol (VRRP), Equal Cost Multipath (ECMP), routing policy services, and client address redistribution (CAR).
- *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 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 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

To print selected technical manuals and release notes for free, directly from the Internet, go to www.nortel.com/documentation, find the product for which you need documentation, then locate the specific category and model or version for your hardware or software product. Use Adobe Reader to open the manuals and release notes, search for the sections you need, and print them on most standard printers. Go to the Adobe Systems Web site at 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 most recent 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 most recent documentation and software for VPN Router, click one of the following links.

Link	Web site
Most recent software	Nortel page for VPN Router software located at support.nortel.com/go/ main.jsp?cscat=SOFTWARE&poid=12325
Most recent documentation	Nortel page for 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 perform the following activities:

- 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 the following Web site:

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's new in *Nortel VPN Router Configuration — Firewalls, Filters, NAT, and QoS* (NN46110-508) for Release 8.0.

- [“Features” on page 19](#)
- [“Other changes” on page 20](#)

Features

For information about feature-related changes, see the following sections:

- [“Interface filters” on page 19](#)
- [“Branch office NAT Traversal” on page 19](#)
- [“QoS information” on page 20](#)

Interface filters

Interface filters do not apply to packets sent to internal circuitless IP (CLIP) addresses. For more information about filters, see [“Filter configuration” on page 61](#).

Branch office NAT Traversal

Release 8.0 introduces Network Address Translation (NAT) traversal for branch office tunnels between VPN Routers when one router is in a private network that uses one or more NAT devices. For more information about NAT Traversal, see [“NAT Traversal” on page 78](#).

QoS information

For more information about existing features, see [“QoS configuration” on page 121](#).

Other changes

For more information about changes that are not feature related, see the following sections:

- [“Document changes” on page 20](#)
- [“Title change” on page 20](#)

Document changes

This document is changed to comply with Nortel writing conventions.

Title change

This document is renamed from *Nortel VPN Router Security — Firewalls, Filters, NAT, and QoS* (NN46110-601).

Chapter 1

Overview of firewalls, filters, and NAT

The VPN Router designs integrated firewall solutions to meet the needs of a variety of customers. The VPN Router provides the following firewall solutions:

- VPN Router Stateful Firewall
- VPN Router Interface Filters

With the VPN Router Stateful Firewall, the VPN Router performs a variety of secure routing functions, which depends on how you configure the routing capabilities. For example, you can configure the VPN Router to securely route nontunneled traffic from its private interface, through the firewall, and out its public interface. With this configuration, users on the private network can access the Internet without requiring a separate, dedicated router. The VPN Router Stateful Firewall achieves optimum performance because of advanced memory management techniques and optimized packet inspection.

The VPN Router Stateful Firewall provides a high level of security, the fastest runtime, and the flexibility to define the rules to fit your environment. The Stateful Firewall delivers full firewall capabilities and assures the highest level of network security. To do this, the Stateful Firewall examines both incoming and outgoing packets and compares them to a common security policy. All service rules are interpreted based on IP conversations (not packets) and are fully stateful. Security rules do not filter packets directly, but the Stateful Firewall services base how to process the packets on the defined security policy.

The VPN Router interface filters provide a cost-effective level of protection. You can disable the interface filters only after you enable the VPN Router Stateful Firewall.

Because no routing protocols run on untrusted interfaces, the IP public address table (PAT) provides the routing information to route packets to the appropriate trusted interfaces. The IP PAT limits unauthorized sources. If you enable either VPN Router Stateful Firewall or VPN Router Interface Filter, the router disables PAT because the former two provide better policy-based security.

After you disable the firewall, PAT applies only to packets received on a public interface. PAT maintains a list of trusted sources that includes the remote client or branch office tunnel end point, Remote Authentication Dial-In User Service (RADIUS), Certificate Management Protocol (CMP), or Certificate Revocation List (CRL) server address (if on the public side). PAT does not limit the packets from those trusted sources. For packets coming from an address that does not exist in the trusted source list, PAT applies a rate limit (6 packets every 10 seconds) based on the source address.

The VPN Router Stateful Firewall public address table information does not relate to Network Address Translation (NAT) or network address port translation (NAPT), which is often referred to as port address translation.

This chapter includes the following topics:

- [“VPN Router Stateful Firewall concepts” on page 22](#)
- [“Filters for access control” on page 26](#)
- [“Network Address Translation” on page 27](#)

VPN Router Stateful Firewall concepts

The VPN Router Stateful Firewall provides a secure access point between an internal network and an external network, such as the Internet. The firewall performs the following actions:

- protects your network and the information on your network from unauthorized intrusion from external networks
- provides a line of defense to allow acceptable traffic, as defined by your organization, and to drop all unacceptable traffic before it enters or leaves the network
- monitors packets and sessions and, based on established rules, determines the appropriate actions to take

In addition, you can configure the firewall to log some or all significant events. This includes all connections over the network, such as all e-mail transactions, firewall status changes, and system failures. You can use the logged information to help enhance network security or track unauthorized use.

Stateful inspection

Some protocols are difficult to securely allow through a firewall using traditional filtering mechanisms. The File Transfer Protocol (FTP), for example, typically uses a known port to create the control connection, but a data connection uses a random port. You need stateful inspection to allow an FTP data connection through a firewall without leaving a large number of open ports. The firewall inspects packets at the application layer to determine the port used by the data connection. Traffic on that port then passes through the firewall for the duration of the FTP session.

Transport-level state inspection provides a number of ways to make Transmission Control Protocol (TCP) traffic more secure and more difficult for hackers to intercept. Stateful inspection of TCP verifies the consistency of the TCP header and prevents some well-known TCP attacks. TCP sequence numbers are randomized to prevent sequence number guessing.

Stateful inspection of each application is unique. Stateful inspection validates and permits nonpredicted ports that an application uses through the firewall. The firewall inspects the following applications:

- FTP
- Trivial FTP (TFTP)
- Remote Command (RCMD)
- Structured Query Language Network (SQLNET)
- VDOLive
- RealAudio

All unique end-to-end communication creates a conversation. For instance, an FTP session between a client and a server can consist of several streams of traffic, with both data and control packets flowing back and forth. All of this traffic is part of the same conversation.

Interfaces

The VPN Router can use many interfaces. Each tunnel (end user or branch office) is a virtual interface, and all VPN Routers use two or more physical interfaces. The interface on which packets arrive at the VPN Router (the source interface) or the interface on which packets leave the VPN Router (the destination interface) classify the packets.

You construct the rules in a policy to either use or ignore this classification. If the rule designates Any as an interface, the rule ignores this classification. If the rule designates an interface or group of interfaces, the rule uses this classification.

Use the following terms to designate an interface for the rules in a policy:

- Any—any physical interface or tunnel
- Trusted—a private physical interface or tunnel
- Untrusted—a public physical interface
- Tunnel:Any—any tunnel
- For tunnels, specify either a group name for user tunnels or the specific branch office tunnel for branch office tunnels:
 - Tunnel:/base—specify the specific branch office tunnel. For example, /base/mktng/tony refers to branch office tony in group /base/mktng.
 - Tunnel:user—specify a group name for user tunnels. For example, /base/engineering refers to all user tunnels in that group.
- Interface name—the value of the Description field assigned to the physical interface on the System, LAN (or System, WAN) window (If the description is blank, the interface name defaults to the value of the Interface field on the same page.)

You can configure a physical interface as private or public on the System, LAN, Interfaces window. By default, the LAN interface (Slot 0) is private and all other interfaces are public.

Filter rules

Filtering uses a set of rules to determine whether to allow a packet through the firewall. Typical options are to accept or drop the packet—these options provide a degree of security for a network.

The rules determine one of the following actions:

- accept the packet
- drop the packet
- reject the packet by sending a reject message to the source address
- log the packet locally (you can use these actions with the previous three actions)

Antispoofing

Antispoofing prevents a packet from forging its source IP address. Typically, antispoofing examines and validates the source address of each packet.

Antispoofing performs the following checks:

- source address is not equal to the destination address
- source address is not equal to 0
- source address from an external network is not one of the directly connected networks

Attack detection rules

The firewall can detect common attacks launched against corporate networks. It also drops packets that result from the attack, which prevents denial-of-service as well as nonauthorized intruders. The VPN Router Stateful Firewall provides a defense against denial of service attacks with well-known prevention methods.

The VPN Router Stateful Firewall protects against the following types of objects:

- Jolt2 is a fragmentation attack that affects Windows PCs by sending the same fragment repetitively.
- Linux Blind Spoof attempts to establish a spoofed connection instead of sending the final ACK with the correct sequence number and with no flag set. Linux does not try to verify if the ACK is not set. The firewall drops a packet if the ACK is not set.
- A SYN flood can disable your network services by flooding them with connection requests. This action fills the SYN queue, which maintains a list of unestablished incoming connections, forcing it to not accept additional connections.

- A User Datagram Protocol (UDP) Bomb sends malformed UDP packets that can crash a remote system.
- Teardrop/Teardrop-2 is a fragmentation attack that sends out invalid fragmented IP packets that trigger a bug in the IP fragment reassembly code of some operating systems.
- Land attack sends a TCP packet to a running service on the target host with a source address of the same host. The TCP packet is a SYN packet that establishes a new connection and sends from the same TCP source port as the destination port. After the target host accepts the packet, the packet causes a loop within the operating system, essentially locking the system.
- Ping of death sends a fragmented packet larger than 65536 bytes, which causes the remote system to incorrectly process this packet. This causes the remote system to either restart or freeze during processing.
- Smurf sends a large number of Internet Control Message Protocol (ICMP) echo (ping) messages to an IP broadcast address with the forged source address of the intended victim. The routing device that forwards traffic to those broadcast addresses performs the IP broadcast to Layer 2 broadcast. This broadcast causes most network hosts to take the ICMP echo request and issue a reply to each, which multiplies the traffic by the number of hosts that respond.
- Fraggle sends a large number of UDP echo messages. On a multiaccess broadcast network, potentially hundreds of machines can reply to each packet.
- ICMP unreachable sends ICMP unreachable packets from a spoofed address to a host, which causes the host to stop all legitimate TCP connections to the spoofed host in the ICMP packet.
- Data flood sends a large amount of data to a system as a denial of service attack, which exhausts available resources and stops responses to other user requests.
- FTP command overflow crashes FTP servers that contain buffer overflows for commands that take arguments. This applies to the user command, which means an attacker does not need a valid account to crash the system.

Filters for access control

As you progressively put in place the components of your VPN Router configuration, access control becomes an important security mechanism. You need complete control over which users can access particular servers and services.

You use filters to fine-tune access to specific hosts and services. All users use custom filter profiles based on their group profiles that describes the resources they can access on the network. The filters are defined by

- protocol ID
- direction
- source and destination IP addresses
- source and destination port
- TCP connection establishment

You create a list of rules for a filter profile to perform precisely the action that you want. The filter tests the rules in order until it finds the first match. Therefore, the order of the rules is very important. The filter mechanism works such that if no rule matches a packet, the router discards the packet (denied); therefore no traffic transmits or receives unless specifically permitted.

Network Address Translation

NAT provides transparent routing between address spaces. If you use NAT in an extranet, multiple private networks can connect dynamically through secure tunnels without requiring address space reconfiguration.

The following two factors increase the use of NAT:

- Shortage of IP addresses—Most Internet service providers (ISP) allocate only one address to a single customer. This address is dynamic, so a client receives a different address each time they connect to the ISP. Because users receive a single IP address, they can use only one computer connected to the Internet at a time. After NAT runs on this single computer, multiple local computers can share that single address to connect them all at the same time. The outside world is unaware of this division and performs all communications as though only a single machine on the local network is accessible.
- Security — NAT automatically provides security without special configuration because it permits only connections that originate on the private network. You can still make some internal servers available to the outside world by statically mapping internal addresses to externally available ones, thus making services such as FTP available in a controlled way.

In the context of virtual private networks, you need NAT to allow multiple intranets with conflicting subnets to communicate. Because you can fix the configuration of branch office or partner networks, a VPN solution must securely route between these networks without requiring unique private addresses across the entire extranet.

Chapter 2

Stateful Firewall configuration

To use the firewall on the VPN Router, you must install a license key and enable the firewall service. Without the firewall enabled, the VPN Router forwards the following traffic patterns:

- private physical interface to private physical interface
- private physical interface to user or branch office tunnel
- tunnel to tunnel (user or branch office)

After you enable the firewall, the VPN Router additionally routes traffic from public to private interfaces.



Note: Shut off all traffic to the VPN Router before you activate the firewall on the Firewall/NAT window. Do this during off hours to prevent inconvenience to the users.

You must create rules for tunnel traffic before the router permits traffic on existing tunnels. The VPN Router Stateful Firewall uses the principle that traffic not specifically allowed is disallowed. The rule set of the active policy applies to all traffic, including tunneled and nontunneled traffic. Therefore, after you first enable the VPN Router Stateful Firewall, the router disallows all traffic until you configure rules that specifically allow certain types of traffic.

This chapter includes the following topics:

- [“Configuration prerequisites” on page 30](#)
- [“Java 2 software installation” on page 31](#)
- [“Enabling firewall options” on page 33](#)
- [“Rule enforcement” on page 36](#)
- [“Log options” on page 36](#)

- [“Configuring antispoofing” on page 38](#)
- [“Configuring malicious scan detection” on page 39](#)
- [“Policy configuration” on page 39](#)
- [“Verifying the configuration” on page 55](#)
- [“Configuring a sample security policy” on page 55](#)
- [“Firewall deployment examples” on page 57](#)

Configuration prerequisites

Before you configure your VPN Router Stateful Firewall, collect the following information:

- The management IP address of your VPN Router. This address is on the System, Identity window.
- The firewall license key. Choose Admin, License Keys, type the key in the Key / Status box to the right of Stateful Firewall and then click OK. You only need to install a key once on each VPN Router. Select and delete the content in the Key / Status box to remove the key.
- The name of the firewall is the name the Domain Name Service (DNS) server uses to identify the management address of the VPN Router. Type this name in the DNS Host Name box of the System Identity window.
- The names and IP addresses of your VPN Router interfaces. This information is on the Statistics > Interfaces window.

You must meet the following system requirements to gain access to the VPN Router Stateful Firewall Manager:

- Supported operating systems and platforms include Solaris (OS 2.8 and 2.9) on an x86 or SPARC platform and Microsoft Windows 2000, or Windows XP.
- Required software includes Java 2 Plug-in Version 1.6.0_u6, available in the Java 2 Runtime Environment (J2RE) Version 1.6.0_u6. The J2RE is available for automatic download on a Windows platform for all VPN Routers except the 1010, 1050 and 1100 (for more information, see the Java 2 Runtime Environment Installation). J2RE installation files for Windows and Solaris are also available on the Nortel CD in the tools/java directory.

- Supported browsers include Internet Explorer 6.0 and 7.0 and Firefox 2.0 and 3.0. The VPN Router does not support the version of the Java 2 Plug-in that comes with Netscape 6.


Java 2 software installation

To access the VPN Router Stateful Firewall Manager, you must install Java 2 Runtime Environment on the computer that administers the VPN Router. Choose from two separate procedures to install the Java 2 software that depend on whether you use Internet Explorer or Firefox to access the VPN Router.

Using Internet Explorer

To install the Java 2 software on Windows 9x, Windows 2000, or Windows NT from Internet Explorer

- 1 Log on to the management IP address of the VPN Router.
- 2 Choose **Services, Firewall/NAT**.
The Firewall/NAT window appears.
- 3 In the **VPN Router Stateful Firewall** row, click **Manage Policies**.
A window appears and tries to load the VPN Router Stateful Firewall Manager.
- 4 If the Security Warning window appears, click **Yes** to install the Java 2 Runtime Environment [“Security Warning window” on page 32](#).

Figure 1 Security Warning window

The installation program downloads the software from the VPN Router. (This option is not available for the 1010, 1050, and 1100 hardware platforms.) The program can take several minutes to load, depending on the speed of your connection to the VPN Router.

- 5 After the installation program displays the Software Licensing Agreement, click **Yes** to accept the agreement.
- 6 After the installation program asks for an installation location, accept the default location or choose another installation location.
- 7 Click **Next** to finish the installation.
- 8 After the installation is complete, close all open Web browsers.
- 9 Restart the computer for the changes to take effect.

Using Firefox

Nortel supports Firefox 2.0 and 3.0. To install the Java 2 software from Firefox

- 1 Go to addons.mozilla.org.
- 2 On the left navigation bar, click **Plugins**.
- 3 Under the Java category, click **Download Now**.
- 4 Complete the instructions on the Web page.
- 5 Log on to the management IP address of the VPN Router.

6 Choose Services, Firewall/NAT.

The Firewall/NAT window appears.

7 In the VPN Router Stateful Firewall row, click Manage Policies.

A window appears and loads the VPN Router Stateful Firewall Manager.

Enabling firewall options

You can select only one firewall choice at a time. The choices are

- **VPN Router Firewall**—enables the VPN Router Stateful Firewall feature. After you enable the firewall, you can run a combination of the following:
 - VPN Router Stateful Firewall
 - VPN Router Interface Filter
 - Interface NAT
 - Antispoofing
 - Malicious Scan Detection
- **No Firewall**—disables all firewall features on the VPN Router. In this configuration, the VPN Router performs VPN routing only.

To enable the VPN Router firewall

1 Choose Services, Firewall/NAT.

The Firewall/NAT window appears.

2 Select VPN Router Firewall. After you enable the VPN Router Firewall, you can run any combination of the following:

- VPN Router Stateful Firewall
- VPN Router Interface Filter
- Interface NAT
- Anti-spoofing
- Malicious Scan Detection

3 Click OK.

- 4 Confirm your selection.
- 5 At the prompt, restart the VPN Router.

You must restart the VPN Router before the firewall becomes active. After you enable firewall support, you must configure the specified firewall.

To enable no firewall

- 1 Choose **Services, Firewall/NAT**.

The Firewall/NAT window appears.

- 2 Select **No Firewall**. This disables all firewall features on the VPN Router. In this configuration, the VPN Router performs VPN routing only.
- 3 Click **OK**.

The configuration procedures assume that you configure the VPN Router (except for the firewall component) and that you obtain the required firewall license. You do not need a license for the VPN Router Interface Filter.

To enable the VPN Router Stateful Firewall

- 1 Choose **System, LAN**.

The LAN Interfaces window appears.

- 2 For each interface, click **Configure**.
- 3 Type a label in the **Description** box. This name identifies interfaces in the security policy rules. You assign an IP address to the LAN, which represents the physical port interface. Slot n Interface n represents an optional LAN card in expansion Slot n using Interface n.

For example, you can make Internet the description for Slot 1 Interface 1 and ServiceNet the description for Slot 2 Interface 1. The description is case sensitive and you cannot abbreviate it when you specify the interface in the rules. If you do not specify a description, the default name for the interface is Slot n Interface 1 (n=1 to 6) and is case sensitive. You cannot abbreviate the name. The available slot numbers are hardware platform specific.

- 4 Choose **Services, Firewall/NAT**.
- 5 Enable **VPN Router Stateful Firewall**.

- 6 Click **Schedule System Reboot** to restart the system now.
- 7 On the system shutdown window, click **OK** and on the confirmation page, click **OK** to indicate the restart.
- 8 After the VPN Router restarts, return to **Services, Firewall/NAT**.
- 9 Click **Manage Policies** to load the VPN Router Stateful Firewall Manager applet. The first time you do this on a workstation, you must load the Java applet. The message `Retrieving policies` appears.
- 10 Select the **System Default** policy, which is read-only.
- 11 Click **View** to review this policy. Every new policy includes the implied rules.
- 12 You can toggle the browser windows between the VPN Router Stateful Firewall Manager applet and the **Firewall/NAT** window. If you use your browser to change other settings on the VPN Router while you run the VPN Router Stateful Firewall Manager applet, the current VPN Router Stateful Firewall Manager applet does not reflect these changes. Click the **Firewall** icon in the VPN Router Stateful Firewall Manager applet to refresh the list of policies and other VPN Router settings. Changes made in the VPN Router Stateful Firewall Manager applet are not evident in the Firewall/NAT window until you save the policy.
- 13 Choose **Manager, Exit SFw/Nat** to exit the VPN Router Stateful Firewall Manager.
- 14 Click **Yes**.
- 15 After you exit the VPN Router Stateful Firewall Manager applet, click **Refresh** on the **Firewall/NAT** window.

The new policies you create do not automatically apply to the firewall. You can apply only one policy at a time to the firewall.



Note: You cannot import or export new policies. However, no restrictions exist to create new policies.

Rule enforcement

ICMP is allowed or disallowed on public and private interfaces. To enable ICMP, you must establish a complete three-way handshake prior to the application of data.

Log options

The following options control the amount of firewall event information recorded in the event log. The router does not save this information in the system log.

- All—includes traffic, policy manager, firewall, and NAT
- Traffic—logs the creation or removal of flows and conversations
- Policy manager—logs firewall processes and the creation of rules and policies
- Firewall—logs how the firewall handles packets within a flow
- NAT—logs NAT-related events
- Debug—creates special log messages intended for use only by Nortel customer-support personnel

You edit these options on the Firewall/NAT > Edit window.

You can also configure a maximum connection number, which reserves memory for a maximum number of connections. Determine the optimum memory allocation to make it easier to configure your system for firewall traffic. In the Maximum Connection Number box, type a number in the indicated range. The range shown varies depending on the model and amount of memory for your VPN Router. Each IPsec tunnel requires two connections. Nortel recommends that you configure the number near the middle of the range displayed unless you must consider specific requirements. You must restart the VPN Router if you change the maximum connection number.



After you disable the syslog server parameter, the VPN Router sends a message to the syslog indicating that the server is disabled.

Application-specific logging

Firewall-specific logging includes application-specific logging, denial of service attack logging, and the ability to send firewall-specific events to a remote syslog server. The application-specific logs for Hypertext Transfer Protocol (HTTP) and File Transfer Protocol (FTP) contain a unique connection identifier so that events are traced to the start and end of a TCP session. You can configure the firewall rules to enable logging in either brief or detail format for rules with FTP and HTTP service.

Configuring remote system logging

The VPN Router can forward firewall-specific events to a remote syslog server. You can select whether to send all events or only firewall-specific events to the remote syslog server.

To configure remote syslog

- 1** Choose **Services, Firewall/NAT, VPN Router Firewall, Edit**.

The Firewall/NAT > Edit window appears.

- 2** Enable **Logging** beside each feature you want to configure for the VPN Router Stateful Firewall. The options are

- All
- Traffic
- Policy Manager
- Firewall
- NAT
- Debug

- 3** Select which type of log you require from the **Implied Rule Log Level** list:

- None
- Brief
- Detail
- Trap

- 4** Click **OK**.

- 5 Choose **Services, Syslog**.
The SysLog Forwarding window appears.
- 6 Type a Hostname or IP address.
- 7 In the **Filter Level** list, select **All**.
- 8 In the **Entity** list, select **Security**.
- 9 In the **Subentity** list, select **Firewall**.
- 10 In the **Tagged Facility** list, select **KERN**.
- 11 Type **514**, the default, for the UDP port.
- 12 Click **Enabled** for the server.
- 13 Click **OK**.
- 14 Start syslog on the remote syslog system.
- 15 To verify that firewall-specific events appear on the remote syslog system, send traffic through the VPN Router that generates firewall events.

Configuring antispoofing

To configure antispoofing

- 1 Choose **Services, Firewall/NAT**.
The Firewall/NAT window appears.
- 2 Select **Anti-spoofing**.
- 3 Click **Edit**.
The Firewall/NAT > Anti-Spoofing window appears.
- 4 Select the public interface on which you want to enable antispoofing.
- 5 Click **OK**.

Configuring malicious scan detection

Scan detection detects port scanning attempts through the VPN Router that are aimed at private resources.

To configure scan detection

- 1 Choose **Services, Firewall/NAT**.

The Firewall/NAT window appears.

- 2 Select **Malicious Scan Detection**.

- 3 Click **Edit**.

The Firewall/NAT > Scan Detection window appears.

- 4 In the **Detection Interval** box, type the interval (1 through 60) over which the number of port scans or host scans are inspected. If the number of scans exceeds the configured threshold during this interval, the security log logs the scan.

- 5 In the **Port Scan Threshold** box, type the number of host-to-host connections (between 1 and 10 000) on the private side to which an attacking machine must send scan packets during the inspection interval to trigger an event in the security log.

- 6 In the **Network Scan Threshold** box, type the number of one-to-many connections (between 1 and 10 000) needed to trigger an event. This value is the number of ports on one host on the private side to which an attacking machine must send scan packets during the inspection interval to trigger an event in the security log.

- 7 Click **OK**.

Policy configuration

Firewall service consists of two primary components:

- service properties
- security policy

Service properties define the offered service and includes a service name, the protocol (TCP, UDP, ICMP), and the port number (or range) on which the service occurs.

Security policies consist of a set of rules that specify what service is allowed or denied. You use service objects to specify all rule fields for service policies. Each rule consists of a combination of network objects, services, actions, and logging mechanisms. You can define custom policies if you need more complex security policies and the standard policies are not sufficient. Customize your policies to further refine the control over what traffic you allow on your internal networks.

The firewall policies use standard actions, which represent the most commonly used policies. A set of rules defines a specific security policy. A rule defines whether the router accepts or rejects (or logs) communication based on the source, destination, and service.

You must create rules for tunnel traffic before the router allows traffic on existing tunnels. The VPN Router Stateful Firewall uses the principle that traffic that is not specifically allowed is disallowed. The rule set of the active policy applies to all traffic, including tunneled and nontunneled traffic. Therefore, after you first enable the VPN Router Stateful Firewall, the firewall disallows all traffic until you configure rules that specifically allow certain types of traffic.

Firewall policy creation and modification

You implement access control parameters through the graphical user interface (GUI) or the command line interface (CLI). You can use either interface to configure the following items:

- network objects
- service objects
- rules

For more information about CLI commands, see *Nortel VPN Router Using the Command Line Interface* (NN46110-507).

Policy creation

Use the Services, Firewall/NAT, VPN Router Stateful Firewall, Manage Policies menu path to create, edit, delete, copy, or rename a firewall policy. The current policy appears in bold and read-only policies appear in italic. The System Default policy always appears. This read-only policy defines the firewall behavior if you do not apply user-defined policies, or if the selected policy is not available.

Adding a policy

To add a new policy

- 1** Choose **Services, Firewall/NAT**.

The Firewall/NAT window appears.

- 2** Click **Manage Policies** beside **VPN Router Stateful Firewall**.

- 3** Click **New**.

The New Policy window appears and prompts you for a name for the new policy.

- 4** Type the policy name. The name must begin with a letter and cannot contain the : + =] , ; " characters.

- 5** Click **OK** to go to the **Policy Edit** window, which provides a blank firewall policy.

- 6 Choose **Manager, Exit SFw/Nat**.
- 7 Click **Yes** to exit the applet.
- 8 Click **Yes** to save the new policy.

Deleting a policy

You cannot delete a read-only policy, or the policy that currently applies to the VPN Router. If you select one of these policies, the Delete option is dimmed. To delete a policy

- 1 Choose **Services, Firewall/NAT**.
The Firewall/NAT window appears.
- 2 In the **VPN Router Stateful Firewall** row, click **Manage Policies**.
- 3 Select the policy that you want to delete, and then click **Delete**.
- 4 Click **OK** to delete the selected policy.

Copying a policy

To copy a firewall policy

- 1 Choose **Services, Firewall/NAT**.
The Firewall/NAT window appears.
- 2 Click **Manage Policies** beside **VPN Router Stateful Firewall**.
- 3 Select the policy that you want to copy.
- 4 Click **Copy**.
- 5 Type a name for the copied policy.
- 6 Click **OK**.

The new policy appears in the list of policies in the firewall policies window. This policy contains the same rules as the original policy.

Renaming a policy

You cannot rename a read-only policy, or the policy that applies to the VPN Router. If you select a read-only policy, the Rename option is dimmed. To rename a firewall policy

- 1 Choose **Services, Firewall/NAT**.
The Firewall/NAT window appears.
- 2 Click **Manage Policies** beside **VPN Router Stateful Firewall**.
- 3 Select the policy that you want to rename.
- 4 Click **Rename**.
- 5 Type the new name of the policy.
- 6 Click **OK**.

Navigating rules

You use the Firewall Policy, Edit window to add, delete, and modify the rules within a policy. This window divides into the following rule tabs:

- Implied Rules
- Override Rules
- Interface Specific Rules
- Default Rules



Note: When you create a firewall rule, under Interface Specific Rules, Slot 7 Interface 1 appears, which is the serial port. The serial port does not appear on versions of the VPN Router prior to Version 4.80.

Implied rules

The firewall processes implied rules first. These rules permit tunnel termination and access to the management interface. They are derived from the configuration you apply by using the Services, Available menu path and other configuration windows, for example, Routing Information Protocol (RIP), Open Shortest Path First (OSPF), and Virtual Router Redundancy Protocol (VRRP). The system

statically generates and defines some rules, which are read-only. **“Implied rules”** on page 44 You cannot modify these rules—they are for display purposes only. Implied rules regulate traffic that originated from or terminated at the VPN Router. You can control routed traffic that is not directed to the VPN Router with Override rules, Interface-specific, or Default rules.

Figure 2 Implied rules

#	Src Interface	Dst Interface	Source	Destination	Service	Action	Log	Status	Remark
1	System	Trusted	_MngtIP	any	any	accept		✓	
2	Tunnel:Any	System	any	_MngtIP	any	accept		✓	
3	Private	System:Private	any	any	I2xservicegroup pptpservicegroup icmpservicegroup ipsecservicegroup	accept		✓	Allow/Deny services to come in from any private interfaces and terminated at any private interfaces
4	Private	System:Private	any	any	tcp8000 tcp1000	drop		✓	Allow/Deny services to come in from any private interfaces and terminated at any private interfaces
5	System	Trusted	any	any	dhcp tcp1000 dhcp-client I2xservicegroup pptpservicegroup	accept		✓	Allow/Deny services originated from the system to leave the system through any trusted interfaces

Static preimplied rules

The first rule in the implied rules section is the only statically generated rule. It always exists in the implied rules section regardless of the configuration. This rule allows the listed services to leave the VPN Router on the private interfaces as long as the services originated from the VPN Router. [“Servers and menu paths” on page 45](#) shows the server type and its corresponding configuration windows.

Table 1 Servers and menu paths

Servers	Menu path	Description
DHCP, DHCP-CLIENT	Servers, DHCP Relay	
DNS	System, Identity	
Remote-RPC		UDP port 17185
Nbdatagram, nbssession		Remote Netbios
Pptp	Services, Available	
IPSEC	Services, Available	
L2TP & L2F	Services, Available	
FWUA	Services, Available	
Radius	Services, Available	
HTTP, HTTPS	Services, Available	
SNMP	Services, Available	
FTP	Services, Available	
TELNET	Services, Available	
CRL	Services, Available	
CMP	Services, Available	
LDAP	Servers, LDAP	
UDP Wrapper	Services, IPsec (Ipsec Settings)	Enable/Disable NAT Traversal UDP, configured port
NTP	System, DATE&TIME, Network Time Protocol	
VRRP	Routing, VRRP	
RIP	Routing, RIP	
OSPF	Routing, OSPF	

Table 1 Servers and menu paths (continued)


Servers	Menu path	Description
SSH Server	Services, SSH Server	
BGP	Services, BGP	Must install the PR or BGP key.

Dynamic implied rules

All of the available services on the Services window generate dynamic implied rules. Implied rules for ports that are not well known use a service name that consists of the protocol and the port number. For example, a tcp10 rule generates from port numbers associated with external Lightweight Directory Access Protocol (LDAP) and Remote Authentication Dial-In User Service (RADIUS) servers and configurable Firewall User Authentication (FWUA) ports.

Override rules

Override rules are the first set of modifiable rules in the policy. [“Override rules” on page 46](#) The purpose of these rules is to quickly override the rest of the rules described later in the policy, possibly for a short period, while you debug a problem. These rules do not specify a specific interface in the source or destination interface column. You can only select from the interface groupings (Any, Trusted, Untrusted, Tunnel:Any, User Tunnel:Any, Branch Tunnel:Any, SSL-VPN).

Figure 3 Override rules

Interface-specific rules

Interface-specific rules apply only to packets that enter or leave the VPN Router through one specific interface (physical or tunnel). Interface-specific rules use two rule types: source and destination. “[Interface-specific rules \(Source rules\)](#)” on page 47 and “[Interface-specific rules \(Destination rules\)](#)” on page 48 Source rules define the selected interface as the source. Destination rules define the selected interface as the destination. Physical interface names correspond to the names configured on either the LAN Interfaces or WAN Interfaces window. Tunnels that are also interfaces correspond either to a group name for user tunnels or the specific branch office tunnel name. The interface-specific rule section shows only one interface at a time. To view all of the interface-specific rules, select All Interfaces.

Figure 4 Interface-specific rules (Source rules)






Figure 5 Interface-specific rules (Destination rules)

Default rules

Default rules “[Default rules](#)” on page 48 apply to all traffic, but are not restricted to a specific interface. These rules specify interface groupings for the source or destination (Any, Trusted, Untrusted, Tunnel:Any, User Tunnel:Any, Branch Tunnel:Any).

Figure 6 Default rules

Rule creation

Menus control actions on rules. Right-click an option to access menus. Each menu controls a different aspect of the rule.

Header row menu

Right-click on a header cell to show the header row menu. This menu contains one item, Add New Rule. You use this menu item to add a new rule to the top of the list. The new rule appears in position one and all existing rules increment by one.

Row menu

Right-click on the number next to a rule to activate the row menu. You use this menu to add a new rule at a particular location, delete the specific rule, and perform cut, copy, or paste operations on a rule.

Cell menus

Right-click on an individual cell to access cell-specific menus. Two types of cell menus exist: option menus and procedure menus. Option menus provide a list of possible values for the cell. The cell displays the selection after you click on one of the items.

Procedure menus provide a list of operations that you can perform on the cell, such as Add and Edit. After you click one of the items, either the operation is performed immediately (such as Copy) or an additional window appears, which prompts you for more information (such as Add).

Rule columns

Each rule within a firewall policy uses the same attributes, which are specified by the column headers. The following sections describe the columns within a firewall rule:

#

This column specifies the order of the rules within the section. The order applies only to the section in which the rule appears and does not apply across the entire policy. If you log a rule, the log information includes this number (#).

Src interface and Dst interface

These columns specify the source and destination interfaces for the rule. Right-click on the cell to display an option menu that contains possible interfaces. What appears in this option menu depends on which section of the Firewall policy the particular column appears in. For the Override and Default rules, the interfaces are only interface groupings.

These groupings are

- Any—any physical interface or tunnel
- System—an NVR interface that is either the source or destination
- Trusted—a private physical interface or tunnel
- Untrusted—a public physical interface
- Tunnel:Any—any tunnel, excluding any physical interfaces
- User Tunnel:Any—any user tunnel
- Branch Tunnel:Any—any branch tunnel
- SSL-VPN—an SSL-VPN tunnel

For interface-specific rules, you can specify the interfaces as either groupings or individual interfaces.

Click on the user tunnel or branch office menu items to display the tunnel selection window. You use this window to select a specific tunnel (branch office or user tunnel).

Source and Destination

These columns specify the source and destination network object for the rule. Right-click on a column in the cell to modify these attributes, which then shows a procedure menu. You can add more than one source or destination address to a rule.

Click Add to display the Network Object Selection window. [“Network Object Selection window” on page 51](#) Use this window to define and apply a new network object. You can create the following network objects: host, network, IP range, and group (a collection of these objects).




Note: You use NOT operand to specify which networks you do not want included.

Figure 7 Network Object Selection window



Italicized objects in the list are read-only—you cannot modify them. You use the New, Edit, and Delete options in this window to create, edit, and delete network objects.

Click Edit to display the Network Object Edit window. [“Network object edit window” on page 52](#) You use this window to modify the attributes for the selected network object.

Figure 8 Network object edit window


Click Delete to remove the selected network object. If the object that you want to delete is the last object, the cell returns to the default value.

Click Copy, Cut, or Paste to perform those operations on the current network object.

Service

This column specifies the service objects handled by the selected rule. Right-click on the cell to display the standard procedure menu (Add or Edit).

Click Add to access the Service Object Selection window [“Service Object Selection window” on page 53](#), where you define and apply a new service object. You can create the following service objects: TCP, UDP, ICMP, IP protocol, and object groups (a collection of these objects). You can add more than one service to a rule.

Figure 9 Service Object Selection window

Italicized objects in the list are read-only—you cannot modify them. You use the New, Edit, and Delete options in this window to create, edit, and delete service objects.

Click Edit to display the Service Object Edit window. You use this window to modify the attributes for the selected service object.

Click Delete to remove the selected service object from the cell. If the object you want to delete is the last object in the cell, the cell returns to its default value.

Click Copy, Cut, or Paste to perform those operations on the current service object.

Action

The Action column specifies the action that occurs after you activate a rule. Right-click on the cell to display an option list that contains four items: Accept, Drop, Reject, and User Authentication. Click one of these items to configure the cell to the selected state.

Log

Use the Log column to specify the logging level for this rule. Right-click on this cell to show an option list that contains the following logging levels: None, Brief, Detail, and Trap.

Status

The Status column specifies the status of the particular rule, either Enabled or Disabled.

Remark

Use the Remark column to attach a remark to a particular rule. Right-click Remark and select Add or Edit remark, and then type a comment in the dialog box that appears.

Creating a new policy

To configure the firewall policies

- 1 Choose **Services, Firewall/NAT**.

The Firewall/NAT window appears.

- 2 In the **VPN Router Stateful Firewall** row, click **Manage Policies**.

- 3 Click **New** to create a new policy.

- 4 Type the policy name, and then click **OK**. The name must begin with a letter and cannot contain the : + =] , ; " characters.

The Firewall, Edit Policy: <polycyname> window appears with no rules defined. In this window, you can add, delete, and modify the rules for the policy.

- 5 You can select the rule group as follows:

- Implied rules (view only)
- Override rules
- Interface-specific rules
- Default rules

- 6 Click the **Interface Specific Rules** tab.

- 7 Select an interface and a subinterface from the lists.

- 8 Select either **Source Interface Rules** or **Destination Interface Rules**.

- 9 Right-click the appropriate cell to add a new rule.

- 10 Repeat these steps to add more rules.
- 11 Select **Policy**, and then click **Save Policy** to save your changes.
- 12 Click **OK**.
- 13 After the policies are saved, choose **Manage, Exit SFw/Nat**.
- 14 Click **Yes** to exit.
- 15 Click **Yes** to save the policy.

Successful completion of these steps indicates that the VPN Router firewall is functioning and that the VPN Router routing patterns are available.

Verifying the configuration

After you complete the configuration tasks for the firewall, you can check the routing patterns. To verify that the firewall functions properly, you can use a procedure similar to the following:

- 1 Make sure the firewall uses a security policy that allows the type of traffic you use for the test (or you can use an Accept All policy for the test).
- 2 Verify public-to-private traffic. Perform an FTP operation from a host on the public side of the VPN Router to a host on the private side.
- 3 Verify private-to-public traffic. Perform an FTP operation from a host on the private side of the VPN Router to a host on the public side.
- 4 Verify tunnel-to-internal network traffic. Connect a remote VPN Router system to the local VPN Router. From the client, access a Web page on the internal network.
- 5 Verify tunnel-to-Internet traffic. Connect a remote Nortel VPN Client system to the VPN Router. From the client, access a Web page on the Internet.

Configuring a sample security policy

In this configuration example, the following configuration exists:

- The public IP address is 192.168.3.22 (Internet Access).

- The private IP address is 10.3.3.102 (VPN Router default is LAN).
- The FTP server IP address is 192.168. 3.20 on the public network.
- The security policy allows users to download files to the FTP server, with no other access to the Internet.

To configure the VPN Router Stateful Firewall to implement a security policy

1 Choose **Services, Firewall/NAT**.

The Firewall/NAT window appears.

2 Click **Manage Policies** for VPN Router Stateful Firewall.

3 Click **New**.

4 Type **AllowFTPAccess** as the policy name, and then click **OK**.

5 Click the **Interface Specific Rules** tab.

Make no changes to the interface or subinterface lists and leave Source Interface Rules selected.

6 On the **Interface Specific Rules** tab, right-click # in the header, and then select **Add New Rule**.

7 Click the **DST Interface** value (*any), right-click to display the selection menu, and then select **SSL-VPN**.

8 Click the **Destination** value (*any), right-click to display the selection menu, and then select **Add**.

a In the **Network Object Selection** window, click **New**.

b In the **Network Object Type Selection** window, select **Host** as the type of object to create, and then click **OK**.

c In the **Network Object Insert** window, type the Host name (**externalFTPserver**) and the IP address (**192.168.3.20**), and then click **OK**.

d In the **Network Object Selection** window, click **OK** to add the externalFTPserver network object into the **Destination** field.

9 Click the **Service** value (*any), right-click, and then click **Add** to display the **Service Object Selection** box, click **FTP**, and then click **OK**.

- 10 Click the **Action** value (drop), right-click to display the **Action** menu, and then click **Accept** to enter it into the **Action** field.
- 11 Click the **Log** value (blank = none), right-click to display the **Log** menu, and then click the required log value to enter it into the **Log** field. In this example, the log value is brief.
- 12 Click the **Status** value (checked means enabled), right-click to display the **Status** menu, and then click the required status value to enter it into the Status field. (Within a policy, you can independently disable each rule in the Override, Interface-Specific, and Default groups.)
- 13 Choose the **Manager** menu, and then click **Exit SFw/Nat**. Click **Yes** to exit.
- 14 In the **Save Changes to this policy** dialog box, click **Yes**.
- 15 Click **Refresh**.
- 16 On the **Services, Firewall/NAT** window, select **AllowFTPAccess** from the policy box, and then click **OK**. (You can apply only a single policy to the VPN Router.)
- 17 Restart the VPN Router to activate the new firewall configuration.

Firewall deployment examples

You can customize security policies and apply them to individual subscribers, or you can create them as templates and apply them to many subscribers.

Some questions to consider before you establish firewall rules include:

- What are the IP addresses for all of your servers (FTP, DNS, Web, e-mail) accessible through this firewall?
- If you configure NAT, what IP addresses can you list that are otherwise not visible?
- What applications, other than HTTP, FTP, mail protocols, and other typical network traffic, run across your firewall?

Residential firewall example

A residential firewall, see [“Example of a basic residential firewall” on page 58](#), is generally a simple firewall that allows user-initiated traffic while it blocks incoming traffic or port scans.

Figure 10 Example of a basic residential firewall




Use the Override Rules tab on the Firewall, Edit Policy window to configure your residential firewall with a single override rule that allows all trusted traffic. Trusted traffic is traffic that comes from either a trusted physical interface or a tunnel.

Alternatively, you can use the Interface Specific Rules tab on the Firewall, Edit Policy window to configure a single interface specific rule that allows traffic sourced from the physical interface LAN (slot 1/0).

Business firewall example

A business firewall, see [“Business firewall” on page 59](#), requires a more complex rule configuration. A business user requires access to internal resources, such as mail servers and Web servers. The choices for service indicate which protocols to accept or reject on the network. Typically, these include HTTP, Simple Mail Transfer Protocol (SMTP), FTP and network protocols, such as some forms of ICMP.

Figure 11 Business firewall

After you configure a business firewall, you must configure override rules to perform the following actions:

- require branch office users to authenticate themselves prior to accessing internal resources
- allow user tunnel traffic to go anywhere
- allow nontunneled FTP and HTTP to gain access to DMZ

You must also configure an interface specific rule to allow all traffic that enters from the private LAN to go anywhere. You configure the override rules in the Override Rules tab on the Firewall, Edit Policy window. You configure the interface specific rule in the Interface Specific tab in the Firewall, Edit Policy window.

Chapter 3

Filter configuration

Two types of filters exist: tunnel filters and interface filters. You use tunnel filters for user groups, and you use interface filters for LAN and WAN interfaces. Interface filters do not apply to packets sent to an internal circuitless IP (CLIP) address. After you change a tunnel filter, it does not affect existing tunnels. You must reestablish the existing tunnels before changes take effect.

A filter consists of one or more inbound rules (for traffic coming into the network) and one or more outbound rules (for traffic leaving the network). Filter names are a convenient way to manage a set of rules.

To view the available filters, choose Profiles, Filters. The Current VPN Router Tunnel Filters and Current VPN Router Interface Filters lists show the currently available filters.

This chapter includes the following topics:

- [“Adding and editing filters” on page 61](#)
- [“Management access restrictions” on page 63](#)
- [“Configuring next-hop traffic filters” on page 65](#)

Adding and editing filters

To add a filter

- 1 Choose **Profiles, Filters**.

The Filters window appears.

- 2 Type a new filter name in the **Create** dialog box.

- 3 Click **Create**.
- 4 To add a rule to the **Rules in Set** list, select a rule from the **Available Rules** list, and then click the left arrow.
- 5 To remove or delete a rule from the **Rules in Set** list, select the rule, and then click the right arrow.
- 6 To move the rule up one place in the **Rules in Set** list, select the rule, and then click the up arrow.
- 7 To move the rule down one place in the **Rules in Set** list, select the rule, and then click the down arrow.

The Available Rules box lists all of the available rules you can add to the filter. They appear in the format of Name: Rule String.

- 8 Click **OK**.

To edit a filter

- 1 From the **Profiles, Filters, Edit** window, click **Manage Rules**.
- 2 Click **Edit**.
- 3 Select the filter action, either **Permit**, **Deny**, or **Nexthop**.
- 4 Select the direction, either **Inbound** or **Outbound**.
- 5 Select an address.
- 6 Select a protocol. The choices are **icmp**, **ip**, **tcp**, or **udp**.
- 7 For the **Source Port**, select options from both lists.
- 8 For the **Destination Port**, select options from both lists.
- 9 For the **TCP Connection**, select either **Established** or **Don't Care**.
- 10 Click **OK**.

Management access restrictions

Use the Allow Management Traffic options from the Profiles, Filters, Edit menu path to restrict management access to the VPN Router through tunnels. Each filter set uses an explicit list of management services. Specify the management services allowed through a tunnel to control which groups of users perform different management tasks while tunneled into the VPN Router.

The default filter is Permit All, and the configuration of this filter allows Hyper Text Transfer Protocol (HTTP), Simple Network Management Protocol (SNMP), and Ping. However, if you create a new filter, all management traffic settings are disabled by default.

The management protocols consist of two groups: Local Services and Remote Servers. The Local Services selections refer to services that reside on the VPN Router. The Remote Servers selections refer to services that reside on other systems that the VPN Router uses. After you enable services, the router allows network traffic for these services through tunnels.

The management services apply to user and branch office connections. These options do not affect HTTP, SNMP, File Transfer Protocol (FTP), Telnet, or Ping protocol traffic that passes through the VPN Router outside a tunnel.

The Local Services options are

- HTTP—enable or disable access to the Web server on the VPN Router
- HTTPS—enable or disable access to the secure Web server on the VPN Router
- SNMP—enable or disable SNMP gets from the VPN Router
- FTP—enable or disable FTP puts or gets to and from the VPN Router
- Telnet—enable or disable Telnet access to the VPN Router
- SSH—enable or disable Secure Socket Shell (SSH) access to the VPN Router
- PING—enable or disable Ping access to the VPN Router
- RADIUS—enable or disable access to the Remote Authentication Dial-In User Service (RADIUS) authentication service
- Identification—Enable or disable access that uses identification services
- DNS—enable or disable access to the Domain Name Server (DNS) service

The Remote Servers options restrict traffic to external services that the VPN Router needs. Specify these services to restrict which VPN Router tunnels can send protocol traffic for the external services it requires.

The Remote Servers options are

- **FTP**—enable or disable FTP access from the VPN Router to external FTP servers on the other end of a tunnel. The FTP back-up and FTP upgrades facilities are examples of external services that this option controls.
- **RADIUS**—enable or disable the ability of the VPN Router to access a remote RADIUS server.
- **DNS**—enable or disable remote users from using the DNS service for the VPN Router.
- **NTP**—enable or disable remote Network Time Protocol (NTP) server access by using the filter.
- **LDAP Proxy**—enable or disable external Lightweight Directory Access Protocol (LDAP) proxy server access by using the filter.
- **CMP**—enable or disable remote Certificate Management Protocol (CMP) server access through this filter to the gateway.
- **TunnelGuard**—enable or disable TunnelGuard traffic on this filter.

Use Copy Filter to copy a filter from one filter set to the other. For example, if you already use a filter for tunnels, you can copy it for use by your VPN Router interfaces.



Note: If you plan to use a filter for both tunnels and interfaces, it must appear in both windows on the Filters window.

To copy a filter

- 1** Choose **Profiles, Filters**.

The Filters window appears.

- 1** Click the filter in one of the **Current Filters** lists.

- 2** Click **Up** or **Down** to copy the filter to the other **Current Filters** list.

The Copy Filters window appears and prompts you to confirm that you want to copy the filter.

If you copy a tunnel filter for use by a VPN Router Stateful Firewall, you need to configure additional steps because the traffic that uses the VPN Router Stateful Firewall traverses two VPN Router interfaces. For example, it can enter through a public interface and exit through a private interface. However, tunnel traffic only enters and exits through a single physical interface.

Configuring next-hop traffic filters

Customers use next-hop traffic filters to control the next-hop selection and route traffic within their domain. If a packet matches filter criteria, the configured next hop performs a forwarding lookup and forwards the packet using that routing table instance. If the lookup fails, traditional destination-based routing occurs using the routing table.

Each IP interface can use inbound and outbound filters that cause an action on a packet if the packet matches the filter criteria. After a filter rule with a configured next hop [“Filter rule with next hop” on page 65](#) matches an incoming packet, the filter accepts the packet and uses the next hop for forwarding.



Note: Interface filters do not apply to packets sent to an internal CLIP address.

Next-hop traffic filters are only applicable for inbound filters for each interface (physical or virtual) for each protocol.

Table 2 Filter rule with next hop

Source address	Destination address	Service	Action	Next-hop address	Comment
10.0.0.0 (255.0.0.0)	47.17.253.0 (255.255.255.0)	IP	Nexthop	192.32.140.216 (255.255.255.0)	Filtered traffic forwards to 192.32.140.216

After you apply a next-hop filter on an interface, all incoming IP traffic to that interface from the 10 network and going to the 47 network forwards to the next-hop address. This assumes that a reachable route exists to the next-hop address. If the next hop is not reachable, the VPN Router uses the destination address in the IP header (as in normal routing) to forward the packet.

For tunnels, make sure the next-hop address is beyond the remote end point of the tunnel and along the path to the actual destination.

To configure next-hop traffic filters

1 Choose **Profiles, Filters**.

The Filters window appears.

2 Click **Manage Rules**.

3 Select the rule that you want to change, and then click **Edit**.

4 Select **Nexthop** for the filter action. You can optionally enter the source and destination address fields

5 Click **OK**.

6 To enable private to tunnel forwarding, select **System, Forwarding**.

The Forwarding window appears.

7 Select **Apply Packet Filter on Private to Tunnel Traffic** in the Next Hop Forwarding section.

8 Click **OK**.

Chapter 4

NAT configuration

Network Address Translation (NAT) uses one or more globally unique IP addresses to give ports on a private network access to the Internet. For virtual private networks (VPN), multiple intranets with conflicting subnets implement NAT to communicate. The configuration of branch office or partner networks must securely route between these networks without requiring unique private addresses across the entire extranet.

NAT contains a pool of continually reused global addresses. A network can use one set of network addresses internally and a different set when it deals with external networks. The internal considerations of the network determine the allocation of internal network addresses. Global addresses must remain unique to distinguish between different hosts. After the router directs a packet, NAT replaces the internal corporate address with a global address. As soon as the application session ends, the global address returns to the pool so that subsequent connections can use the global address. NAT can also modify the source and destination port numbers.

This chapter includes the following topics:

- [“Address translations” on page 68](#)
- [“NAT modes” on page 74](#)
- [“NAT Traversal” on page 78](#)
- [“NAT and VoIP” on page 81](#)
- [“NAT usage” on page 85](#)
- [“NAT policy configuration” on page 89](#)
- [“Sample NAT procedures” on page 95](#)
- [“NAT ALG for SIP” on page 99](#)
- [“Firewall SIP ALG” on page 101](#)
- [“Configuring Firewall Virtual ALG” on page 102](#)

- [“Hairpinning” on page 104](#)
- [“Timeouts” on page 109](#)
- [“NAT statistics” on page 110](#)
- [“Proxy ARP” on page 110](#)

Address translations

You can configure address translation permanently (static) or allocate it dynamically, which permits many devices on an internal network to share a few IP addresses. Static translation allocates one external host address for each internal address and converts it to a different global IP address. Dynamic address translation occurs after a session starts. No guaranteed one-to-one mapping takes place. An example of dynamic translation is port mapping, which uses the TCP or User Datagram Protocol (UDP) source port and source address to allow multiple sessions from many hosts using a single public NAT address.

NAT supports the following address translations:

- Dynamic many-to-one
- Dynamic many-to-many
- Static one-to-one
- Port forwarding
- IPsec-aware NAT
- Double NAT


Dynamic many-to-one—port translation

With network address port translation (NAPT), many internal IP addresses hide behind a single external address. Dynamically-assigned ports distinguish one IP address from another. This method is especially useful if you need to use several IP addresses and your ISP provides only one address. Use dynamic many-to-one translation only for traffic initiated from an internal host.

NAT attempts to assign a port from the corresponding port list. NAT assigns the original port if it is available. If the port is not available, NAT tries to assign a port from the largest port number that is smaller than the original port. If all smaller ports are unavailable, NAT assigns a port greater than the one requested. If all ports are unavailable, the VPN Router drops the packet.

“Port translation” on page 69 shows the private network 10.0.1.0 hidden behind the public address 30.0.1.154. NAT replaces the source IP address of all requests that originate from the private network (10.0.1.0) with the public IP address 30.0.1.154; only the public IP address is visible from the public network. In addition, NAT dynamically translates source ports to unique translated ports.

Figure 12 Port translation




Dynamic many-to-many—pooled translation

In dynamic many-to-many translation, NAT translates only the address (not the port). Usually, the number of externally visible IP addresses is less than the number hidden behind the VPN Router. Each time a host on the private network makes a request, the VPN Router chooses an unused external IP address, and then performs the translation. Use dynamic many-to-many only for traffic that initiates from an internal host.

The following example “[Dynamic pooled address translation](#)” on page 70 illustrates many-to-many dynamic translation. The user configures a pooled NAT rule that converts the internal address range 10.0.1.154–10.0.1.164 to 30.0.1.154–30.0.1.154. Traffic initiates from 10.0.1.154 and 10.0.1.156 destined to a machine (11.1.1.2) on the public Internet. NAT translates both addresses to unique public addresses dynamically.

Figure 13 Dynamic pooled address translation




Static one-to-one translation

Static address translation allocates one external host address for each internal address. This allocation is always the same.

“[Static address translation](#)” on page 71 shows host 10.0.1.154 on the private side statically mapped to an external address 30.0.1.154, which allows Internet host 11.1.1.2 to initiate a session using the translated external address. The host that uses this rule is always bound to the same external address.

Figure 14 Static address translation




Port forwarding

With port forwarding, one externally accessible IP address forwards incoming requests to different addresses behind the NAT device based on the protocol used. You can route incoming Web traffic to a Web server, and you can forward File Transfer Protocol (FTP) traffic destined to the same external IP address to a different device that provides FTP services.

“[Port forwarding example](#)” on page 72 illustrates port forwarding. A host 11.1.1.2 on the Internet needs to access a Web server and an FTP server that run on two separate internal machines that are hidden behind the single externally visible address 30.0.1.154. To do this, you use a port forwarding NAT rule that sends the traffic to the two different machines based on the forwarding ports.

Figure 15 Port forwarding example




Double NAT

You can use double NAT to translate both external and internal networks at the same time. You can modify both the source and destination addresses for each packet that enters and leaves the VPN Router. You use rules to achieve this, one to translate the source address and one to translate the destination address. The destination address translation must use a static rule.

“[Double NAT](#)” on page 73 shows a host 11.1.1.2 on the Internet that initiates a connection to 30.0.1.154, the translated address of the internal host. NAT translates both the source and destination addresses as the packet traverses NAT.

Figure 16 Double NAT




IPsec-aware NAT

IPsec-aware NAT protects against the alteration of TCP/IP headers, usually performed by NAT. Use IPsec-aware NAT if an IPsec tunnel passes through a VPN Router that performs NAT translation, but the tunnel does not terminate at the VPN Router. IPsec-aware NAT provides interoperability with IPsec implementations that do not support the UDP wrapper solution to perform NAT on IPsec traffic. Unlike NAT Traversal, IPsec-aware NAT is always on and you cannot configure it.

“[IPsec-aware NAT example](#)” on page 73 shows an IPsec-aware NAT example.

Figure 17 IPsec-aware NAT example



NAT modes

Based on the handling of UDP packets, you can classify NATs in four different modes:

- Full Cone NAT
- Restricted Cone NAT
- Port Restricted Cone NAT
- Symmetric NAT




Note: The VPN Router supports only Restricted Cone NAT and Symmetric NAT modes. All visible references to Cone NAT in the system refer to Restricted Cone NAT.

Full Cone NAT

A Full Cone NAT maps all requests from the same internal IP address and port to the same external IP address and port. An external host can send a packet to the internal host by sending a packet to the mapped external address.


Figure 18 Full Cone NAT



“Full Cone NAT” on page 75 shows an example of a private client behind a NAT with IP 10.0.0.1 sending and receiving on port 8000 mapped to the external IP/port on the NAT device of 202.123.211.25:12345. Anyone on the public side can send packets to that external IP or port and the internal client IP or port correctly translates those packets.

Restricted Cone NAT


A Restricted Cone NAT maps all requests from the same internal IP address and port to the same external IP address and port. Unlike a Full Cone NAT, an external client can send a packet to the internal client only if the internal client previously sent a packet to the IP address.

Figure 19 Restricted Cone NAT

“Restricted Cone NAT” on page 76 shows an example of a private client sending a packet to an external client (computer A). The NAT device maps 10.0.0.1:8000 to 202.123.211.25:12345, which allows the public client to send back packets to the NAT address of the private client. However, the NAT device blocks all packets from an external client (computer B) until the private client sends a packet to that external IP address. Both external clients can send packets destined to the NAT address, and they are translated correctly to the private address of the client.

Port Restricted Cone NAT

A Port Restricted Cone NAT is similar to a Restricted Cone NAT, but the restriction includes port numbers. An external client can send a packet to the internal client only if the internal client previously sent a packet to the IP address and port.

Figure 20 Port Restricted Cone NAT


“Port Restricted Cone NAT” on page 77 shows an example of a Port Restricted Cone NAT. If an internal client sends a packet to an external client at IP 222.111.99.1 and port 10101, the NAT device only allows packets that come from the same IP and port. If the internal client sent packets to multiple external IP addresses or ports, they can all respond to the client at the same mapped IP address and port, and the NAT device performs the reverse translation to the internal IP address.

Symmetric NAT

A Symmetric NAT maps all requests from the same internal IP address and port to a specific destination IP address, to the same external IP address, and port. If the same host sends a packet with the same source address and port to a different destination, NAT uses a different mapping. Only the external host that receives a packet can send a packet back to the internal host.

The default NAT mode is Symmetric. To change the mode to Restricted Cone NAT, choose Services, Firewall/NAT, and then in the VPN Router Firewall row, click Edit.

Figure 21 Symmetric NAT



“Symmetric NAT” on page 78 shows an example of a Symmetric NAT. If the internal client 10.0.0.1:8000 sends a packet to the external IP 222.111.88.2, it maps to 202.123.211.25:12345 while a packet sent from the same address and port to 222.111.99.1 maps to a different public IP and port (202.123.211.25:45678). The external client on computer B can only send a packet to the mapped source address of the packet it received, and the external client on computer A can only send a packet to the mapped external source IP of its received packets.

NAT Traversal

The VPN client or server user tunnels use NAT Traversal to pass through intermediate routers or gateways, each of which can perform NAT actions on the packet. Most hotels and airports that provide Internet connectivity use NAT to connect to the Internet. VPN Routers use NAT Traversal for branch office tunnels between routers if one router is in a private network that uses one or more NAT devices.

The problem occurs after you place a VPN device behind a NAT device that does not support IPsec. In this situation, the VPN devices on both ends of the tunnel must wrap the IPsec packets in such a manner that the NAT device does not drop the packets. NAT Traversal solves this problem.

The VPN Router supports branch office NAT Traversal in the following two modes:

- draft mode—based on IETF draft *draft-ietf-ipsec-nat-t-ike-00*
- RFC mode—based on IETF RFC 3947

The VPN router always proposes both modes. If the other side supports both modes, the connection uses RFC mode. If the other side supports only one mode, the mode supported by the other side is used. The VPN Router chooses the mode based on the proposals received (only Draft or RFC). If no common mode exists, the IPSEC negotiation continues as if Branch Office NAT Traversal is disabled globally. If the router receives both modes, it chooses RFC mode. The following list describes the differences between these two modes:

- Draft mode uses port 500 for UDP encapsulation of Encapsulating Security Payload (ESP) and Authentication Header (AH) while RFC mode uses port 4500.
- RFC mode switches Internet Key Exchange (IKE) on port 4500 as soon as it discovers NAT translation. The RFC-compliant peer must also accept IKE negotiations initiated on port 4500.
- The UDP-encapsulated ESP header format is different. Draft mode considers IKE as the primary protocol multiplexed on port 500 and defines a non-IKE marker to add when encapsulating ESP. RFC mode considers ESP as the primary protocol on port 4500 and defines a non-ESP marker to add when encapsulating IKE.

To use User Tunnel NAT Traversal, you must also define a UDP port that all client connections use to connect to the VPN Router. This port must be a unique and unused UDP port within the private network (supported range 1025 to 49151). By default, no UDP port is defined.

For branch office NAT Traversal, VPN peers must use the same port for ESP or AH encapsulation in UDP. For more information about IKE NAT Traversal negotiation, see the IETF draft document *draft-ietf-IPSec-udp-encaps-00* or *UDP Encapsulation of IPSec ESP Packets (RFC 3498)*. Peers must be RFC or draft compliant to establish VPN communications.



Note: To allow NAT Traversal with the IPsec client, you must enable NAT Traversal by using the Profiles, Groups, Edit IPsec menu path.

You use the group-level NAT Traversal settings to configure the User Tunnel NAT Traversal mode at the group level. By default, NAT Traversal is Auto-Detect NAT, therefore the client and VPN Router UDP encapsulate ESP data whenever NAT is detected. If you select Not Allowed no UDP encapsulation will occur. The option Auto-Detect IPsec NAT allows the client and VPN Router to UDP encapsulate ESP data, but only if the NAT device detected is non-IPsec aware (when the NAT device does not allow for IPsec pass-through).

Because a variety of NAT devices and IPsec pass-through implementations exist, not all environments function properly in the Auto-Detect IPsec NAT mode. In environments with unknown NAT devices, Nortel recommends that you use the Auto-Detect NAT setting. Nortel recommends that you use the Auto-Detect IPsec NAT setting only for environments with well-known NAT devices.



Note: You can use an unused UDP port for NAT Traversal. Do not use Layer 2 Tunneling Protocol (L2TP)/Layer 2 Forwarding (L2F) port 1701 or General Packet Radio Service (GPRS) port 3386. Make sure that the port you select does not conflict with ports already in use.

To configure NAT Traversal for user tunnels

1 Choose **Services, IPsec**.

The IPsec Settings window appears.

2 In the **NAT Traversal** section, select **User Tunnel**.

3 Select **Disable Client IKE Source Port Switching** to disable the automatic client IKE source port switching.

4 Specify the UDP port to use for NAT Traversal.

5 Click **OK**.

To configure NAT Traversal for branch office tunnels

1 Choose **Services, IPsec**.

The IPsec Settings window appears.

2 In the **NAT Traversal** section, select **Branch Office Tunnel**.

3 Click **OK**.

4 Choose **Profiles, Branch Office**.

The Branch Office window appears.

5 Select a group, and then click **Configure**.

The Branch Office > Edit Group window appears.

6 In the IPsec section, click **Configure**.

7 Configure the interval after which keepalive messages transmit to the NAT device detected between the tunnel peers. The format is HH:MM:SS.

8 Click **OK**.

For information about how to configure branch office NAT Traversal in the CLI, see *Nortel VPN Router Using the Command Line Interface* (NN46110-507).

NAT and VoIP

As traffic traverses between private and public networks, NAT translates IP addresses and port numbers in private address ranges into public addresses. Private addresses are typically assigned to the IP endpoints in a Voice over IP (VoIP) network (IP phones, soft clients) to hide the IP identity from the public network. Voice calls from and to the public network must reach endpoints in the private network and, as a result, proper routing of media to endpoints with private addresses requires network address translation.

VoIP protocols introduce a number of complexities for NAT, because they carry IP address and port information within the body of the message that is not accessible to NAT. NAT cannot conduct translation on private IP addresses within the payload of application layer messages. Therefore, the voice media, which gets directed to the private IP address identified in the signaling message, is not routed to the private address, which results in a one-way speech path.

The challenges for VoIP traversal in NAT occur for the following reasons:

- NATs only look at Layer 3 addressing
- VoIP signaling protocols embed IP addresses at Layer 5
- Real-time Transport Protocol (RTP) and Real-time Transport Control Protocol (RTCP) work at Layer 5

Two of the most common solutions proposed to fix the NAT Traversal issue are

- Application Level Gateways (ALG)
- address and port discovery


The following section focuses on the address and port discovery mechanisms for VoIP. For more information about ALGs, see [“NAT ALG for SIP” on page 99](#).

Address and port discovery

In address and port discovery, the media end points send probe packets to a server to discover the public IP address and port to use for a specific media stream. The server echoes back to the end point the source IP address as seen after the translation.

Applications use Simple Traversal of UDP through NATs (STUN), a lightweight protocol, to discover the presence and types of NATs and firewalls between the application and the public Internet. Applications also use STUN to determine the public IP addresses allocated by the NAT device.

[“STUN” on page 83](#) shows how STUN works.

Figure 22 STUN

STUN inspects exploratory STUN messages that arrive at the STUN server to identify the public-side NAT details. The STUN-enabled client sends an exploratory message to the external STUN server to determine the transmit and receive ports to use. The STUN server examines the incoming message and informs the client about which public IP address and ports the NAT device uses. The client uses this information in the call establishment messages sent to the SIP server. The STUN server does not sit in the signaling or media data flows.


To discover a valid IP address and port, NAT must use the same IP address and port binding, regardless of where the packet is going. This means that Symmetric NAT does not work for peer-to-peer media with address and port discovery. STUN requires a Cone NAT implementation. Restricted Cone NAT makes the VPN Router more secure.

Network address port translation (NAPT)

Network address port translation (NAPT) is a dynamic NAT where many internal IP addresses hide behind a single external IP address, distinguished only by their dynamic port assignment. The Symmetric NAT maps an IP address and port to a unique IP address and port for each session initiated from a private client. With Cone NAT, this mapping changes so that each internal IP address and port maps to the same external IP address and port, regardless of the destination and the session.

“Restricted Cone NAT — NAPT” on page 84 shows the flow of a Restricted Cone NAT.

Figure 23 Restricted Cone NAT — NAPT



Configuring Cone NAT

You can enable or disable Cone NAT with the graphical user interface (GUI) or the Command Line Interface (CLI). For more information about the CLI, see *Nortel VPN Router Using the Command Line Interface* (NN46110-507).

To configure Cone NAT

- 1 Choose **Services, Firewall/NAT**.
The Firewall/NAT window appears.
- 2 Click **Edit** in the **VPN Router Firewall** row.
- 3 Under the **NAT Mode** section, select **Cone NAT**.

4 Click **OK**.



Note: If you change the NAT mode, you clear the NAT flow cache. After you clear the NAT cache flow, you disrupt all active NAT sessions.

NAT usage

NAT applies to routed traffic that passes through the physical interfaces (interface NAT) and branch office interfaces (branch office NAT) using separate NAT policies. Each branch office uses one NAT policy, and one global NAT policy applies to nontunneled traffic.



Note: If you make changes to a branch office parameter, you must disable and then reenable the branch office for the changes to take effect. You can use the flow cache clear capability to force NAT changes to take effect on existing sessions.

Branch office tunnel NAT

In branch offices, two or more branches can use the same private addressing scheme. The branch offices must still communicate with one another.


A typical scenario can include a client on LAN 1 who tries to access the FTP server on LAN 2, and who sends a packet with a source address of 10.0.0.13 and a destination address of 10.0.0.14. Without NAT, the VPN Router looks at the destination address and assumes that the destination is on the same LAN as the source device because the addresses are both on the 10.0.0.0 network and no tunnel connection is brought up. Because you cannot use an Interior Gateway Protocol (IGP) to dynamically learn routes at the remote end of the tunnel to allow the client to access the server on the other LAN, you implement NAT on both sides of the branch office connection. This issue is common for branch office tunnels where the address space overlaps for each end.

To allow the client to access the server on the other LAN, you can implement NAT on both sides of the branch office connection. In this example, VPN Router1 defines a remote accessible network of 12.0.0.0, and VPN Router2 defines a remote accessible network of 11.0.0.0. VPN Router2 uses a static translation of 10.0.0.14 (server) to 12.0.0.1. VPN Router1 uses a translation of 10.0.0.13 (client) to 11.0.0.1. As a result, VPN Router2 must define 11.0.0.0 as the remote accessible network.

With NAT implemented on both sides of the branch office connection, the client can access the FTP server. A packet generated from the client uses a source address of 10.0.0.13 and a destination address of 12.0.0.1. VPN Router1 recognizes that 12.0.0.0 is the remote LAN for the branch office connection. VPN Router1 translates the source address of the packet to 11.0.0.1 based on the NAT table. VPN Router 2 looks at the destination address of the incoming packet and translates it to 10.0.0.14, but the source address remains 11.0.0.1.

“Overlapping address translation” on page 86 shows a simple branch office connection with two LANs, and a branch office tunnel across the Internet. VPN Router1, which connects the local network to the remote network through its branch office tunnel, uses a pooled NAT rule.

Figure 24 Overlapping address translation



Interface NAT


If you apply interface NAT to IP packets that leave or arrive into the VPN Router through the physical interfaces, either the source or destination IP address is translated to another IP address, depending on the NAT policy.



Note: The difference between interface and branch office NAT is when and where you apply the NAT policy.

“Interface NAT” on page 87 shows an example of interface NAT.

Figure 25 Interface NAT



Use the Services, Firewall/NAT window to apply interface NAT. Interface NAT can use one of the following rule types:

- Static—for static mapping, an internal address range maps one-to-one to an external range.
- Port forwarding—for port forwarding mapping, external packets route on a specified port to one of the internal systems.
- Port—for port mapping, the range of internal addresses is hidden behind a single external address. These external addresses are distinguished by using dynamically assigned port numbers.

- Pooled—for pooled mapping, an internal address dynamically maps to the next available address from the external address range.




Note: Interface NAT applies only to clear text traffic (nontunneled, routed through the VPN Router). Branch office NAT only applies to specific branch office tunnel traffic. If you disable interface NAT, it does not impact branch office NAT.

Dynamic routing protocols

You can advertise NAT routes on all interfaces. You use the routing policy list to restrict the route redistribution to only specific interfaces. Whenever you apply a NAT policy to interface or branch office tunnels, the routes to the translated IP addresses are added to the routing table. After you disable NAT, the routes to the translated IP addresses are deleted. Destination NAT adds the original destination address and source NAT adds the translated source address.

In “[NAT with dynamic routing example](#)” on page 88, the VPN Router uses a NAT rule to convert IP addresses in the range of 10.0.1.1 to 10.0.1.10 to 192.168.1.1.

Figure 26 NAT with dynamic routing example



By default, the Routing Information Protocol (RIP) and the Open Shortest Path First (OSPF) protocol distribute NAT routes. However, you can disable the redistribution for a particular protocol on the Routing, Policy, Redistribution Table window.

You can enable NAT on a branch office with dynamic routing. When you configure NAT for a branch office, you do not want it to announce the route to original IP addresses. You can create a routing policy to block the route advertisement to the original IP addresses, but it cannot announce a part of a subnet. Therefore, if you apply NAT to part of a subnet, routes do not advertise to the entire subnet.

You can add the translated address range to the routing table as a single subnet. However, if you choose a nonsubnet IP address range, you can add those addresses as individual host entries or as a group of smaller subnets (summarization). Summarization reduces the number of NAT route entries in the routing table manager (RTM) and thereby the number of entries redistributed. You can either enable or disable the summarization option. By default the option is enabled.

If you configure both NAT and dynamic routing, do not enable a branch office if no routing policy associates with the corresponding branch office interface. You must create a routing policy on the Routing, Policy window.

NAT uses a port mapping table to track the ports of outgoing packets for each client. The port mapping table relates the actual local IP address, source port, and translated source port number of the client to a destination address and port. NAT can then reverse the process for returning packets and route them back to the correct clients. This applies to TCP and UDP traffic only.

NAT policy configuration

A NAT policy consists of service properties and a security policy. Service properties define the service offered and includes a service name, the protocol (TCP, UDP, ICMP), and the port number (or range) on which the service occurs.

Security policies consist of a set of rules that specify what service is allowed or denied. You use service objects to specify all rule fields for service policies. Each rule consists of a combination of network objects, services, actions, and logging mechanisms. You can define custom policies when you need more complex security policies and the standard policies are not sufficient.



Note: Read-only NAT Policies created prior to Version 4.80 work according to the previous translation until you apply a modified copy to the interface. If you reapply the read-only NAT policy after the copy, the read-only policy translates in accordance with the new rules.

NAT policy sets

The VPN Router maintains one set (source and destination address pair) of active global NAT policies for all nontunneled traffic and a configurable NAT policy set for each branch office tunnel definition. To view active NAT policies for interface and branch offices, choose Status, Statistics.

At system startup, NAT obtains a cached policy (if one exists) while the system initializes. If no cached policy exists, it takes the default NAT policy, which is no NAT translation. The default NAT policy for the VPN Router 1010, 1050, and 1100 port maps its private address space to the public IP address.

After the system initialization is complete, the router retrieves the NAT policy from the LDAP database and it becomes the active policy. After you change the policy, the system stores it on the local disk as a cached policy and in the LDAP database. NAT uses the active policy for new sessions. For the existing sessions, it uses the original policy.

Rule creation

Menus control actions on rules. Right-click an option to access menus. Each of the following menus control a different aspect of the rule:

- Header row menus—contain only Add New Rule, which you use to add a new rule to the top of the list. The new rule appears in position one and all existing rules increment by one.

- Row menus—use this menu to add a new rule at a particular location, delete the specific rule, and perform cut, copy, or paste operations on a rule.
- Cell menus—are cell-specific and contain cell option menus and procedure menus.
 - Option menus provide a list of possible values for the cell. After you click one of the items, the selection is displayed in the cell.
 - Procedure menus provide a list of operations that you can perform on the cell, such as Add and Edit. After you click one of the items, either the operation performs immediately (such as Copy) or an additional dialog box appears and prompts you for more information (such as Add).

For rule columns, each rule within a NAT policy uses the same attributes, which are specified by the following column headers:

- # specifies the order of the rules within the section. The order applies only to the section in which the rule appears and does not apply across the entire policy.
- Source and Destination specify the source and destination network object for the rule. You can add more than one source or destination address to a rule. To modify these attributes, right-click on a column in the cell, which shows a procedure menu. Click Add to display the Network Object Selection dialog box. In this dialog box you define and apply a new network object. You can create the following network objects: host, network, IP range, and group (a collection of these objects).



Note: You use the NOT operand to specify which networks you do not want to use NAT.

Italicized objects in the list are read-only. You cannot modify them. Use the New, Edit, and Delete options to create, edit, and delete network objects.

Click Edit to display the Network Object Edit window. Use this window to modify the attributes for the selected network object.

Click Delete to remove the selected network object. If the object that you want to delete is the last object, it returns to the default value.

- Service specifies which service objects are handled by the selected rule. Right-click on the cell to display the standard procedure menu (Add or Edit).

Click **Add** to access the Service Object Selection dialog box, where you define and apply a new service object. You can create the following service objects: TCP, UDP, ICMP, IP protocol, and object groups (a collection of these objects).

Italicized objects in the list are read-only. You cannot modify them. Use the **New**, **Edit**, and **Delete** options in this window to create, edit, and delete service objects. Click **Edit** to display the Service Object Edit window. Use this window to modify the attributes for the selected service object.

Click **Delete** to remove the selected service object from the cell. If the object you want to delete is the last object in the cell, the cell returns to its default value (in this case, **Any**).

Click **Copy**, **Cut**, or **Paste** to perform those operations on the current service object.

- **NAT Action** specifies the action that occurs after you activate the rule. Right-click the cell to display an option list that contains the following items: **None**, **Static**, **Pooled**, **Port Mapping**, and **Port Forwarding**. Click one of these items to configure the cell to the selected state.
- **Translated Source**—specifies the source IP address of the first packet (static, pooled, port). To modify this attribute, right-click a column in the cell. You can add more than one source address to a rule. You can create the following network objects: **host**, **network**, **IP range**, and **group** (a collection of these objects).
- **Translated Destination**—specifies the destination IP address of the first packet of a port forwarding application session. To modify this attribute, right-click a column in the cell, which shows a procedure menu. You can add more than one destination address to a rule.
- **Status**—specifies the status of the particular rule. The status is either **Enabled** or **Disabled**.
- **Remark**—attaches a remark to a particular rule. After you right-click **Remark** and choose **Add** or **Edit remark**, a dialog box appears where you can type a comment.

Creating a new policy

To configure NAT policies

- 1 Choose **Services, Firewall/NAT**.

The Firewall/NAT window appears.

- 2 Enable **Interface NAT**.
- 3 Click **Manage Policies**.

The NAT, Select Policy window appears. Use this window to create, edit, delete, copy, or rename a NAT policy. Bold denotes the policy that is currently applied to the VPN Router, and italics denotes read-only policies.

The System Default policy is always listed. This read-only policy defines the NAT behavior when you do not apply user-defined policies or when the selected policy is not available.



Note: The exception to this rule is the VPN Router 1010, 1050, and 1100 where the default NAT policy is translate everything to the public interface IP (Interface NAT). These VPN Router systems are generally used in a small office environment where you want to apply NAT to everything on the private side of the single global IP address assigned by the ISP.

- 4 Click **New** to create a new policy.
- 5 Type the policy name, and then click **OK**. The name must begin with a letter and cannot contain the : + =] , ; " characters.
The NAT, Edit Policy: <polycyname> window appears with no rules defined. In this window, you can add, delete, and modify the rules for the policy.
- 6 Right-click the appropriate cell to add a new rule.
- 7 Repeat these steps to add more rules.
- 8 Select **Policy**, and then click **Save Policy** to save your changes.
- 9 After the policies are saved, choose **Manager, Exit SFw/Nat**.

Adding a policy

To add a new policy

- 1 Click **New**.

- 2 Type the policy name. The name must begin with a letter and cannot contain the : + =] , ; " characters.
- 3 Click **OK** to go to the **Policy Edit** window, which provides a blank NAT policy, or click **Cancel** to return to the **Policy Selection** window.

Deleting a policy

You cannot delete a read-only policy or the policy that is currently applied to the VPN Router. If you select one of these policies, the Delete option is disabled. To delete a policy

- 1 Select the policy that you want to delete, and then click **Delete**.
- 2 Click **OK** to delete the selected policy.

Copying a policy

To copy a NAT policy

- 1 Select the policy that you want to copy.
- 2 Click **Copy**.
- 3 Enter a name for the copied policy.
- 4 Click **OK**.

The new policy appears in the list of policies in the NAT policies window. This policy contains the same rules as the original policy.

Renaming a policy

You cannot rename a read-only policy or the policy that is applied to the VPN Router. If you select a read-only policy, the Rename option is disabled. To rename a policy

- 1 Select the policy that you want to rename.
- 2 Click **Rename**.

The Rename dialog box appears.

- 3 Enter the new name of the policy.
- 4 Click **OK**.

Sample NAT procedures

The following sections describe the steps for sample NAT procedures.

For the following configuration on the VPN Router, create the NAT policy:

STATIC: 10.0.1.0 - 10.0.1.255 -> 30.0.0.0 - 30.0.0.255

Go to Routing, Access List and create an access list acc1 to permit 30.0.0.0/24 and deny 10.0.1.0/24. Create another access list acc2 to permit 10.0.0.0/16 and deny 30.0.0.0/24.

Configuring interface NAT with RIP

This procedure shows interface NAT with RIP:

- 1 On the VPN Router, enable **Interface NAT**, and then apply the preceding NAT policy to **Interface NAT**.
- 2 Choose **Routing, RIP**, and then enable **RIP**.
- 3 Choose **Routing, Policy**, and then verify the redistribution table for the **RIP** protocol to redistribute NAT routes.
- 4 Create a policy list of type **Announce** on Interface **20.0.9.100** for protocol **RIP** with **acc1** access list.
- 5 Create another policy list of type **Announce** on Interface **10.0.9.100** for protocol **RIP** with **acc2** access list.
- 6 Send a ping request from **10.0.1.1** to **20.0.1.1**. Ping gets the reply back.

Configuring interface NAT with OSPF

This procedure shows interface NAT with OSPF:

- 1 On the VPN Router, enable **Interface NAT**, and then attach the preceding NAT policy to **Interface NAT**.
- 2 Choose **Routing, OSPF**, and then enable **OSPF**.
- 3 Choose **Routing, Policy**, and then verify the redistribution table for the **OSPF** protocol to redistribute NAT routes.
- 4 Create a policy list of type **Announce** on Interface **20.0.9.100** for protocol **OSPF** with an **acc1** access list.
- 5 Create another policy list of type **Announce** on Interface **10.0.9.100** for protocol **OSPF** with an **acc2** access list.
- 6 Send a ping request from **10.0.1.1** to **20.0.1.1**. Ping gets the reply back.

Configuring branch office NAT with RIP

This procedure shows NAT on a branch office with dynamic routing enabled:

- 1 On the VPN Router, select **Profiles, Branch Office**, and then create a branch office with a local end point as **20.0.9.100** and remote end point as **20.0.9.1**.
- 2 Enable **dynamic routing** for that branch office, and then enable **RIP**. Enable **NAT** and create the preceding NAT policy.
- 3 Choose **Routing, RIP**, and then enable **RIP**.
- 4 Choose **Routing, Policy**, and then verify the redistribution table for **RIP** protocol to redistribute NAT routes.
- 5 Create a policy list of type **Announce** on **Branch Office** Interface for protocol **RIP** with an **acc1** access list.
- 6 Create another policy list of type **Announce** on Interface **10.0.9.100** for protocol **RIP** with an **acc2** access list.
- 7 To configure Router-2 (VPN Router), select **Profiles, Branch Office**, and then create a branch office with a local end point as **20.0.9.1** and remote end point as **20.0.9.100**.
- 8 Enable **Dynamic Routing** for that branch office, and then enable **RIP**.
- 9 Choose **Routing, RIP**, and then enable **RIP**.
- 10 Send a ping request from **10.0.1.1** to **20.0.0.1**. Ping gets the reply back.

Configuring branch office NAT with OSPF

This procedure shows NAT on a branch office with dynamic routing enabled:

- 1 On VPN Router-1, select **Profiles, Branch Office**, and then create a branch office with a local end point as **20.0.9.100** and remote end point as **20.0.9.1**.
- 2 Enable **Dynamic Routing** for that Branch Office, and then enable **OSPF**. Enable **NAT** and create the preceding NAT policy.
- 3 Choose **Routing, OSPF**, and then enable **OSPF**.
- 4 Choose **Routing, Policy**, and then verify the redistribution table for **OSPF** protocol to redistribute NAT routes.
- 5 Create a policy list of type **Announce** on the **Branch Office** Interface for protocol **OSPF** with an **acc1** access list.
- 6 Create another policy list of type **Announce** on Interface **10.0.9.100** for protocol **OSPF** with an **acc2** access list.
- 7 To configure the Router-2 (VPN Router), select **Profiles, Branch Office**, and then create a branch office with a local end point as **20.0.9.1** and remote end point as **20.0.9.100**.
- 8 Enable **Dynamic Routing** for that branch office, and then enable **OSPF**.
- 9 Choose **Routing, OSPF**, and then enable **OSPF**.
- 10 Send a ping request from **10.0.1.1** to **20.0.0.1**. Ping gets the reply back.

Configuring branch office NAT

This configuration example “[NAT configuration example](#)” on page 97 adds a NAT static rule with a single host as the source.

Figure 27 NAT configuration example



- 1 Select **Services, Firewall/NAT**, and then in the **Interface NAT** row, click **Manage Policies**.

- 2 Click **New**, type the policy name, and then click **OK**.
- 3 Right-click # and click **Add New Rule**.
- 4 Right-click **Source**, and then click **Add**.

The Network Object Selection window appears. You use this window to create network objects. After you create the network object, you can apply a network object to an Address column of the rule.
- 5 Click **New**, select **Host**, and then click **OK**.
- 6 In the **Host Object Insert** window, type the host name and IP address: **Sqa64; 1.0.0.64**. Click **OK** twice to return to the NAT Translate Action window.
- 7 Right-click **NAT Action** (None), and then select an action.
- 8 Right-click **Translated Source**, and then click **Add**.
- 9 Click **New**, select **Host** and click **OK**.
- 10 In the **Host Object Insert** window, type information for the translated host: Host Name **Sqa64Trans**; IP Address **30.0.0.64**. Click **OK** twice to return to the NAT Translate Action window.
- 11 Click **Policy, Save policy**. A `Please wait ...` message appears to show that you saved the policy.
- 12 Choose **Profiles, Branch Office**, select a branch office tunnel, and click **Configure**.
- 13 From the **NAT** menu, select the policy, and then click **OK**.
- 14 From **SQA64**, use ping, Telnet or another application to pass traffic over the tunnel.

Configuring NAT with the VPN Router Stateful Firewall

To use NAT on the VPN Router with the VPN Router Stateful Firewall, where the NAT address is within the same subnet as the public interface.

- 1 Create a NAT policy with the following specifications:
 - a Type **static** in the name field.
 - b Leave the **Translation** type configured as **static**.

- c Add the internal VPN Router address (for example, 10.4.4.204) as the start and the end internal address.
 - d Add the external address (for example, 192.168.4.204) as the starting external address.
- 2 Choose **System, Forwarding** and enable Proxy ARP for physical interfaces and click **OK**.
 - 3 Enable **Interface NAT** and select the NAT rule created in Step 1.




Note: You must configure an Allow All policy on the VPN Router Stateful Firewall.

NAT ALG for SIP

Traditional NATs do not translate Layer 5 addresses. Therefore, the VoIP signaling, RTP, and RTCP become unreachable after NAT translation (one-way signaling and audio) due to the embedded IP address and port specified within the IP payload.

“[NAT and SIP](#)” on page 99 illustrates the problem caused by NAT for Session Initiation Protocol (SIP) signaling.

Figure 28 NAT and SIP



“[NAT and SIP](#)” on page 99 shows the following process:

- 1 User A sends an invite to User B.
- 2 The NAT device translates the Layer 3 address, but not the Layer 5 (SIP and Session Description Protocol [SDP]) addresses.
- 3 User B receives the invite and responds back to the NAT address. The signaling is complete (for example, 200 OK).
- 4 User A sends RTP to User B SDP c= / m= address: port.
- 5 User B tries to send RTP to User A c= / m= address: port, but this fails because it cannot route to User A (the SDP address and port did not receive the NAT), which results in one-way audio.
- 6 If User A hangs up (because of one-way audio), the BYE is sent to User B correctly.
- 7 If User B hangs up, the BYE does not reach User A because the header address did not receive the NAT. This leaves the state of User A for that session as up until User A hangs up.

Two of the solutions that correct the NAT Traversal issue are

- ALGs
- address and port discovery

For more information about the address and port discovery method, see [“Address and port discovery” on page 82](#).

The following section focuses on NAT ALG for SIP to support VoIP phones that use SIP as the signaling protocol.

Application level gateways

A NAT ALG translates embedded IP addresses and port numbers in application protocol messages. NAT ALG supports FTP, ICMP, Berkeley R commands, NetBIOS, IPsec (ESP only), and the Simple Network Management Protocol (SNMP). For application traffic flows that embed an IP address in the data portion (such as FTP or NetBIOS), you must use an ALG.

Use SNMP ALG support to use SNMP traps with NAT. The data within the SNMP traps is translated, which prevents inconsistencies within the packet. The SNMP ALG applies to SNMP traps that originate from the VPN Router only if NAT rules translate traffic that originates from the VPN Router. You must enable the SNMP management system to send SNMP Gets from the Admin, SNMP window.

The NAT ALG provides support for SIP traffic to and from SIP phones and the SIP Server MCS 5100 because i2004 phones are Unified Networks IP Stimulus (UNISim) devices.

Configuring NAT ALG for SIP

You can enable or disable NAT ALG for SIP with either the GUI or the CLI. For more information about the CLI commands, see *Nortel VPN Router Using the Command Line Interface* (NN46110-507).

To configure NAT ALG for SIP

- 1 Choose **Services, Firewall/NAT**, and then click **Edit** in the **VPN Router Firewall** row.

The Firewall/NAT > Edit window appears.

- 2 Under **NAT Application Level Gateway**, click **SIP**.
- 3 Click **OK**.



Note: If you enable Firewall in the Logging section, you receive a log with Firewall events in it.

Firewall SIP ALG

Firewalls, by default, cannot identify port numbers within the payload of signaling protocols and cannot dynamically open ports for media traversal, which blocks voice traffic. Firewalls operate with Layer 3 or Layer 4 information and cannot access information in higher layer protocols.

The development of ALGs for the VoIP signaling protocols solves this issue. The SIP ALG performs the necessary translation of the IP addresses embedded in the SIP messages and updates the SDP information. The Firewall ALG examines the SDP information, identifies the RTP port number for the call and opens the port in the firewall during call setup. The Firewall ALG also raises a flag to tell NAT to perform an application level translation. The ALG then performs the address to port mapping and state setup to ensure that the data channels map according to the information in the SDP. The ALG closes the port after call termination. This provides a mechanism to dynamically open and close ports in the firewall and restrict the voice traffic to active sessions only to increase network security.

Configuring Firewall Virtual ALG

The Firewall Virtual ALG is a syntax-independent ALG for firewall traversal that works for both encrypted and nonencrypted UNISlim signaling, such as VoIP. A Firewall Virtual ALG works only with UNISlim signaling.

The Firewall Virtual ALG is based on a trust model that assumes that the phone authenticates itself with the call server, and that continuous detection of signaling traffic between the phone and the call server allows media to or from the phone to traverse the firewall. Continuous communication implies that the call server trusts the endpoint and that the call server does not communicate constantly with the endpoint device if the endpoint device is not authorized to send media through the firewall. The controlling entity does not acknowledge requests from unauthorized devices.

The entity that controls the phone in Succession 1000 Call servers is the Terminal Proxy Server (TPS). With TPS, UNISlim phones on the private side can make calls to phones on the public side without explicitly opening holes in the firewall.

To enforce a more stringent and secure protocol, the Firewall Virtual ALG waits until it receives an RTP or RTCP packet from the phone on the private side to open a pinhole in the firewall. The advantage of this late pinhole creation is that the ALG uses the exact 5 tuple for which it needs to open a pinhole. The Firewall Virtual ALG creates the pinhole only for outbound traffic, which prevents unauthorized access from the outside. The Firewall Virtual ALG creates a reverse path in response to the outbound pinhole. The system drops all packets from the outside phone until the internal phone sending packets to the external phone creates the pinhole.

Because the Firewall Virtual ALG cannot interpret and inspect the UNISlim protocol, the Firewall Virtual ALG closes the pinholes only after the default timeout period of the underlying transport protocol.

To enable or disable the Firewall Virtual ALG

1 Choose Services, Firewall/NAT.

The Firewall/NAT window appears.

2 In the VPN Router Firewall row, click Edit.

3 In the FW Application Level Gateway section, click Enable or Disable. The default is disabled.

To configure the Firewall Virtual ALG:

1 Choose Services, Firewall/NAT.

The Firewall/NAT window appears.

2 In the VPN Router Firewall row, click Edit.

3 In the FW Application Level Gateway section, click Configure.

The port number in the Signaling Port and the Media Port dialog boxes depends on the configuration of the server.

4 To add a server, click Add.

a Type the name of the server in the **Server Name** box.

b Type the IP address in the **IP Address** box.

c Type the port number in the **Port** box.

d Select either **TCP** or **UDP** as the protocol.

e Click **Apply**.

5 To edit a call server, click Edit.

6 To delete a call server, click Delete.

To enable the Virtual ALG with the CLI, enter the following command:

```
CES(config)#firewall alg virtual enable
```

To disable the Virtual ALG, enter the following command:

```
CES(config)#no firewall alg virtual enable
```

To configure the Virtual ALG Server, enter the following command:

```
CES(config)#firewall alg virtual server <servername> ip  
<ipaddress> port <portnumber> proto <tcp/udp>
```

The following example shows how to configure ports:

```
CES(config)#firewall alg virtual port-media 5200  
CES(config)#firewall alg virtual port-signaling 5000
```

Hairpinning

You need hairpinning when two IP phones behind the same NAT want to communicate. VPN Router NAT blocks packets from the private side of the NAT device that are destined for the private side for which a NAT binding to a specific port exists. This does not allow peer-to-peer communication between two endpoints behind the same NAT device if they try to use their public address. Hairpinning examines the destination address of a packet, evaluates the destination address NAT binding, and makes a determination on the requirement for hairpinning.

NAT hairpinning performs payload translation on SIP and UNISim messages.


Hairpinning with SIP

Hairpinning solves another special issue that is introduced when voice phones are on one side of a NAT boundary and the call server is on the other side. The SIP NAT ALG translates the IP addresses of the SIP phones from private space to public. After the call server is queried for the IP address of the person you call, it responds with the public IP address. The server also supplies the called person with the public IP address of the caller.

Although both clients are in the same private address space, each thinks the other resides in the public address space. The media traffic between the clients needs to go to and from the public addresses, looping through the NAT device.

“[Hairpinning with SIP](#)” on page 105 shows hairpinning support required for VoIP media. The MCS call server sees both private side phones as using a 47.17.248.1:x address, telling the private side caller that the caller uses a 47.17.248.1:x IP.


Figure 29 Hairpinning with SIP



Hairpinning with a UNISstim call server

When a UNISstim call server sends an Open Audio Stream (OAS) message to an IP phone, it always uses the public address as the far end address for the other IP phone. If both IP phones are behind the same NAT, this creates problems because the media packets are sent to the NAT device, which does not know what these packets are for. However, if the NAT device supports hairpinning, it redirects the packets to the right destination, which generates the voice path.

“[Intrarealm call with hairpinning](#)” on page 106 shows an intrarealm call with hairpinning.

Figure 30 Intrarealm call with hairpinning

In “[Intrarealm call with hairpinning](#)” on page 106, both i2004a and i2004b are behind the same NAT and registered into the same CS 1000 TPS server. The server encrypts the UNISTim messages and the ERouter NAT cannot translate UNISTim messages payload. After successful registration of both IP phones, ERouter NAT generates the following NAT table entries:

Table 3 NAT entries

Internal address	External address	Remote address
192.168.0.2:5000	47.135.152.15:12345	47.135.152.16:7000
192.168.0.2:5200	47.135.152.15:52000	47.135.152.16:10000
192.168.0.2:5201	47.135.152.15:52001	47.135.152.16:10001
192.168.0.3:5000	47.135.152.15:12347	47.135.152.16:7000
192.168.0.3:5200	47.135.152.15:52002	47.135.152.16:10000
192.168.0.3:5201	47.135.152.15:52003	47.135.152.16:10001

After i2004a calls i2004b, TPS sends OAS to i2004b with the following contents:

Far end address = 47.135.152.15:52000

Near end port = 5200

TPS sends OAS to i2004a with the following contents:

Far end address = 47.135.152.15:52002

Near end port = 5200

After i2004a sends media packets to i2004b, the packet header looks like this:

Source Address = 192.168.0.2:5200, Destination = 47.135.152.15:52002.

After i2004b sends media packets to i2004a, the packet header looks like this:

Source = 192.168.0.3:5200, Destination = 47.135.152.15:52000.

After ERouter NAT receives the media packet generated by i2004a, it first compares the destination address in the packet header against its external address entries on its NAT table. It finds a match (47.135.152.15:52002) and translates the destination address from 47.135.152.15:52002 to 192.168.0.3:5200.

The ERouter NAT further compares the source address in the packet header against the internal address entries on the NAT table. The ERouter finds a match (192.168.0.2:5200), translates the source address from 192.168.0.2:5200 to 47.135.152.15:52000, and forwards the translated packet to i2004b.

Similarly, after ERouter NAT receives the media packet generated by i2004b, it first compares the destination address in the packet header against its external address entries on its NAT table. It finds a match (47.135.152.15:52000) and translates the destination address from 47.135.152.15:52000 to 192.168.0.2:5200.

The ERouter NAT further compares the source address in the packet header against the internal address entries on the NAT table. The ERouter finds a match (192.168.0.3:5200), translates the source address from 192.168.0.3:5200 to 47.135.152.15:52002, and forwards the translated packet to i2004a.




Note: Hairpinning support is part of the solution, and can coexist with the other portions of the solution. For example, with nonencrypted UNISTim messages, the hairpinning logic automatically turns off, and a direct media path is achieved.

Hairpinning with a STUN server

When NAT Traversal for phones behind the NAT device is based on STUN, the phones use the port discovery protocol between the phone and the STUN server to discover their public addresses and use the discovered public addresses for peer-to-peer communication.

The diagram in “[NAT Hairpinning](#)” on page 108 describes the hairpinning solution with the STUN server. Phone A and Phone B discover their public addresses. Phone A on the private side of the VPN Router initiates a call to Phone B on the private side. After the call establishes, Phone A starts to send media to Phone B and Phone B to Phone A with public NAT destination addresses in the media packets. VPN Router NAT, unaware that the voice packets need NAT hairpinning, blocks the media packets. After you enable NAT hairpinning, it examines the destination address of a packet, evaluates the destination address NAT binding, and makes a determination on the requirement for hairpinning.

Figure 31 NAT Hairpinning



Hairpinning requirements

Two requirements exist for NAT Hairpinning:

- Because IP phones do not always accept packets from arbitrary IP addresses, the source IP address must be the public IP address of the NAT device.

- If the device performs NAT on a VPN tunnel, packets sent from private devices to the assigned VPN IP are hairpinned back without entering the VPN tunnel. After the packets reach the private endpoint, the source IP address must be the assigned VPN IP address.

Enabling hairpinning

You can use the GUI or the CLI to turn the hairpinning of packets on or off. For more information about the CLI commands, see *Nortel VPN Router Using the Command Line Interface* (NN46110-507).

To configure hairpinning

- 1 Choose **Services, Firewall/NAT**.
The Firewall/NAT window appears.
- 2 In the **VPN Router Firewall** row, click **Edit**.
- 3 For Hair-pinning, select **Enable**.
- 4 Click **OK**.

Hairpinning statistics are shown on the Status, Statistics, NAT Stats window.

Timeouts

After a session terminates, NAT deletes the associated translations. However, if a server goes down unexpectedly, the associated translation must age out so that the available translation addresses are not exhausted. The NAT timeouts are grouped by the following protocol:

- ICMP—3 minutes
- UDP—3 minutes
- TCP—120 minutes

NAT statistics

The following statistics counters are provided for source and destination NAT services:

- Source Translated—number of packets with the source address translated
- Destination Translated—number of packets with the destination address translated
- Flows Translated—number of flows translated by NAT service
- No Action—number of flows for which no translation occurs
- Dropped—number of packets dropped because NAT did not translate the source or destination address
- Pooled Address Translations failed—number of packets dropped because NAT did not map a new address from the available address pool
- Port Translations failed—number of packets dropped because NAT did not map a new port for translation

You can view the NAT statistics by choosing the Status, Statistics menu path.

Proxy ARP


You need Proxy Address Resolution Protocol (ARP) if the translated address NAT assigns to a private host makes it appear as if that private host is on the network of the other host. The other host sends an ARP request and does not receive a response unless you enable Proxy ARP for physical interfaces on the VPN Router.

In [“Proxy ARP example” on page 111](#), the numbers correspond to the following actions:

- 1 Host 20.0.1.150 pings the host 20.0.1.1.
- 2 The ARP request for host 20.0.1.1 broadcasts to the network.
- 3 The VPN Router responds to the ARP request using its own hardware address for the ARP reply.
- 4 The ICMP echo reply is sent directly to the host 20.0.1.1.

- 5 Because the interface NAT policy statically maps 20.0.1.1 to 10.0.1.1, this first packet is translated and sent to 10.0.1.1.
- 6 Host 10.0.1.1 receives the ping.
- 7 It replies with its own ICMP echo reply and sends the packet to the VPN Router.
- 8 The source IP 10.0.1.1 is translated to 20.0.1.1 and sent to 20.0.1.150.
- 9 The target host receives the packet, processes the ICMP, and the ping program reports the results.

Figure 32 Proxy ARP example



Chapter 5

Firewall user authentication configuration

You use firewall user authentication (FWUA) to ensure users log on to the VPN Router Stateful Firewall before they are granted network access. FWUA provides more granular security controls against unauthorized firewall use. Use FWUA for user-level accounting information for firewall users.

FWUA extends and enforces user authentication on traffic between branch office (BO) tunnels. You can also apply FWUA on nontunneled traffic when the VPN Router acts as a router and firewall edge device.

FWUA uses the existing authentication services, with username and passwords supported for both internal authentication services, for example, Lightweight Directory Access Protocol (LDAP). FWUA can also use external authentication services, for example, Remote Authentication Dial-in User Service (RADIUS) or LDAP proxy. [“FWUA example” on page 114](#) shows authentication by internal LDAP and [“FWUA configuration” on page 116](#) shows authentication by an external service (RADIUS and LDAP proxy).

FWUA by SecurID extends the authentication approach of FWUA, which enforces user authentication on traffic between branch office connections in the VPN environment. This authentication method applies to nontunneled traffic FWUA when the VPN Router acts as a router and a firewall edge device.


FWUA with TunnelGuard extends the capabilities of FWUA by downloading the TunnelGuard applet after the user authenticates. How you configure TunnelGuard determines if TunnelGuard verifies that, for example, the PC uses the proper patches and runs antivirus software before it grants access to the network.

For more information about FWUA with TunnelGuard, see *Nortel VPN Router Configuration — Tunnel Guard* (NN46110-307).

Policies within the VPN Router can contain a user authentication specification for a rule. Users must register an active Hyper Text Transfer Protocol Secure (HTTPS) logon session with the User Authentication Table Manager (UATM) before they gain the access granted by the rule. Users without an existing logon session registered with the UATM are not granted access even if the rule explicitly permits the traffic profile. User UATM sessions are mapped to the active session table by source IP address.

The following figure shows an example of FWUA.

Figure 33 FWUA example



HTTPS support provides a secured communication channel for administration traffic to the VPN Router system and for firewall users to provide their authentication credentials to the VPN Router Stateful Firewall. A FWUA user directs their HTTPS-enabled Web browser to a specific Uniform Resource Locator (URL) designated for the FWUA logon on the VPN Router. The VPN Router supports both Secure Socket Layer (SSL) 2.0/3.0 and Transport Layer Security (TLS) 1.0.

The following suites are supported:

- Symmetric Ciphers—RC4, Data Encryption Standard (DES), and Triple DES, Cipher Block Chaining (CBC)
- Public Key Cryptography and Key Agreement Protocols—RSA and Diffie-Hellman
- Authentication Codes and Hash Algorithms—Message Digest 5 (MD5) and Secure Hash Algorithm (SHA)-1

The following combinations of ciphers, key agreement protocols, and hashing algorithms are available:

- EDH-RSA-DES-CBC3-SHA
- DES-CBC3-SHA
- RC4-SHA
- RC4-MD5
- EXP1024-RC4-SHA
- EXP1024-DES-CBC-SHA
- EXP1024-RC4-MD5
- EDH-RSA-DES-CBC-SHA
- DES-CBC-SHA
- EXP-EDH-RSA-DES-CBC-SHA
- EXP-EDH-DSS-DES-CBC-SHA
- EXP-DES-CBC-SHA


The authentication facilities for FWUA use the existing authentication services available on the VPN Router except RADIUS-based tokens and digital certificates. All user-level accounting mechanisms available for VPN users are also available for FWUA users.

The following list identifies prerequisites to FWUA configuration:

- The VPN Router runs the VPN Router Stateful Firewall.
- Enable Secure Socket Layer/Transport Security (SSL/TLS), which requires that you install a valid digital certificate to support HTTPS communication.
- FWUA users must use an HTTPS-enabled Web browser with a compatible SSL/TLS crypto suite.

The following figure shows an example of FWUA configuration.

Figure 34 FWUA configuration



To configure FWUA

- 1 Choose **Services, Available**.
The Services window appears.
- 2 Click **Public** and **Private** for **Firewall User Authentication**.
- 3 Click **OK**.
- 4 Choose **Services, FWUA**.
The Firewall UA Settings window appears.

- 5 Type the text for a welcome banner, the port value (default 8000), and the default maximum session value. You can add RADIUS or LDAP proxy authentication servers to the authentication order later.
- 6 Select **TunnelGuard Checking Only** to enable FWUA for TunnelGuard enforcement only, which removes the need for the user to log on to FWUA. If you select this option, you must provide a user ID and password for the user. Use this user name and password to anonymously log on all FWUA users. For more information about TunnelGuard, see *Nortel VPN Router Configuration — TunnelGuard* (NN46110-307).
- 7 Click **OK**.
- 8 Choose **Services, SSL/TLS**.
The SSL window appears.
- 9 Select the desired **Ciphers** (default all) and select an X.509 digital server certificate for this VPN Router (for example, CN=ces48, O=CSE, C=US). If no certificates appear in the list, you did not define server certificates on your VPN Router, or the server certificate is disabled.
- 10 Click **OK**.
- 11 Choose **Profiles, Users**, and then click **Edit** for a user.
The User Management > Edit User window appears.
- 12 Select a group.
- 13 In the **Firewall User Authentication** boxes, type the user name and password.
- 14 Click **OK**.
- 15 Choose **Services, Firewall/NAT**, and then in the VPN Router Stateful Firewall row, click **Manage Policies**.



Note: The firewall UI requires Java 2 Runtime Environment (JRE) 1.6.0_u6 or later. If you do not use a sufficient JRE you are prompted by the VPN Router to download and install JRE 1.6.0_u6 directly from the VPN Router. A copy of JRE 1.6.0_u6 is also available on the VPN Router server CD.

- a Click **New** and type the name of the policy.

- b** Click the **Default Rules** tab, right-click on the number sign (#), and then select **Add New Rule**.
- c** Right-click the **Action** cell, and then select **User Authentication**.
- d** Select the group that contains the FWUA user. If you select ***any** for the group, it forces all users, regardless of their group association to authenticate to the firewall. Click **OK**.
- e** Choose **Policy, Save Policy**. Click **OK**.
- f** Choose **Manager, Exit SFw/Nat**. Click **Yes** to exit and **Yes** to save.
- g** Verify **VPN Router Stateful Firewall** on the **Firewall/NAT** window is enabled.
- h** Select the new firewall policy (refresh the screen for the new policy to appear in the list), and click **OK**.



Note: You must install a valid VPN Router Stateful Firewall license key. Also you must restart the VPN Router the first time you enable the VPN Router Stateful Firewall. You can disable the VPN Router tunnel filters as they are no longer needed.

To test the FWUA rule, try to communicate through the VPN Router. Communication attempts fail.

- 16** Direct your HTTPS enabled browser to the predefined FWUA logon URL on the VPN Router and log on to the firewall using the FWUA user profile that you created. The FWUA logon URL follows the format of `https://VPNRouterhostname:port/FWUA.htm` or `https://VPNRouterIPAddress:port/FWUA.htm` where `VPNRouterhostname` or `VPNRouterIPAddress` resolves to a VPN Router interface (not management IP). The port is the port number you specify on the **Services, FWUA** window.



Note: If the domain VPN Router digital server certificate is not part of a certificate domain trusted by your Web browser (you do not use a certificate issued by the same CA), or the domain listed on the VPN Router certificate does not match the DNS domain of the VPN Router, your Web browser prompts you with a security alert dialog box. Click **Yes** to trust the certificate and proceed.

After a successful authentication, the browser window must remain open during the entire time that you want to communicate through the firewall. This keeps an active FWUA session in the UATM.

- 17 Try to communicate through the firewall again. Communication attempts succeed.
- 18 To modify the current FWUA configuration to accommodate external authentication methods, go to **Services, FWUA, Add RADIUS** or **Add LDAP Proxy Server**. The Associated Group specifies the group the RADIUS or LDAP proxy authentication users obtain their privileges from as defined on the Server, RADIUS Auth or the Server, LDAP Proxy windows. If you configure the /Base group to authenticate RADIUS or LDAP Proxy Auth users for VPN connections, it also authenticates FWUA users.

Chapter 6

QoS configuration

Quality of Service (QoS) on the VPN Router supports several features such as Admission Control, Bandwidth Management, Call Admission Priority, Forwarding Priority, and Differentiated Services (DiffServ) features like behavior aggregate classifiers, multifield classifiers and traffic conditioning. In addition, the VPN Router uses Resource Reservation Protocol (RSVP) signals to the public network to reserve a portion of the network bandwidth for a specific connection.

You must purchase and install the Advanced Routing license before you can configure and use certain QoS features. The following table indicates the Advanced Routing license requirements.

Table 4 Advanced Routing license requirements

Action	License required
Assigning Call Admission Priority to the user or branch office group parameter	No
Enabling DiffServ Code Point 802.1p mapping	No
Assigning Forwarding Priority to the user or branch office group parameter	No
Globally enabling Bandwidth Management	Yes
Globally enabling Admission Control	Yes
Applying classifiers to an interface	Yes
Changing the Egress Queueing Mode from Legacy to Per-Hop Behavior	Yes
Enabling interface shaping	Yes
Enabling traffic conditioning	Yes

This chapter includes the following topics:

- “Admission control” on page 122
- “Bandwidth Management” on page 124
- “Call Admission Priority” on page 125
- “Forwarding Priority” on page 127
- “NNSC queues” on page 128
- “Differentiated Services” on page 133
- “Traffic conditioning” on page 146
- “RSVP” on page 150

Admission control

Admission control is a traffic control feature that determines if the gateway can supply the requested bandwidth and CPU resources for a new session while it continues to provide the bandwidth and CPU resources for existing sessions. Use Admission Control with bandwidth policies to limit the number of concurrent user and branch office tunnels.

Globally enabling Admission Control

To enable or disable Admission Control globally

- 1 Choose **QoS, Admission Control**.

The Admission Control window appears.

- 2 Select either **Enabled** or **Disabled**.
- 3 Click **OK**.

You can over subscribe an interface to work with Admission Control to enforce bandwidth policies. Use this type of configuration to ensure that existing or high-priority connections do not lose bandwidth to subsequent user or branch office tunnels. This potentially denies a tunnel connection request before the number of tunnel licenses reaches the limit. Configure an over-subscription ratio

to control the total guaranteed bandwidth on the interface. Configure Bandwidth Management policies for user and branch office tunnels on an individual group basis. If you enable Bandwidth Management and Admission Control, the router enforces the bandwidth policies.

To calculate the available bandwidth, multiply the interface committed rate for the tunnel traffic by the over-subscription ratio. For example, if the tunnel traffic committed rate is 9 Mb/s and the over-subscription ratio is 1, then the total available bandwidth that can be used for Admission Control is 9 Mb (9 Mb/s x 1). If the used bandwidth is 5 Mb/s, only 4 Mb/s of available bandwidth remains for the next tunnel connection requests. If the available bandwidth is not enough, the status of the tunnel connection request is pending or denied.

The VPN Router checks the committed rate of the incoming user or branch office tunnel against the available bandwidth. If no available bandwidth remains and new user tunnels try to connect, the router refuses the tunnel connections. Unlike user tunnel connections, branch office connections establish but do not activate (pending) until the specified bandwidth becomes available. After a user tunnel logs off, the router gives the bandwidth to the branch office. If the user attempts to log on again, the router refuses the connection on the assumption that no available bandwidth remains, or the available bandwidth is not enough to service the connection.

Use the total available bandwidth when you assign a bandwidth policy to group of users or branch office tunnels. If the total available bandwidth is 9 Mb/s and the expected tunnel connections is 500, the guaranteed bandwidth for each connection is 18 Kb/s (9 Mb/s / 500). Use 18 Kb/s for the committed rate.

Adjust the over-subscription ratio to adjust the total available bandwidth. In the preceding example, an over-subscription ratio of 1:1 provides 9 Mb/s total available bandwidth. If the over-subscription ratio is 10:1 (default), the total available bandwidth will be 90 Mb/s. This bandwidth provides a high committed rate for x number of tunnels. For 500 tunnels and a total available bandwidth of 90 Mb/s, this provides 180 Kb/s of committed bandwidth (90 Mb/s / 500), which is 10 times higher than the committed rate for 1:1 over-subscription ratio. Use this approach to allocate unused portions of bandwidth to other sessions for optimum network efficiency.

To configure over subscription

1 Choose **QoS, Interfaces**.

The QoS Interfaces window appears.

- 2 Under **Admission Control**, click **Configure**.
- 3 Type a value in the **Over-Subscription Ratio** box.
- 4 Click **OK**.

Over-subscription example

An administrator configures two branch office tunnels in two different groups. Tunnel A is part of a group that uses a 10 Mb bandwidth policy, and Tunnel B is part of a group that uses a 5 Mb bandwidth policy. The public interface is a 10 Mb interface with a 1:1 over-subscription ratio. The administrator enables Admission Control and Bandwidth Management. If Tunnel A connects to the router, Tunnel B cannot connect or use bandwidth until the Tunnel A connection drops.

Bandwidth Management

Bandwidth management forces tunnels to conform to a set of rates. Two rates (committed and excess) and two excess actions (mark or drop) exist. Packets use different drop preferences, depending on whether they are lower than the committed rate (lowest drop preference), between the committed and excess rate (higher drop preference), and higher than the excess rate (highest drop preference if excess action is Mark). After congestion occurs, the VPN Router drops packets according to their drop preference. If the excess action is Drop, the VPN Router drops all the packets higher than the excess action.

Configuring Bandwidth Management

Use Bandwidth Management to manage the VPN Router CPU and interface bandwidth resources to ensure that tunneled sessions get predictable and adequate levels of service. You use Bandwidth Management to configure the VPN Router resources for users, branch offices, and interface-routed traffic. Bandwidth components keep track of and control the level of bandwidth used on the physical interfaces and the tunnels.

You can use Admission Control to guarantee that resources are available to support the committed bandwidth assigned to a user. This potentially denies a client access before the router reaches the license limit. The VPN Router interface speed determines the available bandwidth.

To configure Bandwidth Management

- 1 Choose **Admin, License Keys**, and then type the key for the Advanced Routing license.
- 2 Click **OK**.
- 3 Choose **QoS, Bandwidth Mgmt**.

The Bandwidth Management window appears. You must define bandwidth rates in bits per second (bps). For example, 10 Mb/s equals 10 000 000 bps.

- 4 Choose **Profiles, Groups**, and then click **Edit** for the group.
- 5 From the **Connectivity** section, click **Configure**.
- 6 In the **User Bandwidth Policy** section, click **Configure**, and then define the committed and excess bandwidth rates.
- 7 Click **OK**.
- 8 From **QoS, Bandwidth Mgmt**, enable **Bandwidth Management**.
- 9 Click **OK**.
- 10 Choose **QoS, Interfaces** to configure the over-subscription rate. Use this ratio to adjust for some users that do not use all of their allotted bandwidth simultaneously under normal circumstance. The default is 10:1.

Call Admission Priority

Use Call Admission Priority to assign each user group profile to one of four priority classes (from 1—high to 4—low) for Call Admission. The VPN Router reserves connections for each class of user, which guarantees that a large number of low-priority users do not lock out the high-priority users. The VPN Router does not accept further low-priority connections if it services the maximum number of low-priority sessions. After the router accepts a connection, it never drops the connection.

Because the VPN Router supports a maximum number of sessions, it is important to assign users to the proper priority classes. This configuration ensures that connections are available to the appropriate users during periods of heavy traffic. Although other callers are permitted access to the VPN Router, this access is proportional to the assigned priority level for their group.

By default, the router admits access for any call for the first 50 percent of connections, regardless of the assigned priority. The next 25 percent of calls guarantee access to only priority 1, 2, and 3 callers. The next 15 percent of calls guarantee access to only priority 1 and 2 callers. For the final 10 percent of calls, only priority 1 callers are guaranteed access.

For example, assume a hypothetical maximum of 2000 sessions, “[Call admission priority](#)” on page 126 shows the connections available for each priority based on a percentage of the total capacity.

Table 5 Call admission priority

Capacity	Priority	Available connections
0–50%	All	1000
51–75%	1, 2, 3	500
76–90%	1, 2	300
91–100%	1	200

The following figure shows the maximum number of connections available for each priority.

Table 6 Maximum connections for each priority

Priority	Connections
1	2000
2	1800
3	1500
4	1000

Forwarding Priority

Use Forwarding Priority to assign each user to one of four priority classes. Forwarding Priority guarantees each class different maximum forwarding times between the interfaces of the VPN Router. For example, Forwarding Priority protects high-priority traffic generated by the company CEO from high-bandwidth traffic generated by lower-priority users. Or, you can assign the sales team to priority 1 to ensure they can always place orders, especially during the quarter-end rush.

The technology that supports forwarding priority is weighted fair queuing with random early detection (RED). This queuing mechanism gives each of the four user classes (from 1—high to 4—low) a different weight in the amount of service time they receive by the packet-forwarding process. Each class, however, is guaranteed some level of service so that no traffic through the VPN Router is ever completely stalled. Assign users to the four different class levels to make sure they receive the proper service and performance, especially during heavily congested times. QoS is only effective when all associated lines can service the forwarding demands at the required speeds.

If a group profile uses a forwarding priority of 1 (highest), it receives the highest possible bandwidth guarantee and the lowest level of latency. Packets sent by this group transmit immediately even during heavy traffic flows on the VPN Router. Conversely, if a group profile uses a forwarding priority of 4 (lowest), it receives the least amount of bandwidth allocated and possibly the highest level of latency. Therefore, the VPN Router experiences heavy traffic, fewer packets sent by this group transmit when higher-level priority packets exist in the queue.

To illustrate how Forwarding Priority works, the example in [“Bandwidth allocation for each priority level” on page 127](#) assumes heavy traffic and a queue of packets. Packets transmit according to the approximate rates for each pass that are cited in the following table.

Table 7 Bandwidth allocation for each priority level

Priority 1	Priority 2	Priority 3	Priority 4
60% pass	25% pass	10% pass	5% pass

Of the total packets transmitted in a hypothetical pass, 60 percent come from the Priority 1 queue; 25 percent from the Priority 2 queue; 10 percent from the Priority 3 queue; and 5 percent from the Priority 4 queue.

NNSC queues

Nortel Networks Service Classes (NNSC) define a set of service class behaviors that can help construct various levels of QoS-aware services. Think of NNSCs as default QoS policies built into the VPN Router. You can match the behavior provided by the NNSC to the traffic type for which you apply QoS. Then configure the various devices in the network to place specific traffic into the appropriate NNSC that provides the closest performance behavior required for that service.

Nortel uses a standard default set of eight NNSCs that provide eight levels of service class behaviors. The router associates incoming traffic with an NNSC based on a DiffServ Code Point (DSCP) to NNSC mapping function. The VPN Router supports the following seven NNSCs queues.

The NNSCs are

- Critical and Network
- Premium
- Platinum
- Gold
- Silver
- Bronze
- Standard

Critical and Network service classes

The Critical NNSC services traffic within a single administrative network domain. If such traffic does not get through, the network cannot function. Examples of such types of traffic are heartbeats between core network switches or routers.

The Network service class includes network control traffic, for example, Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), and routing table updates.

Premium service class

The Premium NNSC services traffic that is marked as EF or CS5, predominantly interactive IP telephony services, and provides the low latency required to support such services. For more information about the mapping between NNSC queues and DSCP values, see [“NNSC to DSCP mapping” on page 136](#).

Metal service classes

The Platinum, Gold, Silver, and Bronze NNSCs are collectively referred to as the metal classes. The metal NNSCs provide a minimum bandwidth guarantee and are used for variable bit rate. Normally, applications that use the metal NNSCs support mechanisms that dynamically adjust their transmit rate and burst size based on congestion (packet loss) detected from the network.

The Platinum NNSC services applications that require low latency; for example, real-time services such as video conferencing and interactive gaming. The Platinum NNSC provides a minimum bandwidth assurance for flows marked Assured Forwarding—AF4x- and Class Selector—CS4. In a congestion situation, DiffServ nodes use drop precedence to control variable bit rates that exceed the minimum assured bandwidth.

The Gold NNSC services applications that require near-real-time service and are not as delay sensitive as applications used in the Platinum service. Such applications include streaming audio and video, video (movies) on demand, and surveillance video.

In general, the Gold NNSC is based on the assumption that the traffic is buffered at the source or destination and, therefore, the traffic is less sensitive to delay and jitter. As a default, the Gold NNSC provides a minimum bandwidth assurance for AF31-, AF32-, AF33-, and CS3-marked flows. Under congestion, DiffServ nodes use drop precedence to control variable bit rates and burst sizes that exceed the minimum assured bandwidth.

The Silver NNSC services responsive (typically client and server-based) applications. Such applications include a Systems Network Architecture (SNA) terminal (PC or Automatic Teller Machine) to mainframe (host) transactions that use Data Link Switching (DLSw—SNA over IP), Web-based ordering and credit card processing, financial wire transfers, and Enterprise Resource Planning (ERP) applications such as SAP and BaaN.

Silver NNSC applications are those that require a fast response. Typically, Silver NNSC applications need asymmetrical bandwidth. In other words, the client typically sends a short message to the server, and the server responds with a large data flow back to the client. The most common example occurs after a user clicks a hyperlink (a few dozen bytes) on a Web page, and a new Web page appears (KB of data). The Silver NNSC provides a minimum bandwidth assurance for AF2x- and CS2-marked flows.

The Silver NNSC favors short-lived low bandwidth TCP-based flows. Under congestion, DiffServ nodes use drop precedence to control variable bit rates and burst sizes that exceed the minimum assured bandwidth.

The Bronze NNSC services long-lived, TCP-based flows, such as file transfers, Telnet sessions, e-mail, or noncritical Operations and Maintenance (OAM) traffic. The Bronze NNSC provides a minimum bandwidth assurance for AF1x- and CS1-marked flows and favors long-lived, high bandwidth Transport Control Protocol (TCP)-like flows. Under congestion, DiffServ nodes use drop precedence to control variable bit rates and burst sizes that exceed the minimum assured bandwidth. Nortel recommends that you use the Bronze NNSC for noncritical operations, administration, maintenance, and provisioning (OAMP) traffic with Nortel recommended CS1 DSCP marking.

Standard service class

Best effort services use the Standard NNSC. This NNSC does not specify delay, loss, or jitter.

Queuing mechanisms

Queuing mechanisms protect queuing memory if the queue becomes congested by intelligently discarding packets before they enter the queue. Packet discard ensures that individual queues do not grow to unreasonable lengths. The VPN Router uses specific queuing mechanisms to service queues at egress, which also manages queue congestion.

The following table provides details about the queue configuration if you configure the Egress Queuing mode to Diffserv Per-Hop Behavior (PHB).



Note: Ensure that you configure the data rate on the WAN interfaces to match the actual line rate (choose System, WAN, Configure) to correctly calculate the PHB queues

Table 8 Queue configuration

NNSC	Queue	Queuing mechanism	Bandwidth allocated
Critical and Network	1	Strict priority	15% of total bandwidth
Premium	2	Strict priority	15% of total bandwidth
Platinum	3	Weighted fair queuing	60% of remaining bandwidth after priority queues
Gold	4	Weighted fair queuing	24% of remaining bandwidth after priority queues
Silver	5	Weighted fair queuing	10% of remaining bandwidth after priority queues
Bronze	6	Weighted fair queuing	4% of remaining bandwidth after priority queues
Standard	7	Weighted fair queuing	2% of remaining bandwidth after priority queues

You can override the default bandwidth allocation for the Premium queue if you enable and configure a rate for Egress Expedited Forwarding Shaping Rate under DiffServ Edge. For more information, see [“Configuring DiffServ” on page 141](#).

After traffic packets flow to the interface, they are taken from the head of the queues for transmission. The router selects which interfaces to service according to a weighted round-robin (WRR) algorithm based on the speed of the interface.

Seven active egress queues exist for each physical interface. Each time a device driver requests a packet for transmission, the router serves the network control and EF queues (strict priority queues) before the weighted queues. PHB queuing applies only on the egress of an interface. You can configure and enforce the EF egress (outbound) shaping rate with the PHB queue. The queue delays nonconforming EF traffic, but does not drop it.

Weighted fair queuing

In a weighted fair queuing (WFQ) scenario, the VPN Router uses the allocated bandwidth to determine how often the queues are serviced.

Strict priority

Strict priority ensures that specific queues are treated as priority queues. Packets in the priority queues transmit before packets in the lower queues. On the VPN Router, the Critical and Premium queues use a defined portion of the egress line rate.

Congestion avoidance

All queues use Multi-level Random Early Detection (MRED). MRED is an extension of the original RED algorithm.

RED detects congestion as it begins by computing the average queue size. After the average queue size exceeds a preset threshold, packets are dropped probabilistically. Some of the advantages of RED are

- RED keeps the average queue size low while allowing an occasional burst of packets in the queue.
- RED does not process bursty traffic differently than nonbursty traffic.
- RED helps avoid global synchronization, which occurs with Tail Drop gateways when multiple TCP connections drop packets at approximately the same time. The TCP connections reduce their congestion windows at the same time, which results in lowered link utilization of the congested link.

Differentiated Services

DiffServ settings classify and mark packets to receive specified per-hop forwarding behavior on each node along their path. Diffserv settings also implement sophisticated classification, marking, policing, and shaping operations at network boundaries or hosts and allocate resources to traffic streams by service provisioning policies that govern how the router marks and conditions traffic on entry to a differentiated services-capable network, and how that traffic is forwarded within that network. DSCPs not recognized are forwarded as if marked for the default behavior, Best Effort (BE).



Note: You can activate either DiffServ or Legacy Forwarding Priority as the egress queuing mode for an interface. You cannot activate both at the same time.

The following table defines common DiffServ terms and concepts.

Table 9 DiffServ terms and concepts


Term	Definition
DS boundary or access point	The DS boundary is the edge of a DS domain where classifiers and traffic conditioners are likely deployed.
DS field	The DiffServ field was formerly the IPv4 Type of Service octet or the IPv6 Traffic Class octet. The first six bits of the DS field are called the DSCP, and the value of the DSCP determines the PHB.
Microflow	A microflow is a single instance of an application-to-application flow of packets, identified by source address, destination address, protocol ID, and source port.
Marking	Marking is the process of configuring the DSCP in a packet based on defined rules.
PHB	The PHB is the forwarding treatment a DiffServ node applies to a packet in a DiffServ network.
Policing	Policing ensures that a traffic stream performs in accordance with the domain service provisioning policy or service level agreement.
Remarking	Remarking changes the DSCP of a packet, usually in accordance with a service level agreement.

Table 9 DiffServ terms and concepts

Term	Definition
Service level agreement	A service level agreement is a service contract that specifies the forwarding service that traffic must receive.
Shaping	Shaping delays the nonconforming packets to bring the stream into compliance with the agreed-on traffic profiles. A shaper is a strong form of policing; it ensures that excessive packets are not allowed into the network.
Traffic profile	A traffic profile represents the temporal properties of a traffic stream such as rate.

After traffic enters the DiffServ network, the router places packets in a queue according to their marking, which in turn determines the PHB of that packet. For example, if a video stream is marked so that it receives the highest priority, it is placed in a high-priority queue. As these packets traverse the DiffServ network, the video stream forwards before other packets.

To differentiate between classes of service, the DiffServ field in the IP packet header, as defined in RFC 2474 and RFC 2475, is marked. The Differentiated Services (DS) field in the IP header is an octet, and the first six bits, the DSCP, are used in the DiffServ architecture. The following figure illustrates the DSCP position in the IP header.

Figure 35 IPv4 packet header with DSCP field

Two standard PHBs are defined in RFC 2597 and RFC 3246: the Assured Forwarding (AF) PHB group and the Expedited Forwarding (EF) PHB group. The VPN Router also uses the default (DF) and Class Selector (CS) groups as defined in RFC 2474. A DiffServ network uses CS for backward compatibility with IP precedence; the CS group supports legacy routers. Best-effort services use the DF group.

IP precedence classifies packets in a non-DiffServ network. Whereas the DSCP uses the first six bits of the TOS field, IP precedence uses the first three. Use IP precedence to provide eight possible precedence values (CS0 to CS7). IP precedence is appropriate for simple QoS solutions. If you require more customization, configure a custom DSCP mapping for a DiffServ network.

Assured Forwarding PHB group

RFC 2597 describes the Assured Forwarding PHB group, which further divides the delivery of IP packets into four independent classes. The Assured Forwarding PHB group offers different levels of forwarding resources in each DiffServ node. Within each Assured Forwarding PHB group, IP packets are marked with one of three possible drop precedence values. In case of network congestion, the drop precedence of a packet determines its relative importance within the Assured Forwarding group.

Both real-time and nonreal-time bursty traffic use the AF group.

Expedited Forwarding PHB group

RFC 3246 describes the Expedited Forwarding PHB group as the Premium service: the best service your network can offer. Expedited Forwarding PHB is a forwarding treatment for a DiffServ microflow when the rate of flow transmission ensures that it is the highest priority and experiences no packet loss for in-profile traffic.

Voice services use the EF group. The EF group emulates a leased line, and EF packets preempt lower priority traffic.

The following table provides the mapping between NNSC queues and the corresponding DSCP.

Table 10 NNSC to DSCP mapping

NNSC	Queue	DSCP
Critical and Network	1	CS7, CS6
Premium	2	EF, CS5
Platinum	3	AF4, CS4
Gold	4	AF3, CS3
Silver	5	AF2, CS2
Bronze	6	AF1, CS1
Standard	7	DF, CS0

The following table summarizes the defined 21 PHBs and their codepoint mappings.

Table 11 DSCP mapping to PHB

DSCP	Defined PHB
111 000	CS7
110 000	CS6
101 000	CS5
101 110	EF
100 000	CS4

Table 11 DSCP mapping to PHB

DSCP	Defined PHB
100 010	AF41
100 100	AF42
100 110	AF43
011 000	CS3
011 010	AF31
011 100	AF32
011 110	AF33
010 000	CS2
010 010	AF21
010 100	AF22
010 110	AF23
001 000	CS1
001 010	AF11
001 100	AF12
001 110	AF13
000 0000	CS0

Classifier configuration

Traffic classification includes any function where a packet is examined to determine what further action to take according to defined rules. Classification identifies the flow to which a particular packet belongs so that the router can possibly modify the packet contents or PHB, apply conditioning treatments to the packet, and determine how to forward the packet to the egress interface. The VPN Router supports two types of classifiers: multifield (MF) classifier and behavior aggregate (BA) classifier.

A multifield classifier selects packets based on the content of some header fields, typically a combination of source address, destination address, DS field, protocol ID, source port, and destination port. Based on the selection, a marker configures the DSCP in the packet. Configure rules to specify how the router classifies

packets based on the contents of the header fields. The router examines, classifies, and marks all incoming packets for each session and physically routed traffic based on the rules you configure. You can enforce PHB by using multifield classifiers.


The BA classifier categorizes traffic based solely on the existing DSCP value in the packet. The classifier uses the DSCP value to apply a PHB to the packet if it exits an interface and the interface uses Diffserv Per-Hop Behavior as the egress queuing mode.

The router forwards packets that contain an unrecognized or unsupported DSCP as default best-effort as if the packet is marked for the default forwarding behavior. The default marking of DSCP is best-effort. If you do not define or apply a rule, the packet is marked as best-effort and forwarded to the best-effort queue.


You can define an MF classifier for an interface. After you create a classifier, you apply it to an interface and enable the Multi-Field Classifier option under Diffserv Edge. The interface MF classifier applies to routing traffic that passes through that interface.

When traffic flows through a VPN Router, it is subject to two MF Classifiers: ingress MF classifier on the ingress interface and egress MF classifier on the egress interface. To use MF classifiers for packet DSCP marking, you need a clear picture of which interfaces traffic flows through and understand which MF classifier governs the packet marking result. The following list identifies rules to help you decide how to use MF classifiers.

- You can enable or disable ingress and egress MF classifiers on an individual interface basis.
- If you disable MF classifiers, the DiffServ marking process does not occur on the interface.
- If you enable MF classifiers, but do not apply a rule, the interface marks the packets as BE.
- If you enable both ingress and egress MF classifiers, the ingress MF classifier governs the marking result. The following figure shows the packets in traffic flow 1 from PC 1 to PC 2 are marked as BE (x00), while packets in traffic Flow 2 from PC 2 to PC1 are marked as EF (xB8).

Figure 36 Ingress classifier marks packets

- If you disable ingress MF classifiers, egress MF classifiers control the DSCP marking results. The following figure shows the packets in traffic flow 1 from PC 2 to PC 1 are marked as AF3(x68) because ingress MFC is disabled on interface 2. The egress MF classifier on interface 1 controls the marking.

Figure 37 Egress classifier marks packets

Configuring an MF classifier

To configure an MF classifier:

- 1 Choose **QoS, Classifiers**.
- 2 On the right pane, type new classifier name and click **Create**.

The Classifiers----> Edit Classifier screen appears.

- 3 Click on **Manage Rules** and then click **Create** to create a rule.
- 4 Type the Rule Name, select the options based on the criteria to filter the packet for the Source Address, Destination Address, Protocol, TCP/UDP Source Port, TCP/UDP Destination Port, Current DSCP Value or Diffserv Marking options, and click **OK**.

The Classifiers---> Edit--->Rules screen appears.

- 5 Click **Close** to exit from the Rules screen.

The Classifiers---> Edit Classifier screen appears.

- 6 Select the rule displayed in the Available Rules window that was created, click << button to add the rule to the classifier, and click **OK**.
- 7 Click **Update Classifier** to apply the new classifier to an interface.

Using a BA classifier and the current DSCP

To use a BA classifier and honor the current DSCP

- 1 Choose **QoS, Interfaces**.
The QoS Interfaces window appears.
- 2 Click **Configure** for **DiffServ Edge**.

- 3 Ensure that MF classifiers are disabled on the ingress and egress interfaces.



Note: If they disable MF classifiers on the ingress interface but have it enabled on the egress interface, then NVR will not honor the BA classifier.

- 4 Click **OK**.
- 5 Click **Configure** for **Egress (Outbound) Queuing Mode**.
- 6 Select **DiffServ Per-Hop Behavior**.
- 7 Click **OK**.

Configuring DiffServ

To configure DiffServ

- 1 Choose **QoS, Interfaces**, and then click **Configure** in the **DiffServ Edge** section.

The QoS Interfaces > DiffServ Edge window appears.

- 2 In the **Multi-Field Classifier State** list, select **Enabled** or **Disabled**.
- 3 In the **Ingress** (Inbound) list, select the MF classifier to apply.
- 4 In the **Egress** (Outbound) list, select the MF classifier to apply.
- 5 In the **Traffic Conditioning State** list, select **Enabled** or **Disabled**.

Traffic conditioning drops and remarks a traffic stream to shape it into compliance with a traffic metering profile. For Expedited Forwarding (EF) and Assured Forwarding 1 to Assured Forwarding 4 (AF1 to AF4), you can configure a Traffic Conditioning Meter (in bps).

- For EF, the rate is an average rate, although at times traffic can burst as much as twice the configured rate. Traffic lower than the rate forwards; traffic higher than the rate drops.

- For AF1 to AF4, packets under the rate are marked as low drop precedence. Packets under two times the configured rate are marked as medium drop precedence. Packets higher than two times the configured rate are marked as high drop precedence.



Note: Enter values for EF and AF1 to AF4 greater than 512 bps. Traffic conditioning does not work with configured rates smaller than 512 bps or with packets smaller than 64 bytes.

- 6 Type a value, in bps, in the **Expedited Forwarding (EF) Rates** box. Nonconforming traffic drops.
- 7 Type values, in bps, for the Assured Forwarding rates (AF4—AF1).
- 8 Select the **Excess Action** for each AF rate to either drop traffic that exceeds the configured rate or to mark the traffic.
- 9 Type a value, in bps, in the **Best Effort Forwarding Rate** box. Nonconforming traffic drops.
- 10 For Egress (Outbound) traffic conditioning, type a value, in bps, for **Expedited Forwarding Shaping Rate**. Shaping delays the packets in a stream to conform to a defined traffic profile (the EF Shaping value). Nonconforming traffic delays; it does not drop.

You must disable Anti-Replay if you use IPsec tunnels over LANs or WANs (the typical usage). DiffServ sorting is incorrect if you enable Anti-Replay. Anti-Replay does not acknowledge DiffServ and uses different methods to discard packets, which adversely affects the DiffServ sorting.

DSCP to 802.1p mapping

The 802.1p specification prioritizes network traffic at the data link layer. The 802.1p uses the User Priority field of the 802.1Q header. This priority extension tags Ethernet frames with 1 of 8 different classes of service to provide service differentiation at the Ethernet layer. The 802.1p to DSCP markings are static and are set according to the Nortel standard.

DiffServ provides QoS at the IP level by redefining the 8 bit Type of Service (TOS) field of the IPv4 header Type of Service field as a DS field.


The DSCP uses six bits of the DS field to select the PHB a packet experiences at each node. The DSCP identifies the priority of service a packet receives in the network.

The VPN Router uses DSCP to 802.1p mapping to tag frames for prioritization over public and private physical interfaces. The router supports mapping DSCP to 802.1p marking on ingress to or egress from the VPN Router and can separately enable or disable 802.1p to DSCP mapping on ingress or egress.

The 802.1p tag often does not remain with the packet as it travels from source to destination. However, the DSCP marker in an IP header does remain with the packet. Although some Ethernet switches cannot interpret the DSCP, they can interpret the 802.1p tag. Provide a consistent mapping between DSCP and 802.1p so that Ethernet networks achieve the required end-to-end QoS behavior.

In “[Example 802.1p to DSCP mapping](#)” on page 143, the Layer 2 switches are DSCP-unaware and the Layer 3 switch and router are DSCP-aware. If a packet that transmits from the Layer 2 switch to the router uses the 802.1p tag as it enters the Layer 3 switch, the Layer 3 switch performs a 802.1p to DSCP mapping and forwards the packet to the router. After the router sends a packet back to one of the DSCP-unaware switches, the Layer 3 switch performs a DSCP to 802.1p mapping and forwards a packet to the Layer 2 switch.

Figure 38 Example 802.1p to DSCP mapping



After you enable the mappings and the router receives an incoming packet with an 802.1p marking, the VPN Router uses the default 802.1p to DSCP mappings shown in the following table.

Table 12 Default incoming 802.1p mappings

802.1p user priority	Maps to DSCP
7	CS7
6	EF
5	AF41
4	AF31
3	AF21
2	AF11
1	DF
0	DF

If you enable the mappings and a packet exits the router, VPN Router uses the default DSCP to 802.1p mappings shown in the following table.

Table 13 Default outgoing 802.1p mappings

DSCP	Maps to 802.1p user priority
CS7	7
CS6	7
EF, CS5	6
AF41, AF42, AF43, CS4	5
AF31, AF32, AF33, CS3	4
AF21, AF22, AD23, CS2	3
AF11, AF12, AF13, CS1	2
DF, CS0, All undefined DSCPs	0

If you disable the mappings, the router ignores the 802.1p tag value and applies normal multifield classifier (MFC) action to all packets.

Configuring DSCP to 802.1p mapping

To configure DSCP to 802.1p mapping:

1 Choose **QoS, Interfaces**.

The QoS Interfaces window appears.

2 From the **Current Interface** list, select the interface to which you want to apply the mappings.

3 Click **Display** to display the selected interface (Fast Ethernet appears by default).

4 In the **DSCP 802.1p Mapping** section, click **Configure**.

5 Select either **Custom** or **Standard** for the Egress (outbound) and for Ingress (inbound).

6 If you select the **Custom** setting, click **configure custom mappings**.

7 Configure the **DSCP Class to 802.1p precedence mapping** and the **802.1 precedence to DSCP mapping** sections. For more information, see [“Default incoming 802.1p mappings” on page 144](#) and [“Default outgoing 802.1p mappings” on page 144](#).

8 Click **OK**.

Router-generated packets

The VPN Router must process various types of network traffic with various degrees of delay, loss, and jitter requirements. Network traffic characteristics range from network critical traffic such as routing updates or loss-sensitive streaming audio and video to best-effort traffic like file transfers and e-mail.

The router manages the behavior of traffic flows based on DSCP markings in the IP packet. With Release 4.75 and later, the VPN Router includes DSCP markings in specific router-generated packets. The following table identifies the type of packets and their markings.

Table 14 DSCP markings in router-generated packets

Packet type	DSCP marking
Internet Security Association and Key Management Protocol keepalives	CS7
OSPF Hello	CS7
OSPF link state advertisement updates, LS acknowledge	CS6
Routing Information Protocol	CS6
VRRP	CS6
Dynamic Host Configuration Protocol	AF21
DNS	AF21
Internet Control Message Protocol requests	AF21
Network Time Protocol	AF21

Point-to-Point Protocol (PPP) control packets and Frame Relay (FR) Local Management Interface (LMI) packets transmit out the highest priority egress queue on the router, regardless of the configured egress queuing mode.

Traffic conditioning

QoS provides the option to drop data that exceeds configured AF rates. This option provides guaranteed bandwidth based on Diffserv code points that guarantees a fixed percentage of total bandwidth to each of several applications. The VPN Router conditions traffic based on the bytes in the packet, excluding the Data Link header, that arrives on the ingress interface. For example, if a 200 byte packet enters the router, the router conditions the traffic based on 186 bytes—Layer 3 and above as the router removes the Data Link Layer (14 bytes).

Traffic conditioning by DSCP provides a method to limit traffic at ingress to the VPN Router based on DSCP value. This method ensures that particular DSCP values obtain the desired amount of outbound bandwidth. Traffic that exceeds the configured rate for a particular DSCP drops at inbound to the VPN Router. You can configure traffic conditioning for the following DSCP values:

- EF, inbound and outbound
- AF4, inbound only
- AF3, inbound only
- AF2, inbound only
- AF1, inbound only
- Best Effort, inbound only

Traffic conditioning drops or remarks a traffic stream to shape it into compliance with a traffic metering profile. For Expedited Forwarding (EF) and Assured Forwarding 1 to Assured Forwarding 4 (AF1 to AF4), you can configure a Traffic Conditioning Meter in bps.

For EF, the rate is an average rate, although at times traffic can burst as much as twice the configured rate. Traffic below the rate forwards; traffic above the rate drops.

For AF1 to AF4, packets under the rate are marked as low drop precedence. Packets under two times the configured rate are marked as medium drop precedence. Packets above two times the configured rate are marked as high drop precedence.

You can configure the assured forwarding queues option to drop data that exceeds the configured rate. (EF excess data always drops.) This data drops on ingress and never enters a queue. If the configured data rates for the assured forwarding queues are based on the interface shaping rate, which is based on the downstream data rate, the queues are the appropriate size.

You can show traffic conditioning statistics on the interface, however, traffic conditioning for the outbound interface does not show under the traffic conditioning statistics. Only ingress statistics show under traffic conditioning statistics.

EF outbound traffic conditioning

You can configure egress traffic conditioning (shaping) for EF traffic only. The queue delays nonconforming traffic, but it does not drop or down-mark the traffic.

Configure the EF shaping rate on egress (outbound) to determine the EF queue size on an interface. You cannot directly configure the queues used by AF and Best-Effort traffic.

An example of when to configure the EF shaping rate on egress is if you pass voice traffic through the router. By default, if you choose DiffServ Per-Hop Behavior as the egress queuing method for an outbound interface, EF traffic receives a Per-Hop Behavior 15 percent of the interface. However, you can enable traffic conditioning on the outbound interface and configure the egress outbound expedited forward rate to the interface speed to override the bandwidth limitation.

Configuring traffic conditioning

To configure traffic conditioning

- 1 Choose **QoS, Interfaces**, then and click **Configure** in the **DiffServ Edge** section.

The QoS Interfaces > DiffServ Edge window appears.

- 2 In the **Traffic Conditioning State** list, select **Enabled** or **Disabled**.



Note: Enter values for EF and AF1 to AF4 greater than 512 bps. Traffic conditioning does not work with configured rates smaller than 512 bps or with packets smaller than 64 bytes.

- 3 Type a value, in bps, for the **Expedited Forwarding (EF) Rate**. Nonconforming traffic drops.
- 4 Enter values, in bps, for the Assured Forwarding rates (AF4 to AF1).
- 5 Configure the **Excess Action** list for each AF rate to either drop traffic that exceeds the configured rate or to mark the traffic.
- 6 For Egress (Outbound) traffic conditioning, enter a value, in bps, for **Expedited Forwarding Shaping Rate**. Shaping delays the packets in a

stream to conform to a defined traffic profile (the EF Shaping value). The queue delays nonconforming traffic, but it does not drop the traffic.

- 7 Click **OK**.

Configuring interface shaping

Interface shaping shapes or delays the outgoing packet flow through an interface to better match the throughput of a downstream device. The VPN Router supports interface shaping for Ethernet interfaces only and the shaping is not DiffServ aware. Interface shaping uses the egress port statistics for measurement.

Unlike EF egress shaping, interface shaping applies to all traffic regardless of protocols or markings. If you enable and configure both shaping features for an Ethernet interface, the smallest shaping rate controls the EF egress shaping.

To configure interface shaping

- 1 Choose **QoS, Interfaces**.

The QoS Interfaces window appears. The current interface shows the current QoS configuration, which includes interface shaping.

- 2 Under **Current Interface**, select the Ethernet interface to configure, and then click **Display**.

- 3 Under **Interface Shaping**, click **Configure**.

The QoS Interfaces > Interface Shaping window appears.

- 4 Under **Interface Shaping State**, select **Enabled**. The default value is disabled.

- 5 Type the shaping rate, in bps.

- 6 Click **OK**.

RSVP

The VPN Router supports Resource Reservation protocol (RSVP) QoS for the Internet. Successful external network-level QoS requires the cooperation of all the devices on the network (between the user and either the access point to the private network or the ultimate destination host). Currently, RSVP is the best-defined technology for resource reservation. However, only a few service providers offer a service that uses RSVP.

The VPN Router signals to the other devices on the public network and describes the level of bandwidth that it needs to ensure adequate performance. Both the data rate of the user connection to the Internet and the data rate of the link between the Internet and the VPN Router determines this amount of bandwidth.

The two key components of RSVP are

- PATH messages, which are constant announcements by the host system or the VPN Router that a certain amount of bandwidth must remain available.
- RESV messages, which are responses from the client that it wants to reserve the requested bandwidth.

If the client responds to the PATH messages with RESV messages, the RSVP-ready routers attempt the resource reservation. These routers actually reserve the resources requested if they are RSVP-compliant.

Index

A

- access control filters 26
- actions on rules 53, 92
- Admission Control
 - enforcing bandwidth policies 122
- anti-spoofing 25, 33
- application layer gateway 100
- attack detection 25
- available rules 62

B

- bandwidth management 124

C

- Call Admission
 - guarantees 125
 - Priority 125
- cell menu 49
- columns
 - Dst interface 50
 - Src interface 50
- configuration
 - initial 29
- conversation 23

D

- default rules 48
- Differentiated Services (DiffServ) 133
- DSCP to 802.1p mapping 143
- dynamic many-to-one 68

E

- egress (outbound) queueing mode 142

F

- filter
 - rules 24
- filters
 - copy 64
 - edit current 61
 - storing 64
- firewall
 - imbedded 21
 - installation prerequisites 30
 - integrated 21
 - options 33, 36
- Firewall Virtual ALG 102
- forwarding priority 127
 - quality of service 127
- FTP 63, 64

H

- header row menu 49

I

- ICMP rule enforcement 36
- implied rules 43
- installation prerequisites 30
- interface
 - classifiers 138
 - over-subscription ratio 123
- interfaces 24

interface-specific rules 47

J

Java 2 Runtime Environment
Internet Explorer 31

L

log
column 53
levels 53

logging
application-specific 37
HTTP 37
remote system 37

logging FTP 37

M

menus
cell 49
header row 49
row 49

N

NAPT 22, 68, 83

NAT 22
branch office 85
creating policies 92
double 72
dynamic routing protocol 88
interface NAT 87
IPsec-aware 73
pooled translation 69
port forwarding 71
statistics 110

NAT SIP ALG 101

NAT Traversal 78
branch office, configuring 81
user tunnel, configuring 80

Network Address Translation 67

Network Address Translation (NAT) 27
network objects 51, 91

NNSC
description 128
mapping to DSCP 136

O

override rules 46

P

policies
actions 40
adding 41, 93
components 39
copying 42, 94
creating 41
deleting 42, 94
editing 41
renaming 43, 94
selecting 41, 93

pooled translation type 88

port mapping 87

port translation (NAPT) 68

proxy ARP 110

publications
hard copy 15

Q

quality of service
forwarding priority 127
RSVP 150

queues
bandwidth allocation 131
NNSC 128

queuing
strict priority 132
weighted fair queuing 132

R

- remarks 54, 92
- remote system logging 37
- row menu 49
- RSVP quality of service 150
- rule column 49, 91
- rules
 - default 48
 - implied 43
 - in policies 24
 - interface-specific 47
 - navigating 43, 49, 90
 - override 46
- Rules in Set 62

S

- service objects 52, 91
- SNMP 63
- stateful inspection 23
 - application 23
 - TCP 23
- static address
 - NAT 70
- static translation type 87
- status 54, 92
- strict priority, see queuing 132
- syslog 37
- system requirements 30

T

- technical publications 15
- traffic conditioning 141

V

- VoIP 81

W

- weighted fair queuing, see queuing 132

