

Version 6.00

Part No. 315896-E Rev 00
August 2005

600 Technology Park Drive
Billerica, MA 01821-4130

Configuring Firewalls, Filters, NAT, and QoS for the Contivity Secure IP Services Gateway

NORTEL

Copyright © Nortel Limited 2005. All rights reserved.

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

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

Trademarks

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

Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.

AXENT and OmiGuard Defender are trademarks of AXENT Technologies, Inc.

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

Entrust and Entrust Authority are trademarks of Entrust Technologies, Incorporated.

Java and Solaris are trademarks of Sun Microsystems.

Linux and Linux FreeS/WAN are trademarks of Linus Torvalds.

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 and registered trademarks are the property of their respective owners.

The asterisk after a name denotes a trademarked item.

Restricted rights legend

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

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

Statement of conditions

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

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

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

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

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

Nortel Networks Inc. software license agreement

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

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

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

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

3. Limitation of Remedies. IN NO EVENT SHALL NORTEL NETWORKS OR ITS AGENTS OR SUPPLIERS BE LIABLE FOR ANY OF THE FOLLOWING: a) DAMAGES BASED ON ANY THIRD PARTY CLAIM; b) LOSS OF, OR DAMAGE TO, CUSTOMER’S RECORDS, FILES OR DATA; OR c) DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS OR SAVINGS), WHETHER IN CONTRACT, TORT OR OTHERWISE (INCLUDING NEGLIGENCE) ARISING OUT OF YOUR USE OF THE SOFTWARE, EVEN IF NORTEL NETWORKS, ITS AGENTS OR SUPPLIERS HAVE BEEN

ADVISED OF THEIR POSSIBILITY. The forgoing limitations of remedies also apply to any developer and/or supplier of the Software. Such developer and/or supplier is an intended beneficiary of this Section. Some jurisdictions do not allow these limitations or exclusions and, in such event, they may not apply.

4. General

- a. If Customer is the United States Government, the following paragraph shall apply: All Nortel Networks Software available under this License Agreement is commercial computer software and commercial computer software documentation and, in the event Software is licensed for or on behalf of the United States Government, the respective rights to the software and software documentation are governed by Nortel Networks standard commercial license in accordance with U.S. Federal Regulations at 48 C.F.R. Sections 12.212 (for non-DoD entities) and 48 C.F.R. 227.7202 (for DoD entities).
- b. Customer may terminate the license at any time. Nortel Networks may terminate the license if Customer fails to comply with the terms and conditions of this license. In either event, upon termination, Customer must either return the Software to Nortel Networks or certify its destruction.
- c. Customer is responsible for payment of any taxes, including personal property taxes, resulting from Customer's use of the Software. Customer agrees to comply with all applicable laws including all applicable export and import laws and regulations.
- d. Neither party may bring an action, regardless of form, more than two years after the cause of the action arose.
- e. The terms and conditions of this License Agreement form the complete and exclusive agreement between Customer and Nortel Networks.
- f. This License Agreement is governed by the laws of the country in which Customer acquires the Software. If the Software is acquired in the United States, then this License Agreement is governed by the laws of the state of New York.

Contents

Preface	13
Before you begin	13
Text conventions	13
Acronyms	15
Related publications	17
Hard-copy technical manuals	18
Chapter 1	
Overview of firewalls, filters and NAT	21
Contivity Stateful Firewall concepts	22
Stateful inspection	23
Interfaces	24
Filter rules	25
Anti-spoofing	25
Attack detection rules	25
Filters for access control	27
Network address translation	27
Chapter 2	
Configuring the Contivity Stateful Firewall	29
Configuring prerequisites	30
Installing Java 2 software	31
Using Internet Explorer	31
Using Netscape	32
Using Netscape 6	33
Using Netscape on Solaris	34
Enabling firewall options	35
Rule enforcement	37

Selecting logging options	37
Application-specific logging	38
Remote system logging	39
Configuring anti-spoofing	40
Configuring malicious scan detection	41
Setting up policies	43
Creating and editing firewall policies	43
Creating policies	44
Adding a policy	44
Deleting an existing policy	45
Copying an existing policy	45
Renaming an existing policy	45
Navigating rules	46
Implied rules	46
Override rules	48
Interface-specific rules	49
Default rules	51
Creating rules	51
Header row menu	51
Row menu	52
Cell menus	52
Rule columns	53
Creating a new policy	60
Verifying your configuration	61
Configuring a sample security policy	61
Firewall deployment examples	64
Residential firewall example	64
Business firewall example	66
Chapter 3	
Configuring filters	69
Adding and editing filters	69
Next hop traffic filters	73

Chapter 4	
Configuring NAT	77
Address translations	77
Dynamic many-to-one — port translation	78
Dynamic many-to-many — pooled translation	79
Static one-to-one translation	80
Port forwarding	81
Double NAT	82
IPsec-aware NAT	83
NAT modes	84
Full Cone NAT	84
Restricted Cone NAT	85
Port restricted Cone NAT	86
Symmetric NAT	87
NAT traversal	88
NAT and VoIP	89
Address/Port discovery	90
Network address port translation (NAPT)	91
Configuring Cone NAT	92
NAT Usage	94
Branch office tunnel NAT	95
Interface NAT	96
Dynamic routing protocols	97
Configuring NAT policy	99
NAT policy sets	99
Creating rules	100
Creating a new policy	105
Adding a policy	107
Deleting an existing policy	107
Copying an existing policy	107
Renaming an existing policy	108
Sample NAT procedures	108
Interface NAT with RIP	108
Interface NAT with OSPF	109
Branch Office NAT with RIP	109

Branch Office NAT with OSPF	110
Sample branch office NAT configuration	111
Configuring NAT with the Contivity Stateful Firewall	112
NAT ALG for SIP	112
Application level gateways (ALG)	114
Configuring NAT ALG for SIP	114
Firewall SIP ALG	115
Hairpinning	116
Hairpinning with SIP	116
Hairpinning with a UNISTIM call server	117
Hairpinning with a STUN server	119
Hairpinning requirements	120
Configuring hairpinning	121
Time-outs	122
NAT statistics	122
Proxy ARP	123
Chapter 5	
Configuring firewall user authentication	125
Chapter 6	
Configuring QoS	133
Configuring classifiers	133
Configuring bandwidth management	135
Configuring Differentiated Services (DiffServ)	136
Using forwarding priority	138
Using call admission priority	139
Using RSVP	140
DSCP to 802.1p mapping	141
Index	145

Figures

Figure 1	Security Warning screen	31
Figure 2	Download Java Runtime screen	32
Figure 3	TCP rule enforcement	37
Figure 4	FTP and HTTP logging	39
Figure 5	Syslog forwarding screen	40
Figure 6	Anti-spoofing configuration screen	41
Figure 7	Scan detection configuration screen	42
Figure 8	Firewall - Select Policy screen	44
Figure 9	Implied rules	47
Figure 10	Override rules	49
Figure 11	Interface-specific rules (Source rules)	50
Figure 12	Interface-specific rules (Destination rules)	50
Figure 13	Default rules	51
Figure 14	Row menu	52
Figure 15	Cell menu (option)	52
Figure 16	Cell menu (procedure)	53
Figure 17	Rule column	54
Figure 18	Rule column for interface-specific rules	54
Figure 19	Tunnel Selection box for a user tunnel	55
Figure 20	Tunnel Selection box for a branch office tunnel	55
Figure 21	Network Object Selection box	56
Figure 22	Network object edit box	56
Figure 23	Service Object Selection box	57
Figure 24	Tcp object insert box	58
Figure 25	Action column options	58
Figure 26	Log column options	59
Figure 27	Status column options	59
Figure 28	Firewall - Select Policy screen	62
Figure 29	Firewall - Edit Policy screen with new policy	63

Figure 30	Example of a basic residential firewall	64
Figure 31	Single override rule	65
Figure 32	Single interface specific rule	66
Figure 33	Business firewall	66
Figure 34	Business override rules	67
Figure 35	Edit Filters screen	70
Figure 36	Next hop filter action	74
Figure 37	Next hop forwarding	75
Figure 38	Port translation	79
Figure 39	Dynamic pooled address translation	80
Figure 40	Static address translation	81
Figure 41	Port forwarding example	82
Figure 42	Double NAT	83
Figure 43	IPsec-aware NAT example	83
Figure 44	Full Cone NAT	85
Figure 45	Restricted Cone NAT	86
Figure 46	Port Restricted Cone NAT	87
Figure 47	Symmetric NAT	88
Figure 48	STUN	91
Figure 49	Restricted Cone NAT — NAPT	92
Figure 50	Firewall/NAT page	93
Figure 51	Firewall/NAT Edit page	94
Figure 52	Overlapping address translation	96
Figure 53	Interface NAT	97
Figure 54	NAT with dynamic routing example	98
Figure 55	Network Object Selection box	101
Figure 56	Network object edit box	102
Figure 57	Service Object Selection box	102
Figure 58	Tcp object insert box	103
Figure 59	NAT actions	104
Figure 60	Status column options	104
Figure 61	NAT policy	106
Figure 62	Edit Policy screen	108
Figure 63	NAT configuration example	111
Figure 64	NAT and SIP	113

Figure 65	SIP enabled	115
Figure 66	Hairpinning with SIP	117
Figure 67	Intra-realm call with hairpinning	118
Figure 68	NAT Hairpinning	120
Figure 69	Hairpinning page	121
Figure 70	Proxy ARP example	123
Figure 71	FWUA example	126
Figure 72	FWUA configuration	128
Figure 73	Firewall policy screen	130
Figure 74	Firewall/Nat screen	131
Figure 75	Example 802.1p to DSCP mapping	142

Preface

This guide describes overview and configuration information for the Nortel * Contivity* Secure IP Services Gateway Contivity Stateful Firewall and Contivity filters.

Before you begin

This guide is for network managers who are responsible for setting up and configuring the Contivity Secure IP Services Gateway. This guide assumes that you have experience with windowing systems or graphical user interfaces (GUI) and familiarity with the 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 of the choices. 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.

Acronyms

This guide uses the following acronyms:

ACK	acknowledgement
ALG	application level gateway
BCM	business communications manager
FTP	File Transfer Protocol
H.323	ITU-T specification for multimedia over IP networks of non-guaranteed QoS
LAN	local area network
MCS	multimedia communications server
NAPT	network address port translation
NAT	network address translation
RTCP	RTP control protocol
RTP	real time transport protocol
SDP	session description protocol
SIP	session initiation protocol
STUN	simple traversal of UDP through NAT
TCP	transmission control protocol
UDP	user datagram protocol

VOIP	voice over IP
VPN	virtual private networks
WAN	wide area network

Related publications

For more information about the Contivity Secure IP Services Gateway, refer to the following publications:

- Release notes provide the latest information, including brief descriptions of the new features, problems fixed in this release, and known problems and workarounds.
- *Configuring Basic Features for the Contivity Secure IP Services Gateway* introduces the product and provides information about initial setup and configuration.
- *Configuring SSL VPN Services on the Contivity Secure IP Services Gateway* provides instructions for configuring services on the SSL VPN Module 1000, including authentication, networks, user groups, and portal links.
- *Configuring Servers, Authentication, and Certificates for the Contivity Secure IP Services Gateway* provides instructions for configuring authentication services and digital certificates.
- *Configuring Advanced Features for the Contivity Secure IP Services Gateway* provides instructions for configuring advanced LAN and WAN settings, PPP, frame relay, PPPoE, ADSL and ATM, T1CSU/DSU, dial services and demand services, DLSw, IPX, and SSL VPN.
- *Configuring Tunneling Protocols for the Contivity Secure IP Services Gateway* configuration information for the tunneling protocols IPsec, L2TP, PPTP, and L2F.
- *Configuring Routing for the Contivity Secure IP Services Gateway* provides instructions for configuring BGP, RIP, OSPF, and VRRP, as well as instructions for configuring ECMP, routing policy services, and client address redistribution (CAR).
- *Managing and Troubleshooting the Contivity Secure IP Services Gateway* provides information about system administrator tasks such as backup and recovery, file management, and upgrading software, and instructions for monitoring gateway status and performance. Also, provides troubleshooting information and inter operability considerations.
- *Reference for the Contivity Secure IP Services Gateway Command Line Interface* provides syntax, descriptions, and examples for the commands that you can use from the command line interface.
- *Configuring TunnelGuard for the Contivity Secure IP Services Gateway* provides information about configuring and using the TunnelGuard feature.

Hard-copy technical manuals

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

How to get Help

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

Getting Help from the Nortel Web site

The best source of support for Nortel products is the Nortel Support Web site:

<http://www.nortel.com/support>

This site enables customers to:

- download software and related tools
- download technical documents, release notes, and product bulletins
- sign up for automatic notification of new software and documentation
- search the Support Web site and Nortel Knowledge Base
- open and manage technical support cases

Getting Help over the phone from a Nortel Solutions Center

If you have a Nortel support contract and cannot find the information you require on the Nortel Support Web site, you can 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 Web site below and look up the phone number that applies in your region:

<http://www.nortel.com/callus>

When you speak to the phone agent, you can reference an Express Routing Code (ERC) to more quickly route your call to the appropriate support specialist. To locate the ERC for your product or service, go to:

<http://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, you can contact the technical support staff for that distributor or reseller.

Chapter 1

Overview of firewalls, filters and NAT

The Contivity Secure IP Services Gateway includes integrated firewall solutions that are designed to meet the needs of a variety of customers. The Contivity gateway provides the following firewall solutions: the Contivity Stateful Firewall and Contivity Filters.

With the Contivity Stateful Firewall, the Contivity gateway can perform a variety of secure routing functions, depending upon how you set up the Contivity gateway's routing capabilities. For example, you can configure the Contivity gateway to securely route non-tunneled traffic from its private interface, through the firewall, and out its public interface. This configuration enables users on the Contivity gateway's private network to access the Internet without requiring a separate, dedicated router. The Contivity Stateful Firewall achieves optimum performance as a result of advanced memory management techniques and optimized packet inspection.

The Contivity Stateful Firewall provides a high level of security, the fastest runtime, and the flexibility to define the rules to fit your environment. The firewall delivers full firewall capabilities, assuring the highest level of network security. To do this, the 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 firewall services determine how to process the packets based on the defined security policy.

The Contivity interface filters provide a cost-effective level of protection. The interface filters can only be disabled when the Contivity Stateful Firewall is enabled.

Because there are no routing protocols (such as RIP) running 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. PAT is disabled when either the Contivity Stateful Firewall or filter firewall is enabled as these provide better policy-based security.

When the firewall is disabled, PAT applies only to packets received on a public interface. PAT has a list of trusted sources that includes the remote client or branch office tunnel end point, remote Radius/CMP/CRL server address (if on the public side). PAT does not limit the packets from any of those trusted sources. For packets coming from any address that is not in the trusted source list, a rate limit (6 packet/10 second) is applied based on the source address.

The Contivity Stateful Firewall public address table information is not related to network address translation (NAT) or network address port translation (NAPT), which is often referred to as port address translation (PAT).

Contivity Stateful Firewall concepts

The Contivity Stateful Firewall provides a secure access point between an internal network and an external network, such as the Internet. The firewall allows you to protect your network and the information on it from unauthorized intrusion from external networks. The firewall 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. It monitors packets and sessions to make decisions based on established rules to determine 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 allow through a firewall securely using traditional filtering mechanisms. In FTP, for example, the control connection is typically created using a known port, but the data connection is over a random port. To allow an FTP data connection through a firewall without leaving a large number of open ports requires stateful inspection: packets are inspected at the application layer to determine which port the data connection is using. Traffic on that port can then be allowed to pass through the firewall for the duration of the FTP session.

Transport-level state inspection provides a number of ways to make TCP traffic more secure and more difficult for hackers to intercept. Stateful inspection of TCP consists of verifying the consistency of the TCP header as well as preventing some well-known TCP attacks. TCP sequence numbers are randomized to prevent sequence number guessing.

Stateful inspection of an application is unique for each application. Any non-predicted ports used by an application are validated and allowed through the firewall using stateful inspection. The following applications are inspected:

- FTP
- TFTP
- RCMD
- SQLNET
- VDOLive
- RealAudio

A *conversation* is created for all unique end-to-end communication. 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 Contivity Secure IP Services Gateway can have many interfaces. Each tunnel (end user or branch office) is a virtual interface, and all Contivity gateways have two or more physical interfaces. Packets can be classified by the interface on which they arrive at the Contivity gateway (the source interface) or the interface on which they leave the Contivity gateway (the destination interface).

The rules in a policy can be constructed 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.

The rules in any policy can use the following terms to designate an interface:

- Any — any physical interface or tunnel
- Trusted — any private physical interface or tunnel
- Untrusted — any public physical interface
- Tunnel:Any — any tunnel
- For tunnels, you can specify either a group name for user tunnels or the specific branch office tunnel for branch office tunnels
 - Tunnel:/base — you can specify the specific branch office tunnel. For example, /base/mktng/tony refers to branch office tony in group /base/mktng.
 - Tunnel:user — you can 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) screen. If the description is blank, the interface name defaults to the value of the Interface field on the same screen.

Any physical interface can be configured as private or public on the System > LAN > Interfaces screen. 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 a packet should be allowed 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 to the source address.
- Log the packet locally (these actions can be use with any of the previous three actions).

Anti-spoofing

Anti-spoofing is a method used to prevent a packet from forging its source IP address. Typically, the source address of each packet is examined and validated. Anti-spoofing 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

When a common attack is launched against corporate networks, the firewall should be able to detect these attacks. It should also drop any packets resulting from the attack, preventing denial-of-service as well as non-authorized intruders. The Contivity Stateful Firewall provides a defense against denial of service attacks with well-known prevention methods.

The Contivity Stateful Firewall protects against the following types of objects:

- Jolt2 is a fragmentation attack affecting Windows PCs by sending the same fragment repetitively.
- Linux* Blind Spoof attempts to establish a spoofed connection instead of sending final ACK with correct sequence number and with no flag set. Linux does not try to verify if the ACK is not set. The firewall drops any packet that does not have the ACK set.
- SYN flood can disable your network services by flooding them with connection requests. This fills the SYN queue, which maintains a list of unestablished incoming connections, forcing it not to accept additional connections.
- 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 used to establish a new connection and is sent from the same TCP source port as the destination port. When accepted by the target host, this packet causes a loop within the operating system, essentially locking the system.
- Ping of death sends a fragmented packet larger than 65536 bytes, causing the remote system to incorrectly process this packet. This causes the remote system to either reboot or panic during processing.
- Smurf sends a large number of ICMP echo (ping) messages to an IP broadcast address with the forged source address of the intended victim. The routing device forwarding traffic to those broadcast addresses performs the IP broadcast to layer 2 broadcast, causing most network hosts to take the ICMP echo request and issue a reply to each, multiplying the traffic by the number of hosts responding.
- Fraggle sends a large number of UDP echo messages. On a multi-access broadcast network, there could potentially be hundreds of machines to reply to each packet.
- ICMP unreachable sends ICMP unreachable packets from a spoofed address to a host, causing the host to stop all legitimate TCP connections to the host that is being spoofed in the ICMP packet.

- Data flood sends a large amount of data to a system that can be used as a denial of service attack, which exhausts available resources and stops response 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 Contivity gateway configuration, access control becomes an important security mechanism. You need complete control over which users have access to particular servers and services.

Filtering is the mechanism that fine tunes access to specific hosts and services. All users have 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

A filter profile consists of a list of rules that you create to perform precisely the action that you want. These rules are tested in order until the first match is found. Therefore, the order of the rules is very important. The filtering mechanism works such that if no rule matches a packet, the packet is discarded (denied); therefore no traffic is transmitted or received unless it is specifically permitted.

Network address translation

Network address translation (NAT) enables transparent routing between address spaces. Using NAT in an extranet allows the dynamic connection of multiple private networks via secure tunnels without requiring any address space reconfiguration.

Increasing use of NAT comes from two major factors:

- Shortage of IP addresses — Most Internet service providers (ISPs) allocate only one address to a single customer and this address is assigned dynamically, so every time a client connects to the ISP a different address is provided. Because users are given a single IP address, they can have only one computer connected to the Internet at a time. With NAT running on this single computer, it is possible to share that single address between multiple local computers and 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 were accessible.
- Security — NAT automatically provides security without any special set-up because it allows only connections that originate on the private network. It is still possible to 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, NAT is necessary to allow multiple intranets with conflicting subnets to communicate. Because the configuration of branch office or partner networks may be fixed, a VPN solution must be able to securely route between these networks without requiring all the private addresses to be unique across the entire extranet.

Chapter 2

Configuring the Contivity Stateful Firewall

To use the firewall on the Contivity Secure IP Services Gateway, you must install the license key and enable the firewall service. Without the firewall enabled, the Contivity gateway allows forwarding of 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)

When the firewall is enabled, the Contivity gateway additionally allows routing of traffic from public to private interfaces.



Note: Before you activate the firewall on the Services > Firewall/NAT screen, be sure to shut off all traffic to the Contivity gateway. This should be done off hours to prevent inconvenience to your users.

You must create rules for tunnel traffic before traffic on existing tunnels is allowed. The Contivity Stateful Firewall uses the principle that whatever traffic is not specifically allowed is disallowed. The rule set of the active policy is applied to all traffic, including tunneled and non-tunneled traffic. Therefore, when the Contivity Stateful Firewall is first enabled, all traffic is disallowed until you configure rules specifically allowing certain types of traffic.

Configuring prerequisites

Before you set up your Contivity Stateful Firewall, be sure you have the following information:

- The management IP address of your Contivity gateway. This address is found on the Contivity gateway's System > Identity screen.
- The firewall license key. Go to the Admin > Install Keys screen. Type the key that you obtained from Nortel Networks in the box to the right of Contivity Stateful Firewall and click on the Install button. It is only necessary to install a key once on each Contivity gateway. Click on the Delete button to remove the key.
- The name of the firewall, which is the name that is used by the Domain Name Service (DNS) server to identify the management address of the Contivity gateway. This name is entered in the DNS Host Name field of the Contivity gateway's System > Identity screen.
- The names and IP addresses of your Contivity gateway's interfaces. These are found on the Contivity gateway's Status > Statistics: Interfaces screen.

System requirements necessary to access the Contivity Stateful Firewall Manager:

- Supported operating systems and platforms include Solaris* (OS 2.6, 7, or 8) on an x86 or SPARC* platform and Microsoft Windows 95, 98, 2000, or Windows NT 4.0.
- Required software includes Java* 2 Plug-in Version 1.4.1_02, available in the Java 2 Runtime Environment Version 1.4.1_02. The J2RE is available for automatic download on a Windows platform for all Contivity gateways except the Contivity 1010, 1050 and 1100 (refer to the Java 2 Runtime Environment Installation). J2RE installation files for Windows and Solaris are also available on the Nortel Networks CD in the tools/java directory.
- Supported browsers include Internet Explorer 4 and higher (IE 5.01 service pack 2 is not supported) and Netscape Navigator* and Netscape Communicator* 4 and higher. Netscape 6 currently comes with a version of the Java 2 Plug-in that is not supported. If you wish to use Netscape 6, refer to the Netscape section of the Java 2 Runtime Environment Installation.

Installing Java 2 software

To access the Contivity Stateful Firewall Manager, the computer used to administer your Contivity must have the Java 2 Runtime Environment installed. There are two separate procedures that you can use to install the Java 2 software, depending on whether you use Internet Explorer or Netscape Navigator to access the Contivity.

Using Internet Explorer

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

- 1 Connect to the management IP address of the Contivity and log in.
- 2 Go to the Services > Firewall/NAT screen.
- 3 Click on the Manage Policies button. A popup window appears and tries to load the Contivity Stateful Firewall Manager.
- 4 When the Security Warning box appears, click on Yes to install the Java 2 Runtime Environment (Figure 1).

Figure 1 Security Warning screen



The installation program begins to download the software from the Contivity. (This is not available for the Contivity 1010, 1050, and 1100 hardware platforms.) This may take several minutes to load, depending on the speed of your connection to the Contivity.

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

Using Netscape

To install the Java 2 software on Windows 9x, Windows 2000, or Windows NT from Netscape Navigator:

- 1 Connect to the management IP address of the Contivity and log in.
- 2 Navigate to the Services > Firewall screen.
- 3 Click on the Manage Policies button. A popup window appears and tries to load the Contivity Stateful Firewall Manager. The Plug-in Not Loaded box appears. (If this box does not appear, click on the white or gray box that appears on the browser popup window.)
- 4 Click on the Get the Plug-in button to download the Java 2 Runtime Environment. The Java Plugin Download screen appears (Figure 2).

Figure 2 Download Java Runtime screen



- 5 Click on the Download now link next to the Windows version of the Java Runtime Environment.

- 6 When the browser prompts you for a location to save the file, choose a download location and click on OK to continue. (This may take several minutes to load, depending on the speed of your connection to the Contivity.)
- 7 When the download finishes, go to the download location and double-click on the icon for the Java Runtime Environment.
- 8 When the installation program displays the Software Licensing Agreement, click on Yes to accept the agreement.
- 9 When the installation program asks for an installation location, accept the default location or choose an alternate installation location.
- 10 Click on Next to finish the installation.
- 11 When the installation is complete, close all open Web browsers.
- 12 Reboot the computer for the changes to take effect.

Using Netscape 6

Netscape 6 currently includes a version of the Java 2 Plug-in that is not supported (Version 1.4.1_02). To successfully load the Contivity Stateful Firewall Manager, you must use Version 1.4.1_02. The following steps change the default plug-in to Version 1.4.1_02.

- 1 Install the Java 2 Runtime Environment as described in the previous Netscape section and be sure to restart the computer.
- 2 Load the Java Plug-in Properties from Start > Settings > Control Panel > Java Plug-in.
- 3 Click on the Advanced tab.
- 4 Choose JRE V 1.4.1_02 from the list.
- 5 Click on Apply.
- 6 Close the window.
- 7 Close all open instances of Netscape.
- 8 Restart Netscape. The correct plug-in should be available.

Using Netscape on Solaris

The Java 2 Runtime Environment for Solaris is available on the Nortel Networks CD. The installation files and instructions are available for x86 and SPARC platforms.

To install the Java 2 software on Solaris (OS 2.6, 7, 8) from Netscape Navigator:

- 1 Ensure that a version of Netscape is installed on the computer.
- 2 Close all instances of Netscape if any are opened.
- 3 Go to the tools/java/solaris directory on the Nortel Networks CD.
- 4 Choose the subdirectory for the installed platform, either intel for x86 or sparc for SPARC.
- 5 Copy the binary (.bin) and the README files to the computer.
- 6 Follow the platform-specific installation instructions contained in the README file.
- 7 Set the NPX_PLUGIN_PATH environment variable to the directory containing the javaplugin.so file.

For example, if the J2RE was installed in the /usr/j2re1.4.1_02 directory on a SPARC, the command to set the NPX_PLUGIN_PATH from the C shell would be:

```
setenv NPX_PLUGIN_PATH "/usr/j2re1.4.1_02/plugin/sparc"
```

- 8 Start Netscape and then close it.
- 9 Start Netscape again; the plug-in should now be available.

Enabling firewall options

To enable your Contivity firewall, select one of the following choices on the Services > Firewall/NAT screen and click on the OK button. You are then prompted to reboot your Contivity. You can select only one firewall choice at any one time.

- Contivity Firewall enables the Contivity Stateful Firewall feature. When you enable the Contivity Firewall, you can run any combination of the following:
 - Contivity Stateful Firewall
 - Contivity Interface Filter
 - Interface NAT
 - Anti-spoofing
- No Firewall disables all firewall features on the Contivity. In this configuration, the Contivity performs VPN routing only.

After you change your Firewall selection on the Services > Firewall/NAT screen and click on the OK button, you are prompted to confirm your selection. If you selected No Firewall, click on the OK button. If you selected the Contivity Firewall option, you must restart your Contivity gateway before the firewall becomes active. After you enable firewall support, you must configure the specified firewall.

The configuration procedures assume that you have already configured your Contivity Secure IP Services Gateway (except for the firewall component) and that you have obtained the required firewall license. You do not need a license for the Contivity Interface Filter.

To turn on and enable the Contivity Stateful Firewall:

- 1 Go to the System > LAN screen. For each interface, click on Configure and enter a label in the Description field. You use this name to identify 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 could make “Internet” the description for Slot 1 Interface 1 and “ServiceNet” the description for Slot 2 Interface 1. The description is case sensitive and cannot be abbreviated when specifying 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), is case sensitive, and cannot be abbreviated. The available slot numbers are hardware platform specific.

- 2 Go to the Services > Firewall/NAT screen.
- 3 Check the box next to Contivity Stateful Firewall.
- 4 On the system shutdown screen, click on OK and on the confirmation screen, click on OK to indicate the reboot.
- 5 After the Contivity reboots, return to the Services > Firewall/NAT screen.
- 6 Click on Manage policies to load the Contivity Stateful Firewall Manager applet. The first time you do this on any workstation, you need to load the Java applet. You see the message “Retrieving policies.”
- 7 Select the System Default policy, which is read-only.
- 8 Click on the View button to review this policy. The implied rules are included with every new policy.
- 9 You can toggle the browser windows between the Contivity Stateful Firewall Manager applet and the Services > Firewall/NAT screen. If you use your browser to change other settings on the Contivity while the Contivity Stateful Firewall Manager applet is running, these changes are not reflected in the current Contivity Stateful Firewall Manager applet. Click on the Firewall icon in the Contivity Stateful Firewall Manager applet to refresh the list of policies and other Contivity settings. Any changes made in the Contivity Stateful Firewall Manager applet are not evident in the Services > Firewall/NAT screen until you save the policy.
- 10 To exit the Contivity Stateful Firewall Manager, click on Manager > Exit.
- 11 After you exit the Contivity Stateful Firewall Manager applet, click on Refresh on the Services > Firewall/NAT screen.

The new policies you create are not automatically applied to the firewall. Only one policy at a time can be in effect on the firewall.

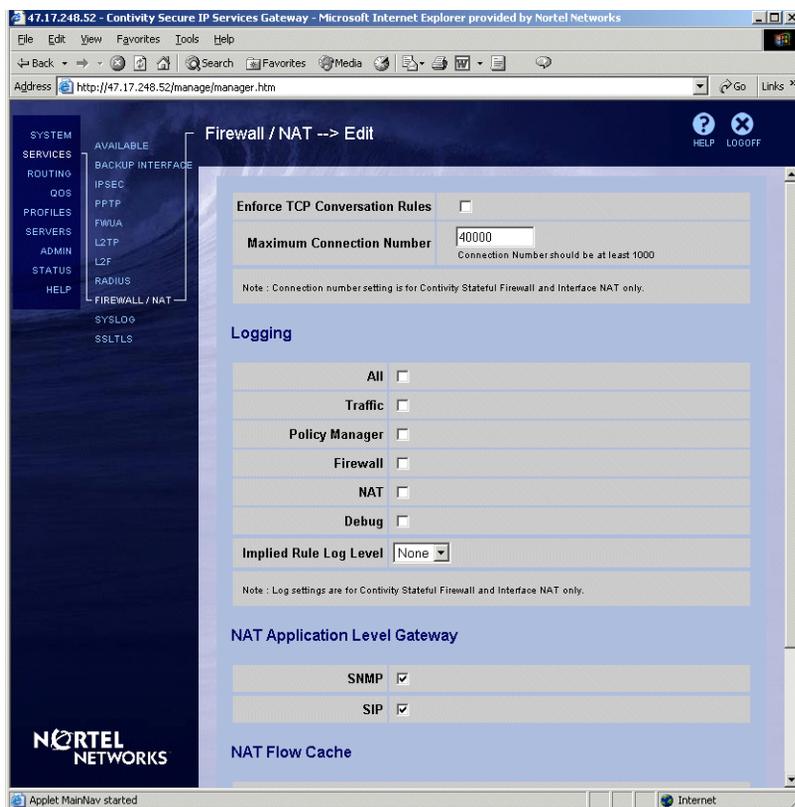


Note: You cannot import or export new policies. However, there are no restrictions on creating new policies.

Rule enforcement

ICMP can now be allowed or disallowed on public and private interfaces (Figure 3). To enable this, you must have a complete three-way handshake prior to the application of data.

Figure 3 TCP rule enforcement



Selecting logging options

The following options control the amount of firewall event information recorded in the event log. This information is not saved in the system log.

- All – includes traffic, policy manager, firewall, and NAT
- Traffic – logs when flows and conversations are created or removed

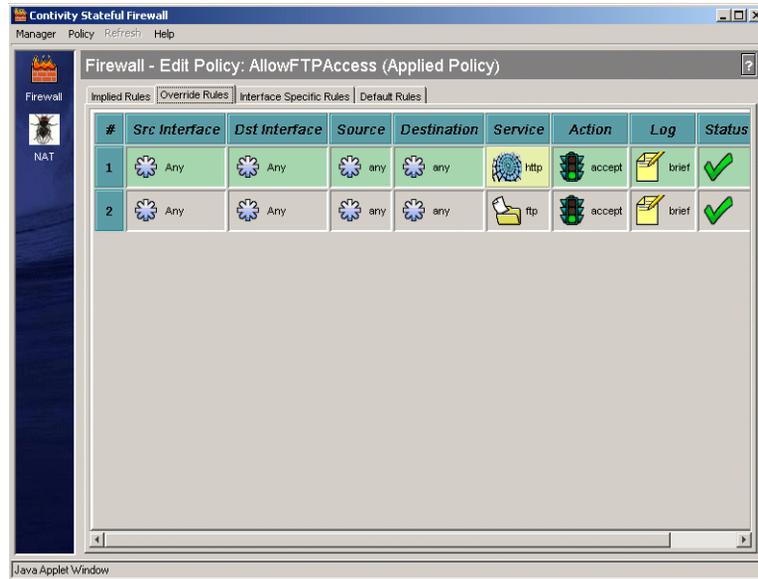
- Policy manager – logs firewall processes and when rules and policies are created
- Firewall – logs how the firewall handles packets within a flow
- NAT – logs NAT-related events
- Debug – creates special log messages intended for use only by Nortel Networks customer support personnel

Go to the Contivity Firewall > Edit screen to edit these options.

You can also set a maximum connection number, which allows you to reserve memory for a maximum number of connections. Determining the optimum memory allocation makes it easier to tune your system for firewall traffic. Under the Maximum Connection Number section, enter a number in the indicated range. The range displayed varies depending on the model and amount of memory for your Contivity. Each IPsec tunnel requires two connections. Nortel Networks recommends that you set the number near the middle of the range displayed unless you have specific requirements that you need to consider. You must reboot your Contivity gateway if you change the maximum connection number.

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 HTTP and FTP contain a unique connection identifier so events can be traced to the start and end of a TCP session. Configure the firewall rules to enable logging in either brief or detail format for rules with FTP and HTTP service, as shown in [Figure 4 on page 39](#).

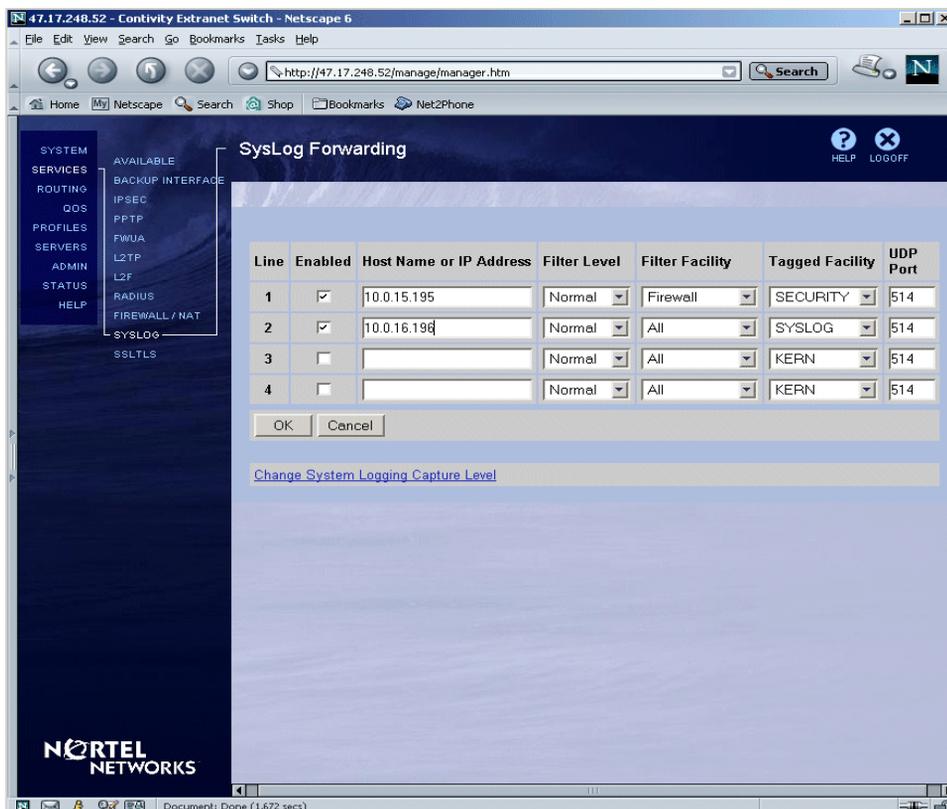
Figure 4 FTP and HTTP logging

Remote system logging

The Contivity syslog capability allows firewall-specific events to be forwarded to a remote syslog server. You can select whether to send all events to the remote syslog server or only firewall-specific events.

To configure remote syslog:

- 1 Enable firewall logging either by enabling it from the rules screen or from the Services > Firewall/NAT screen.
- 2 Configure a remote syslog server from the Services > Syslog screen, as shown in [Figure 5](#) on page 40.

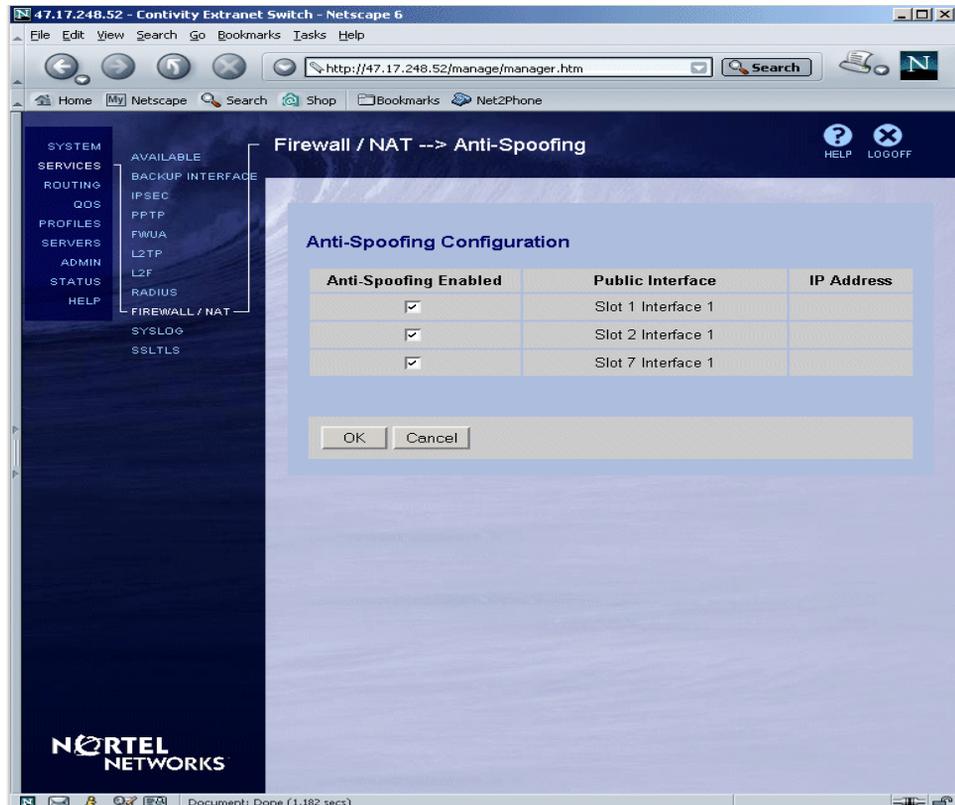
Figure 5 Syslog forwarding screen

- 3 Select Firewall for the Filter Facility.
- 4 Start syslog on the remote syslog system.
- 5 To verify that firewall-specific events appear on the remote syslog system, send traffic through the Contivity that generates firewall events.

Configuring anti-spoofing

To configure anti-spoofing:

- 1 Go to the Firewall/NAT screen and select the Anti-spoofing option.
- 2 Click on the Edit button. The anti-spoofing screen appears, as shown in [Figure 6 on page 41](#).

Figure 6 Anti-spoofing configuration screen

- 3 Click on the check box next to the public interface on which you want to enable anti-spoofing.
- 4 Click on OK.

Configuring malicious scan detection

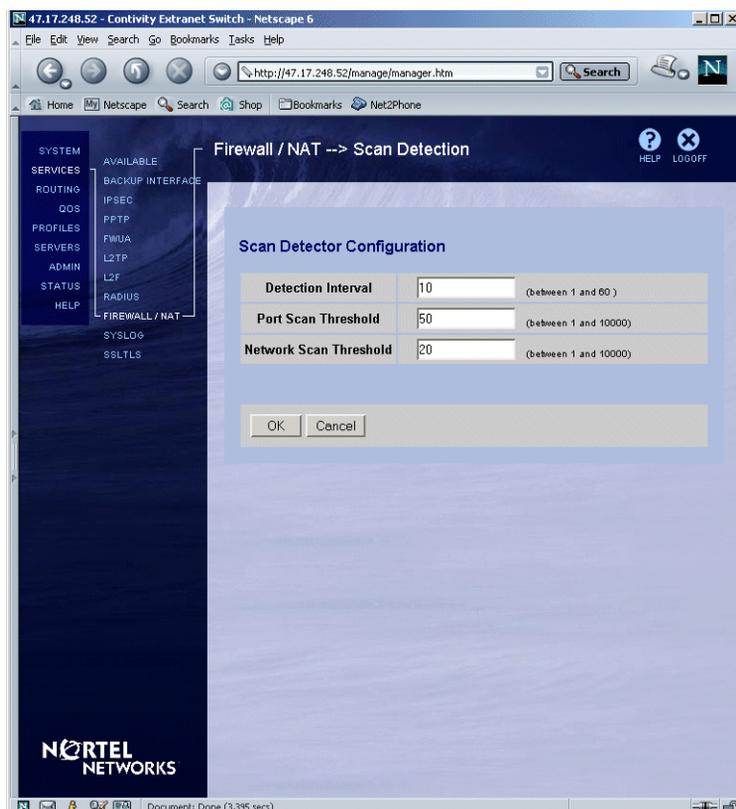
Scan detection detects port scanning attempts through the Contivity that are aimed at private resources.

To configure scan detection:

- 1 Go to the Services > Firewall/NAT screen and select the Malicious Scan Detection box.

- 2 Click on the Edit button. The Scan Detection screen appears (Figure 7).

Figure 7 Scan detection configuration screen



- 3 Specify the interval (1 through 60) over which the number of port scans or host scans will be inspected. If the number of scans exceeds the configured threshold during this interval, the scan is logged in the security log.
- 4 Specify the number of host-to-host connections (between 1 and 10000) 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.
- 5 Enter the number of one-to-many connections (between 1 and 10000) needed to trigger an event in the Network Scan Threshold field. This 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.
- 6 Click on OK.

Setting up policies

Firewall service consists of two primary components: the service properties and the security policy. Service properties define what service is 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 specify all rule fields for service policies using service objects. Each rule consists of a combination of network objects, services, actions, and logging mechanisms. You can define custom policies when more complex security policies are needed and the standard policies are not sufficient. By customizing your policies, you can further refine the control over what traffic is allowed 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 communication should be accepted or rejected (or logged) based on its source, destination, and service.

You must create rules for tunnel traffic before traffic on existing tunnel definitions will be allowed. The Contivity Stateful Firewall uses the principle that whatever traffic is not specifically allowed is disallowed. The rule set of the active policy is applied to all traffic, including tunneled and non-tunneled traffic. Therefore, when the Contivity Stateful Firewall is first enabled, all traffic is disallowed until you configure rules specifically allowing certain types of traffic.

Creating and editing firewall policies

Access control parameters are implemented through the graphical user interface or the command line interface (CLI). Using either interface, you can configure the following:

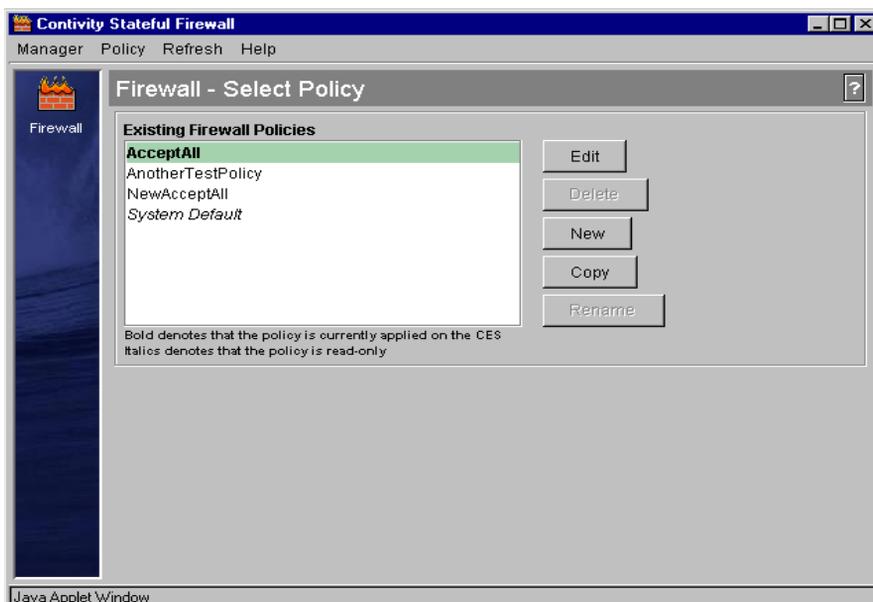
- Network objects
- Service objects
- Rules

See Reference for the *Contivity Secure IP Services Gateway Command Line Interface* for a list of commands you can use on the CLI.

Creating policies

The Firewall - Select Policy screen (Figure 8) allows you to create, edit, delete, copy, or rename a firewall policy. Bold denotes the policy that is currently applied to the Contivity Secure IP Services Gateway and italics denotes read-only policies. The System Default policy is always listed. This read-only policy defines the firewall behavior when no user-defined policies have been applied or when the selected policy is not available.

Figure 8 Firewall - Select Policy screen



Adding a policy

To add a new policy:

- 1 Click on the New button. The New Policy box appears and prompts you for a name for the new policy.
- 2 Enter the policy name. The name must begin with a letter and may not contain the : + =] , ; " characters.
- 3 Click on OK to go to the Policy Edit screen, which has a blank firewall policy, or click on Cancel to return to the policy selection screen.

Deleting an existing policy

You cannot delete a read-only policy or the policy that is currently applied to the Contivity. If you select one of these policies, the Delete button is not enabled. To delete an existing policy:

- 1 Select the policy that you want to delete and click on the Delete button. The delete policy confirmation box appears.
- 2 Click on OK to delete the selected policy.

Copying an existing policy

To copy a firewall policy:

- 1 Select the policy that you want to copy.
- 2 Click on the Copy button. The copy box appears.
- 3 Enter a name for the copied policy.
- 4 Click on OK.

The new policy appears in the list of policies in the firewall policies screen. This policy contains the same rules as the policy from which it was copied.

Renaming an existing policy

You cannot rename a read-only policy or the policy that is applied to the Contivity. If you select a read-only policy, the Rename button is not enabled. To rename an existing firewall policy:

- 1 Select the policy that you want to rename.
- 2 Click on the Rename button. The Rename box appears.
- 3 Enter the new name of the policy.
- 4 Click on OK.

Navigating rules

The firewall policy edit screen allows you to add, delete, and modify the rules within a policy. This screen is divided into the following rule groups:

- Implied rules
- Override rules
- Interface-specific rules
- Default rules



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

Implied rules

Implied rules are processed first by the firewall. These rules permit tunnel termination and access to the management interface. They are derived from the Services > Available screen and other configuration screens (such as RIP, OSPF, and VRRP). Some rules are statically generated, as shown in [Figure 9 on page 47](#), and are defined by the system and are read-only. They cannot be modified and are for display purposes only. Implied rules regulate traffic that originated from or terminated at the Contivity. You should control any routed traffic that is not directed to the Contivity with Override rules, Interface-specific or Default rules.

Figure 9 Implied rules

#	Src Interface	Dst Interface	Source	Destination	Service	Action
1	System	Trusted	_MngtIP	any	any	accept
2	Tunnel: Any	System	any	_MngtIP	any	accept
3	Any	CLIP	any	any	any	accept
4	Private	System:Private	any	any	l2xservicegroup pptpservicegroup icmpservicegroup ipsecservicegroup	accept
5	Private	System:Private	any	any	tcp8000	drop

Static pre-implied rules

The first rule in the implied rules section is the only rule that is statically generated. It always exists in the implied rules section regardless of the configuration. This rule allows the listed services to leave the Contivity on any of the private interfaces as long as the services are originated from the Contivity. Table 1 shows the server type and its corresponding configuration screens.

Table 1 Servers and corresponding configuration screens

Servers	Configuration Screen	Description
DHCP, DHCP-CLIENT	Servers > DHCP Relay	
DNS	System > Identity	
Remote-RPC		UDP port 17185
Nbdatagram, nbssession		Remote Netbios
Pptp	Service > Available	
IPSEC	Service > Available	
L2TP & L2F	Service > Available	
FWUA	Service > Available	

Table 1 Servers and corresponding configuration screens

Servers	Configuration Screen	Description
Radius	Service > Available	
HTTP, HTTPS	Service > Available	
SNMP	Service > Available	
FTP	Service > Available	
TELNET	Service > Available	
CRL	Service > Available	
CMP	Service > 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	

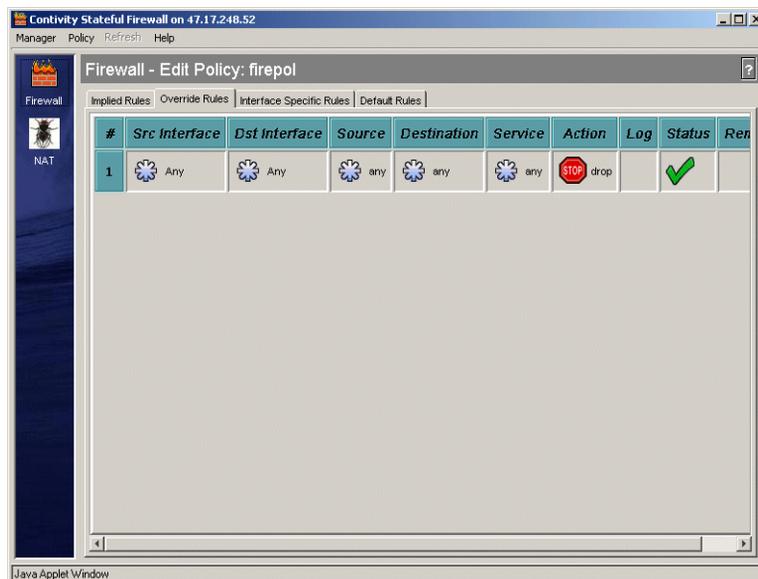
Dynamic implied rules

Dynamic implied rules are generated from all the available services on the Services > Available screen. Implied rules for ports that are not well known have a service name that consists of the protocol and the port number. For example, a tcp10 rule is generated from port numbers associated with external LDAP and RADIUS servers and configurable FWUA ports.

Override rules

Override rules, as shown in [Figure 10 on page 49](#), are the first set of modifiable rules in the policy. 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 debugging 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).

Figure 10 Override rules



Interface-specific rules

Interface-specific rules only apply to packets that enter or leave the Contivity through one specific interface (physical or tunnel). Interface-specific rules, as shown in [Figure 11 on page 50](#) and [Figure 12 on page 50](#), are divided into two rule types: source and destination. Source rules define the selected interface as the source and destination rules define the selected interface as the destination. Physical interface names correspond to the names configured on either the System > LAN or System > WAN screen. 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 only displays one interface at a time. To view all of the interface-specific rules, select All Interfaces.

Figure 11 Interface-specific rules (Source rules)

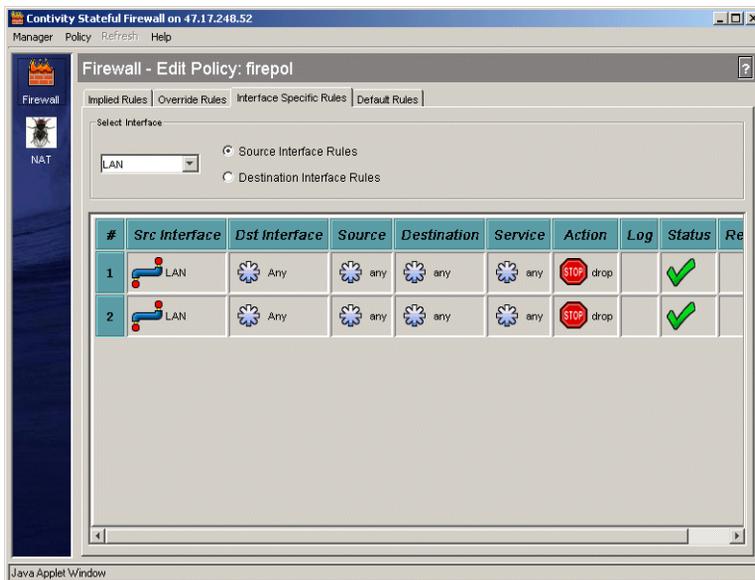
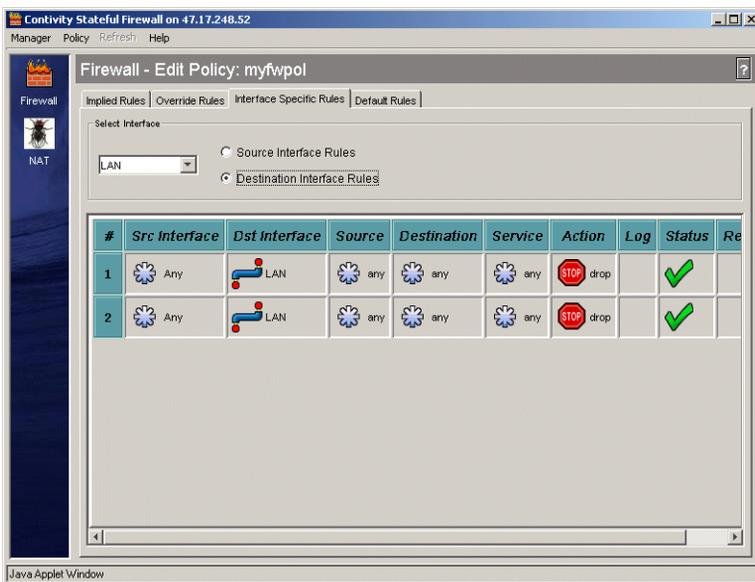


Figure 12 Interface-specific rules (Destination rules)



Default rules

Default rules (Figure 13) are applied to all traffic, but not restricted to a specific interface. These rules specify interface groupings for the source or destination (Any, Trusted, Untrusted, Tunnel:Any).

Figure 13 Default rules

The screenshot shows the 'Contivity Stateful Firewall on 47.17.248.52' configuration window. The title bar includes 'Manager Policy Refresh Help'. The main window is titled 'Firewall - Edit Policy: myfwpol' and has tabs for 'Implied Rules', 'Override Rules', 'Interface Specific Rules', and 'Default Rules'. The 'Default Rules' tab is active, displaying a table with the following data:

#	Src Interface	Dst Interface	Source	Destination	Service	Action	Log	Status	Ren
1	Any	Any	any	any	any	STOP drop		✓	
2	Any	Any	any	any	any	STOP drop		✓	

The interface also shows a sidebar with 'Firewall' and 'NAT' options, and a 'Java Applet Window' at the bottom.

Creating rules

Actions on rules are controlled by menus that you access by right-clicking on an option. Each of these menus controls a different aspect of the rule.

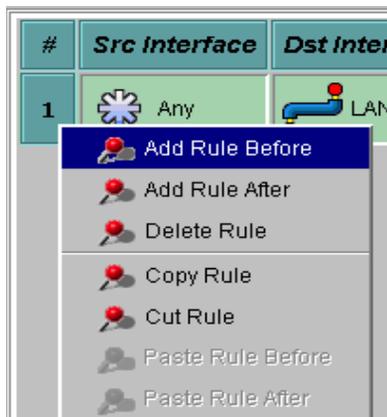
Header row menu

Right-clicking on any of the header cells brings up the Header row menu. This menu contains one item, Add New Rule. This menu item allows you 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-clicking on the number next to an existing rule activates the row menu (Figure 14). This menu allows you to add a new rule at a particular location, delete the specific rule, and perform cut/copy/paste operations on a rule.

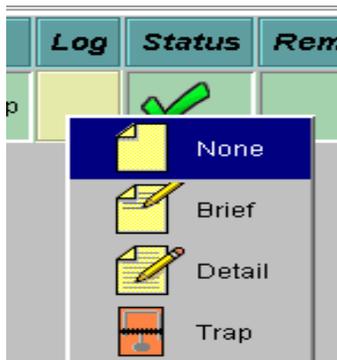
Figure 14 Row menu



Cell menus

Cell menus (Figure 15) are cell specific and accessed by right-clicking on an individual cell. There are two types of cell menus: option menus and procedure menus. Option menus provide a list of possible values for the cell. These menus are similar to a drop-down list box. When you click on one of the items, the selection is displayed in the cell.

Figure 15 Cell menu (option)



Procedure menus (Figure 16) provide a list of operations that you can perform on the cell, such as Add and Edit. When you click on one of the items, either the operation is performed immediately (such as Copy) or an additional box appears, prompting you for more information (such as Add).

Figure 16 Cell menu (procedure)



Rule columns

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

#

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

Src interface and Dst interface

These columns specify the source and destination interfaces for the rule. Right-clicking on the cell displays an option menu containing 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 may only be interface groupings.

These groupings are:

- Any — any physical interface or tunnel
- Trusted — any private physical interface or tunnel
- Untrusted — any public physical interface
- Tunnel:Any — any tunnel, excluding any physical interfaces

Figure 17 shows an example of a rule column.

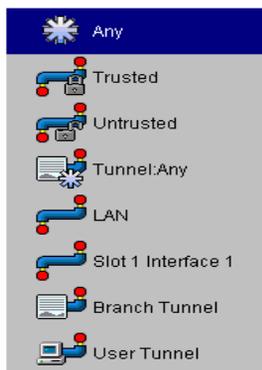
Figure 17 Rule column



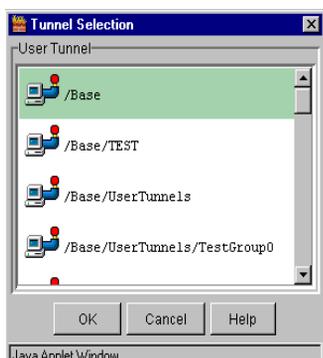
For the interface-specific rules, you can specify the interfaces as either groupings or individual interfaces.

Figure 18 shows a rule column for interface-specific rules.

Figure 18 Rule column for interface-specific rules



Clicking on the user tunnel, as shown in [Figure 19 on page 55](#), or branch office, as shown in [Figure 20 on page 55](#), menu items displays the tunnel selection box. This box allows you to select a specific tunnel (branch office or user tunnel).

Figure 19 Tunnel Selection box for a user tunnel**Figure 20** Tunnel Selection box for a branch office tunnel

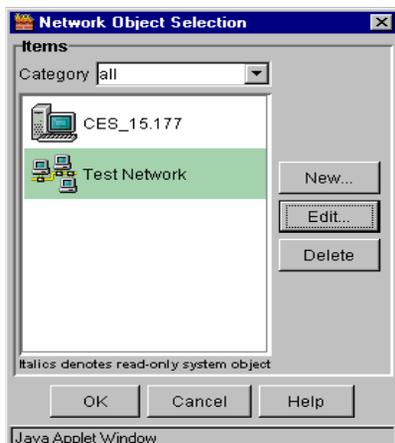
Source and Destination

These columns specify the source and destination network object for the rule. You can modify these attributes by right-clicking on a column in the cell, which then brings up a procedure menu. It is possible to add more than one source or destination address to a rule.

Clicking on Add displays the Network Object Selection box, as shown in [Figure 21 on page 56](#). In this 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).

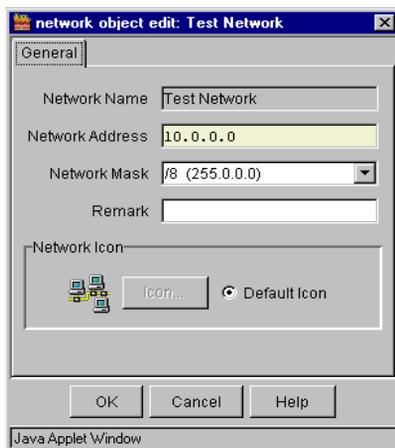


Note: The NOT operand allows you to specify which networks you do not want to be included.

Figure 21 Network Object Selection box

Italicized objects in the list are read-only and cannot be modified. The New, Edit, and Delete buttons in this box allow you to create, edit and delete network objects.

Clicking on Edit displays the Network Object Edit box ([Figure 22](#)). This box allows you to modify the attributes for the selected network object.

Figure 22 Network object edit box

Clicking on Delete removes the selected network object. If the object that you want to delete is the last object, it returns to the default value.

Clicking on Copy, Cut, or Paste performs those operations on the current network object.

Service

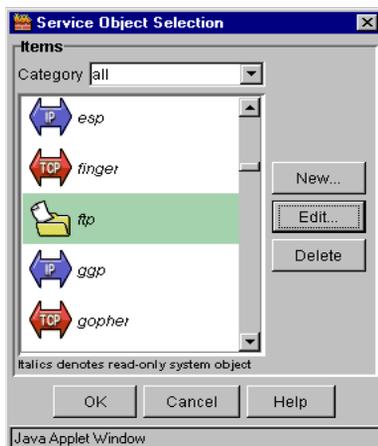
This column specifies which service objects are handled by the selected rule. Right-clicking on the cell displays the standard procedure menu (Add or Edit).

Clicking on Add accesses the Service Object Selection box (Figure 23), which allows you to 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).

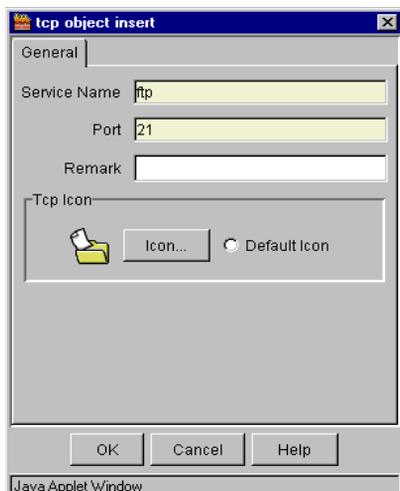


Note: It is possible to add more than one service to a rule.

Figure 23 Service Object Selection box



Italicized objects in the list are read-only and cannot be modified. The New, Edit, and Delete buttons in this box allow you to create, edit, and delete service objects. Clicking on Edit displays the Service Object Edit box, as shown in Figure 24 on page 58. This box allows you to modify the attributes for the selected service object.

Figure 24 Tcp object insert box

Clicking on Delete removes the selected service object from the cell. If the object to be deleted is the last object in the cell, the cell returns to its default value (in this case, Any).

Clicking on Copy, Cut, or Paste performs those operations on the current service object.

Action

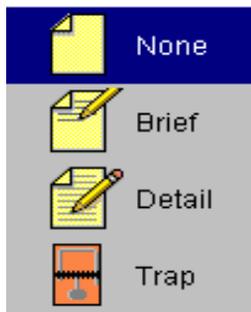
This column ([Figure 25](#)) specifies the action that occurs when the rule is activated. Right-clicking on the cell displays an option list containing three items: Accept, Drop, Reject, and User Authentication. Clicking on one of these items sets the cell to the selected state.

Figure 25 Action column options

Log

The Log column (Figure 26) allows you to specify the logging level for this rule. Right-clicking on this cell brings up an option list containing the following logging levels: None, Brief, Detail, and Trap.

Figure 26 Log column options



Status

This column (Figure 27) specifies the status of the particular rule. The status can be either Enabled or Disabled.

Figure 27 Status column options



Remark

This column allows you to attach a remark to a particular rule. When you right-click on Remark and choose Add or Edit remark, a box appears where you can type a comment.

Creating a new policy

To configure your firewall policies:

- 1 Log in to the Contivity. Select Services > Firewall/NAT. The Firewall/NAT screen appears.
- 2 Under Configuration, click on the Enabled radio button next to Contivity Firewall.
- 3 Click on the Manage Policies button. The Firewall - Select Policy screen appears.
- 4 Click on New to create a new policy. The New Policy box appears.
- 5 Enter the policy name and click on OK. The name must begin with a letter and may not contain the : + =] , ; " characters. The Firewall - Edit Policy: <policyname> screen appears with no rules defined. In this screen, you can add, delete, and modify the rules for the policy.
- 6 You can select the rule group as follows:
 - Implied rules (view only)
 - Override rules
 - Interface-specific rules
 - Default rules
- 7 Select the Interface Specific Rules tab.
- 8 Select an interface and a subinterface from the drop-down lists.
- 9 Select either Source Interface Rules or Destination Interface Rules.
- 10 Right-click on the appropriate cell to add a new rule.
- 11 Repeat these steps to add more rules.
- 12 Select Policy and click on Save Policy to save your changes.
- 13 When the policies are saved, go to the Manage menu and click on Close Manager.

Successful completion of these steps indicates that your Contivity's firewall is functioning and that the Contivity's routing patterns are available.

Verifying your configuration

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

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

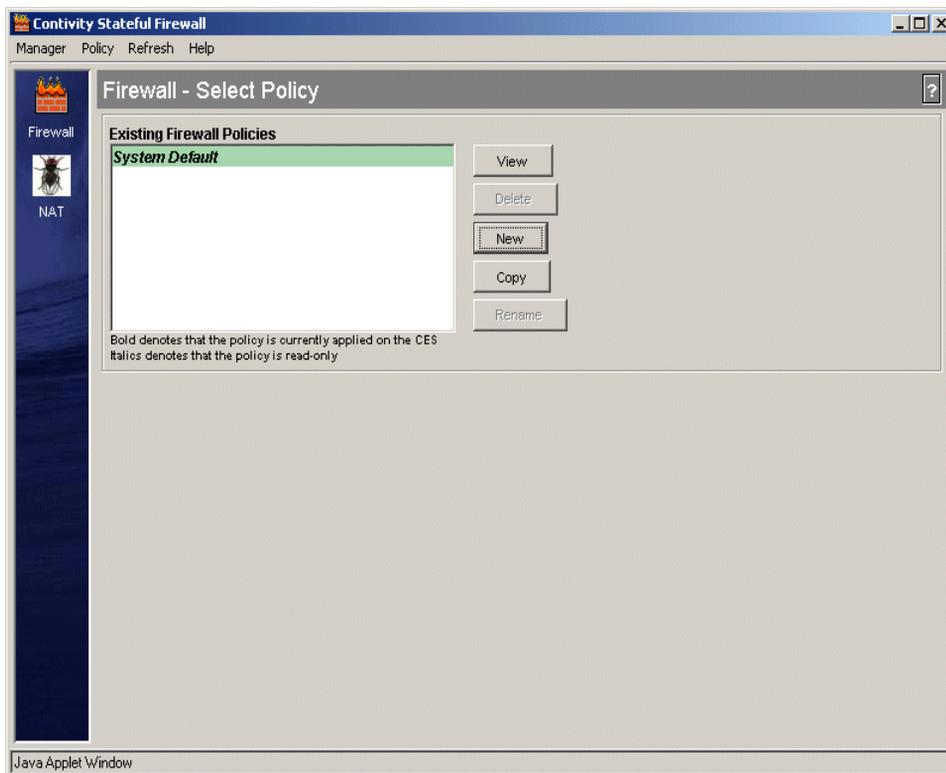
Configuring a sample security policy

In this configuration example, the following setup exists:

- Public IP address 192.168.3.22 (Internet Access)
- Private IP address 10.3.3.102 (Contivity default is LAN)
- FTP server IP address 192.168.3.20 on the public network
- Security policy allows users to FTP to the FTP server to download files, with no other access to the Internet permitted

To configure the stateful firewall to implement a security policy:

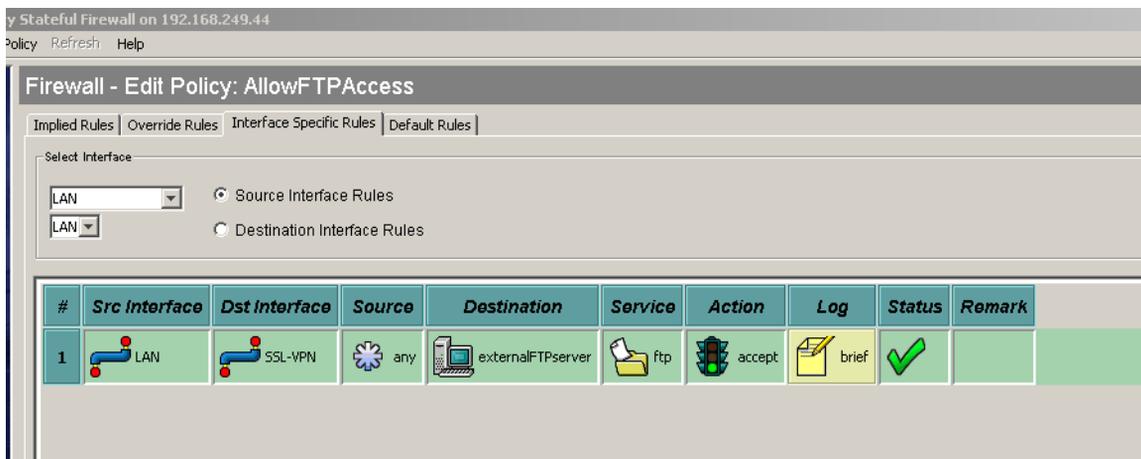
- 1 From the Services > Firewall/NAT screen, click on the Manage Policies button for Contivity Stateful Firewall.
- 2 On the Firewall - Select Policy screen, click on the New button, enter AllowFTPAccess as the policy name and click on OK. The Firewall - Select Policy screen is shown in [Figure 28 on page 62](#).

Figure 28 Firewall - Select Policy screen

- 3** On the Firewall - Edit Policy screen, click on the Interface Specific Rules tab. The Firewall - Edit Policy screen with Interface Specific Rules tab selected is shown in [Figure 29 on page 63](#).
 - a** Make no changes to the interface or subinterface drop-down boxes and leave Source Interface Rules selected.
 - b** Right click on the # box in the header and select Add New Rule.
- 4** On the Firewall - Edit Policy (Interface Specific Rules) screen, click on the DST Interface value (*any), right-click to display the selection menu, and select SSL-VPN.
- 5** On the Firewall - Edit Policy (Interface Specific Rules) screen, click on the Destination value (*any), right-click to display the selection menu, and select Add.
 - a** In the Network Object Selection box, click on the New button.

- b** In the Network Object Type Selection box, select Host as the type of object to create.
 - c** In the Network Object Insert box, enter the Host name (externalFTPserver) and the IP address (192.168.3.20), and click on OK.
 - d** In the Network Object Selection box, click on OK to add the externalFTPserver network object into the Destination field.
- 6** On the Firewall - Edit Policy (Interface Specific Rules) screen, click on the Service value (*any), right-click to display the Service Object Selection box, scroll down to and click on FTP, and click OK.
 - 7** On the Firewall - Edit Policy (Interface Specific Rules) screen, click on the Action value (drop), right-click to display the Action menu, and click on Accept to enter it into the Action field.
 - 8** On the Firewall - Edit Policy (Interface Specific Rules) screen, click on the Log value (blank = none), right click to display the Log menu, click on the required log value to enter it into the Log field. In this example, the log value is brief.
 - 9** On the Firewall - Edit Policy (Interface Specific Rules) screen, click on the Status value (checked means enabled), right-click to display the Status menu, click on the required status value to enter it into the Status field. (Within a policy, each rule in the Override, Interface-Specific, and Default groups can be independently disabled.)

Figure 29 Firewall - Edit Policy screen with new policy



- 10 On the Firewall - Edit Policy (Interface Specific Rules) screen, click on the Manager menu at the top left of the screen and click on Exit CSF/NAT. In the Save Changes to this policy box, click on Yes.
- 11 On the Services > Firewall/NAT screen, select AllowFTPAccess from the policy box, and click on OK. (Only a single policy can be applied to the Contivity.)
- 12 Click on the Contivity Firewall button, check Contivity Stateful Firewall, and click on OK. You will be prompted to reboot the Contivity to activate the new firewall configuration.

Firewall deployment examples

Security policies can be customized and applied to individual subscribers, or created as templates to be applied to many subscribers.

Some questions you should consider when establishing firewall rules include:

- What are the IP addresses for all of your servers (FTP, DNS, Web, mail) accessible through this firewall?
- If you are setting up NAT, what IP addresses should be listed that are otherwise not visible?
- What applications, other than HTTP, FTP, mail protocols, and other typical network traffic, will be running across your firewall?

Residential firewall example

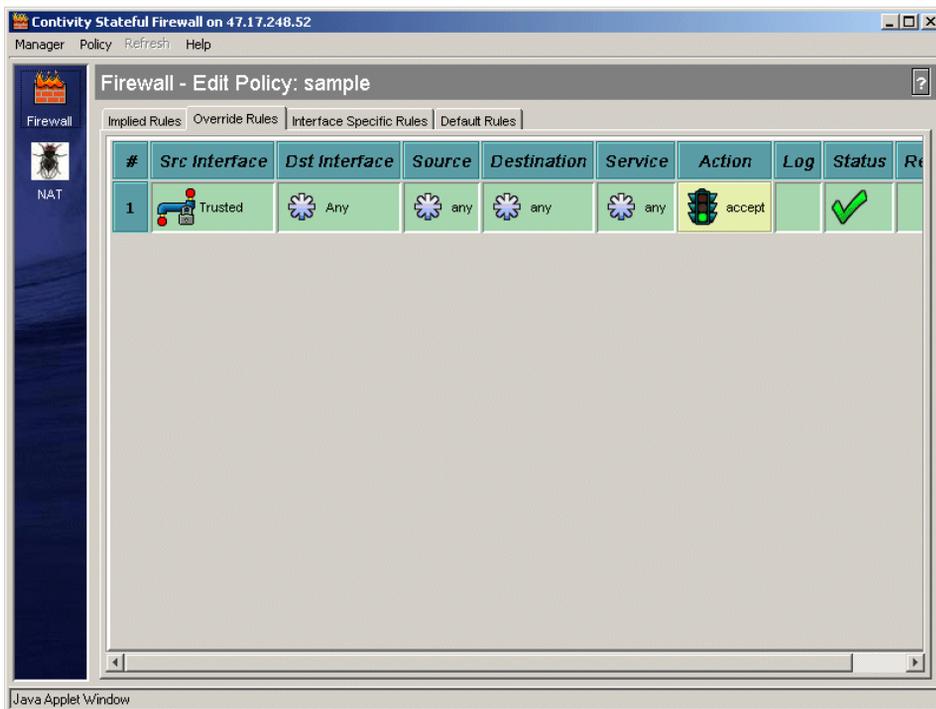
A residential firewall ([Figure 30](#)) is generally a simple firewall designed to allow user-initiated traffic while blocking any incoming traffic or port scans.

Figure 30 Example of a basic residential firewall

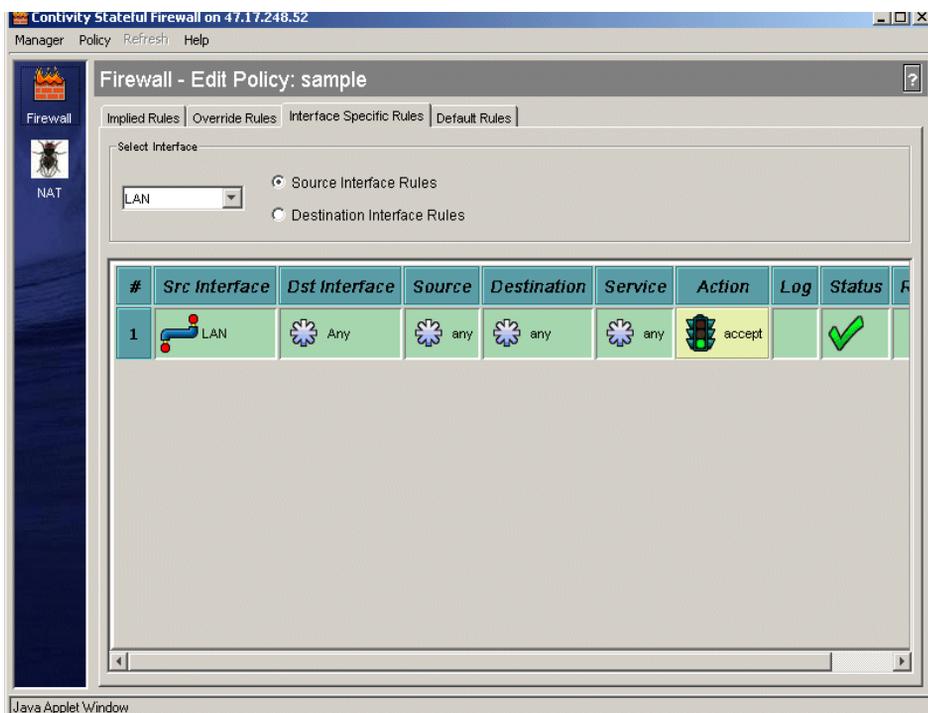


You can configure your residential firewall with a single override rule (Figure 31) that allows all trusted traffic. Trusted traffic is traffic that is sourced from either a trusted physical interface or a tunnel.

Figure 31 Single override rule

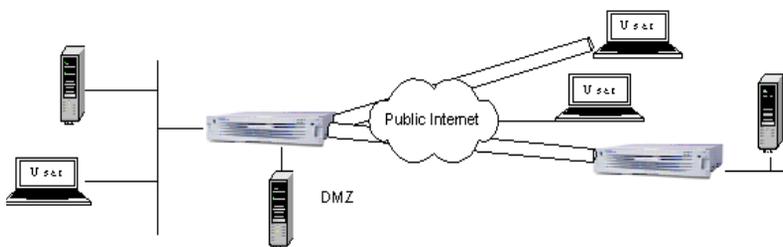


Alternatively, you can configure a single interface specific rule, as shown in Figure 32 on page 66, that allows traffic sourced from the physical interface LAN (slot 1/0).

Figure 32 Single interface specific rule

Business firewall example

A business firewall (Figure 33) requires a more complex rule configuration. A business user must have 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 will include HTTP, SMTP, FTP and network protocols, such as some forms of ICMP.

Figure 33 Business firewall

When configuring a business firewall, you must set override rules (Figure 34) that:

- Require branch office user to authenticate themselves prior to accessing internal resources
- Allows user tunnel traffic to go anywhere
- Allows non-tunneled FTP and HTTP to gain access to the DMZ

You must also set an interface specific rule (Figure 33 on page 66) that allows all traffic that enters from the private (LAN) to go anywhere.

Figure 34 Business override rules

The screenshot shows the 'Firewall - Edit Policy: sample' window in the Contivity Stateful Firewall management console. The 'Override Rules' tab is selected, displaying a table of three rules. The table columns are: #, Src Interface, Dst Interface, Source, Destination, Service, Action, and Log. Rule 1 is 'Branch Tunnel: Any' with 'user authentication' action. Rule 2 is 'User Tunnel: Any' with 'accept' action. Rule 3 is 'Any' with 'DMZ' destination and 'ftp' and 'http' services, with 'accept' action.

#	Src Interface	Dst Interface	Source	Destination	Service	Action	Log
1	Branch Tunnel: Any	Any	any	any	any	user authentication	
2	User Tunnel: Any	Any	any	any	any	accept	
3	Any	Any	any	DMZ	ftp http	accept	

Chapter 3

Configuring filters

You can manage the Contivity Secure IP Services Gateway tunnel filters (for user groups) or interface filters (for LAN and WAN interfaces). Changing a tunnel filter does not affect any existing tunnels. You must re-establish the existing tunnels for changes to take effect.

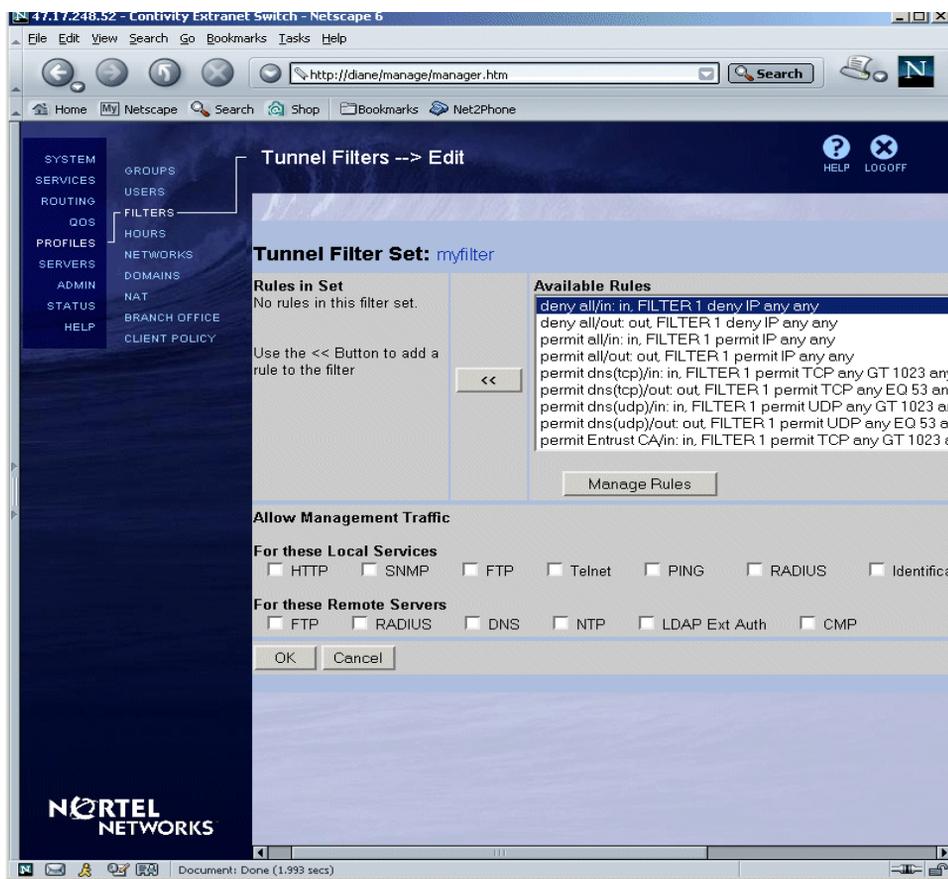
The Current Tunnel Filters and Current Interface Filters screens show the currently available filters (Profiles > Filters). A filter usually consists of one or more inbound rules (for traffic coming into the network) and one or more outbound rules (for traffic leaving the network). Filter names are a convenient way of managing a set of rules.

Adding and editing filters

To add or edit a filter:

- 1 On the Services > Firewall/NAT screen, enter the new filter name and click on the Create button or click on the name of the filter set and on the Edit button. The Edit Filter screen appears, as shown in [Figure 35 on page 70](#).

Figure 35 Edit Filters screen



- 2 For the Filter Set, enter the name of the filter. For existing filters, Rules in Set lists the rules that are already contained in the filter that you are editing.
- 3 Click on a rule from the Available Rules list on the right of the screen, then click on the left arrow to add the rule. This adds the selected rule to the current rules list. The new rule is added after the rule currently selected in the Rules in Set list.
- 4 Click on a rule, then click the right arrow to remove or delete it from the Rules in Set list.
- 5 Click on a rule in the Rules in Set list, then click the up arrow to move the rule up one place in the list.
- 6 Click on a rule in the Rules in Set list, then click the down arrow to move the rule down one place in the list.

- 7 The Available Rules field lists all of the rules that are available on the Contivity gateway to add to the filter. They appear in the format of Name: Rule String. Click to view the Current Rules screen, from which you can Create, Edit, Copy, or Delete a Rule.



Note: The Allow Management Traffic section applies only to tunnel filters, and does not appear on the screens for interface filters.

- 8 By manipulating the Allow Management Traffic options, you can restrict management access to the Contivity gateway through tunnels. Each filter set has an explicit list of management services. By specifying the management services allowed through a tunnel, you can control which groups of users are able to perform different management tasks while tunneled into the Contivity gateway.

The Contivity gateway's default filter is Permit All, and the settings for this filter are to allow HTTP, SNMP, and PING. But if you create a new filter, all management traffic settings are disabled by default.

The management protocols are divided into two groups. The Local Services selections refer to services that reside on the Contivity gateway. The Remote Servers selections refer to services that reside on other systems that are used by the Contivity gateway. When enabled, network traffic for these services is allowed through tunnels.

The management services apply to user and branch office connections. These options do not affect HTTP, SNMP, FTP, Telnet, or PING protocol traffic that is passing through the Contivity gateway outside a tunnel.

When you select the boxes for these Local Services, you:

- HTTP--Enable or disable access to the Web server on the Contivity gateway.
- SNMP--Enable or disable SNMP “gets” to the Contivity gateway.
- FTP--Enable or disable FTP “puts” or “gets” to the Contivity gateway.
- Telnet--Enable or disable Telnet access to the Contivity gateway.
- PING--Enable or disable PING access to the Contivity gateway.
- RADIUS--Enable or disable access to the Contivity gateway’s RADIUS authentication service.

When you select the boxes for these Remote Servers, the Remote Servers options restrict traffic to external services that are required by the Contivity gateway. By specifying these services, you can restrict which tunnels on a Contivity gateway can send protocol traffic for external services it requires.

- FTP — enable or disable FTP access from the Contivity gateway 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 are controlled by this option.
- DHCP — enable or disable access to dynamic host configuration protocol (DHCP) servers from the Contivity gateway.
- RADIUS — enable or disable the Contivity gateway's ability to access a remote RADIUS server.
- DNS — enable or disable remote users from using the Domain Name Server (DNS) service for the Contivity gateway.

Use the Copy Filter buttons to copy an existing filter from one filter set to the other. For example, if you have already created a filter for tunnels, you can copy it for use by your Contivity gateway's interfaces.



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

To copy a filter, click on the existing filter in one Current Filters window, then click the appropriate Up or Down button to move the filter to the other Current Filters window. The Copy Filters screen appears, asking you to confirm that you want to copy the filter.



Note: If you copy a tunnel filter for use by a Contivity Stateful Firewall, you might need additional set up steps because the traffic that uses the Contivity Stateful Firewall traverses two Contivity gateway interfaces. For example, it might enter via a public interface and exit through a private interface. However, tunnel traffic only enters and exits through a single physical interface.

Next hop traffic filters

Next hop traffic filters allow customers to control the next hop selection and route traffic within their domain. If a packet matches filter criteria, a forwarding lookup is performed using the configured next hop and is forwarded using that routing table instance. If the lookup fails, then traditional destination-based routing occurs using the routing table.

Each IP interface can have inbound and/or outbound filters that cause an action to be taken on a packet if the packet matches the filter criteria. When a filter rule with next hop (shown in [Table 2](#)) configured matches an incoming packet, the filter accepts the packet and uses the next hop for forwarding.

Next hop traffic filters are only applicable for inbound filters per interface (physical or virtual) per 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 will be forwarded to 192.32.140.216

When this filter is applied on an interface, all incoming IP traffic coming to that interface from 10 network and going to the 47 network is forwarded to the next hop address. This assumes that there is a reachable route to the next hop address. If the next hop is not reachable, than the destination address in the IP header is used (as in normal routing) to forward the packet.

For tunnels, the +next hop address should be beyond the remote end point of the tunnel and should be along the path to the actual destination.

To configure next hop traffic filters:

- 1 Go to Profiles > Filters and click on Manage Rules.
- 2 Select the rule that you want to change and click on Edit.
- 3 Select Nexthop for the filter action. You can optionally enter the source and destination address fields, as shown in [Figure 36 on page 74](#).

Figure 36 Nexthop filter action

SYSTEM SERVICES ROUTING QOS PROFILES SERVERS ADMIN STATUS HELP

GROUPS USERS FILTERS HOURS NETWORKS DOMAINS BRANCH OFFICE CLIENT POLICY MAP CLASS

Tunnel Filters --> Rules --> Edit

HELP LOGOFF

Filter Rule for:	deny all/in	
Filter Action	<input type="radio"/> Permit <input type="radio"/> Deny <input checked="" type="radio"/> Nexthop	
Direction	Inbound	
Nexthop Address	any, 0.0.0.0, 255.255.255.255	Modify
Source Address	any, 0.0.0.0, 255.255.255.255	Modify
Destination Address	any, 0.0.0.0, 255.255.255.255	Modify
Protocol	ip	Modify
Source Port	Equals any, 0	Modify
Destination Port	Equals any, 0	Modify
TCP Connection	<input type="radio"/> Established <input checked="" type="radio"/> Don't Care	
OK Cancel		

NORTEL NETWORKS

- 4 To enable private to tunnel forwarding, go to System > Forwarding. The Forwarding screen appears, as shown in [Figure 37 on page 75](#).
- 5 Select Next Hop Forwarding and click on OK.

Figure 37 Next hop forwarding



Chapter 4

Configuring NAT

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

NAT contains a pool of available global addresses that are continually reused. It allows a network to use one set of network addresses internally and a different set when dealing with external networks. Internal network addresses are allocated according to internal considerations of the network. Global addresses must remain unique to distinguish between different hosts. When a packet is routed, NAT replaces the internal corporate address with a global address. As soon as the application session is over, the global address is returned to the pool and can be used by subsequent connections. NAT can also modify the source and destination port numbers.

Address translations

Address translation can be set up permanently (static) or allocated dynamically, allowing many devices on an internal network to share a few IP addresses. Static translation allocates one external host address for each internal address and is converted to a different global IP address. Dynamic address translation occurs when a session is started. No guaranteed one-to-one mapping takes place. An example of dynamic translation is port mapping, which uses the TCP/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:

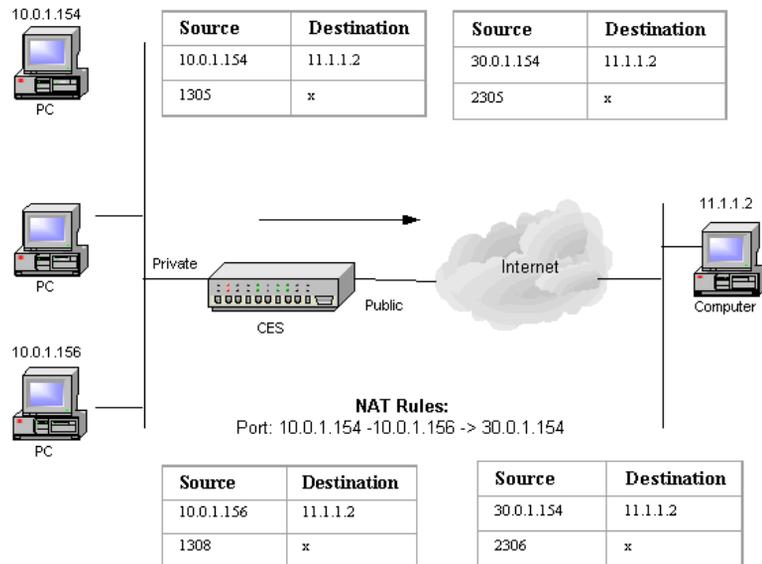
- Dynamic many-to-one
- Dynamic many-to-many
- Static one-to-one
- Port forwarding
- IPsec-aware NAT
- Double NAT

Dynamic many-to-one — port translation

With Network address port translation (NAPT), many internal IP addresses are hidden behind a single external address using dynamically-assigned ports to distinguish between them. This is especially useful if you need to use several IP addresses and have only one address available from your ISP. Dynamic many-to-one translation can only be used for traffic initiated from an internal host.

NAT attempts to assign a port from the corresponding port list. If the original port is available, it is assigned. If not, 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 packet is dropped.

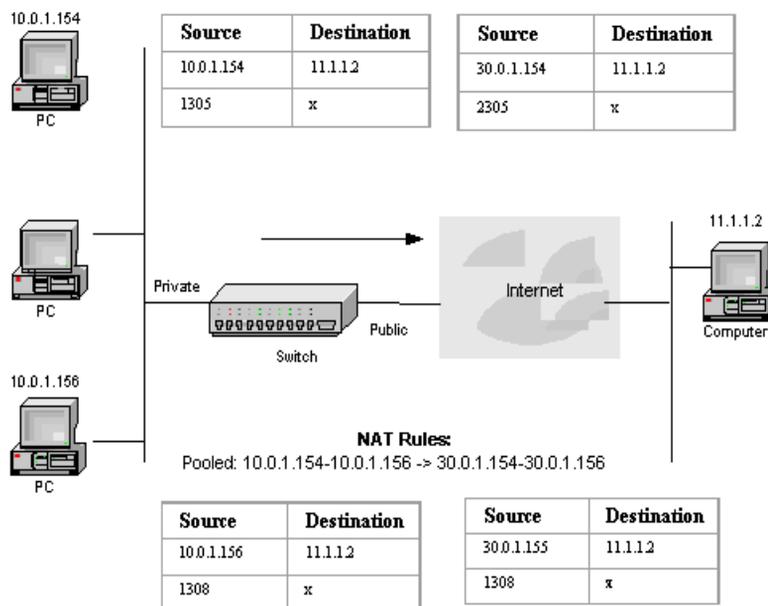
[Figure 38 on page 79](#) shows the private network 10.0.1.0 is hidden behind the public address 30.0.1.154. All requests originating from the private network (10.0.1.0) have their source IP addresses replaced with the public IP address 30.0.1.154; only the public IP address is visible from the public network. In addition, source ports are dynamically translated to unique translated ports.

Figure 38 Port translation

Dynamic many-to-many — pooled translation

In dynamic many-to-many NAT, only the address (not the port) is translated. Usually, the number of externally visible IP addresses is less than the number being hidden behind the Contivity gateway. Each time a request is made from a host on the private network, the Contivity gateway chooses an external IP address that is currently unused, and then performs the translation. Dynamic many-to-many can only be used for traffic initiated from an internal host.

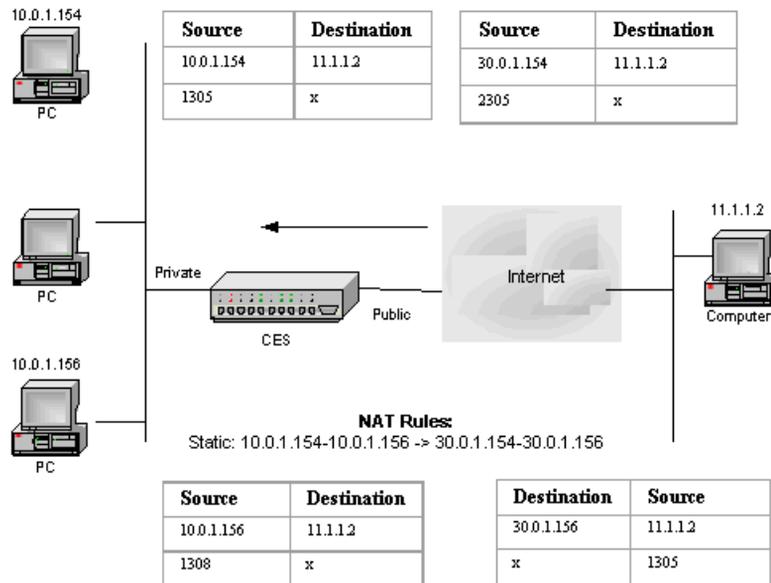
The following example ([Figure 39 on page 80](#)) illustrates the use of many-to-many dynamic translation. The user configures a pooled NAT rule converting the internal address range 10.0.1.154-10.0.1.164 to 30.0.1.154-30.0.1.154. Traffic is initiated from 10.0.1.154 and 10.0.156 destined to a machine (11.1.1.2) on the public Internet. Both addresses are translated to unique public addresses dynamically.

Figure 39 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.

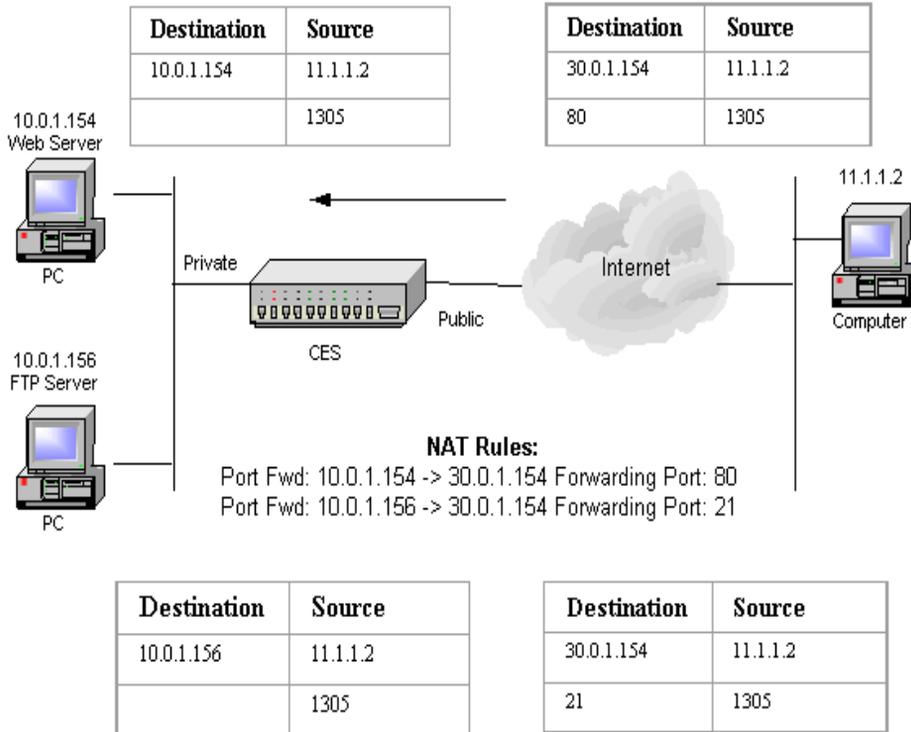
[Figure 40 on page 81](#) shows host 10.0.1.154 on the private side statically mapped to an external address 30.0.1.154, which allows Internet host 11.1.1.2 to initiate a session using the translated external address. The host using this rule is always bound to the same external address.

Figure 40 Static address translation

Port forwarding

Port Forwarding allows one externally accessible IP address to forward incoming requests to different addresses behind the NAT device based on the protocol used. Incoming Web traffic could be routing to a Web server, while FTP traffic destined to the same external IP address could be forwarded to a different device capable of providing FTP services.

[Figure 41 on page 82](#) illustrates the use of Port Forwarding. A host 11.1.1.2 on the Internet needs to access a Web server and an FTP server running on two separate internal machines that are hidden behind a single externally visible address 30.0.1.154. This is achieved using a port forwarding NAT rule that sends the traffic to the two different machines based on the forwarding ports.

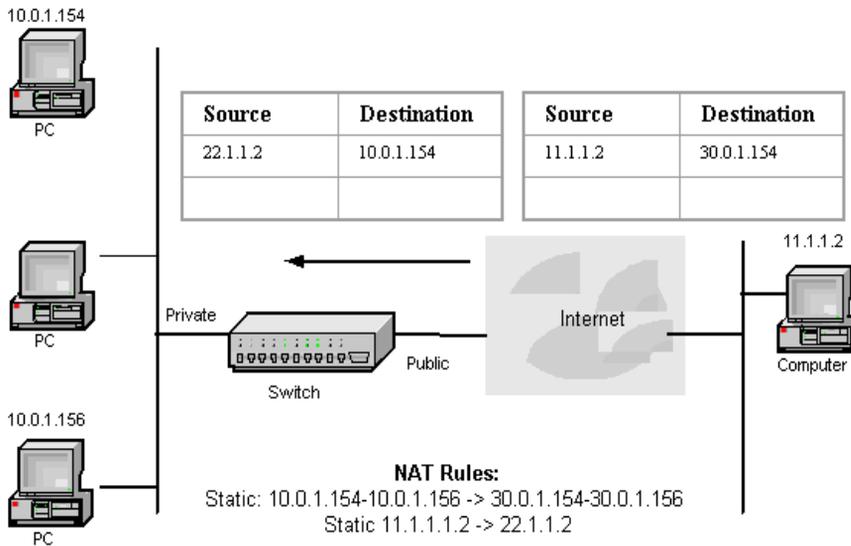
Figure 41 Port forwarding example

Double NAT

Double NAT allows both external and internal networks to be translated at the same time. For each packet entering and leaving the Contivity gateway, both the source and destination addresses could be modified. For Contivity gateway NAT, this is achieved using rules, one for translating the source address and one for translating the destination address. The destination address translation must use a static rule.

[Figure 42 on page 83](#) shows a host 11.1.1.2 on the Internet initiating a connection to 30.0.1.154, the translated address of the internal host. As the packet traverses NAT, both the source and destination addresses are translated.

Figure 42 Double NAT

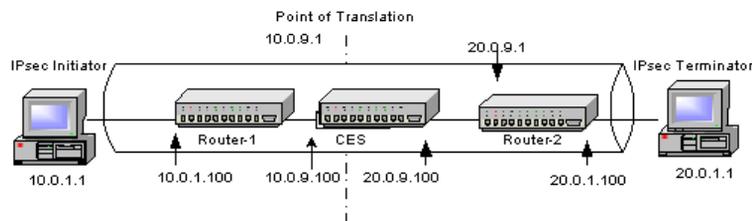


IPsec-aware NAT

IPsec-aware NAT provides a means of protecting against the alteration of TCP/IP headers, usually performed by NAT. IPsec-aware NAT is used when an IPsec tunnel passes through a Contivity gateway performing NAT translation, but does not terminate at the Contivity gateway. This allows inter-operability 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 cannot be configured.

Figure 43 shows an IPsec-aware NAT example.

Figure 43 IPsec-aware NAT example



NAT modes

Based on the handling of UDP packets, NATs can be classified as four different modes:

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

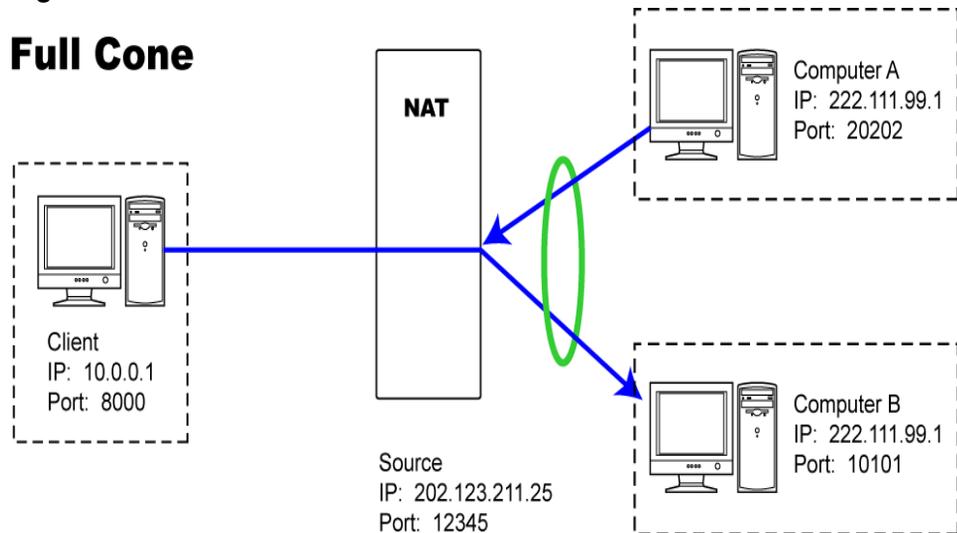


Note: Only Restricted Cone NAT and Symmetric NAT modes are supported. All visible references to Cone NAT in the system refer to Restricted cone NAT.

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. Any external host can send a packet to the internal host by sending a packet to the mapped external address.

Figure 44 Full Cone NAT

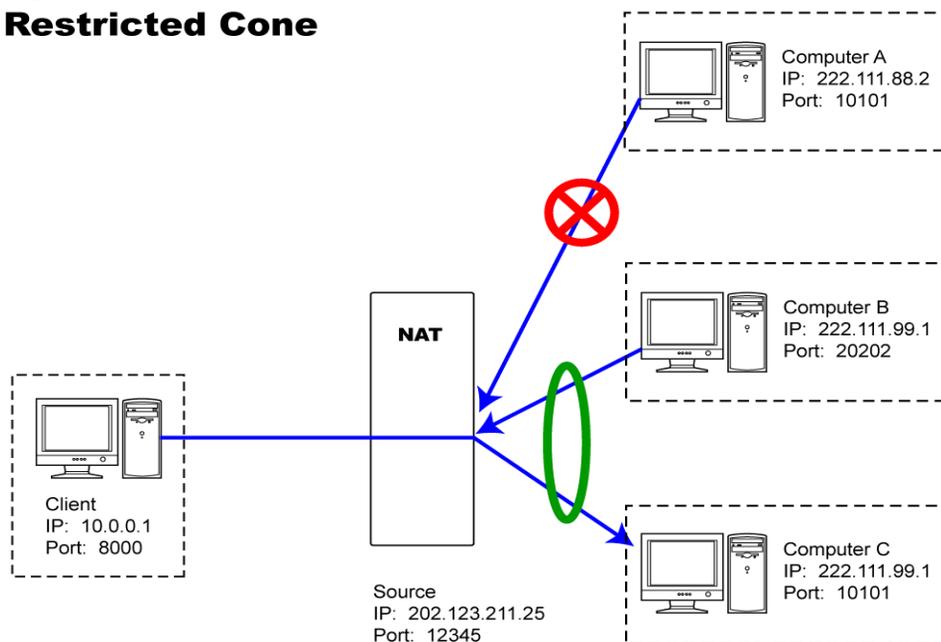


553-AAA3064

Figure 44 is 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 of 202.123.211.25:12345. Anyone on the public side can send packets to that external IP/port and those packets will be correctly translated to the client's internal IP/port.

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 has previously sent a packet to the IP address.

Figure 45 Restricted Cone NAT**Restricted Cone**

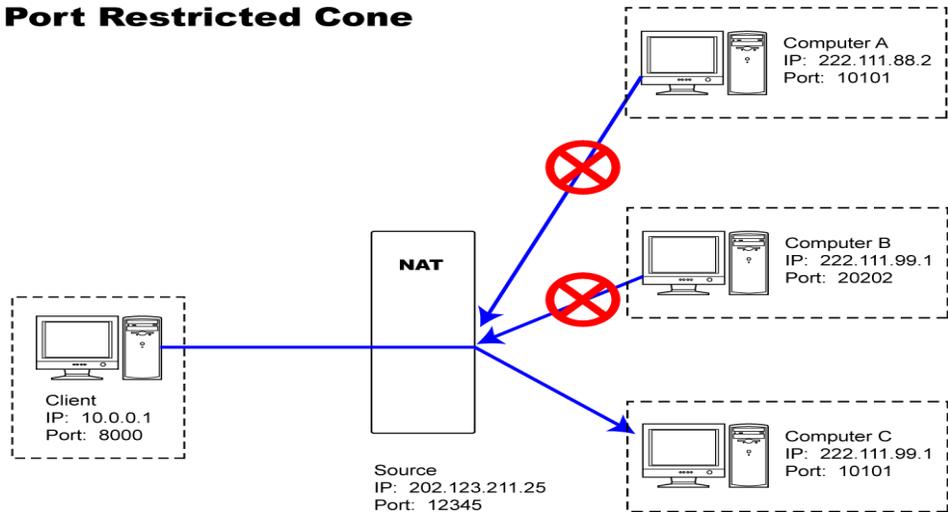
553-AA3066

Figure 45 shows an example of a private client sending a packet to an external client (computer A). The NAT 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 blocks all packets coming from an external client (computer B) until the private client sends a packet to that external IP address. Once that is done, both external clients can send packets destined to the NAT address and they are translated correctly to the clients' private address.

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 has previously sent a packet to the IP address and port.

Figure 46 Port Restricted Cone NAT
Port Restricted Cone



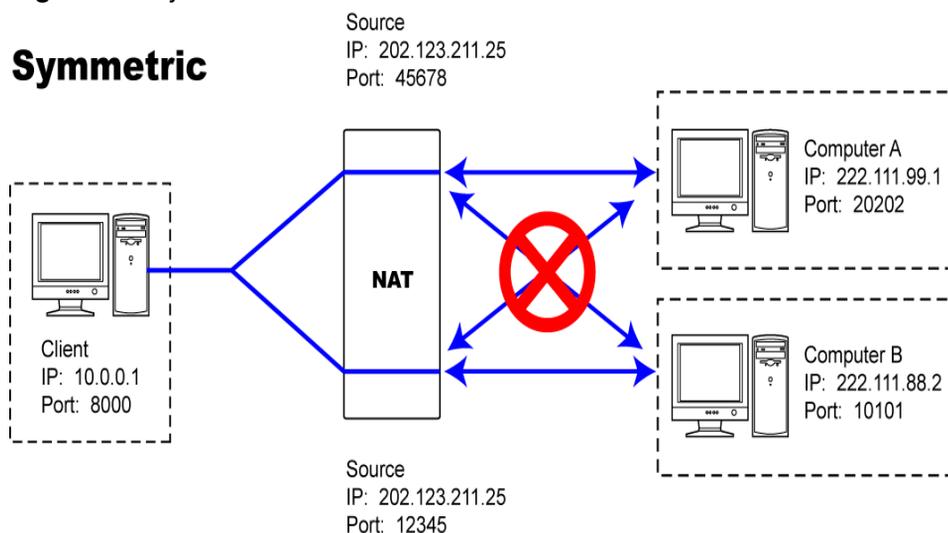
553-AAA3065

Figure 46 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 will only allow packets that come from the same IP and port. If the internal client has sent out packets to multiple external IP address/ports, they can all respond to the client at the same mapped IP address and port and the NAT will do 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, a different mapping is used. Only the external host that receives a packet can send a packet back to the internal host.

The default NAT mode is Symmetric and can be changed to restricted Cone NAT on the Services > Firewall > NAT > Edit page.

Figure 47 Symmetric NAT

553-AAA3067

Figure 47 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 may be mapped to 202.123.211.25:12345 while a packet sent from the same address and port to 222.111.99.1 may be mapped 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

NAT traversal enables the Contivity VPN client or server user tunnels to pass through intermediate routers or gateways, each of which may NAT the packet. Most hotels and airports that provide Internet connectivity use NAT to connect to the Internet. You can enable NAT traversal on the Services > IPsec screen. By default, NAT traversal is disabled.

NAT traversal solves the user tunnel case where the IPsec-aware NAT may not always work because other NATs are between the source and destination PC hosts.

To use NAT traversal, you must also define a UDP port that is used for all client connections to the Contivity. This port must be a unique and unused UDP port within the private network (supported range 1025 - 49151). By default, no UDP port is defined.



Note: To allow NAT traversal with the IPsec client, you must enable the NAT traversal setting on the Profiles > Groups > Edit IPsec screen.

The group-level NAT traversal setting allows you to configure the NAT traversal mode at the group level. By default, NAT traversal is Not Allowed. Therefore, even if NAT is detected between the client and the Contivity, UDP encapsulation of ESP data will not occur. Selecting Auto-Detect NAT allows the client and Contivity to UDP encapsulate ESP data whenever NAT is detected. It also allows the client and switch to UDP encapsulate ESP data but only if the NAT detected is non-IPSec aware (when the NAT box does not allow for IPsec pass-through).

Because there are a variety of NAT devices and varying IPsec pass-through implementations, not all environments can be guaranteed to function properly using the Auto-Detect IPsec NAT mode. In environments where the NAT devices may be unknown the Auto-Detect NAT setting is recommended. The Auto-Detect IPsec NAT setting is only recommended for environments where the NAT devices are well known.



Note: Nortel Networks recommends that you use port 10001 for NAT traversal. Do not use L2TP/L2F port 1701 or General Packet Radio Service (GPRS) port 3386. Make sure that any port that you select does not conflict with any ports that you are already using.

NAT and VoIP

NAT translates IP addresses and port numbers in private address ranges into public addresses when traffic traverses between private and public networks. The IP endpoints in a VoIP network (IP Phones, Soft Clients) are typically assigned private addresses to hide their identity from the public network. Voice calls from and to the public network have to reach endpoints in the private network and, as a result, network address translation is required to allow proper routing of media to endpoints with private addresses.

VoIP protocols introduce a number of complexities for NAT, since 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, cannot be routed to the private address, resulting in a one way speech path.

The challenges for VoIP traversal in NAT occur because:

- NATs only look at Layer 3 addressing
- VoIP signaling protocols embed IP addresses at Layer 5
- RTP and RTCP work at Layer 5

Two of the most common solutions that have been proposed to fix the NAT traversal issue are:

- Application Level Gateways (ALG)
- Address/port discovery

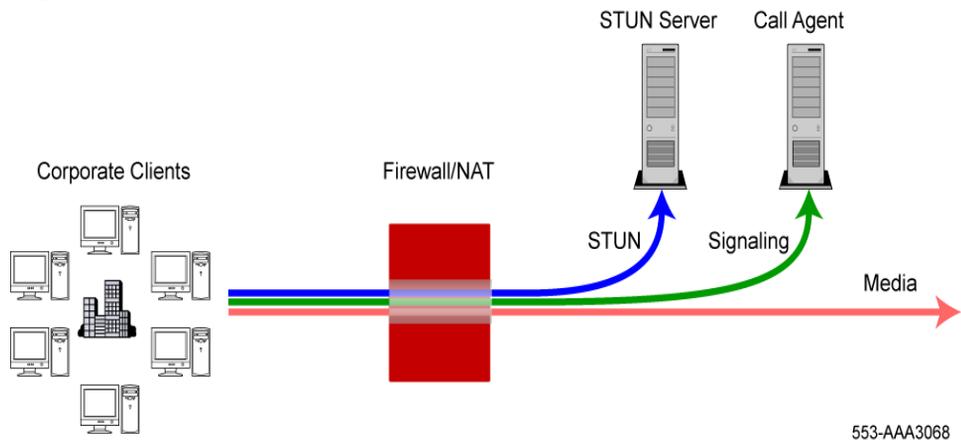
This section focuses on the address/port discovery mechanisms for VoIP. ALGs are discussed in [“NAT ALG for SIP” on page 112](#).

Address/Port discovery

In address/port discovery, the media end points dynamically discover the public IP address and port to be used for a specific media stream by sending probe packets to a server that echo back to the end point its source IP address as seen after the NAT Translation.

Simple Traversal of UDP through NATs (STUN) is a lightweight protocol that allows applications to discover the presence and types of NATs and firewalls between the application and the public Internet. It also allows the applications to determine the public IP addresses allocated by the NAT.

[Figure 48 on page 91](#) shows how STUN works.

Figure 48 STUN

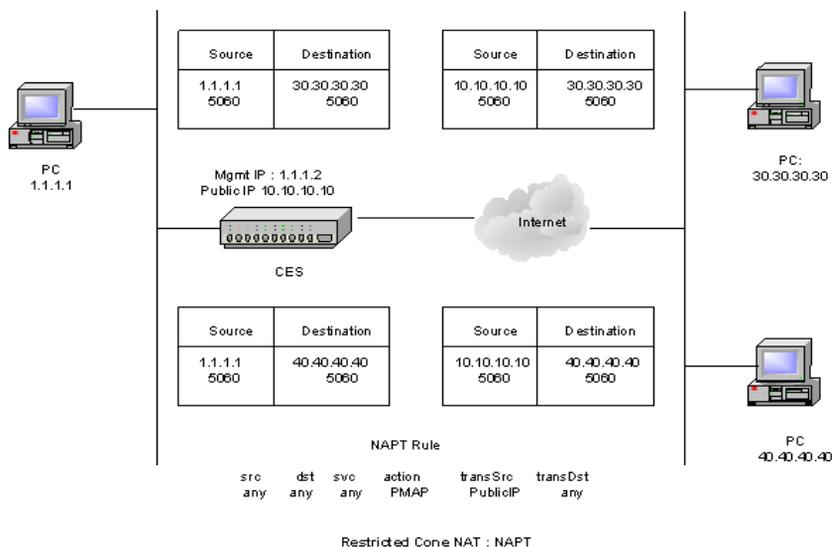
STUN identifies the public-side NAT details by inspecting exploratory STUN messages that arrive at the STUN server. The STUN-enabled client sends an exploratory message to the external STUN server to determine the transmit and receive ports to use. The STUN server examines the incoming message and informs the client which public IP address and ports the NAT used. These are then used in the call establishment messages sent to the SIP server. Note that the STUN server does not sit in the signaling or media data flows.

In order for the discovered IP address and port to be valid, it is imperative that the same IP address and port binding be used by the NAT, regardless of where the packet is going. This means that Symmetric NAT does not work for peer-to-peer media with address/port discovery. STUN requires any Cone NAT implementation. Restricted Cone NAT makes the box more secure.

Network address port translation (NAPT)

Network address port translation (NAPT) is a dynamic NAT where many internal IP addresses are hidden behind a single external IP address using dynamic port assignment to distinguish between them. 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 is changed so that each internal IP address and port is mapped to the same external IP address and port, irrespective of the destination and the session.

[Figure 49 on page 92](#) shows the flow of a Restricted Cone NAT.

Figure 49 Restricted Cone NAT — NAPT

Configuring Cone NAT

Cone NAT can be enabled or disabled with the GUI or the Command Line Interface. To learn the CLI, go to *Reference for the Contivity Secure IP Services Gateway Command Line Interface (CLI)*.

To configure Cone NAT:

- 1 Select Services > Firewall/NAT. The Firewall/NAT page opens.

[Figure 50 on page 93](#) shows the Firewall/NAT page.

Figure 50 Firewall/NAT page

The screenshot shows the Firewall/NAT configuration page. The sidebar menu includes: SYSTEM, SERVICES, ROUTING, QOS, PROFILES, SERVERS, ADMIN, STATUS, HELP, AVAILABLE, AOT, DEMAND, IPSEC, PPTP, FWUA, L2TP, L2F, RADIUS, FIREWALL / NAT, SYSLOG, and SSLTLS. The main content area is titled 'Firewall / NAT' and contains a 'Configuration' section with the following table:

Enabled	Firewall / NAT Type	Firewall / NAT Policy	Action
<input checked="" type="radio"/>	Contivity Firewall *		Edit
	Contivity Stateful Firewall <small>Not available until a license key is installed. See Admin->License Keys page.</small>		
<input checked="" type="checkbox"/>	Contivity Interface Filter		
<input type="checkbox"/>	Interface NAT	NAT Policy: (None) ▾	Manage Policies
<input type="checkbox"/>	Anti-Spoofing		Edit
<input type="checkbox"/>	Malicious Scan Detection		Edit
<input type="radio"/>	No Firewall		

Below the table, there are two footnotes:

- * A firewall must be enabled to allow interface traffic, select either interface(packet) filters, or the Contivity Stateful Firewall
- * Enable Contivity Firewall requires the Interface Filter to be "permit all" in order for OSPF to function properly. [System->LAN](#)

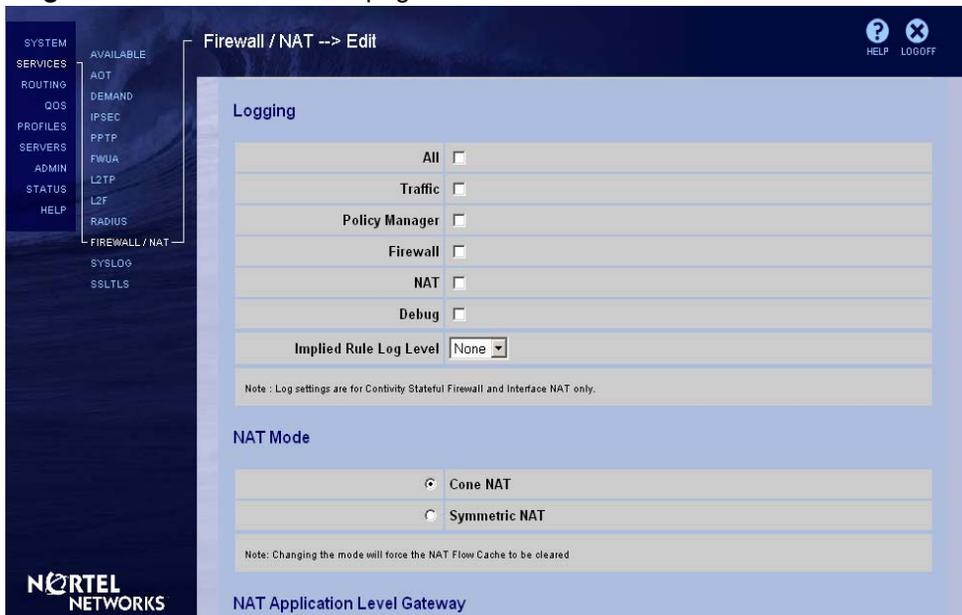
At the bottom of the configuration area, there are two filter settings:

- Contivity Tunnel Filter Enable
- Contivity Tunnel Management Filter Enable

- 2 Click Edit in the Contivity Firewall row. The Firewall/NAT > Edit page opens.
- 3 Under NAT Mode, select Cone NAT.

Figure 51 on page 94 shows the Firewall/NAT > Edit page where you select Cone NAT.

Figure 51 Firewall/NAT Edit page



4 Click OK.

The Firewall/NAT page re-opens. Cone NAT has been applied.



Note: Changing the mode will force the NAT flow cache to be cleared. Clearing the NAT cache flow will result in a disruption of all active NAT sessions.

NAT Usage

NAT is applied to routed traffic passing through its physical interfaces (interface NAT) and branch office interfaces (branch office NAT) using separate NAT policies. Each branch office has one NAT policy, and there is one global NAT policy applied to non-tunneled traffic.



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

Branch office tunnel NAT

In branch offices, you could have two or more branches that use the same private addressing scheme. Nonetheless, the branch offices still have to communicate with one another.

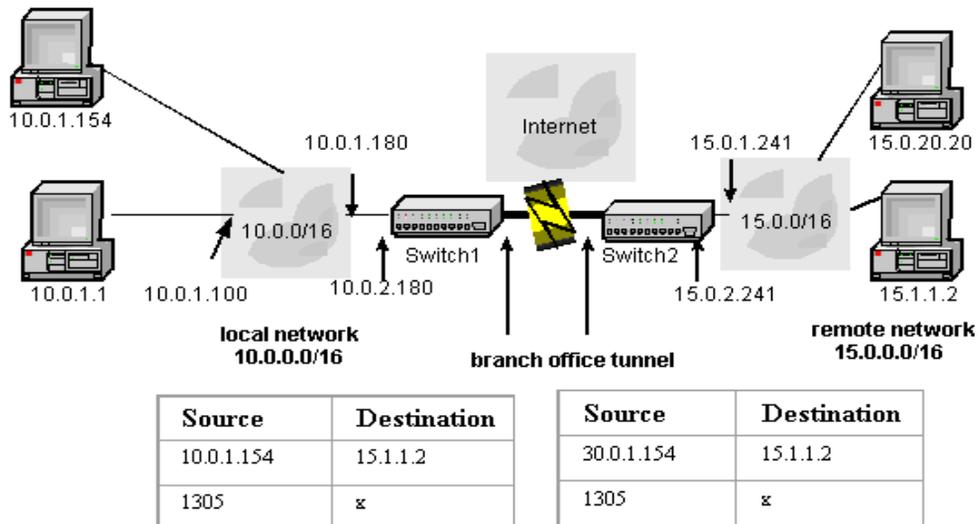
A typical scenario might include a client on LAN 1 who tries to access the FTP server on LAN 2, and who sends a packet with a source address of 10.0.0.13 and a destination address of 10.0.0.14. Without NAT, the Contivity gateway would look at the destination address and assume that the destination is on the same LAN as the source device because the addresses are both on the 10.0.0.0 network and no tunnel connection would be 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 is a common issue 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, Contivity1 defines a remote accessible network of 12.0.0.0, and Contivity2 defines a remote accessible network of 11.0.0.0. Contivity2 uses a static translation of 10.0.0.14 (server) to 12.0.0.1. Contivity1 uses a translation of 10.0.0.13 (client) to 11.0.0.1. As a result, Contivity2 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 has a source address of 10.0.0.13 and a destination address of 12.0.0.1. Contivity1 recognizes that 12.0.0.0 is the remote LAN for the branch office connection. Contivity1 translates the source address of the packet to 11.0.0.1 based on the NAT table. Contivity 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.

[Figure 52 on page 96](#) shows a simple branch office connection with two LANs, and a branch office tunnel across the internet. A pooled NAT rule is applied to Contivity1, which connects the local network to the remote network via its branch office tunnel.

Figure 52 Overlapping address translation



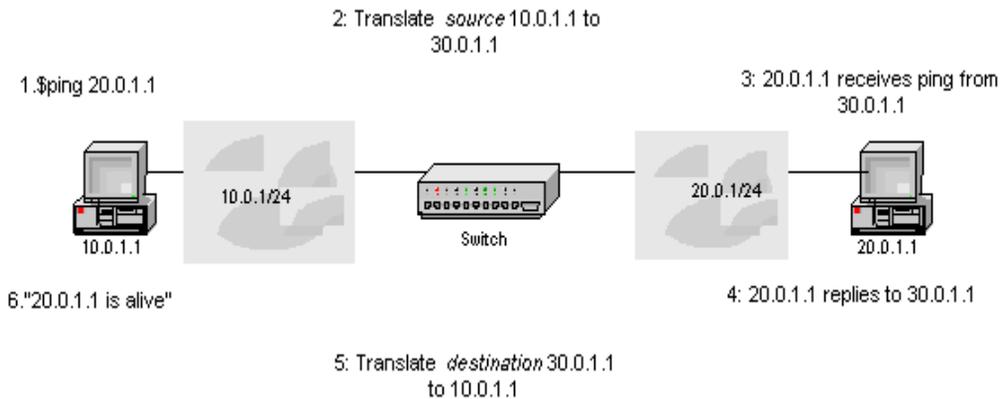
Interface NAT

Interface NAT is applied to IP packets going out from or coming into the Contivity gateway via its physical interfaces causing either the source or destination IP address to be translated to another IP address, depending upon the NAT policy.



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

Figure 53 on page 97 shows an example of interface NAT.

Figure 53 Interface NAT

NAT is applied to interface NAT using the Services > Firewall/NAT screen. Interface NAT rules can be one of the following types:

- Static – for static mapping, an internal address range is mapped “one to one” to an external range.
- Port Forwarding – for port forwarding mapping, external packets are routed on a specified port onto 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 is dynamically mapped to the next available address from the external address range.



Note: Interface NAT only applies to clear text traffic (non-tunneled, routed through the Contivity gateway). Branch office NAT only applies to specific branch office tunnel traffic. If you disable interface NAT, it does not impact branch office NAT.

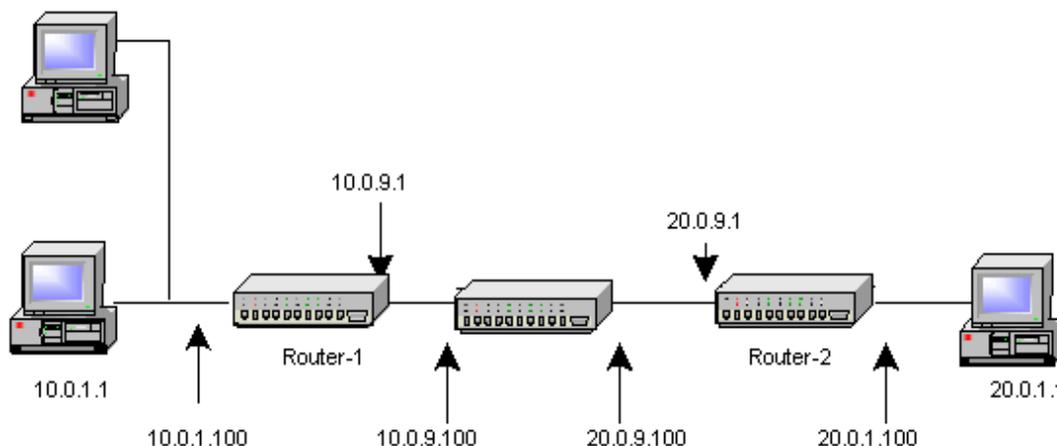
Dynamic routing protocols

NAT routes are advertised on all interfaces. You can restrict the route redistribution to only specific interfaces using the routing policy list. Routing protocols, such as RIP and OSPF, could previously only send packets to the translated addresses if they were configured statically. Now, whenever a NAT

policy is applied to interface or branch office tunnels, the routes to the translated IP addresses are added to the routing table. When NAT is disabled, 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 [Figure 54](#), the Contivity gateway has a NAT rule to convert IP addresses in the range of 10.0.1.1 - 10.0.1.10 to 192.168.1.1.

Figure 54 NAT with dynamic routing example



By default, NAT routes are distributed by RIP and OSPF protocols. However, you have the option to disable the redistribution for a particular protocol from the Routing > Policy > Redistribution Table screen.

You can enable NAT on a branch office with dynamic routing. When NAT is configured for a branch office, you do not want it to announce the route to original IP addresses. You can have 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 subnet, there will not be a route advertisement to the entire subnet.

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

A branch office should not be enabled when there is no routing policy associated with the corresponding branch office interface if both NAT and dynamic routing are configured. You must create a routing policy on the Routing > Policy screen.

NAT uses a port mapping table to track the ports for each client's outgoing packets. The port mapping table relates the client's actual local IP address, source port, and translated source port number 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.

Configuring NAT policy

A NAT policy consists of service properties and a security policy. Service properties define what service is 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 specify all rule fields for service policies using service objects. Each rule consists of a combination of network objects, services, actions, and logging mechanisms. You can define custom policies when more complex security policies are needed 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 a modified copy is applied to the interface. If the read-only NAT policy is applied again after the copy, then the read-only policy translates according to the new rules. This means that it is possible to modify and apply the policies created prior to Version 4.80.

NAT policy sets

The Contivity gateway maintains one set (source and destination address pair) of active global NAT policies for all non-tunneled traffic and a configurable NAT policy set for each branch office tunnel definition. To view active NAT policies for interface and branch offices, go to the Status > Statistics screen.

At system startup, NAT obtains a cached policy if one exists while the system initialization is in process. If there is no cached policy, it takes the default NAT policy, which is no NAT translation. The default NAT policy for the Contivity gateway 1010, 1050, and 1100 port maps its private address space to the public IP address.

Once the system initialization is complete, the NAT policy is retrieved from the LDAP database and that becomes the active policy. When you change the policy, it is stored 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.

Creating rules

Actions on rules are controlled by menus that you access by right-clicking on an option. Each of these menus controls a different aspect of the rule:

- Header row menus – contain only Add New Rule, which allows you 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 – allow you to add a new rule at a particular location, delete the specific rule, and perform cut/copy/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. These menus are similar to a list box. When you click on 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. When you click on one of the items, either the operation is performed immediately (such as Copy) or an additional dialog box appears, prompting you for more information (such as Add).

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

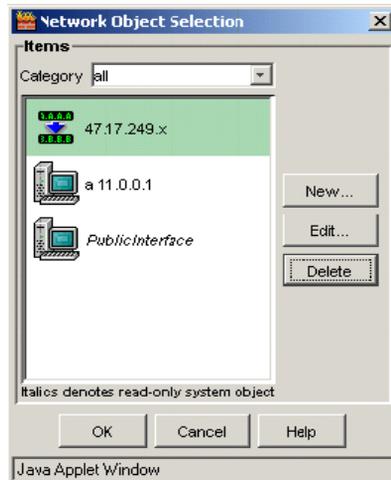
- # specifies the ordering of the rules within the section. The order only applies to the section in which the rule appears and does not have meaning across the entire policy.

- Source and Destination specify the source and destination network object for the rule. It is possible to add more than one source or destination address to a rule. You can modify these attributes by right-clicking on a column in the cell, which then brings up a procedure menu. Clicking on Add displays the Network Object Selection dialog box (Figure 55). 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: The NOT operand allows you to specify which networks you do not want use NAT.

Figure 55 Network Object Selection box



Italicized objects in the list are read-only and cannot be modified. The New, Edit, and Delete buttons in this dialog box allow you to create, edit and delete network objects.

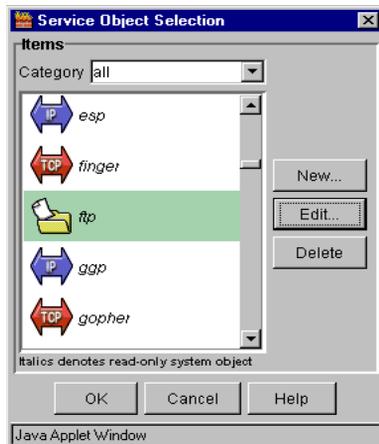
Clicking on Edit displays the Network Object Edit dialog box, as shown in [Figure 56 on page 102](#). This dialog box allows you to modify the attributes for the selected network object.

Figure 56 Network object edit box

Clicking on Delete removes the selected network object. If the object that you want to delete is the last object, it returns to the default value.

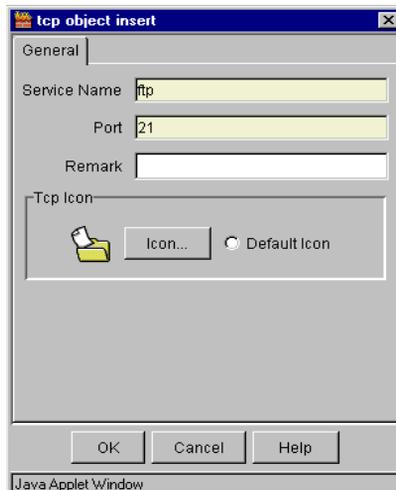
- Service specifies which service objects are handled by the selected rule. Right-clicking on the cell displays the standard procedure menu (Add or Edit).

Clicking on Add accesses the Service Object Selection dialog box (Figure 57), which allows you to 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).

Figure 57 Service Object Selection box

Italicized objects in the list are read-only and cannot be modified. The New, Edit, and Delete buttons in this dialog box allow you to create, edit and delete service objects. Clicking on Edit displays the Service Object Edit dialog box (Figure 58). This dialog box allows you to modify the attributes for the selected service object.

Figure 58 Tcp object insert box



Clicking on Delete removes the selected service object from the cell. If the object to be deleted is the last object in the cell, the cell returns to its default value (in this case, Any).

Clicking on Copy, Cut, or Paste performs those operations on the current service object.

- NAT Action specifies the action that occurs when the rule is activated. Right-clicking on the cell displays an option list containing three items: None, Static, Pooled, Port Mapping, and Port Forwarding. Clicking on one of these items sets the cell to the selected state, as shown in [Figure 59 on page 104](#).

Figure 59 NAT actions

- Translated Source – specifies the source IP address of the first packet (static, pooled, port). You can modify this attribute by right-clicking on a column in the cell. It is possible to 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. You can modify this attribute by right-clicking on a column in the cell, which then brings up a procedure menu. It is possible to add more than one destination address to a rule.
- Status – specifies the status of the particular rule. The status can be either Enabled or Disabled. (Figure 60)

Figure 60 Status column options

- Remark allows you to attach a remark to a particular rule. When you right click on Remark and choose Add or Edit remark, a dialog box appears where you can type a comment.

Creating a new policy

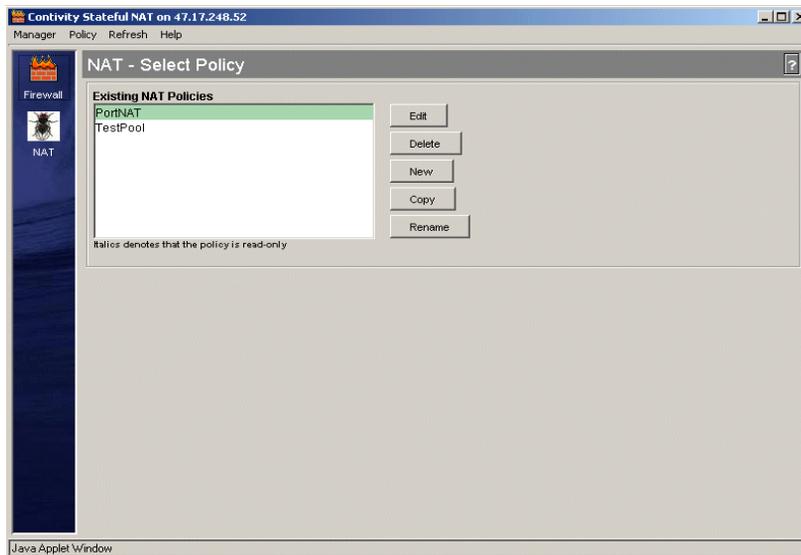
To configure NAT policies:

- 1 Go to the Services > Firewall/NAT screen.
- 2 Click on the Interface NAT check box under the Enabled column.
- 3 Select the NAT Policy from the list.
- 4 Click on the Manage Policies button. The NAT - Select Policy screen appears, as shown in [Figure 61 on page 106](#). This screen allows you to create, edit, delete, copy, or rename a NAT policy. Bold denotes the policy that is currently applied to the Contivity Secure IP Services Gateway and italics denotes read-only policies.

The System Default policy is always listed. This read-only policy defines the NAT behavior when no user-defined policies have been applied or when the selected policy is not available.



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

Figure 61 NAT policy

- 5 Click on New to create a new policy. The New Policy dialog box appears.
- 6 Enter the policy name and click on OK. The name must begin with a letter and may not contain the : + =] , ; " characters. The NAT - Edit Policy: *<policyname>* screen appears with no rules defined. In this screen, you can add, delete, and modify the rules for the policy.
- 7 You can select the rule group as follows:
 - Implied rules (view only)
 - Override rules
 - Interface-specific rules
 - Default rules
- 8 Select either Source Interface Rules or Destination Interface Rules.
- 9 Right-click on the appropriate cell to add a new rule.
- 10 Repeat these steps to add more rules.
- 11 Select Policy and click on Save Policy to save your changes.
- 12 When the policies are saved, go to the Manage menu and click on Close Manager.

Adding a policy

To add a new policy:

- 1 Click on the New button. The New Policy dialog box appears and prompts you for a name for the new policy.
- 2 Enter the policy name. The name must begin with a letter and may not contain the : + =] , ; " characters.
- 3 Click on OK to go to the Policy Edit screen, which has a blank NAT policy, or click on Cancel to return to the policy selection screen.

Deleting an existing policy

You cannot delete a read-only policy or the policy that is currently applied to the Contivity gateway. If you select one of these policies, the Delete button is not enabled. To delete an existing policy:

- 1 Select the policy that you want to delete and click on the Delete button. The delete policy confirmation dialog box appears.
- 2 Click on OK to delete the selected policy.

Copying an existing policy

To copy a NAT policy:

- 1 Select the policy that you want to copy.
- 2 Click on the Copy button. The copy dialog box appears.
- 3 Enter a name for the copied policy.
- 4 Click on OK.

The new policy appears in the list of policies in the NAT policies screen. This policy contains the same rules as the policy from which it was copied.

Renaming an existing policy

You cannot rename a read-only policy or the policy that is applied to the Contivity gateway. If you select a read-only policy, the Rename button is not enabled. To rename an existing policy:

- 1 Select the policy that you want to rename.
- 2 Click on the Rename button. The Rename dialog box appears.
- 3 Enter the new name of the policy.
- 4 Click on OK.

Sample NAT procedures

The following sections describe the steps for sample NAT procedures.

Figure 62 Edit Policy screen

#	Source	Destination	Service	NAT Action	Translated Source	Translated Destination	Status	Remark
1	a 11.0.0.1	a 12.0.0.1	http	port_fwd	any	a 192.32.250.205	✓	

For the following configuration on the Contivity gateway, 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.

Interface NAT with RIP

This sample shows interface NAT with RIP:

- 1 On the Contivity gateway, enable Interface NAT and attach the above NAT policy to Interface NAT.

- 2 Go to the Routing > RIP screen and enable RIP.
- 3 Go to the Routing > Policy screen and 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 should get the reply back.

Interface NAT with OSPF

This sample shows interface NAT with OSPF:

- 1 On the Contivity gateway, enable Interface NAT and attach the above NAT policy to Interface NAT.
- 2 Go to the Routing > OSPF screen and enable OSPF.
- 3 Go to the Routing > policy screen and 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 should get the reply back.

Branch Office NAT with RIP

This sample shows NAT on a branch office with dynamic routing enabled.

- 1 On the Contivity gateway, go to the Profiles > Branch Office screen and 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 enable RIP. Enable NAT and create the above NAT policy.
- 3 Go to the Routing > RIP screen and enable RIP.

- 4 Go to Routing > policy screen and 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 (Contivity), go to the Profiles > Branch Office screen and 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 enable RIP.
- 9 Go to the Routing > RIP screen and enable RIP.
- 10 Send a ping request from 10.0.1.1 to 20.0.0.1. Ping should get the reply back.

Branch Office NAT with OSPF

This sample shows NAT on a branch office with dynamic routing enabled.

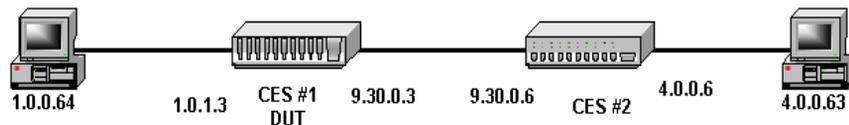
- 1 On Contivity-1, go to the Profiles > Branch Office screen and create a branch office with a local end point as 20.0.9.100 and remote end point as 20.0.9.1.
- 2 Enable Dynamic Routing for that Branch Office and Enable OSPF. Enable NAT and create the above NAT policy.
- 3 Go to the Routing > OSPF screen and enable OSPF.
- 4 Go to the Routing > policy screen and 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 (Contivity), go to the Profiles > Branch Office screen and 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 enable OSPF.

- 9 Go to Routing > OSPF and enable OSPF.
- 10 Send a ping request from 10.0.1.1 to 20.0.0.1. Ping should get the reply back.

Sample branch office NAT configuration

This configuration example (Figure 63) adds a NAT static rule with a single host as the source.

Figure 63 NAT configuration example



- 1 Using a browser with valid JRE (1.4.1_02), go to Services > Firewall/NAT and click on Manage Policies.
- 2 Log in to Contivity Stateful NAT.
- 3 Click on New, enter the policy name, and click on OK.
- 4 Right Click on the # sign click on Add New Rule.
- 5 Right click on Orig Src. The Network Object Selection window opens. This is used to create network objects. Once created, a network object may be applied to any Address column of the rule.
- 6 Click on New, select Host, and click on OK.
- 7 In Host Object Insert window, enter the host name and IP address. Sqa64; 1.0.0.64. Click OK twice to return to the NAT Translate Action window.
- 8 Right click on Trans Src.
- 9 Click New, select Host and click on OK.
- 10 In Host Object Insert window, enter information for the translated host: Host Name = Sqa64Trans; IP Address 30.0.0.64. Click on OK twice to return to the NAT Translate Action window.
- 11 Click on Policy > Save policy. A popup advising you to “Please wait ...” must appear to show that the policy was saved.
- 12 Go to the Profiles > Branch Office screen and select a working branch office tunnel and click on Configure.

- 13 From the NAT menu, select the policy you added and click on OK.
- 14 From SQA64, use ping, Telnet or another application to pass traffic over the tunnel.

Configuring NAT with the Contivity Stateful Firewall

To use NAT on the Contivity gateway with the Contivity Stateful Firewall where the NAT address is within the same subnet as the public interface:

- 1 Go to Profiles > NAT and create a NAT policy called Static by entering static in the name field and click create.
- 2 Click on Add to add a NAT rule.
 - a Leave the Translation type set to static.
 - b Add the internal Contivity gateway address (for example, 10.4.4.204) as the start and the end internal address.
 - c Add the external address (for example, 192.168.4.204) as the starting external address.
- 3 Go to System > Forwarding and enable Proxy ARP for Physical Interfaces and click on OK.
- 4 Enable Interface NAT and select the NAT rule created in Steps 1 and 2.



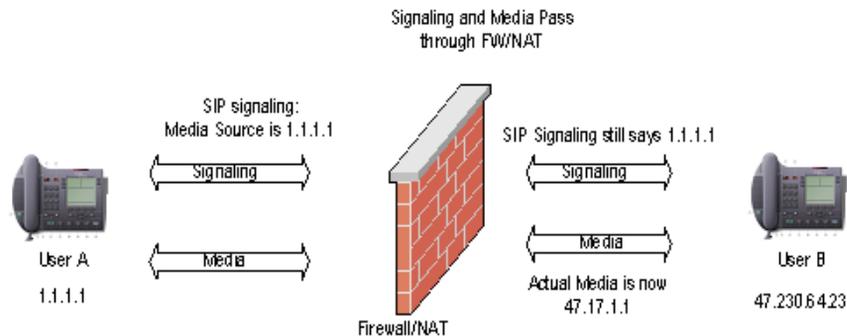
Note: The Contivity Stateful Firewall must have an Allow All policy set.

NAT ALG for SIP

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

[Figure 64 on page 113](#) illustrates the problem caused by NAT for Session Initiation Protocol (SIP) signaling.

Figure 64 NAT and SIP



In Figure 64:

- 1 User A sends an invite to User B.
- 2 The NAT translates the Layer 3 address, but not the Layer 5 (SIP/Session Description Protocol [SDP]) addresses.
- 3 User B receives the invite and responds back to the NAT address. The signaling gets completed (for example, 200 OK).
- 4 User A sends RTP to User B's SDP c= / m= address: port.
- 5 User B tries to send RTP to User A's c= / m= address: port, but this fails since it cannot route to User A (the SDP address and port did not receive the NAT) resulting 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 will not get to User A because the header address did not receive the NAT. This leaves the state of User A for that session to be up until User A hangs up.

Two of the solutions that correct the NAT traversal issue are:

- Application level gateways (ALG)
- Address/port discovery

The address/port discovery method was discussed in [“Address/Port discovery” on page 90](#). This section focuses on NAT ALG for SIP to support VoIP phones that use SIP as their signaling protocol.

Application level gateways (ALG)

NAT ALG translates any embedded IP addresses and port numbers contained in an application’s protocol messages. Support is provided for FTP, ICMP, Berkeley R commands, NetBIOS, IPsec (ESP only), and SNMP. For application traffic flows that embed an IP address in the data portion (such as FTP or NetBIOS), you must have an ALG.

SNMP ALG support allows you to use SNMP traps with NAT. The data within the SNMP traps are translated, preventing inconsistencies within the packet. The SNMP ALG is applied to SNMP traps originating from the Contivity gateway only if there are NAT rules that translate traffic originating from the gateway. You must enable the SNMP management system to send SNMP Gets from the Admin > SNMP screen.

The NAT ALG provides support for SIP traffic to and from Nortel i2004 phones model NTEX00 and the SIP Server MCS 5100.

Configuring NAT ALG for SIP

NAT ALG for SIP can be enabled or disabled with either the GUI or the CLI. To learn the CLI commands, go to *Reference for the Contivity Secure IP Services Gateway Command Line Interface (CLI)*.

To configure NAT ALG for SIP:

- 1 From the Services > Firewall/NAT page, click Edit in the Contivity Firewall row.

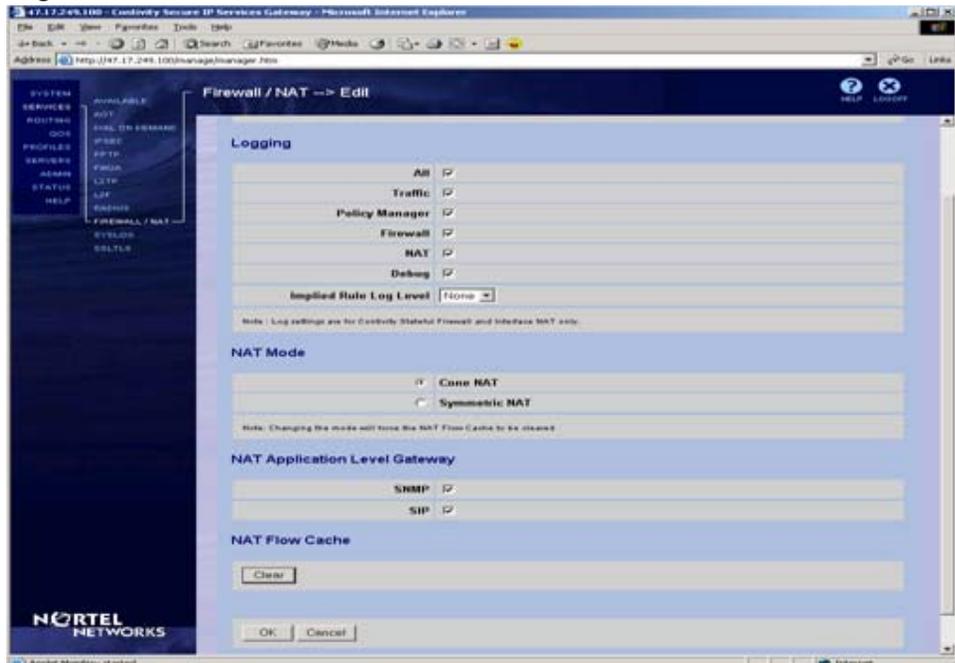
The Firewall/NAT > Edit page opens.

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

The Firewall/NAT page re-opens with the new configuration applied.

Figure 65 shows the interface where you enable NAT ALG for SIP. Note that the box beside SIP is checked.

Figure 65 SIP enabled



Note: If Firewall is enabled in the Logging section, the user will receive a log with Firewall events in it.

Firewall SIP ALG

Firewalls, by default, do not have the intelligence to identify port numbers within the payload of signaling protocols and cannot dynamically open ports for media traversal, resulting in blocking of voice traffic. Firewalls are designed to operate with layer 3/layer 4 information and cannot access information in higher layer protocols.

The development of ALGs for the VoIP signaling protocols solved 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 indicate that NAT should perform an application level translation. The ALG then performs the address/port mapping and state setup to ensure that the data channels get mapped 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 increases network security by restricting the voice traffic to active sessions only.

Hairpinning

Hairpinning is required when two IP phones behind the same NAT want to communicate. Contivity NAT blocks packets coming from the private side of the NAT that are destined for the private side for which a NAT binding to a specific port already exists. This does not allow peer-to-peer communication between two endpoints behind the same NAT if they try to use their public address. Hairpinning corrects this problem by examining the destination address of a packet, evaluating the destination address NAT binding, and making a determination on the requirement for hairpinning.

NAT hairpinning does payload translation on SIP and UNISTIM messages.

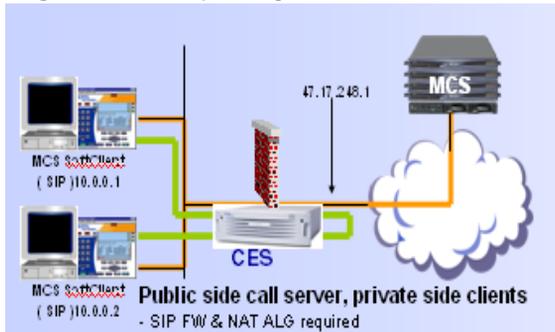
Hairpinning with SIP

A special problem 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. When the call server is queried for the IP address of the person that is being called, it responds with the public IP address. It 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.

Figure 66 shows hairpinning support required for VOIP Media. The MCS call server sees both private side phones as having a 47.17.248.1:x address, telling the private side caller that the called has a 47.17.248.1:x IP, and vice-versa.

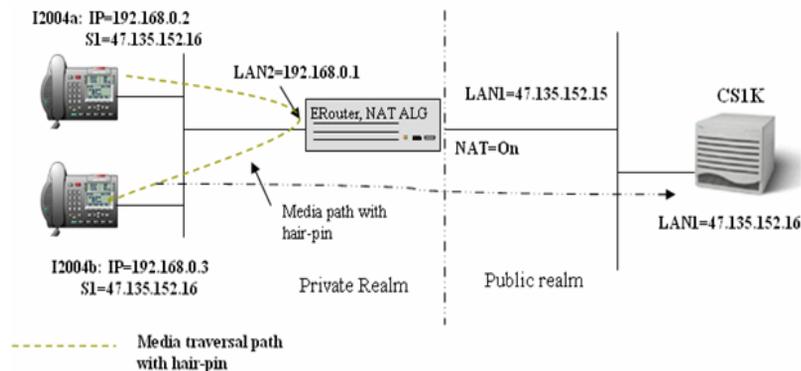
Figure 66 Hairpinning with SIP



Hairpinning with a UNISTIM call server

When a UNISTIM 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 has no idea what these packets are for. However, if the NAT device supports hairpinning, it redirects the packets to the right destination, helping generate the voice path.

[Figure 67 on page 118](#) shows an intra-realm call with hairpinning.

Figure 67 Intra-realm call with hairpinning

In Figure 67, both i2004a and i2004b are behind the same NAT, and registered into the same CS1K TPS server. UNISTIM messages are encrypted and the ERouter NAT is not able to translate UNISTIM messages payload. Upon successful registration of both IP phones, ERouter NAT generates the following NAT table entries:

Table 3

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

When 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

When 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.

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

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

When 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). It translates the destination address from 47.135.152.15:52002 to 192.168.0.3:5200.

It further compares the source address in the packet header against the Internal Address entries on its NAT table. It finds a match (192.168.0.2:5200). It translates the source address from 192.168.0.2:5200 to 47.135.152.15:52000. It forwards the translated packet to i2004b.

Similarly, when 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). It translates the destination address from 47.135.152.15:52000 to 192.168.0.2:5200.

It further compares the source address in the packet header against the Internal Address entries on its NAT table. It finds a match (192.168.0.3:5200). It translates the source address from 192.168.0.3:5200 to 47.135.152.15:52002. It forwards the translated packet to i2004a.



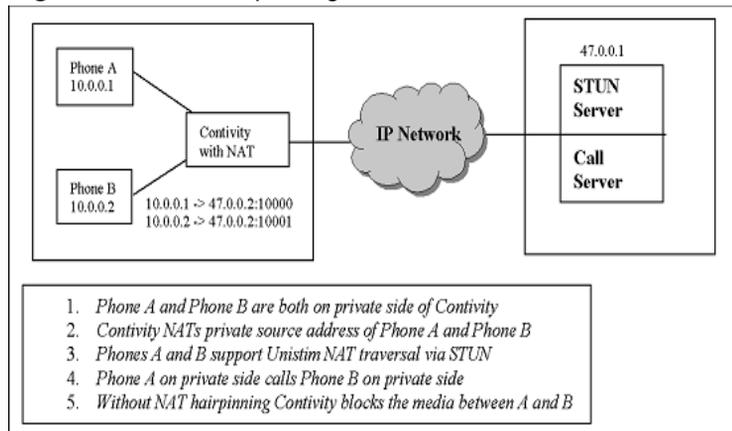
Note: The hairpinning support is part of the solution, and can co-exist with the other portions of the solution. For example, with non-encrypted UNISTIM messages, the hair-pinning logic automatically turns off, and a direct media path can be achieved.

Hairpinning with a STUN server

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

The diagram in Figure 68 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 Contivity initiates a call to Phone B on the private side. The call gets established and then Phone A starts to send media to Phone B and vice versa with public NAT'd destination addresses in the media packets. Contivity NAT, unaware that the voice packets need NAT hairpinning, will block the media packets. The enhancement makes the NAT hairpinning aware so that it examines the destination address of a packet, evaluates the destination address NAT binding, and makes a determination on the requirement for hairpinning.

Figure 68 NAT Hairpinning



Hairpinning requirements

NAT Hairpinning has two requirements:

- Since IP phones may not accept packets from arbitrary IP addresses, the source IP address must be the public IP address of the NAT.
- If the device is performing NAT on a VPN tunnel, packets sent from private devices to the assigned VPN IP are hairpinned back without entering the VPN tunnel. When the packets reach the private endpoint, the source IP address MUST be the assigned VPN IP address.

Configuring hairpinning

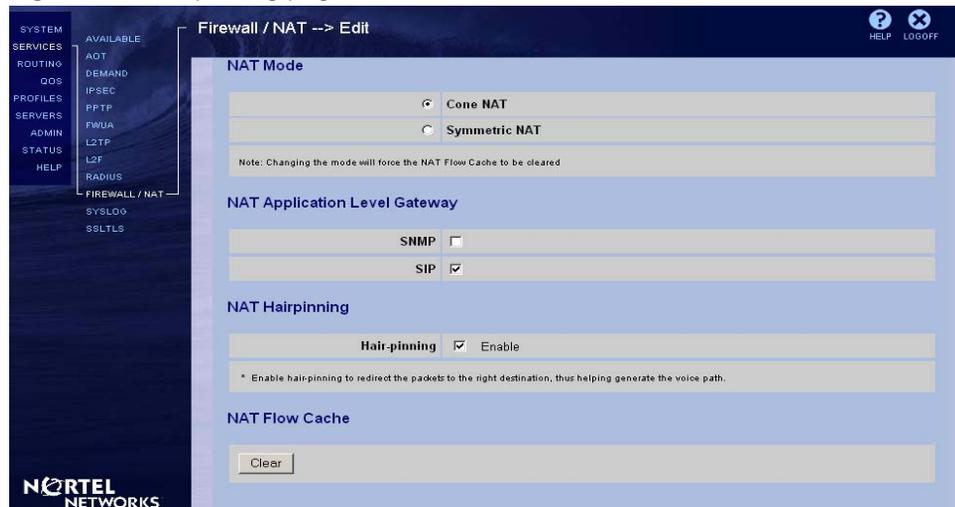
The hairpinning of packets can be turned on/off by the user from the GUI or using the CLI. To learn the CLI commands, go to *Reference for the Contivity Secure IP Services Gateway Command Line Interface (CLI)*.

To configure hairpinning:

- 1 Select Services > Firewall/NAT.
- 2 Click Edit beside Contivity Firewall.
- 3 In the NAT Hairpinning section of the Firewall/NAT > Edit page, check the hairpinning box.
- 4 Click OK.

Figure 69 shows hairpinning enabled.

Figure 69 Hairpinning page



Hairpinning statistics are shown on Status -> Statistics -> NAT Stats page.

Time-outs

When a session terminates, it causes NAT to delete the associated translations. However, if a server goes down unexpectedly, the associated translation must age out so that the available translation addresses will not be exhausted. The NAT time-outs are grouped by protocol as follows:

- 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 that have the source address translated
- Destination Translated – Number of packets that have the destination address translated
- Flows Translated – Number of flows that are translated by NAT service
- No Action – Number of flows for which no translation has been done
- Dropped – Number of packets dropped because NAT could not translate the source/destination address
- Pooled Address Translations failed – Number of packets dropped because NAT could not map a new address from available address pool
- Port Translations failed – Number of packets dropped because NAT could not map a new port for translation

You can view the NAT statistics on the Status > Statistics screen.

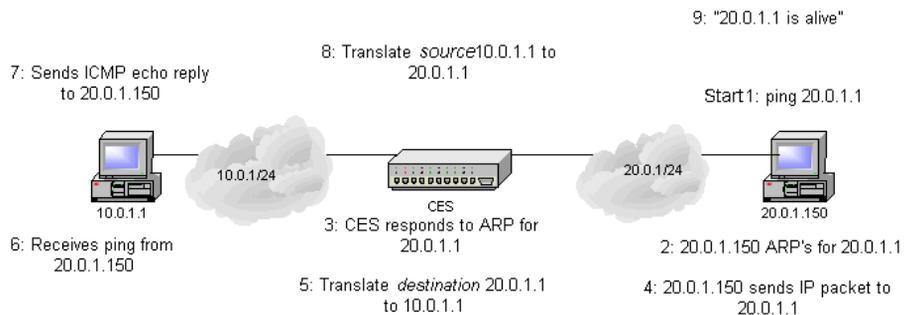
Proxy ARP

Proxy ARP is needed if the translated address assigned by NAT to a private host makes it appear as if that private host is on the other hosts network. The other host will ARP and not get a response unless Proxy ARP is enabled for physical interfaces on the Contivity gateway.

In [Figure 70](#), 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 is broadcast to the network.
- 3 The Contivity gateway 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 Contivity gateway.
- 8 The packet's 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 70 Proxy ARP example



Chapter 5

Configuring firewall user authentication

Firewall user authentication (FWUA) allows you to require users to log in to the Contivity Stateful Firewall before they are granted network access. FWUA provides more granular security controls against unauthorized firewall use and can be used for user-level accounting information for firewall users.

FFWUA extends and enforces user authentication on traffic between branch office (BO) tunnels. It can also be applied on non-tunneled traffic when the gateway acts as a router and firewall edge device.

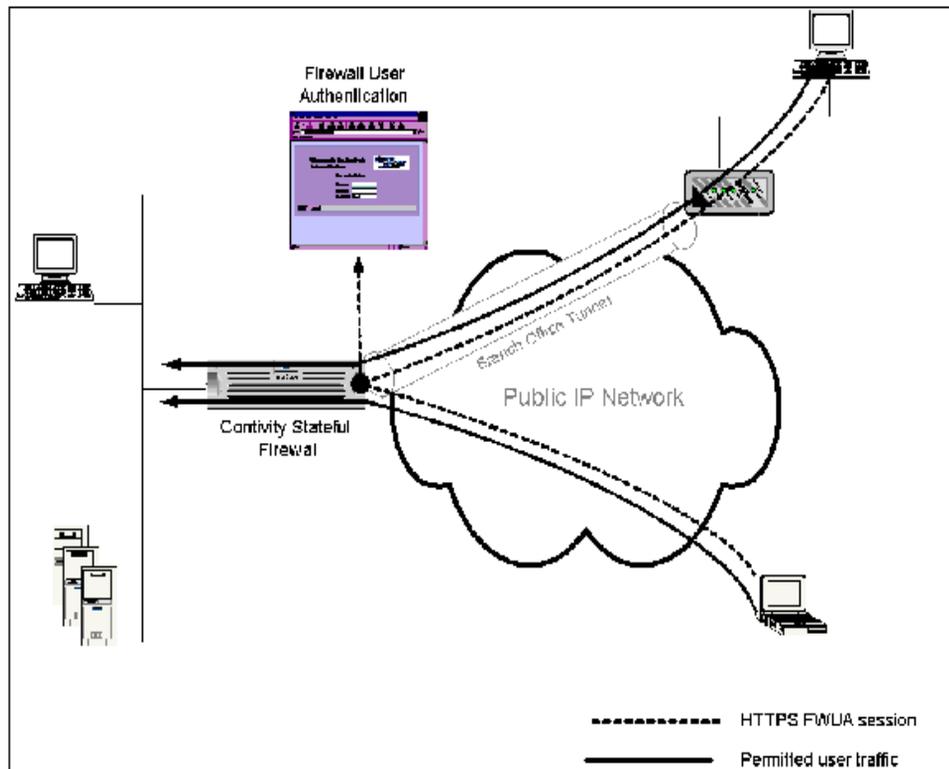
FWUA uses the existing authentication services, with username and passwords supported for both internal authentication services (LDAP) or external authentication services (RADIUS or LDAP proxy). Example 1 is based on authentication by internal LDAP and Example 2 is based on authentication by an external 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 also can be applied to non-tunneled traffic FWUA when the Contivity gateway acts as a router and firewall edge device.

Policies within the Contivity gateway can contain a User Authentication specification for any rule. Users must have an active HTTPS login session registered with the User Authentication Table Manager (UATM) before being permitted access granted by the rule. Users who do not have an existing login session registered with the UATM will not be granted access even if the traffic profile is explicitly permitted by the rule. User UATM sessions are mapped to the active session table by source IP address.

[Figure 71 on page 126](#) is an example of FWUA.

Figure 71 FWUA example



Secure HTTP (HTTPS) support provides a secured communication channel for administration traffic to the Contivity gateway system and for firewall users to provide their authentication credentials to the Contivity Stateful Firewall. A FWUA user directs their HTTPS-enabled Web browser to a specific Uniform Resource Locator (URL) designated for the FWUA login on the Contivity gateway. Both Secure Socket Layer (SSL) 2.0/3.0 and Transport Layer Security (TLS) 1.0 are supported. The following suites are supported:

- Symmetric Ciphers -- RC4, DES, and Triple DES (Cipher Block Chaining or CBC)
- Public Key Cryptography and Key Agreement Protocols -- RSA and Diffie-Hellman
- Authentication Codes and Hash Algorithms -- MD5 and SHA-1

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

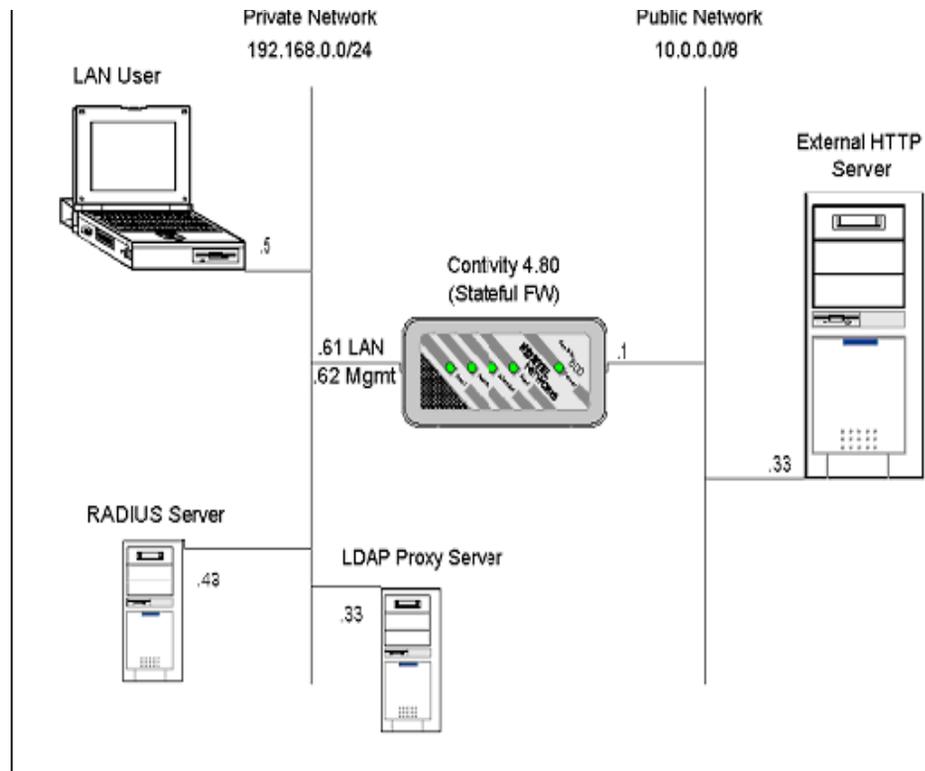
- EDH-RSA-DES-CBC3-SHA
- DES-CBC3-SHA
- RC4-SHA
- RC4-MD5
- EXP1024-RC4-SHA
- EXP1024-DES-CBC-SHA
- EXP1024-RC4-MD5
- EDH-RSA-DES-CBC-SHA
- DES-CBC-SHA
- EXP-EDH-RSA-DES-CBC-SHA
- EXP-EDH-DSS-DES-CBC-SHA
- EXP-DES-CBC-SHA

The authentication facilities for FWUA use the existing authentication services currently available on the Contivity gateway with the exception of RADIUS based tokens and digital certificates. By using the existing authentication services, all user-level accounting mechanisms that are available for VPN users are also available for FWUA users.

Prerequisites for using FWUA are:

- The Contivity Stateful Firewall must be running to configure and process FWUA sessions.
- SSL/TLS must be enabled, which also requires the Contivity gateway to have a valid digital certificate installed to support HTTPS communication.
- FWUA users must have an HTTPS-enabled Web browser with a compatible SSL/TLS crypto suite.

[Figure 72 on page 128](#) is an example of FWUA configuration.

Figure 72 FWUA configuration

To configure FWUA:

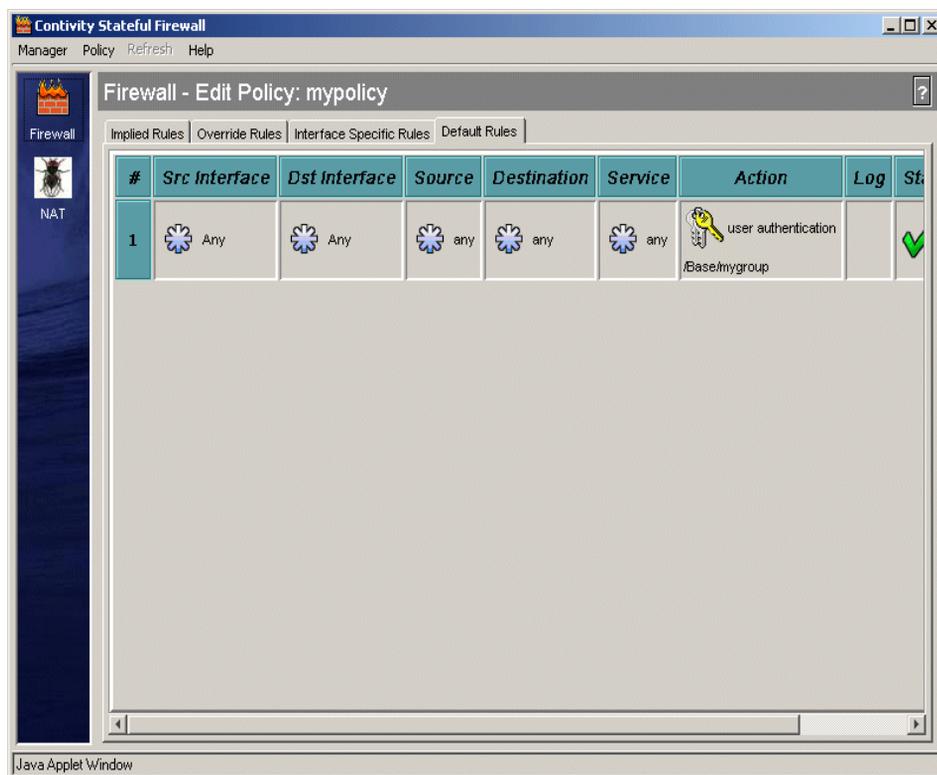
- 1** On the Services > Available screen, check Public and Private for Firewall User Authentication.
- 2** On the Services > Firewall UA Settings screen, enter the text for a welcome banner, the port value (default 8000), and the default max session value (default 1000). You will add RADIUS or LDAP proxy authentication servers to the authentication order later.
- 3** On the Services > SSL/TLS screen, check the desired Ciphers (default all) and enter an existing X.509 digital server certificate pre configured for this Contivity gateway (for example, CN=ces48, O=CSE, C=US). If no available certificates appear in the list, no server certificates are defined on your Contivity gateway or the existing server certificate is disabled.

- 4 On the Profiles > Users > User Management > Edit User screen, create a FWUA user profile in internal LDAP. Enter the user name, select the Group and create a password.
- 5 On the Services > Firewall/NAT > Manage Policies screen, create a firewall policy.



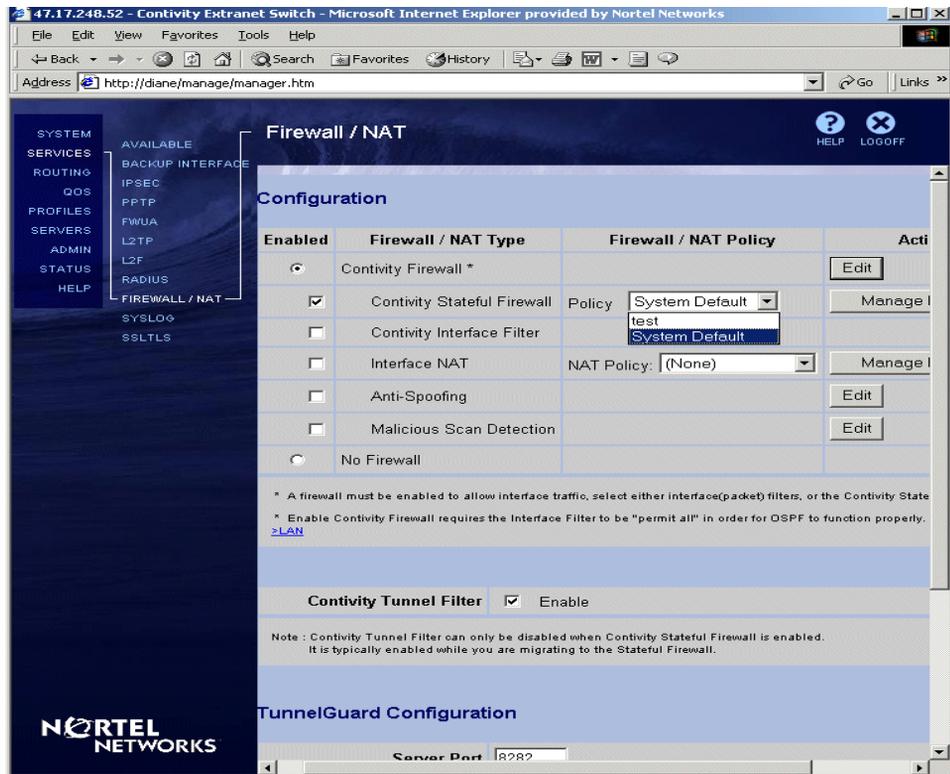
Note: The firewall UI requires JRE 1.4.1_02 or later. If you do not have a sufficient JRE you are prompted by the Contivity gateway to download and install JRE 1.4.1_02 directly from the Contivity gateway. A copy of JRE 1.4.1_02 is also available on the Contivity gateway server CD.

- a After you log in, click on New and enter the name of the policy.
- b Select the Default Rules tab, right-click on the # sign and select Add New Rule. [Figure 73 on page 130](#) shows the Firewall policy screen with Default Rules tab selected.

Figure 73 Firewall policy screen

- c** Right-click on the Action cell and 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.
- e** Select Policy > Save Policy and Manager > Exit CSF. Check Contivity Stateful Firewall to be sure it is enabled. [Figure 74 on page 131](#) shows Contivity Stateful Firewall enabled.

Figure 74 Firewall/Nat screen



- f Select the new firewall policy (refresh the screen for the new policy to appear in the list), and click on OK.



Note: You must have a valid Contivity Stateful Firewall license key installed. Also you must reboot the Contivity gateway if you are enabling the Stateful Firewall for the first time. You should disable the Contivity gateway tunnel filters as they are no longer needed.

You can test the FWUA rule by trying to communicate through the Contivity gateway. Communication attempts should fail.

- 6 Direct your HTTPS enabled browser to the predefined FWUA login URL on the Contivity gateway and log in to the firewall using the FWUA user profile that you created. The FWUA login URL follows the format of `https://Contivityhostname:port/FWUA.htm` or `https://ContivityIPaddress:port/FWUA.htm` where Contivityhostname or ContivityIPaddress resolves to a

Contivity gateway interface (not management IP). The port is the port number you specified on the Services > FWUA menu.



Note: If the domain Contivity gateway digital server certificate is not part of a certificate domain trusted by the users Web browser (the user does not have a certificate issued by the same CA) or the domain listed on the Contivity gateway certificate does not match the DNS domain of the Contivity gateway, you will be prompted by your Web browser with a security alert dialog box. Click on 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 will keep an active FWUA session in the UATM.

- 7 Attempt to communicate through the firewall again. Communication attempts should be successful.
- 8 To modify the current FWUA configuration to accommodate external authentication methods, go to Services > FWUA > Add RADIUS or Add LDAP Authentication Server. The Associated Group specifies what group the RADIUS or LDAP Proxy Authentication that users will obtain their privileges as defined on the Server > RADIUS Auth or the Server > LDAP Proxy menus. If the /Base group is configured to authenticate RADIUS or LDAP Proxy Auth users for VPN connections, it will also be used to authenticate FWUA users.

Chapter 6

Configuring QoS

The Contivity Secure IP Services Gateway supports two internal quality of service (QoS) mechanisms and can also participate in external network signaling to enhance performance. Forwarding priority allows for prioritized traffic, and Call admission priority allows you to reserve connection resources for high-priority users. In addition, external QoS using Resource ReSerVation Protocol (RSVP) signals the public network to reserve a portion of the network's bandwidth for a specific connection.

QoS provides the option of dropping data that exceeds configured traffic conditioning assured forwarding rates. This allows for guaranteed bandwidth based on Diffserv code points that guarantees a fixed percentage of total bandwidth to each of several applications.

Traffic conditioning by DSCP provides a method to limit traffic at ingress to the Contivity based on Diffserv Code Point (DSCP) value. This ensures that particular DSCP values obtain the desired amount of egress bandwidth. Traffic that exceeds the configured rate for a particular DSCP is dropped in ingress to the Contivity gateway.

Configuring classifiers

An MF Classifier can be defined for an interface (interface MF). The interface MF-Classifier is applied to routing traffic going through that interface.

To configure an MF classifier:

- 1 Go to the QoS > Classifiers screen. The Current Multi-Field (MF) Classifiers list includes all existing MF classifiers.

- 2 Select from the Current Multi-Field (MF) Classifiers and click Edit to edit the rules for that MF Classifier. The Edit Rule screen appears.

The Rules in Classifier list shows all of the rules that are applied to the MF Classifier.

The Available Rules list shows all of the existing rules. You can select rules from this list to move them into the Rules in Classifier list and apply them to the MF Classifier.

- 3 Click on a rule from the Available Rules list on the right of the screen, then click on the left arrow to add the rule. This adds the selected rule to the current rules list. The new rule is added after the rule currently selected in the Rules in Classifier list.
- 4 Click on the Edit button to edit an existing rule. The Edit/Create Rules screen appears. The Classifier Rule for field shows the name of the rule.
- 5 Enter the source and destination addresses to limit the rule to acting on packets from and to these addresses. Source and destination are relative to the direction of the rule.
- 6 Click on the Modify button next to the Source and Destination Address fields to edit either of these fields. The DiffServ Rules Definition Address screen appears.
- 7 Select the appropriate protocol from the list. The default list of protocols include:
 - ICMP — Internet Control Message Protocol is a Network protocol layer. The PING utility generates ICMP packets. PING is often used to see if a system's network is available.
 - IP — Internet Protocol is a Network layer protocol in the TCP/IP stack that offers a connectionless internetwork service. IP packets that are encapsulated within other packets create "IP over IP." Multicast IP packets (packets that have multicast destinations), carried between networks that support multicasting over intermediate networks that do not, are the most common implementation. Conferences and other services that are offered through Multicast Backbone (MBONE) are examples.
 - TCP — Transmission Control Protocol is a transport layer protocol in the TCP/IP protocol stack. This is a connection-oriented protocol that provides reliable full-duplex data transmission. Web browsers using HTTP and FTP are examples.

- UDP — User Datagram Protocol is a transport layer protocol in the UDP/IP protocol stack. UDP is a connectionless service that exchanges datagrams without acknowledgment or delivery guarantees, and therefore requires that error handling and retransmissions are handled by other protocols. DNS and WINS are examples.
- 8** Click on the Modify button next to the Protocol field to edit it.
 - 9** Click on the Modify button to the right of the TCP/UDP Source and Destination Port fields to edit them. You can filter packets to or from the Source and Destination ports to permit or deny any packets transferred by the Contivity gateway. The source or destination is relative to the direction of the rule.
 - 10** Click on the Modify button to the right of the Current DSCP Value field to create and edit the DSCP value and mask. The DSCP value and mask assignments allow packets that are already marked to retain their settings or to be remarked based on their previous DSCP value.
 - 11** Select the DSCP to be marked on the next meter, either expedited forwarding (EF) or an assured forwarding (AF) level, that this rule applies to data.

You can configure the assured forwarding queues option to drop data exceeding the configured rate. (EF excess data is always dropped.) This data is dropped on ingress and never enqueued. 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 will be the appropriate size.

Configuring bandwidth management

Bandwidth management capabilities let you manage the Contivity gateway CPU and interface bandwidth resources to ensure that tunneled sessions get predictable and adequate levels of service. Bandwidth management allows you to configure the Contivity gateway resources for users, branch offices, and interface-routed traffic. Bandwidth components keep track of and control the level of bandwidth being used on the physical interfaces and the tunnels.

Bandwidth management forces tunnels to conform to a set of rates. There are two rates (committed and excess) and excess action (mark or drop). Packets are given different drop preferences, depending on whether they are below committed rate (lowest drop preference), between committed and excess rate (higher drop

preference), and above excess rate (highest drop preference if excess action is Mark). When there is congestion on the Contivity gateway, packets are dropped according to their drop preference. When excess action is Drop, all the packets above excess action are dropped.

You can add call admission to guarantee that resources are available to support the committed bandwidth assigned to a user. This potentially denies a client access before the licensed limit of a Contivity gateway is reached. The available bandwidth is based on the Contivity gateway interface speed.

To configure bandwidth management:

- 1 Enable the advanced routing license on the Admin > Install Keys screen.
- 2 Go to the QoS > Bandwidth Management screen to define the bandwidth rates. The maximum rate that you can create is 10 Mbps. This must be defined in bits per second (10 Mbps=10000000).
- 3 Go to the Profiles > Groups > Groups > Edit > Connectivity screen. In the User Bandwidth Policy section, define the committed and excess bandwidth rates.
- 4 Set Bandwidth Management to Enabled.
- 5 Go to the QOS > Interfaces screen to set the over-subscription rate. Adjusting this ratio allows you to adjust for some users not using all of their allotted bandwidth simultaneously under normal circumstance. The default is 10:1.

Configuring Differentiated Services (DiffServ)

DiffServ settings classify and mark packets to receive specified per-hop forwarding behavior on each node along their path. Sophisticated classification, marking, policing, and shaping operations are implemented at network boundaries or hosts. Network resources are allocated to traffic streams by service provisioning policies that govern how traffic is marked and conditioned upon entry to a differentiated services-capable network, and how that traffic is forwarded within that network. Any DiffServ code points (DSCPs) not recognized are forwarded as if marked for the default behavior, Best Effort (BE).



Note: You can only have either DiffServ or Forwarding Priority active at any one time.

Anti-Replay must be disabled when using IPsec tunnels over LANs or WANs (the typical usage). If it is enabled, it causes DiffServ sorting to be incorrect. Anti-Replay does not acknowledge DiffServ and has its own methods of discarding packets, which adversely affects the DiffServ sorting.

To configure DiffServ:

- 1 Go to QOS > Interfaces and click on the Configure button in the DiffServ Edge section.
- 2 In the Multi-Field Classifier State field, select Enabled or Disabled to enable or disable the application of MF Classifiers on this interface.
- 3 In the Ingress (Inbound) field, select from the list the MF Classifier that you want to apply when packets are coming into this interface.
- 4 In the Egress (Outbound) field, select from the list the MF Classifier that you want to apply when packets are going out this interface.
- 5 In the Traffic Conditioning State field, select Enabled or Disabled to enable or disable traffic conditioning on this interface.

Traffic conditioning is the process of dropping and remarking a traffic stream in order to shape it into compliance with a traffic metering profile. For Expedited Forwarding (EF) and Assured Forwarding 1--Assured Forwarding 4 (AF1-AF4), you can configure a Traffic Conditioning Meter (in bps).

- For EF, the rate is used as an average rate, though at times traffic can burst as much as twice the configured rate. Traffic below the rate is forwarded; traffic above the rate is dropped.
- For AF1-AF4, any packets under the rate are marked as low drop precedence. Any packets under two times the configured rate are marked as medium drop precedence. Any packets above two times the configured rate are marked as high drop precedence.



Note: Enter values for EF and AF1--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 Enter a value, in bps, for the Expedited Forwarding (EF) Rates field. Nonconforming traffic is dropped.

- 7 Enter values, in bps, for the Assured Forwarding Rate fields (AF4--AF1). Also, configure the Excess Action field for each AF rate to either drop traffic exceeding the configured rate or to mark the traffic.
- 8 For Egress (Outbound) traffic conditioning, enter a value, in bps, for Expedited Forwarding Shaping Rate. Shaping is a process of delaying the packets in a stream to in order to conform to a defined traffic profile (the EF Shaping value. Nonconforming traffic is delayed, not dropped.

Using forwarding priority

Forwarding priority quality of service allows you to assign each user to one of four priority classes. Each class is guaranteed different maximum forwarding times between the interfaces of the Contivity gateway. For example, high-priority traffic generated by the company CEO would be protected from high-bandwidth traffic generated by lower-priority users. Or, you might assign the sales team to Priority 1 to make sure that they could always place orders, especially during the quarter-end rush.

The technology that supports forwarding priority is called *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 Contivity gateway is ever completely stalled. It is important to assign users to the four different class levels to make sure that they get the proper service and performance, especially during heavily congested times. QoS is only effective when all associated lines are capable of servicing the forwarding demands at the required speeds.

If a group profile has a forwarding priority of 1 (highest), it has the highest possible bandwidth guarantee and the lowest level of latency. Packets sent by this group are transmitted immediately even if there is heavy traffic on the Contivity gateway. Conversely, if a group profile has a forwarding priority of 4 (lowest), it has the least amount of bandwidth allocated and possibly the highest level of latency. Therefore, fewer packets sent by this group are transmitted while there are higher-level priority packets to be sent when the Contivity gateway traffic is heavy.

To illustrate how the Forwarding priority works, the example in Table 4 assumes heavy traffic and a queue of packets. Therefore, packets would be transmitted according to the approximate rates per “pass” that are cited in the table.

Table 4 Bandwidth allocation per 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 would 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.

Using call admission priority

Call admission priority quality of service allows you to assign each user group profile to one of four priority classes (from 1-high to 4-low) for call admission. The Contivity gateway can reserve connections for each class of user, guaranteeing that a large number of low-priority users do not lock out the high-priority users. When the Contivity gateway is servicing the maximum number of low-priority sessions, no further low-priority connections are accepted. Once a connection is accepted it is never dropped.

Since there is a maximum number of sessions supported on the Contivity gateway, it is important to assign users to the proper call admission priority classes. This ensures that connections are available to the appropriate users when there is heavy utilization. Although other callers may be permitted access to the Contivity gateway, this access is proportional to the assigned priority level for their group.

By default, any call is admitted access for the first 50 percent of connections, regardless of the assigned call admission 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, assuming a hypothetical maximum of 2000 sessions, Table 5 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 to 50%	All	1000
51 to 75%	1, 2, 3	500
76 to 90%	1, 2	300
91 to 100%	1	200

Table 6 shows the maximum number of connections available for each priority.

Table 6 Maximum connections per priority

Priority	Connections
1	2000
2	1800
3	1500
4	1000

Using RSVP

The Contivity gateway supports Resource ReSerVation protocol (RSVP) quality of service for the Internet. Successful external network-level quality of service 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 Contivity gateway is ready to take part in the RSVP signaling that is available in some network backbones and that will increase in the future. The Contivity gateway signals to the other devices on the public network and describes the level of bandwidth that is needed to ensure adequate performance. This amount of bandwidth is determined by both the data rate that the user has to the Internet, and

by the data rate of the link between the Internet and the Contivity gateway. This beginning stage of RSVP will be expanded over time to take advantage of advances in the technology. Meanwhile, you can build RSVP networks using the Contivity gateway to gain experience in this environment.

The two key components of RSVP are:

- PATH messages, which are constant announcements by the host system or the Contivity gateway that a certain amount of bandwidth must be kept 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, then RSVP-ready routers attempt the resource reservation. These routers actually reserve the resources requested if they are RSVP-compliant.

DSCP to 802.1p mapping

802.1p is a specification for prioritizing network traffic at the data link layer. 802.1p utilizes the User Priority field of the 802.1Q header. This priority extension allows Ethernet frames to be tagged 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.

Differentiated Services (DiffServ) provides Quality of Service (QoS) at the IP level by utilizing the 8 bit Type of Service field of the IPv4 header. Type of Service field is redefined as Differentiated Services (DS) field.

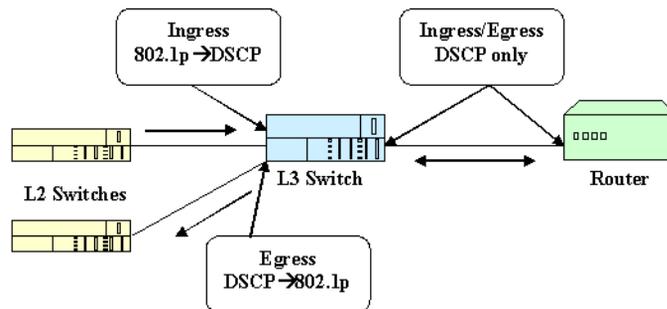
Six bits of the DS field are used as a Differentiated Services Code Point (DSCP) to select the Per Hop Behavior (PHB) a packet experiences at each node. DSCP identifies the priority of service a packet receives in the network. When a packet is being transmitted, the DSCP value of the inner header is copied to the outer IP header.

Support for DSCP to 802.1p mapping allows the Contivity gateway to tag frames for prioritization over public and private physical interfaces. It supports the ability to map DiffServ code point (DSCP) to 802.1p marking on ingress to or egress from the Contivity gateway. It provides the ability to 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. Some Ethernet switches may not be able to interpret the DSCP, but can interpret the 802.1p tag. By providing a consistent mapping between DSCP and 802.1p, the required end-to-end QoS behavior can be achieved over Ethernet networks.

In [Figure 75](#), the layer 2 switches are DSCP-unaware and the layer 3 switch and router are DSCP-aware. If a packet traveling from the layer 2 switch to the router has 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. When 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 75 Example 802.1p to DSCP mapping



When mappings are enabled and an incoming packet with 802.1p marking is received, the Contivity uses the default 802.1p to DSCP mappings shown in [Table 7](#).

Table 7 Default incoming 802.1p mappings

802.1p user priority	Maps to DSCP
7	CS7
6	EF
5	AF41

Table 7 Default incoming 802.1p mappings

802.1p user priority	Maps to DSCP
4	AF31
3	AF21
2	AF11
1	DF
0	DF

When mappings are enabled and an outgoing packet is to be sent out, Contivity uses the default DSCP to 802.1p mappings shown in Table 8.

Table 8 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

When mappings are disabled, the 802.1p tag value is ignored and normal multi-field classifier (MFC) action is applied to all packets.

To configure DSCP to 802.1p mapping:

- 1** Go to the QoS > Interfaces screen.
- 2** Select the interface the mappings to be applied from the list next to Current Interface.
- 3** Click on Display to display the selected interface (Fast Ethernet is displayed by default).
- 4** In the DSCP 802.1p mapping section, click on Configure.

- 5 On Dscp 802.1p Mapping screen, select either Custom or Standard for the Egress (outbound) and for Ingress (inbound).
- 6 If Custom setting selected, click on the “configure custom mappings link”.
- 7 Configure the “DSCP Class to 802.1p precedence mapping” and the “802.1p precedence to DSCP mapping” sections. See [Table 7 on page 142](#) and [Table 8 on page 143](#).
- 8 Click on OK.

Index

A

- access control filters 27
- actions on rules 58, 103
- anti-spoofing 25, 35
- application layer gateway 114
- attack detection 25
- available rules 70, 134

B

- bandwidth guarantees 138
- bandwidth management 135

C

- call admission
 - guarantees 139
 - priority 139
- cell menu 52
- columns
 - Dst interface 53
 - Src interface 53
- configuration
 - initial 29
 - verifying 61
- conversation 23

D

- default rules 51
- Differentiated Services (DiffServ) 136
- DSCP to 802.1p mapping 141
- dynamic many-to-one 78

E

- egress (outbound) queueing mode 138

F

- filter
 - rules 25
- filters
 - copy 72
 - edit current 69
 - storing 72
- firewall
 - imbedded 21
 - installation prerequisites 30
 - integrated 21
 - options 35, 37
- forwarding priority 138
 - quality of service 138
- FTP 71, 72

H

- header row menu 51

I

- ICMP
 - filter 134
- ICMP rule enforcement 37
- implied rules 46
- installation prerequisites 30
- interface
 - classifiers 133
- interfaces 24

interface-specific rules 49

IP packets 134

J

Java 2 Runtime Environment

Internet Explorer 31

Netscape 6 33

Netscape on Solaris 34

L

log

column 59

levels 59

logging

application-specific 38

HTTP 38

remote system 39

logging FTP 38

M

MBONE 134

menus

cell 52

header row 51

row 52

N

NAPT 78

NAT

branch office 95

creating policies 105

double 82

dynamic routing protocol 97

interface NAT 96

IPsec-aware 83

pooled translation 79

port forwarding 81

statistics 122

with Contivity Stateful Firewall 112

NAT SIP ALG 114

NAT Traversal 88

Network Address Translation 77

Network Address Translation (NAT) 27

network objects 55, 101

O

override rules 48

P

policies

actions 43

adding 44, 107

components 43

copying 45, 107

creating 43, 44

deleting 45, 107

editing 43

renaming 45, 108

selecting 44, 105

pooled translation type 97

port mapping 97

port translation (NAPT) 78

proxy ARP 123

publications

hard copy 18

Q

quality of service 133

forwarding priority 138

RSVP 140

R

remarks 59, 104

remote system logging 39

row menu 52

RSVP quality of service 140

rule column 53, 100

rules

- default 51
- implied 46
- in policies 24
- interface-specific 49
- navigating 46, 51, 100
- override 48

S

service objects 57, 102

SNMP 71

stateful inspection 23

- application 23
- TCP 23

static address

- NAT 80

static translation type 97

status 59, 104

syslog 39

system requirements 30

T

TCP

- filter 134

technical publications 18

traffic conditioning 137

U

UDP

- filter 135

