

Version 6.00

Part No. 315898-D Rev 00  
August 2005

600 Technology Park Drive  
Billerica, MA 01821-4130

# **Configuring Routing for the Contivity Secure IP Services Gateway**

**NORTEL**

## **Copyright © Nortel Networks 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.

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

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

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

---

<b>Preface</b> .....	<b>13</b>
Before you begin .....	13
Text conventions .....	13
Acronyms .....	15
Related publications .....	18
Hard-copy technical manuals .....	19
<b>Chapter 1</b>	
<b>Routing overview</b> .....	<b>21</b>
Integrated firewall and routing .....	22
Dynamic routing .....	22
VPN routing .....	22
Static routes .....	23
Route table .....	23
Routing status .....	24
<b>Chapter 2</b>	
<b>Route table</b> .....	<b>27</b>
Route table lookup .....	29
Route selection based on destination .....	30
Route selection based on precedence in route table .....	30
Route table options .....	31
<b>Chapter 3</b>	
<b>Configuring RIP</b> .....	<b>33</b>
Protecting against routing loops .....	34
Configuring RIP on the Contivity gateway .....	35
<b>Chapter 4</b>	
<b>Configuring OSPF</b> .....	<b>41</b>
Installing the Advanced Routing key .....	42
Installing the Premium Routing key .....	42

---

Virtual link support .....	43
Configuring OSPF on the Contivity gateway .....	43
Configuring equal-cost multipath .....	49
Configuring default routes .....	50
<b>Chapter 5</b>	
<b>Border Gateway Protocol .....</b>	<b>53</b>
RFCs .....	54
Installing the Border Gateway key .....	54
EBGP/IBGP peers .....	55
BGP peering and connection processing .....	55
BGP update processing .....	56
Unfeasible route processing .....	56
Feasible route processing .....	56
Path attribute processing .....	56
Keep Alive processing .....	58
BGP policies .....	58
Accept/Announce policies .....	59
Access (Prefix) lists .....	59
AS-Path regular expressions .....	61
Route maps .....	62
Configuring route maps .....	63
Multi-Hop BGP .....	65
Route Reflector .....	65
BGP communities .....	66
Configuring BGP on the Contivity gateway .....	68
Configuring Neighbors .....	70
Adding a Network .....	72
Configuring the Route Reflector .....	73
Configuring AS Path Access Lists .....	74
Configuring Community Lists .....	75
Health Check Support .....	77

---

<b>Chapter 6</b>	
<b>Configuring static routes</b> .....	<b>79</b>
Adding and editing static routes .....	79
Using ping to validate public default route .....	81
<b>Chapter 7</b>	
<b>Route policy service</b> .....	<b>83</b>
Redistribution of routes .....	85
Creating a policy list .....	86
Configuring route policy services (RPS) .....	87
<b>Chapter 8</b>	
<b>Client address redistribution</b> .....	<b>89</b>
<b>Chapter 9</b>	
<b>Configuring multicast relay</b> .....	<b>95</b>
<b>Chapter 10</b>	
<b>Configuring the Virtual Router Redundancy Protocol (VRRP)</b> .....	<b>99</b>
VRRP and dynamic routing for high availability .....	100
Configuring VRRP on the Contivity gateway .....	104
Configuring IP addresses for backups .....	105
Interface groups and critical interface failover .....	106
<b>Index</b> .....	<b>109</b>



---

## Figures

---

Figure 1	Interaction of OSPF, BGP, and RIP with the routing table . . . . .	28
Figure 2	Route maps . . . . .	64
Figure 3	BGP communities . . . . .	68
Figure 4	BGP window . . . . .	69
Figure 5	Neighbor configuration . . . . .	71
Figure 6	Route reflector configuration window . . . . .	73
Figure 7	AS Path access list window . . . . .	74
Figure 8	Community list . . . . .	76
Figure 9	Accept and announce policies . . . . .	84
Figure 10	Client address redistribution . . . . .	90
Figure 11	Aggregation for client address redistribution . . . . .	91
Figure 12	Sample high-availability environment . . . . .	102
Figure 13	VRRP and static tunnels . . . . .	103



---

## Tables

---

Table 1	Forwarding capabilities	21
Table 2	Routing status screen options	24
Table 3	IP Forward Table screen	32
Table 4	IP Route Table screen	32
Table 5	RIP Statistics screen	37
Table 6	RIP Database screen	37
Table 7	RIP Interfaces screen	38
Table 8	LSDB screen	46
Table 9	OSPF Dynamic Neighbors screen	46
Table 10	OSPF Interfaces screen	47
Table 11	OSPF Summary screen	47
Table 12	OSPF Statistics screen	48
Table 13	RFCs	54
Table 14	Path attribute types	57
Table 15	Redistribution rules	85
Table 16	Show user tunnel routes	93
Table 17	Multicast interface-specific rules example	96
Table 18	Multicast Statistics screen	97
Table 19	Multicast Interfaces screen	97



---

## Preface

---

This guide describes the Nortels\* Contivity\* Secure IP Services Gateway routing. It also provides information to help you configure routing.

### 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 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 <b>ldap-server source {external   internal}</b>, you must enter either <b>ldap-server source external</b> or <b>ldap-server source internal</b>, 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 <b>show ntp [associations]</b>, you can enter either <b>show ntp</b> or <b>show ntp associations</b>.</p> <p>Example: If the command syntax is <b>default rsvp [token-bucket {depth   rate}]</b>, you can enter <b>default rsvp</b>, <b>default rsvp token-bucket depth</b>, or <b>default rsvp token-bucket rate</b>.</p>
ellipsis points (. . . )	<p>Indicate that you repeat the last element of the command as needed.</p> <p>Example: If the command syntax is <b>more diskn:&lt;directory&gt;/...&lt;file_name&gt;</b>, you enter <b>more</b> 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 <b>ping &lt;ip_address&gt;</b>, <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 <b>terminal paging {off   on}</b> , you enter either <b>terminal paging off</b> or <b>terminal paging on</b> , but not both.

## Acronyms

This guide uses the following acronyms:

ABOT	asynchronous branch office tunnel
ABR	autonomous boundary router
ABSR	autonomous system boundary router
AS	autonomous system
ASBR	autonomous system border router
BGP	border gateway protocol
BOT	bisync over TCP transport service
CAR	client address redistribution
CMS	circuit mapping service
DN	distinguished name
DNS	domain name system
DR	designated router
EBGP	exterior border gateway protocol
ECMP	equal cost multipath
FEM	forwarding engine mapper
FTP	File Transfer Protocol
IBGP	interior border gateway protocol

IGP	interior gateway protocol
IP	Internet Protocol
IR	information retrieval
ISP	Internet service provider
L2TP	Layer 2 Tunneling Protocol
LAN	local area network
LDAP	lightweight directory access protocol
LSA	link state advertisement
LSDB	link state database
MBGP	multiprotocol BGP
MED	multi-exit discriminator
MD5	message digest
MIB	management information base
NAT	Network Address Translation
NLRE	network layer routing entries
NLRI	network layer reachability information
OSPF	Open Shortest Path First
PACE	packet context engine
PDN	public data network
POP	point-of-presence
PPP	Point-to-Point Protocol
PPTP	Point-to-Point Tunneling Protocol
RIB	Routing Information Base
RIP	Routing Information Protocol
RPA	routing protocol application
RPS	routing policy server
RR	route reflector
RTM	route table manager
SNMP	Simple Network Management Protocol

TCP	transmission control protocol
TTM	time to market
UDP	User Datagram Protocol
URL	uniform resource locator
VLSM	variable-length subnet masks
VPN	virtual private network
VRRP	Virtual Router Redundancy Protocol
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 Firewalls, Filters, NAT, and QoS for the Contivity Secure IP Services Gateway* provides instructions for configuring the Contivity Stateful Firewall and Contivity interface and tunnel filters.
- *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.
- *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 the [www.nortel.com/documentation](http://www.nortel.com/documentation) URL. Find the product for which you need documentation. Then locate the specific category and model or version for your hardware or software product. Use Adobe\* Acrobat Reader\* to open the manuals and release notes, search for the sections you need, and print them on most standard printers. Go to Adobe Systems website [www.adobe.com](http://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

## Routing overview

---

The Contivity Secure IP Services Gateway utilizes Secure Route Technology (SRT) to forward network traffic. SRT operates on the premise that there are trusted and untrusted portions within the network. Trusted interfaces are placed on secure network segments (such as the private LAN) and behave like traditional routed interfaces. Untrusted interfaces are placed on unsecure network segments (such as the Internet), where all insecure services are disabled. Only services considered secure are permitted to run on, or are accessible through, untrusted interfaces.

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

[Table 1](#) is a matrix of Contivity gateway forwarding capabilities between the source interface and destination interfaces.

**Table 1** Forwarding capabilities

	<b>Private</b>	<b>Public</b>	<b>Client tunnel</b>	<b>Branch office tunnel</b>	<b>System management</b>
<b>Private</b>	yes (1)	yes (1)	yes	yes	yes
<b>Public</b>	yes (1)	yes (1)	yes (1)	yes (1)	yes (3)
<b>Client tunnel</b>	yes	yes (1)	yes (2)	yes (2)	yes
<b>Branch office tunnel</b>	yes	yes (1)	yes (2)	yes (2)	yes
<b>System management</b>	yes	yes (3)	yes	yes	not applicable

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

## Integrated firewall and routing

The Contivity gateway is a security device. Therefore, the routing configuration takes effect as it relates to the integrated firewall configuration of the Contivity gateway. In all of the following sections, when there is a reference to integrated firewall, it means the Contivity Firewall option on the Services > Firewall screen. Use this option by selecting either Contivity Stateful Firewall or Contivity interface filter. However, if you use the Contivity interface filter option, you do not need a firewall license.

## Dynamic routing

Dynamic routing protocols are available for private physical interfaces or branch office tunnel interfaces. Public interfaces are not trusted and therefore cannot be configured to run a dynamic routing protocol. The only exception is Border Gateway Protocol (BGP), which can be enabled on public interfaces on request. All physical LAN and WAN interfaces can be configured as either a private or public interface with the exception of slot 0 interface 1, which is always a LAN and private.



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

---

## VPN routing

VPN routing forwards traffic between tunnels or between tunnels and private interfaces. VPN routing enables traffic to enter or exit the Contivity gateway through a tunnel.

Enhanced routing provides additional traffic patterns beyond traditional VPN routing. Either the Contivity Stateful Firewall or Contivity filter must be enabled to support the enhanced routing feature.

## Static routes

You can configure static routes between Contivity gateways when you do not have any dynamic routing protocol, such as OSPF, RIP, or BGP. Even if you do have dynamic routing protocols, you may want to use static routes because they provide stronger security. The Contivity gateway supports multiple default and static routes.

## Route table

The route table contains the routes submitted by the routing protocols and the static route application and dynamic protocols, such as OSPF, RIP, and BGP. The route table manager (RTM) chooses the best routes from the route table to populate the IP forward table. The IP forward table is used by the Contivity gateway during forwarding decisions. The best routes are selected based on the following order of protocol preference:

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

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

## Routing status

The Routing > Status screen provides access to information about each routing protocol. It also provides access to the route table and route table manager (RTM) statistics. [Table 2](#) shows routing status screen options.

**Table 2** Routing status screen options

Button	Description
BGP Summary	Overall summary of BGP running on the Contivity gateway, including the router ID, Local AS, Admin state (enabled or disabled), Hold Interval, Keep Alive Interval, Local Preference, Default Metric, Route Reflector, Client Reflection, Cluster ID, Always Compare MED, Auto summary, Redistribute Internal, Synchronization, Max paths, and Number of Peers.
BGP Routes	Search Type, IP Address, Mask, and Mask Type.
BGP Redistributed Routes	Includes IP Address, IP Mask, and Origin Type.
BGP Neighbors Routes	Includes Routes Type and Neighbor.
BGP Neighbors Summary	Overall summary of Foreign Host, Remote AS, External Link, Remote Router ID, BGP state, Up For, Hold Time, KeepAlive Interval, Advertisement Runs, Received, Received Notifications, Sent, Community Attribute, Accepted Prefixes, Prefix Advertised, Local Host, Local Port, Foreign Host, Foreign Port, Connections Established, Elapsed Time Between Updated Msg, MinASOriginationInterval Timer.
OSPF LSDB	Link state databases in all areas that are known to OSPF, including information on the link state type, ID, advertising router address, metric, ASE, forward address, age, and sequence number for each area.
OSPF Neighbor	Neighbors on all the interfaces running OSPF, including the IP interface address, router ID, neighbor IP address, state, and dead time priority.
OSPF Interfaces	Interfaces configured for OSPF, including the IP address of the interface, the area to which the interface belongs, the type of interface, the state, cost and the designated router in the area to which the interface belongs.
OSPF Summary	Overall summary of OSPF running on the Contivity gateway, including the router ID, global state (up or down), whether an area border router or autonomous system border router.
OSPF Statistics	System-wide OSPF statistics.

**Table 2** Routing status screen options

<b>Button</b>	<b>Description</b>
RIP Database	Contains all routes that can be distributed by RIP (based on routing priorities).
RIP Interfaces	Interfaces that you configured for RIP.
RIP Statistics	System-wide RIP statistics.
VRRP Config	VRRP configuration information.
VRRP Errors	System-wide VRRP errors that have occurred.
VRRP Statistics	System-wide VRRP statistics.
Route Table	Full routing for all routes, including next hops and best routes.
Next Hop Table	Next hop address for each route.
Best Route Table	Used by the forwarding table to determine the best route.
Route Table Stats	Statistics about route table management that provides information about Contivity gateway traffic.
IP Forward Table	Information about the IP routes used to forward traffic.



---

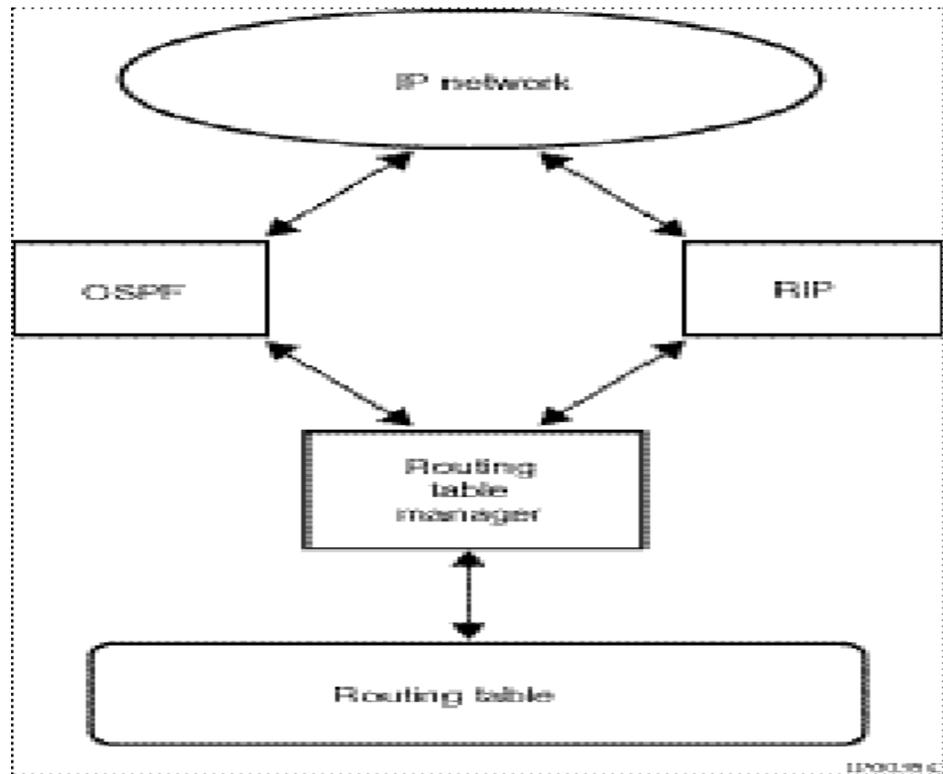
## Chapter 2

# Route table

---

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

[Figure 1 on page 28](#) shows how different routing protocols interact with the route table manager.

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

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

Public traffic has the following public routes:

- Static routes to public interfaces
- Dynamic (BGP only) routes to public interfaces
- Default route to public interface

Private traffic has the following private routes:

- Static routes to private interfaces
- Dynamic routes to private interfaces

- Static routes to branch office tunnel interfaces
- Dynamic routes to branch office tunnel interfaces
- Default route to private interface
- Routes used for tunnels

When a packet arrives, the Contivity gateway performs a full lookup in its IP forwarding table to determine which route to use:

- If firewall support is enabled, all public and private routes in the IP forwarding table are available to the traffic.
- If firewall support is not enabled, only the private portion of the IP forwarding table is available.
- If the traffic's destination route is not found in the table, the table's public or private default route is invoked as described in the following section.

## Route table lookup

The route table has two separate parts. One part contains the routes for traffic that uses the Contivity gateway's public interfaces (untrusted network), and a second part has the routes for traffic using the private interfaces (trusted network). Tunnels are virtual interfaces and are treated as private interfaces.

The following list shows the types of routes in the Contivity gateway's route table:

- Static routes
  - To public interfaces
  - To private interfaces
  - To branch office tunnel interfaces
- Dynamic routes
  - To private interfaces
  - To branch office tunnel interfaces
  - To public interfaces (BGP only)

- Default routes
  - To public interfaces
  - To private interfaces
- Host routes
  - Routes added for VPN users (for example, Contivity VPN Clients or PPTP clients)
- Utunnel routes
  - Host/network routes for clients that log in using the client address redistribution feature

## Route selection based on destination

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

## Route selection based on precedence in route table

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

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

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

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

## Route table options

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

- 1** To view the route table, go to the Routing > Route Table screen.
- 2** To search the route table, select the All, Host, or Network option from the destination field.

If you select Host or Network:

- a** From the interface list, select All or the address.
- b** From the Protocol list, select All or the protocol (BGP, OSPF, RIP, Static, or Direct). You must enter the IP address in the edit box.
- c** Type the IP address.

If you select Network:

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

- 3** Click Search.
- 4** To save the route table to a file:
  - a** Enter the file name in the Filename edit box. You can save the route table as a text file in the directory `ide0/system/xxx`, where `xxx` is the name of the file that you specify.
  - b** Under Route Filter, select Best Routes to view all routes to a single or All Routes to view all destinations. The default is Best Routes.
- 5** Click Save.

- 6 To check the route table status, click the IP Forwarding Table button on the Route Table screen to display the IP Route Network Table, the IP Route Host Table, and the IP Public Address Table.

Table 3 describes the fields on the IP Forward Table screen.

**Table 3** IP Forward Table screen

Column	Description
Destination/Mast	Network address and mask
Contivity	IP address of next-hop Contivity gateway
Flags	Internal use flags
Refcnt	Reference count
Use	How many times used
Interface	Interface identifier
MTU	Size of packet
OuterCtxt	(For internal use only)
CircMap	(For internal use only)
RtEntryP	(For internal use only)

- 7 Click the Route Table button on the Route Table screen to display the full internal route table.

Table 4 describes the fields on the IP Route Table screen.

**Table 4** IP Route Table screen

Column	Description
Seq	Sequence number that shows the best route
Proto	Protocol
IP Address/Netmask	IP address and network mask
Weight	Combination of cost and priority for the best route
NextHop	IP address of the next-hop
NextHopInterface	IP address of the next-hop interface
CId	Circuit ID

---

## Chapter 3

# Configuring RIP

---

Routing Information Protocol (RIP) is a distance-vector routing protocol that allows routers to exchange routing information by means of periodic RIP updates. Routers transmit their own RIP updates to neighboring subnets, and listen for RIP updates from the routers on those neighboring subnets. Routers use the information in the RIP updates to keep their internal routes current.

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

RIP is one of the most common interior Contivity Secure IP Services Gateway protocols used. RIP Version 2 is backward compatible with RIP Version 1 and corrects many RIP Version 1 shortcomings, such as subnet routing, authentication, and multicast support for route messages.

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



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

---

The Contivity gateway supports RIP Version 1 and Version 2. For additional information on RIP, refer to the RFCs located on the Internet Engineering Task Force (IETF) Web site at [www.ietf.org](http://www.ietf.org).

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

## Protecting against routing loops

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

The Contivity gateway supports the following methods used by RIP for minimizing loops and for speeding up the convergence that is caused by the normal correction of a loop:

- Split horizon, where the Contivity gateway does not send routes that it learns from a neighboring router back to that same neighbor.
- Split horizon with poison reverse, where the Contivity gateway *does* send back the routes that it learns from a neighboring router, but it sets the metric for that route to infinity.

- Triggered updates, where an update is sent almost immediately after a routing change has been made on the Contivity gateway. By default, RIP updates routes at regular intervals.

## Configuring RIP on the Contivity gateway

To enable RIP interfaces:

- 1** After you globally enable RIP, you must also enable it on the Routing > Interfaces screen.
- 2** Click Configure. The Interfaces > Routing Interfaces > Configure RIP screen appears.
  - a** The enabled check box indicates that you globally enabled RIP.
  - b** Select V2, V1, or Off as the transmit mode. Transmit mode enables you to specify which version of RIP to use when routing traffic from this Contivity gateway. The default is V2. Selecting OFF specifies that RIP is not used.
  - c** Select V2, V1, Both, or Off as the receive mode. Receive mode enables you to specify which version of RIP accepts incoming traffic. The default is V2. Selecting OFF specifies that RIP is not used. Selecting BOTH specifies that incoming transmissions using either version of RIP are accepted.
  - d** Select None, Simple, or MD5 as the authentication type that is used as part of the RIP transmission. This authentication is specific to RIP and has no bearing on the authentication done as part of the connection to the Contivity gateway. The default is None, which specifies that no authentication is required. Simple indicates that authentication uses a simple password. MD5 specifies that authentication uses an MD5 secret. If you select either Simple or MD5, password and password confirmation fields display.
  - e** Enter a metric value for the cost. This is the cost of sending a packet on the interface expressed in the link state metric.
  - f** Select Enabled or Disabled for poison reverse. Poison reverse updates routing loops in large networks.

- g** If no default route has been set, you can check the Import Default Route box to use the default route learned during RIP updates. Typically, you specify a default route in the route table on the Routing > Static Routes screen. The default is disabled.
- h** Select Enabled to specify that the default route is exported during RIP updates or enter a metric value (1 through 15) to the default route.
- i** Select Enabled to specify that static routes are exported during RIP updates or enter a metric value (1 through 15) to the default route.
- j** Select Enabled to specify that OSPF routes are exported during RIP updates or enter a metric value (1 through 15) to the default route.
- k** Select a metric value (1 through 15) to export the static routes metric if you have a branch office connection. This informs the remote branch office connection of the routes that are used and provides the assigned metric value. The default is 1 and the maximum value is 15.

To globally enable RIP:

- 1** Go to the Routing > RIP screen and click Enable.
- 2** Enter the amount of time in seconds that you want RIP to update the routes. The default is 30 seconds and the range of values is from 5 through 65535 seconds. The hold-down timer is six times the update timer.

To configure RIP interfaces:

- 1** Enable RIP interfaces by clicking Configure on the Routing > Interfaces screen for private interfaces or Profiles > Branch Office > <Group> Edit for branch office tunnel interfaces.
- 2** On the Routing > RIP screen, check the Enabled check box to globally enable RIP. By default RIP is globally disabled.
- 3** Enter the interval of time in seconds for RIP to update the routes. The supported range is from 5 seconds to 65535 seconds, with the default setting at 30 seconds. The RIP hold down timer is automatically 6 times the update timer.
- 4** Configured Physical Interfaces lists the IP address and RIP configuration state (enabled or disabled) of each physical interface.

- 5 Click on the Statistics button to display statistics about RIP on the Contivity gateway.

Table 5 describes the fields on the RIP Statistics screen.

**Table 5** RIP Statistics screen

Column	Description
Global RIP Status	Enabled or disabled
Update interval	Interval in seconds
Trusted Neighbor	Enabled or disabled
Rip Domain	Set or reset
Triggered Update	Set or reset
Route Change	Number of routes changed
Query	Number of queries sent

- 6 Click on the Database button to display information for all of the RIP interfaces.

Table 6 describes the fields on the RIP Database screen.

**Table 6** RIP Database screen

Column	Description
Circuit	Circuit ID
Address	IP address
Mask	Network mask of IP address
Owner	Protocol
Cost	Import cost of RIP routes
Metric	Export metric of RIP routes
Gw	Contivity gateway IP address

- 7 Click on the Interfaces button to display information for all RIP interfaces, including tunnels that are running RIP.

Table 7 describes the fields on the RIP Interfaces screen.

**Table 7** RIP Interfaces screen

Column	Description
Ip	RIP interface IP address
Subnet	Network mask of IP address
RipEnabled	Whether RIP is enabled or disabled
IntfState	Whether up or down
Auth	Authentication type
Type	Interface type
Cid	Circuit ID
RxMode	RIP receive version supported
TxMode	RIP transmit version supported
PoisonRev	Whether enabled or disabled
ImpDRoute	Whether enabled or disabled
ExpTSMetric	Disabled or metric (1-15) export tunnel static route
ExpSMetric	Disabled or metric (1-15) export static route
ExpDMetric	Disabled or metric (1-15) export default route
ExpOspfMetric	Disabled or metric (1-15) export OSPF route

To configure RIP for branch office tunnels:

- 1** Go to the Profiles > Branch Office > Group > Edit screen.
- 2** Click Configure in the RIP section. The list of RIP settings appears.
- 3** Click the Configure button next to each field to change these values.
  - a** Select V2, V1, or Off as the transmit mode. Transmit mode enables you to specify which version of RIP to use when routing traffic from this Contivity gateway. The default is V2. Selecting OFF specifies that RIP is not used.
  - b** Select V2, V1, Both, or Off as the receive mode. Receive mode enables you to specify which version of RIP accepts incoming traffic. The default is V2. Selecting OFF specifies that RIP is not used. Selecting BOTH specifies that incoming transmissions using either version of RIP are accepted.

- c** If no default route has been set, you can check the Import Default Route box to use the default route learned during RIP updates. Typically, you specify a default route in the route table on the Routing > Static Routes screen. The default is Disabled.
  - d** Select Enabled to specify that the default route is exported during RIP updates or enter a metric value (1 through 15) to the default route.
  - e** Select Enabled to specify that static routes are exported during RIP updates or enter a metric value (1 through 15) to the default route.
  - f** Select a metric value (1 through 15) to export the static routes metric if you have a branch office connection. This informs the remote branch office connection of the routes that are used and provides the assigned metric value. The default is 1 and the map value is 15.
  - g** Select Enabled to specify that OSPF routes are exported during RIP updates or enter a metric value (1 through 15) to the default route.
  - h** Enter a metric value for the cost. This is the cost of local RIP interface through the Branch Tunnel.
  - i** Select Enabled or Disabled for poison reverse. Poison reverse updates routing loops in large networks.
  - j** Select None, Simple, or MD5 as the authentication type that is used as part of the RIP transmission. This authentication is specific to RIP and has no bearing on the authentication done as part of the connection to the Contivity gateway. The default is None, which specifies that no authentication is required. Simple indicates that authentication uses a simple password. MD5 specifies that authentication uses an MD5 secret. If you select either Simple or MD5, password and password confirmation fields display.
- 4** Click OK.



---

## Chapter 4

# Configuring OSPF

---

Open Shortest Path First (OSPF) is a link-state routing protocol. With the link state information, a device running OSPF builds a shortest-path tree with itself as the root of the tree. The device can then identify the shortest path from itself to each destination and build its route table. Some of the benefits of OSPF are:

- Fast convergence with minimal routing protocol-related traffic after convergence
- Variable-length subnet masks (VLSM)
- Hierarchical segmentation
- Area routing to provide additional routing protection and a reduction in routing protocol traffic
- Authentication
- Virtual link support
- Equal cost multipath (ECMP) support
- Multicast- or unicast-based route advertisement messages instead of broadcast-based advertisements

The Contivity Secure IP Services Gateway OSPF support allows you to enable or disable OSPF on the Contivity gateway's private and tunneled interfaces. It supports broadcast and point-to-point network types and can act as autonomous boundary router (ABR), information retrieval (IR), autonomous system boundary router (ASBR), designated router (DR), and system designated router (SDR) router types. The Contivity gateway OSPF implementation conforms to OSPF 2 (RFC 2178).

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

## Installing the Advanced Routing key

The Advanced Routing License key must be installed to enable OSPF on the Contivity gateway. (The Firewall License Key is required only when the redistribution capabilities of RIP and OSPF are necessary).

To install a software license key:

- 1 Go to Admin > Install Keys screen.
- 2 Type the key that you obtained from Nortel Customer Support in the box to the right of Advanced Routing.
- 3 Click Install.

After the key is installed, the label Key Installed is displayed. It is only necessary to install a key once on each Contivity gateway. Click Delete to remove the key. A confirmation message appears and, if you click Yes, the key is removed.



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

---

## Installing the Premium Routing key

The Premium Routing key enables the same as the Advanced Routing, BGP, and Data Link Switching (DLSW) keys enable. The procedure for installing the Premium Routing key is the same as the procedure for installing the Advanced Routing key, as described in [“Installing the Advanced Routing key”](#).

---

## Virtual link support

OSPF requires that all non-backbone areas have at least one connection to the backbone area (area 0). If an area does not have a physical connection to the backbone, a virtual link can be used to traverse an intermediate area to connect to the backbone area. The Contivity gateway must be an area border router for the automatic configuration of virtual links to operate properly.

## Configuring OSPF on the Contivity gateway

To configure OSPF interfaces:

- 1** Go to the Routing > Interfaces screen.
- 2** Click Configure. The Interfaces > Routing Interfaces > Configure OSPF screen appears. Interface indicates the type of interface. IP address indicates the IP address of the interface.
  - a** Select Enabled to enable the OSPF state. It is enabled by default.
  - b** Enter the OSPF area to which the attached network belongs. Click the Add an Area link to add an area.
  - c** Select Broadcast or Point to Point for the OSPF network type. The default is Broadcast.
  - d** Select None, Simple, or MD5 as the authentication type that is used as part of the RIP transmission. Simple indicates that authentication uses a simple password. MD5 specifies that authentication uses an MD5 secret. If you select either Simple or MD5, password and password confirmation fields display.
  - e** Enter a metric value for the cost. This is the cost of sending a packet on the interface expressed in the link state metric. The value must always be greater than 0 and the default is 10.
  - f** Enter the priority level of the routers on this interface. The router with the highest priority takes precedence and is the designated router (DR). If there is a tie, the router with the highest Router ID takes precedence. A priority setting of 0 is ineligible to become a designated router on the attached network. Router priority only applies to broadcast networks. The default is 1.

- g** Enter the Length of time in seconds between the Hello packets that the router sends on the interface. It must be the same for all routers attached to a common network. The default is 10.
  - h** Enter the number of seconds after a router ceases to hear Hello packets before declaring that the router is down. The number must be the same for all routers attached to a common network. The default is 40.
  - i** Enter the number of seconds when, if a neighboring router becomes inactive, the router sends packets at a reduced rate in seconds. The default is 120.
  - j** Enter the number of seconds between link state advertisement (LSA) retransmission for adjacencies belonging to this interface. It is also used for retransmitting Database Description and Link State Request packets. This setting should be considerably over the expected round trip delay between any two routers on the attached network. The default is 5.
- 3** Click OK.

To configure OSPF globally:

- 1** Click Routing > OSPF to configure OSPF global parameters. Enabled indicates that OSPF is enabled on this screen. The default setting is Disabled.
- 2** In the Router ID field, type in the IP address used to uniquely identify the OSPF router in the OSPF network. The default address is the lowest IP address of the management or physical interfaces defined on the Contivity gateway. You can change this address provided that it is unique within the area.
- 3** If this Contivity gateway is an autonomous system (AS) boundary router, select True from the AS-Boundary-Router list. This parameter must be set to True to enable the redistribution of non-OSPF routes into OSPF. An AS boundary router is a router that exchanges routing information with routers belonging to other autonomous systems and advertises AS external routing information throughout the AS. The default is False.
- 4** To automatically create virtual links to the backbone network, select True from the list. The default is False.
- 5** Select metric Type 1 or Type 2. Type 1 is the default. Type 1 external metrics are expressed in the same units as OSPF interface cost (in terms of the link state metric). Type 2 external metrics are an order of magnitude larger; any Type 2 metric is considered greater than the cost of any path internal to the AS

boundary router. Use of Type 2 external metrics assumes that routing between AS boundary routers is the major cost of routing a packet, and eliminates the need for conversion of external costs to internal link state metrics.

- 6** Select the maximum number of ECMP paths (1-4). Equal cost multipath provides load balancing of packets to a destination that is reachable over more than one physical interface.
- 7** The Known OSPF Areas section displays all OSPF areas defined locally to the Contivity gateway. The area information is not shared among Contivity gateways. If you want two Contivity gateways to have one of their interfaces in a common area, you must configure both Contivity gateways to define the area information. Area IDs are used as representations of parts of the OSPF network. They help to manage large numbers of networks so that they can exchange information within an area. Each Area ID must be unique for OSPF. By default, all Contivity gateways have an area named 0.0.0.0.

To add an OSPF area, click the Add button. The Routing Protocols > Add Area screen appears.

- a** Enter the IP address in the Area ID field.
  - b** For Stub, select True or False from the list. The default is False.
  - c** For Stub Metric, enter the number of the stub metric. The default is 1.
- 8** The Configured Physical Interfaces section lists:
    - a** IP address of the configured OSPF interfaces.
    - b** Area ID of the configured OSPF interfaces.
    - c** Type is either Broadcast or Point-to-Point.
    - d** State is either Enabled or Disabled.
  - 9** In the Save LSDB Table section, type in the name of the LSDB table that you want to save as a text file in the ide0/system/routing directory.
  - 10** In the Status section, you can display LSDB (link state database), Neighbor, Interfaces, Summary, or Statistics. Click on the LSDB button to display the link state databases in all areas configured for the Contivity gateway.

[Table 8](#) describes the information on the OSPF LSDB screen.

**Table 8** LSDB screen

Column	Description
Link State ID	Link state address
Adv Router	Advertising router address
Age	Age in seconds
Seq Nbr	Sequence number
Checksum	Checksum
Links	Number of links

- 11 Click the Neighbor button to display a list of neighbors for all the interfaces running OSPF.

[Table 9](#) describes information on the OSPF Neighbors screen.

**Table 9** OSPF Dynamic Neighbors screen

Column	Description
Router ID	OSPF ID of neighbor
P	Priority number
State	State of neighbor connection
Dead Time	Time until neighbor is declared dead
Address	Neighbor IP address
Interface	Local IP interface address

- 12 Click the Interfaces button to display the list of interfaces that you configured for OSPF.

Table 10 describes the fields on the OSPF Interfaces screen.

**Table 10** OSPF Interfaces screen

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

**13** Click the Summary button to display the overall summary of OSPF running on the Contivity gateway.

Table 11 describes fields on the OSPF Summary screen.

**Table 11** OSPF Summary screen

Column	Description
Router ID	Unique OSPF ID of router
Router State	OSPF global configured state (up or down)
Supports TOS	Type of service support
SPF schedule delay	Shows delay time before calculating changes to SPF
Hold time between two SPF's	Time between shortest path first calls
Minimum LSA interval	Link state advertisement interval
Minimum LSA arrival	Link state advertisement arrival minimum
Number of external LSA	Number of link state advertisements
Link State Update Interval	Time between link state updates

**Table 11** OSPF Summary screen (continued)

Column	Description
Link State Age Interval	Time between link state aging intervals
Number of Areas in this router	Number of areas
RTM Stats	Route table manager changes for route table changes
Area	Area ID
Number of interfaces in this area	Number of interfaces in this area
SPF algorithm has executed	Number of times shortest path algorithm has been executed

**14** Click the Statistics button to display statistical information about OSPF.

[Table 12](#) describes the fields on the OSPF Statistics screen.

**Table 12** OSPF Statistics screen

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

To configure OSPF for branch offices:

- 1** Click Configure in the OSPF section of the Edit Group screen to configure the OSPF routing attributes of the group.
- 2** Enter the priority level of the routers on this interface. The router with the highest priority takes precedence and is the designated router (DR). If there is a tie, the router with the highest Router ID takes precedence. A priority setting of 0 is ineligible to become a designated router on the attached network. Router priority only applies to broadcast networks. The default is 1.
- 3** Enter the time in seconds until neighbor is declared dead.

- 4 Enter the length of time in seconds between the Hello packets that the router sends on the interface. It must be the same for all routers attached to a common network. The default is 10.
- 5 Enter the number of seconds between LSA retransmission for adjacencies belonging to this interface. It is also used for retransmitting Database Description and Link State Request packets. This setting should be considerably over the expected round trip delay between any two routers on the attached network, and should be conservative. The default is 5.
- 6 Enter the number of seconds for the transmission delay. The default is 1.
- 7 Select None, Simple, or MD5 as the authentication type that is used as part of the OSPF transmission. Simple indicates that authentication uses a simple password. MD5 specifies that authentication uses an MD5 secret. If you select either Simple or MD5, password and password confirmation fields appear.

## Configuring equal-cost multipath

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

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

To configure equal-cost multipath:

- 1 Go to the Routing > Configuration screen.
- 2 Select the maximum equal-cost paths allowed globally by the Contivity gateway (Maximum Paths).

- 3 If you are using OSPF, you must also set the maximum equal-cost paths for OSPF (OSPF Maximum Paths).
- 4 If you are using BGP, you must also set the maximum equal-cost paths for BGP (BGP Maximum Paths).
- 5 If you are using RIP, you must also set the maximum equal-cost paths for RIP (RIP Maximum Paths).
- 6 Select the Forwarding Algorithm. You can change the forwarding algorithm to per-packet, per-destination, or per-source without affecting route or forwarding tables. The load balancing and resource sharing is controlled by the following forwarding algorithms:
  - Per-packet - packets are forwarded in a round-robin fashion. If the Contivity Stateful Firewall is enabled, this policy may cause some overhead in switching the firewall context.
  - Per-destination - packets are forwarded based on source and destination address pair.
  - Per-source - packets are forwarded based on source address.
- 7 Click OK.

## Configuring default routes

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

A default Contivity gateway is the address of the next-hop router. Packets are routed through the default Contivity gateway onto the private or public network when the route table does not have a specific route to the destination.

- 1 Go to the Routing > Configuration screen.
- 2 Under Default Route Preference, Source indicates whether the source is private or public.
- 3 Select or Public or Private as the Outbound Routing Preference:

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

**4** Click OK.



---

## Chapter 5

# Border Gateway Protocol

---

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

BGP runs over a reliable transport protocol. Any authentication scheme used by the transport protocol may be used in addition to BGP's own authentication mechanisms. The error notification mechanism used in BGP assumes that the transport protocol supports a graceful close, that all outstanding data has been delivered before the connection is closed.

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

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

## RFCs

Table 13 shows the RFCs that have been added to those supported on VPN Router.

**Table 13** RFCs

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

## Installing the Border Gateway key

The BGP-4 (Border Gateway Protocol Version 4) License key must be installed to enable BGP on the Contivity gateway.



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

To install a software license key:

- 1 Go to Admin > Install Keys page.
- 2 Type the key that you obtained from Nortel Customer Support in the box to the right of Advanced Routing.
- 3 Click on the Install button.

After the key is installed, the label Key Installed is displayed. It is only necessary to install a key once on each Contivity gateway.

To delete a software license key:

- 1 Click on the Delete button to remove the key.
- 2 A confirmation message appears. Click Yes. The key is removed.



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

---

## EBGP/IBGP peers

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

## BGP peering and connection processing

To begin a BGP peering session, one or both BGP speakers initiate a TCP connection. It is possible for both speakers to initiate a connection simultaneously, resulting in two active TCP sessions between peers. The BGP protocol provides negotiation rules to determine which connection remains and which is deleted. Once the TCP connection is established, the BGP protocol negotiates with its peer using the OPEN message to move into BGP Established State. At this time, each BGP speaker sends BGP update messages to distribute routing information between the speakers.

## BGP update processing

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

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

### Unfeasible route processing

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

### Feasible route processing

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

#### Path attribute processing

Path attributes fall into four separate categories:

- well-known mandatory
- well-known discretionary
- optional transitive
- optional non-transitive

Well-known attributes are recognized by all BGP implementations. Some of these attributes are mandatory and must be included in every UPDATE message. Others are discretionary and may or may not be sent in a particular UPDATE message. Attribute values can be modified using route filters, thus influencing the best path selection. The path attribute information applies to all prefix destinations listed in the NLRI.

Path attribute types are listed in [Table 14](#).

**Table 14** Path attribute types

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

**Table 14** Path attribute types

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

## Keep Alive processing

BGP speakers use a KEEPALIVE message to determine if their peers are reachable. The KEEPALIVE message can be disabled, but when in use, it must be configured so that it is not sent more frequently than once per second.

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

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

Hold Time Interval between BGP peers is negotiable. If the two peers negotiate a Hold Time Interval of zero, then KEEPALIVE messages must not be sent. The Hold Time Interval is configured on BGP > Configure page.

## BGP policies

Policy rules are applied to either permit or deny a route. Policies provide a way of filtering information based on IP prefixes, AS path information, BGP attributes, or source and destination addresses.

There are two types of policies:

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

Outbound policies are just the reverse.

## Accept/Announce policies

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

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

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

## Access (Prefix) lists

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

- Access list example 1:

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

This rule says that any route update that is in the range of 55.1.0.0 -> 55.1.255.255 matches the rule.

- Access list example 2:

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

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

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

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

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

- AS path regular expression example:

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

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

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

## AS-Path regular expressions

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

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



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

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

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

You can create, delete, and modify AS path access lists. You can also apply access lists directly to neighbors for filtering. To configure AS path access lists, go to [“Configuring AS Path Access Lists” on page 74](#).

## Route maps

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

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

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

The route maps can be matched on:

- as-path
- community-list
- ip address

The route maps can set:

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

The following example illustrates how route maps are used:

- **Route map example:**

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

In the above example, any route updates received from neighbor 55.1.1.1 are checked against this route map. First, the sequence number 10 rule states that any route matching ip access list 1 sets the metric to 8. If that check fails to match, then the sequence number 20 rule is checked. This states that any route matching ip access list 2, set the metric to 12.

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

## Configuring route maps

To configure route maps:

- 1 Select Routing > Route Map.

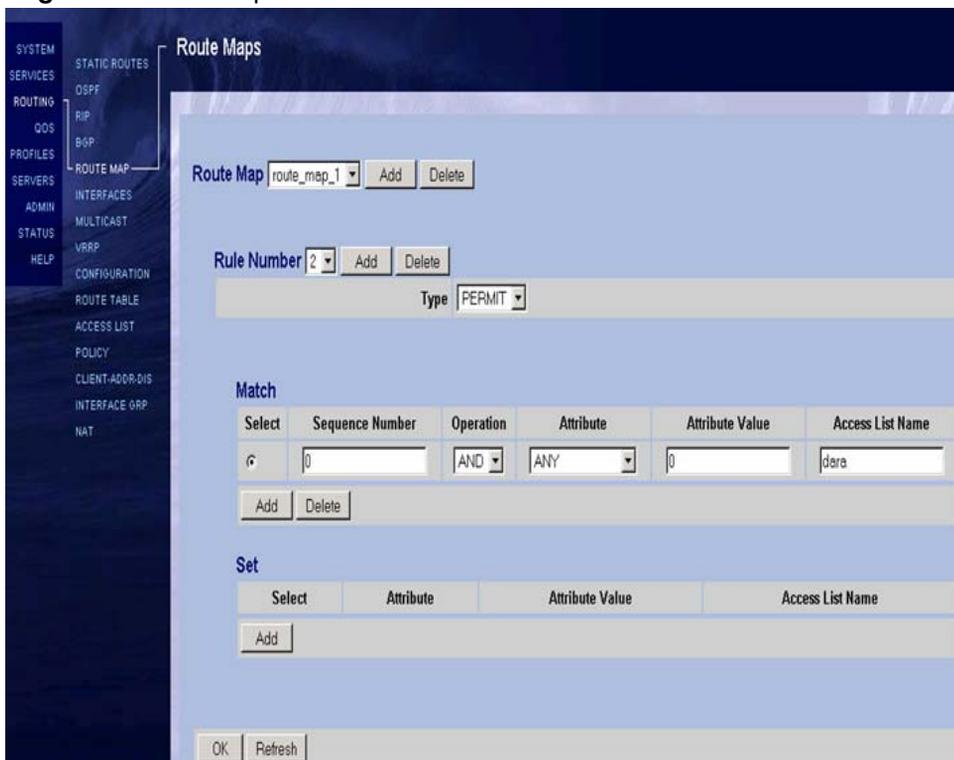


**Note:** If there are no route maps created, the Route Map Creation window opens. If there are route maps created and you want to add a route map to the list, follow the same procedure as creating a route map.

---

- 2 Enter a name in the Name text box.
- 3 Click OK. The Route Maps window opens. The name you entered appears in the Route Map menu.

[Figure 2 on page 64](#) shows the Route Maps window.

**Figure 2** Route maps

- 4 To add a rule number, click Add beside Rule Number. The Route Map Rule Add window opens.
- 5 Enter a number in the Number text box.
- 6 Click OK. The Route Maps window reopens. The number you entered appears in the Number menu.
- 7 Select a type from the Type menu.
- 8 To add a Match, click Add below Match. The Rule Match Add window opens.
- 9 Select an attribute from the Attribute menu.
- 10 Select a value from the Value menu.
- 11 Click OK. The Route Maps window reopens with the information you selected showing under Match.
- 12 To add a set, click Add below Set. The Rule Set Add window opens.

- 13 Select an attribute from the Attribute menu.
- 14 Enter a value in the Value text box.
- 15 Click OK. The Route Maps window reopens with the information you selected showing under Set.
- 16 Click OK.

## Multi-Hop BGP

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

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

Multihop is configured on the BGP > Neighbor > Configuration page. By default, multi-hop BGP is disabled.

## Route Reflector

Using a route reflector, BGP peers are organized into clusters. Each cluster is assigned an ID. Each member of the cluster advertises its routes only to the route reflector. The route reflector, in turn, collects all of the routes from all of the cluster members and advertises them to each of the IBGP peers in its cluster, as well as to any other route reflectors within the AS. Routes learned by the route reflector from other route reflectors are also forwarded to each of its cluster members.

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

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

Route reflector is disabled by default. To configure the route reflector, refer to [“Configuring the Route Reflector” on page 73](#).

## BGP communities

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

BGP communities were developed as a method of simplifying the route distribution based on membership to the community. A set of destination addresses is assigned a community identifier. Network administrators establish a policy for a community instead of a separate policy for each individual prefix. All route updates that are received for members of a community have the same route redistribution characteristics. Control over the distribution of routing information is based on:

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

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

- internet — the Internet community
- no-export — routes with this community are sent to peers in other sub-autonomous systems within a confederation. Do not advertise this route to an EBGp peer.

- local-as — do not advertise this route to an external system
- no-advertise — do not advertise this route to any peer (internal or external)

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

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

[Figure 3 on page 68](#) illustrates the following example.

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

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

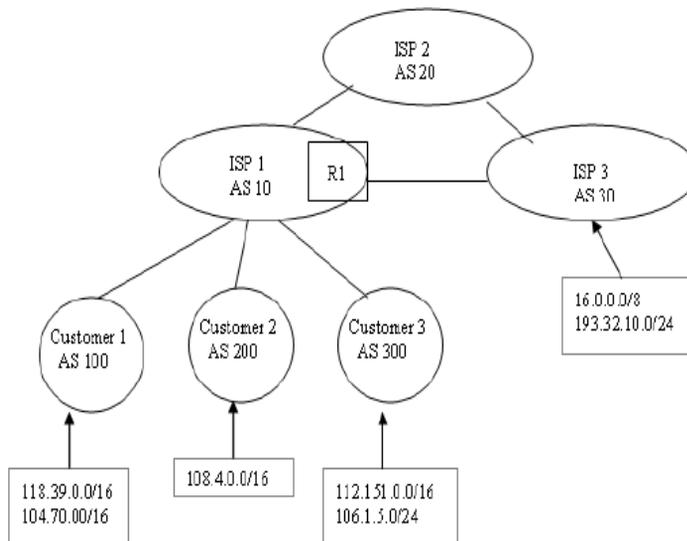
ISP 1 (AS 10) belongs to community 888.

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

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

AS 10 will label all routes learned from AS20 and AS30 as 10:444. This community label represents routes that will receive non-transit service.

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

**Figure 3** BGP communities

To configure a BGP community list, refer to [“Configuring Community Lists” on page 75](#).

## Configuring BGP on the Contivity gateway

BGP is not enabled by default over public interfaces. Enable BGP in Services > Available window.

To enable BGP interfaces:

- 1 Select Routing > BGP. The BGP window appears.

[Figure 4 on page 69](#) shows the BGP window.

Figure 4 BGP window

The screenshot shows the BGP configuration window in a Nortel Networks management interface. On the left is a navigation menu with categories like SYSTEM, SERVICES, ROUTING, PROFILES, SERVERS, ADMIN, STATUS, and HELP. The BGP configuration area includes a 'State' dropdown menu set to 'Disabled'. Below it are input fields for 'Router ID' (192.168.249.44), 'Local AS' (0), 'Hold Timer' (90), 'Keep Alive Timer' (30), 'Synchronization' (checkbox), 'Local Preference' (100), 'Default Metric (MED)' (0), 'Always Compare MED' (checkbox), and 'Maximum Paths' (1). At the bottom, there are 'Configuration' and 'Status' sections with buttons for 'Neighbors', 'Networks', 'Route Reflector', 'AS-Path Access List', 'Community List', 'Summary', 'BGP Routes', 'Redistributed Routes', 'Neighbors Routes', 'OK', and 'Refresh'.

- 2 Select Enabled or Disabled for State. If enabled, BGP enables all neighbors that are in enabled state. If disabled, BGP disables all neighbors that are in enabled state.
- 3 Enter the Router ID.
- 4 Enter a value in Local AS. If you are globally enabling BGP, you should configure Local AS entry.
- 5 Enter the Hold Timer value. The default value is 90 seconds.
- 6 Enter the Keep Alive Timer value. The default value is 30 seconds.
- 7 Synchronization allows routers within an AS to access a route before BGP makes it available to other ASs. Enable Synchronization if there are routers in the AS not speaking BGP.
- 8 Enter the Local Preference value. The default is 100.
- 9 Enter the Default Metric value. Default metric value specifies the appropriate metric for the specified routing protocol. The default metric command is used in conjunction with the redistribute router configuration command to cause

the current routing protocol to use the same metric value for all redistributed routes. This sets the Multi Exit Discriminator (MED) metric as a hint to external neighbors about preferred paths. The MED value can also be set using a route map. By default, during the best-path selection process, MED comparison is done only among paths from the same AS.

- 10** Check the Always Compare MED option if you want to allow the comparison of the MED for paths from neighbors in different ASs.
- 11** Enter the Maximum Paths value. This configuration controls the number of paths allowed. By default, only one path is installed in the IP routing table. If BGP multi-path support is enabled and the EBGp paths are learned from the same neighboring AS, instead of picking one best path, multiple paths are installed in the IP routing table. A maximum of six paths is supported and load balancing is performed among multiple paths.

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

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

## Configuring Neighbors

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

To configure neighbors:

- 1** Click Neighbors from the Routing > BGP window.

[Figure 5 on page 71](#) is the interface used for neighbor configuration.

**Figure 5** Neighbor configuration

The screenshot shows a web-based configuration interface for a BGP neighbor. At the top, there is a 'Neighbor' dropdown menu with the value '10.77.10.161', and 'Add' and 'Delete' buttons. Below this is the 'Configuration' section, which includes a 'State' dropdown set to 'Enabled'. There are input fields for 'Password' and 'Confirm'. Other fields include 'Remote AS' (100), 'Hold Timer' (90), 'Keep Alive Timer' (30), 'Advertisement Interval' (5), 'Retry Interval' (30), 'Source IP Address' (10.77.20.151), 'Weight' (0), 'NH Self' (checkbox), 'EBGP MultiHop' (checkbox), and 'Send Community' (checkbox). The 'Policy' section has a table with columns for 'In' and 'Out' directions, and rows for 'Filter List', 'Distribute List', and 'Route Map', each with a '(None)' dropdown. At the bottom, there is a 'Status' section with 'Summary' and 'Details' buttons, and a footer with 'OK', 'Cancel', and 'Refresh' buttons.

- 2 To add or delete a Neighbor, click the Add or Delete button beside Neighbor at the top of the page.
- 3 Select Enabled or Disabled for State.
- 4 Enter your password and confirm your password.
- 5 Enter a value in Remote AS. At a minimum, remote-AS should be configured for neighbors to be enabled.
- 6 Enter the Hold Timer value. The default value is 90 seconds.
- 7 Enter the Keep Alive Timer value. The default value is 30 seconds.
- 8 Enter the Advertisement Interval value. The minimum advertisement interval is 30 seconds.
- 9 Enter the Retry Interval value. The default is 30 seconds.

10 Enter the Source IP Address.



---

**Note:** The source IP address typically comes from the route table, but the administrator has the option of entering it in the Source IP Address text box.

---

- 11 Enter the Weight value. The administrative weight is local to the router. Any path that a VPN router originates will have a default weight of 32768 and other paths have a weight of 0. You can also assign the weight through filter-lists and route maps.
- 12 Disable NH Self when BGP neighbors do not have direct access to all neighbors on the same IP subnet. You can also specify the next-hop address to be used by route maps.
- 13 Enable EBGP to allow BGP sessions, even when the neighbor is not on a directly connected segment.
- 14 Enable Send Community if you want to include the community parameters in the message when the BGP route is announced to a neighbor.

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

## Adding a Network

To add a network:

- 1 Click Networks on the Routing > BGP page. The BGP > Networks window opens.
- 2 Click Add. The BGP > Networks Add window opens.
- 3 Enter an IP address in the IP Address field.
- 4 Enter a Mask in the Mask field.
- 5 Click OK.

## Configuring the Route Reflector

To configure the Route Reflector:

- 1 From the Routing > BGP page, click Route Reflector. The Route Reflector window opens.

Figure 6 shows the interface used to configure the route reflector.

**Figure 6** Route reflector configuration window

- 2 Select the Status of the route reflector. The status globally enables or disables the feature.
- 3 Enter the Cluster ID. The router ID of the route reflector identifies the cluster.
- 4 Select the Client to Client Route Reflector value. The default is Enabled. However, if the clients are fully meshed, route reflection is not required and the route reflector should be disabled.

To add or remove members from Route Reflector Client lists:

- 1 Under Clients, select a Non Member from the Non Member RR Client List. Click Make RR Client. The Non Member becomes a member of the Member RR Client List.
- 2 Select a member from the Member RR Client List. Click Remove RR Client. The member is removed from the list.

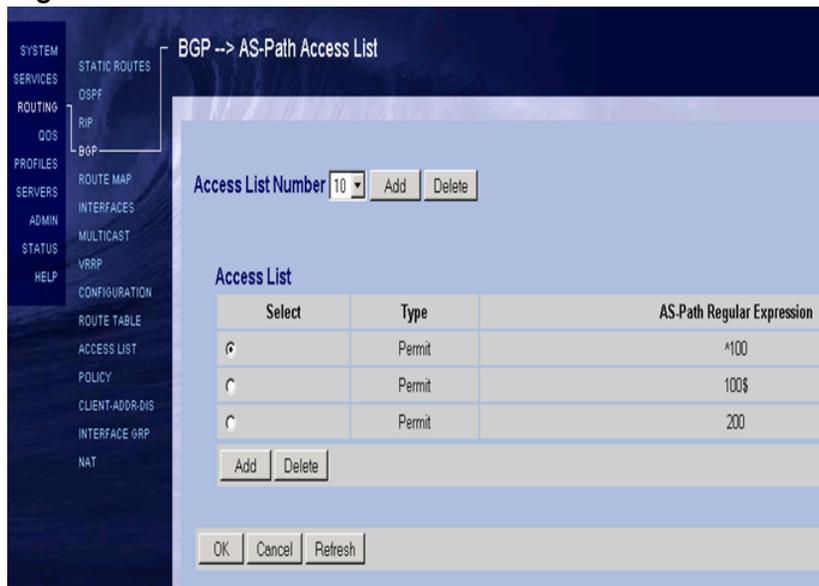
## Configuring AS Path Access Lists

To configure the AS-Path access list:

- 1 Select Routing > BGP. The Routing > BGP window opens.
- 2 Click the AS-Path Access List button. The AS-Path Access List window appears.

Figure 7 shows the interface used to configure AS path access lists.

**Figure 7** AS Path access list window



- 3 To add an Access List number, click Add beside Access List Number. The AS-Path Access List > Add window opens.
- 4 Type a number in the Number text box.
- 5 Click OK. The BGP AS-Path Access List Number window reopens with the number you typed in the Number text box showing in the Access List.
- 6 To create an Access List entry, click Add below Access List. The BGP > AS-Path Access List > Add Entry window opens.
- 7 Select an option from the Type menu.
- 8 Type an entry in AS-Path Regular Expression.

- 9 Click OK. The BGP > AS-Path Access List window reopens with your information showing on the page. At the top of the page is a statement saying Add operation completed successfully.
- 10 To delete an Access List, select the list that you want to delete and click Delete. A new window opens asking if you are sure you want to delete the as-path access list number.
- 11 Click OK. The BGP > AS-Path Access List window reopens with the number you deleted removed from the list. At the top of the window is a note stating Delete operation completed successfully.
- 12 To delete an Access List Entry, click the radio button to select the entry you want to delete. A new window opens asking if you are sure you want to delete the as-path access list entry.
- 13 Click OK. The BGP > AS-Path Access List window reopens with the entry you deleted removed. At the top of the window is a note stating Delete operation completed successfully.

## Configuring Community Lists

To configure a community list:

- 1 From the Routing > BGP page, click Community List. The Community List window opens.

[Figure 8 on page 76](#) shows the interface used to configure a community list.

**Figure 8** Community list

- 2 To add a community list number, click Add. The Community List > Add window opens.
- 3 Enter a number in the Number text box.
- 4 Click OK. The Community List window reopens.
- 5 To add a community entry, click Add. The Community List > Add Entry window opens.
- 6 Select a type from the Type menu.
- 7 Enter a name in the Name text box.
- 8 Click OK. The Community List window reopens with the information you entered showing.
- 9 To delete a community list number, select a number from the Community List menu.

- 10 Click Delete. A new window opens with a warning asking if you are sure you want to delete the community list number.
- 11 Click OK. The Community List window reopens with the community list number deleted.

## Health Check Support

A basic health check support is provided for the BGP-4 protocol.

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



---

## Chapter 6

# Configuring static routes

---

Available routes can be statically defined rather than learned by a dynamic routing protocol, such as Open Shortest Path First (OSPF) or Routing Information Protocol (RIP). Even if you use dynamic routing protocols, you may want to use static routes in certain situations where stronger security is required. The Contivity Secure IP Services Gateway supports multiple default and static routes.

## Adding and editing static routes

To add, edit, or delete static routes:

- 1 Go to Routing > Static Routes and check the Enabled box.

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

- 2 To add a public default route, click the Add Public Route button. The Add Public Default Route window appears.
  - a Click Enabled or Disabled to select the Admin State.
  - b Type the relative cost for the Contivity gateway. You use a lower cost number, such as 1, for the least expensive route. When there are multiple default paths, the Contivity gateway chooses the route with the least cost as the preferred route. The default cost is 10.
  - c Enter the IP address for the next-hop default router in the Contivity gateway address field.
  - d Click OK.

- 3** To add a private default route, click the Add Private Route button. The Add Private Default Route window appears.
  - a** Click Enabled or Disabled to select the Admin State.
  - b** Type the relative cost for the Contivity gateway. You use a lower cost number, such as 1, for the least expensive route. When there are multiple default paths, the Contivity gateway chooses the route with the least cost as the preferred route. The default cost is 10.
  - c** Enter the IP address for the next-hop default router in the Gateway Address field.
  - d** Click OK.
- 4** Click the Add button to add static routes to the route table. The Static Routes > Add window appears. When a static route is added, the Contivity gateway checks whether the next-hop interface address belongs to an attached network. If it does not, the Contivity gateway does not allow the static route.
  - a** Select Enabled or Disabled for the Admin state. The default is Enabled.
  - b** Select the relative cost for the Contivity gateway. You use a lower cost number (for example, 1) for the least expensive route. When there are multiple paths, the Contivity gateway chooses the route with the least cost as the preferred route. The default is 10.
  - c** Enter the network address for the static route to the destination network.
  - d** Enter the subnet mask for the static route to the destination network.
  - e** Enter the Contivity gateway address to the next-hop router to reach the destination network.
  - f** Click OK.
- 5** Click the Show Branch Office Routes button to display the configured branch office tunnels that are set up as static routes. By default, a tunnel is configured as a static route between the tunnel endpoints.
- 6** Click the Adjacent Hosts button to display adjacent host routes.

## Using ping to validate public default route

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

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

If validation is globally enabled, at the expiration of each ping interval it pings the ping address of each public default route for which per-route validation is enabled. If the ping address is not the route's gateway address, a static route is configured and enabled. If a static route with that address already exists, the route is used for validation and its state is saved.

Static routes used to validate public default routes cannot be edited or deleted. They are deleted or returned to their original state when any one of the following conditions occurs:

- Validation is globally disabled.
- The public default route is disabled or deleted.
- Validation is disabled for the public default route.
- The address to ping for validation of the public default route is changed.

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

To configure ping to validate a public default route:

- 1** Go to Routing > Static Routes.
- 2** Click Edit Public Default Route. The default is Disabled.
- 3** Select Validate at Ping Interval. The minimum (and default) is 30 seconds and the maximum is five minutes.
- 4** Enter the address in the Ping Address field.
- 5** Click OK.

---

## Chapter 7

# Route policy service

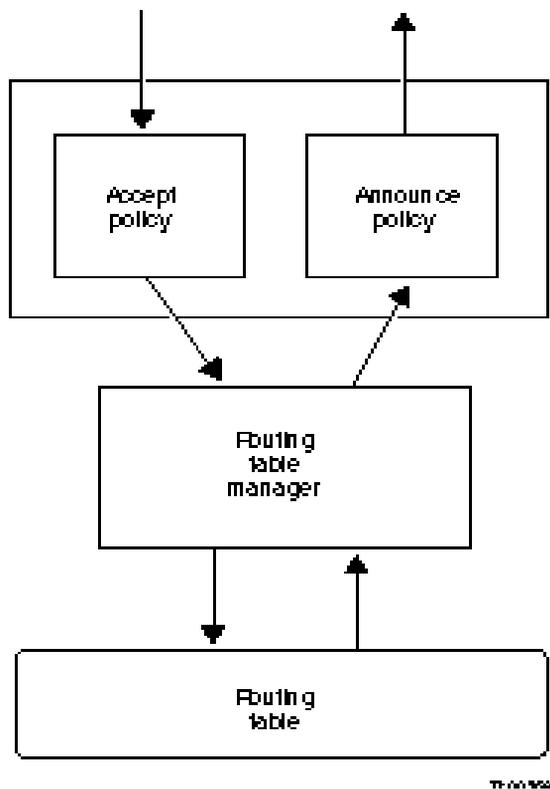
---

The route policy service allows you to control the flow of routing data to and from the route tables. The route policy service provides IP accept and announce policies that you enable or disable as needed.

Accept policies govern the addition of new RIP- or OSPF-derived routes to the route tables. When Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) receives a new routing update, it consults its accept policies to validate the information before entering the update into the route tables. Accept policies contain search information (to match fields in incoming routing updates) and action information (to specify the action to take with matching routes).

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

[Figure 9 on page 84](#) shows the interaction between the route table manager and accept/announce policies.

**Figure 9** Accept and announce policies

The route table manager forwards a route for advertisement to the protocol. The protocol consults an announce policies to determine whether or not to advertise the route to the network.

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

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

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

## Redistribution of routes

The Contivity Secure IP Services Gateway can redistribute static, direct, BGP, and RIP routes into OSPF. It can redistribute static, direct, BGP, and OSPF routes into RIP. It can also redistribute static, direct, OSPF, and RIP routes into BGP. The redistribution of routes from BGP to OSPF is controlled through access lists. Such a redistribution can be further controlled on a per-interface basis in RIP. Route redistribution is also based on security configurations.

[Table 15](#) describes the rules of redistribution for RIP, OSPF, and BGP with the firewall enabled or disabled.

**Table 15** Redistribution rules

Redistributed Route	Firewall ON	Firewall OFF
Public direct route	Yes	No
Public default route	Yes	No
Public static route	Yes	No
Private direct route	Yes	Out physical - No; out tunnel - Yes
Private default route	Yes	Out physical - No; out tunnel - Yes
Private static route	Yes	Out physical - No; out tunnel - Yes
Tunnel static route	Yes	OSPF - Always Yes RIP - In general, Yes, but can be controlled on a per-interface basis
Tunnel dynamic route	Yes	Yes
Utunnel routes	Yes	Yes

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

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

## Creating a policy list

To create a policy list:

- 1** Go to the Routing > Access List window.
- 2** Enter a new access list name. Use any name or number that you choose to a maximum length of 64 characters.
- 3** Click Create. The Access List > Policy window appears.
  - a** Under Action, the options are Permit, Deny, Permit All or Deny All. Permit or Deny is the action applied to a route update when the subnet and mask matches the route update. If you choose Permit All or Deny All, you cannot enter anything in the Subnet, Mask or Mask Type fields.
  - b** If you choose Permit or Deny, type in the subnet mask, mask and mask type (Exact or Range).
  - c** Click Add.
  - d** Click Close to have the new rule go into effect.
- 4** Click Edit to change an existing rule for the selected policy. The current information appears for each policy. You can use either an exact network address or a range of network addresses.
- 5** If you want to move the position of an existing rule, enter a number in the edit box. For example, if you select the third rule and enter 2 in the edit box, this moves the third rule to the second position. The order of the rules is important because the first match causes the action to occur. If there are no matches,

then all traffic is denied. Therefore, build your filter rules by first permitting the services that you want to allow. You can also add a Deny rule early in the rules sequence so that an unwanted packet is dropped before all of the rules are processed.

- 6 Click Close.

## Configuring route policy services (RPS)

To configure route policy services:

- 1 Go to the Routing > Policy window and check the Enabled box to enable RPS. The default setting is Disabled.
- 2 Under Redistribution Table, select the source of the route for each protocol, Static, Direct Nets, Direct Hosts, Utunnel. For correct operation, they should not be enabled at the same time.
- 3 Under Policy List, click the Add button to add a policy.
  - a Enter the Access Name/Number.



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

---

- b Select either OSPF, RIP, or BGP.
  - c Enter the Interface IP address, which is the IP address of the physical interface where you want to apply the policy. Select Global if you want to apply the policy to all interfaces. If the interface is a branch office, select the group name and type the connection name.
  - d Select the policy type, either the accept or announce. You can configure only one accept or announce policy for each protocol per interface.
- 4 Click OK.



---

## Chapter 8

# Client address redistribution

---

When a client initiates a user tunnel, the Contivity Secure IP Services Gateway assigns an inner address to the client. Sources for these addresses can be:

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

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

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

When inner addresses are allocated from an address pool with a range that does not belong to a locally attached network, use the aggregation option to reduce the number of entries in the route table and the route redistribution overhead. Aggregation creates and advertises a single Utunnel network route covering the address pool range when a client tunnel is created using an inner address from this

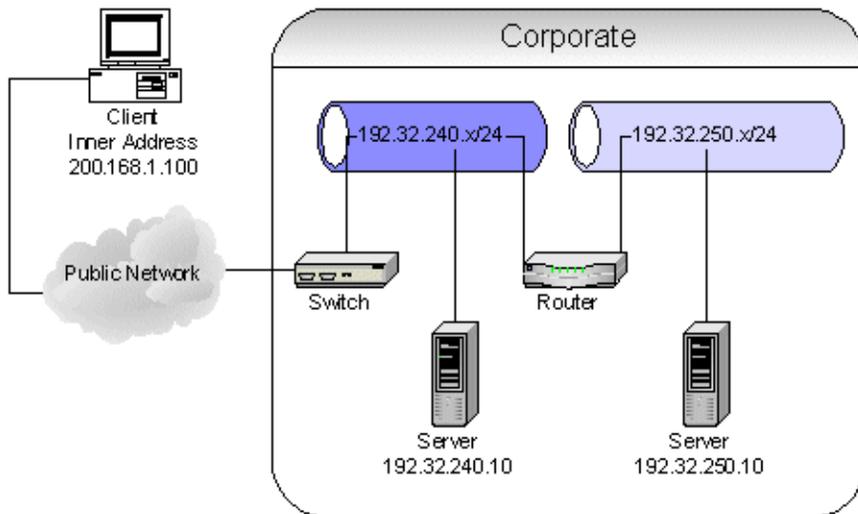
address pool. In Dynamic Aggregation mode, the network route remains in the route table until the last tunnel using an inner address from this address pool is taken down. In Static Aggregation mode, the network route remains in the route table until the user address pool is deleted.



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

Figure 10 shows an example of client address redistribution where the client has an inner address that is not within the local subnet of the private network. The Contivity gateway creates a Utunnel route that is then propagated over the network. The Utunnel route allows the router on the private network to recognize the 200.168.1.100 address and route responses back to it properly.

**Figure 10** Client address redistribution

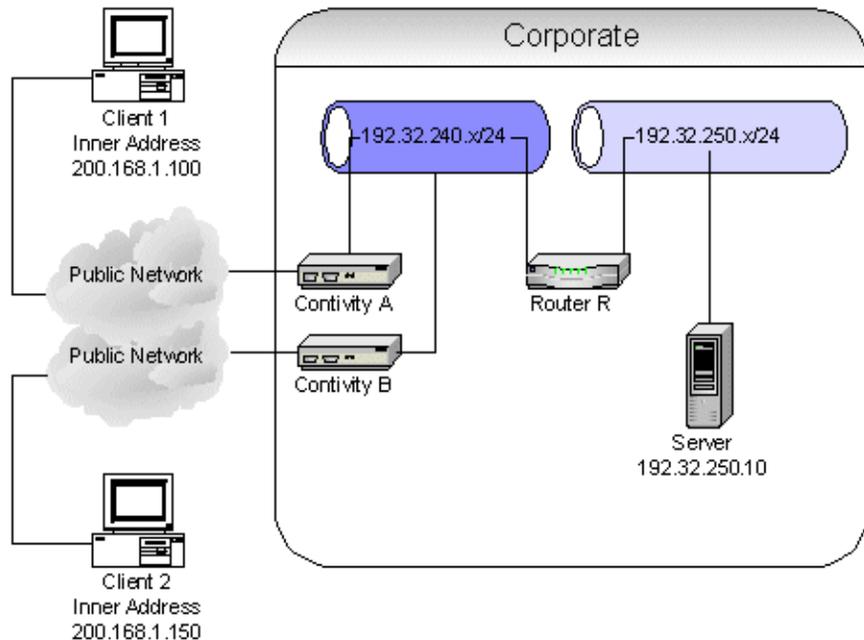


If you enable aggregation, the Contivity gateway identifies the subnet from the address pools where this address belongs and inserts a user tunnel network route for this subnet into the route table manager.

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

For example, in [Figure 11](#), Contivity A has an address range of 200.168.1.100 through 200.168.1.120 and Contivity B has an address range of 200.168.1.150 through 200.168.1.170. Both of these ranges are part of Class C subnet 200.168.1.x/24. Client 1 logs in to Contivity A and Client 2 logs in to Contivity B. Both clients have inner addresses that are not within the local subnet of the private network, but are in the same IP subnet. Contivity A and Contivity B running client address redistribution create Utunnel host routes. These routes are propagated over the network. The router on the private network recognizes addresses 200.168.1.100 and 200.168.1.150 and route responses back to them through the designated Contivity.

**Figure 11** Aggregation for client address redistribution



If you enable aggregation on both Contivity gateways, both gateways will advertise routes to 200.168.1.x/24. Router R will use one of these routes, causing either Client 1 or Client 2 to have communication problems.

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

To configure client address redistribution, go to the Routing > Client-Addr-Dis window:

- 1** On the Routing > Client Address Redistribution (CAR) window, select one of the following CAR modes:
  - **Disable** — CAR is disabled and redistribution of client routes does not take place.
  - **Host Mode** — CAR is enabled and redistribution of client routes is limited only to host routes. Host routes are added to both the forwarding table and the routing table. RIP and OSPF advertise the host routes of the VPN clients to their peers.
  - **Dynamic Aggregation** — CAR is enabled and the client host addresses are added only to the forwarding table. The subnet of the user address pool from which the client address was assigned is added to the routing table. RIP and OSPF only advertise the subnet of the address pool and not the client host addresses. When the last client using this user address pool disconnects, the subnet route is removed from the routing table. RIP and OSPF propagate the route deletion to the surrounding networks.
  - **Static Aggregation** — CAR is enabled and the client host addresses only are added to the forwarding table. The subnet of the user address pool from which the client address was assigned is added to the routing table. RIP and OSPF advertise only the subnet of the address pool and not the client host addresses. When the last client using this user address pool disconnects, the subnet route remains in the routing table. The subnet of the user address pool remains in the routing table as long as the user address pool remains valid. If you delete the user address pool, the subnet for the pool is then deleted from the routing table.
- 2** **Maximum Number of UTunnel Host Routes** allows you to limit the maximum number of user tunnel host routes advertised by the system. The default value is 200.

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

- 3 Click Show User Tunnel Routes to display the user tunnel routes. [Table 16](#) describes the fields.

**Table 16** Show user tunnel routes

Column	Description
IP address	IP address
Mask	IP network mask
Next Hop	Next hop address
Interface	IP interface address
Cost	Relative cost for the Contivity gateway

- 4 Click Statistics to display the configuration of client address redistribution, including mode, the UTunnel limit, and current UTunnel count.
- 5 Click Refresh to view any changes.



---

## Chapter 9

# Configuring multicast relay

---

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

A multicast relay listens to incoming multicast traffic and forwards it out one or more interfaces in the absence of multicast routing. Multicast relay is not supported on public interfaces.

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

The congestion threshold is configured relative to the amount of network processing memory buffers available to process the multicast traffic. The allowable range is 1 to 3000, where 3000 is the default value. If forwarding performance for unicast traffic decreases due to the multicast traffic burden, it is recommended that the threshold be reduced. To view network processing buffer statistics, go to Status > Statistic > snpbufStats.



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

---

Forward multicast packets over a tunnel using the default filter (permit all). For example, to allow multicast packets received over the interface to be relayed over tunnel B01 and not over tunnel B02, define the interface-specific rules as shown in [Table 17](#).

**Table 17** Multicast interface-specific rules example

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

To configure multicast relay:

- 1 On the Routing > Multicast window, check the Enabled check box to enable multicast relay on the Contivity gateway.

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



**Note:** Multicast requires use of the permit all interface filter.

---

- 2 Enter the congestion threshold value.
- 3 To add an interface to the multicast boundary list, click Add to go to the Multicast > Add window.
  - a Enter the Access Name/Number in the edit box.
  - b Select the IP address for the interface.
  - c Select Enabled for the State.
  - d Click the New Access List link to view the existing access screen.
- 4 Click Statistics to display the global multicast relay status and the statistics of the configured multicast interfaces, including branch office interfaces.

[Table 18](#) describes fields on the Multicast Statistics screen.

**Table 18** Multicast Statistics screen

Column	Description
Interface	IP address of interface
CID	Circuit ID
PktsRcvd	Number of packets received
PktsSent	Number of packets sent
PktsDropped	Number of packets dropped

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

[Table 19](#) describes fields on the Multicast Interfaces screen.

**Table 19** Multicast Interfaces screen

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



---

## Chapter 10

# Configuring the Virtual Router Redundancy Protocol (VRRP)

---

The Virtual Router Redundancy Protocol (VRRP) is a standard protocol used by the Contivity Secure IP Services Gateway to handle private interface failures. It is one method to help maintain a state of high availability. Hosts that are configured with static or default Contivity gateways obtain a resilient next-hop address. VRRP provides gateway-level failover in case a private physical interface fails. Using VRRP and dynamic routing provides a high degree of redundancy.

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

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

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

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

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

An external Lightweight Directory Access Protocol (LDAP) server is not a requirement, but may make VRRP easier to use. The LDAP server provides a common location where information for each Contivity gateway can be maintained. It enables each Contivity gateway to see the virtual router settings of other Contivity gateways in the system.

To configure VRRP, the virtual router ID (VRID) for the virtual router group must be identical to all Contivity gateways. If you use the internal LDAP server, the Contivity gateways must have the virtual router parameters configured the same way.

## VRRP and dynamic routing for high availability

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

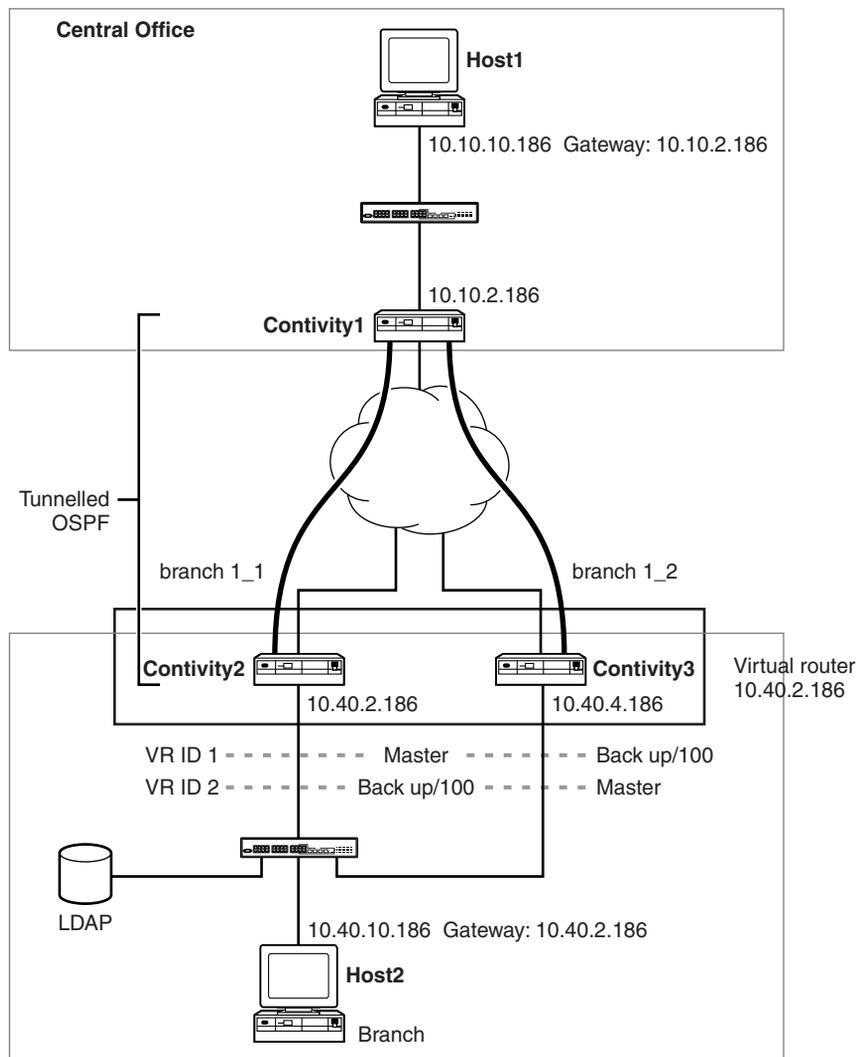
[Figure 12 on page 102](#) shows a deployment where the central office is configured with one Contivity gateway, Contivity1, and Host1 with the default Contivity gateway pointing to Contivity1. The branch is configured with two Contivity gateways: Contivity2 and Contivity3. VRRP is configured on the private side of Contivity2 and Contivity3 backing up each other's physical interface. Host2 in the branch has its default gateway pointing to Contivity2. There are also two branch office tunnels, as indicated, connecting the gateways.

Consider the traffic flow between Host2 and Host1. If Contivity2 fails or the private interface of Contivity2 fails, Contivity3 will assume the mastership of the private interface of Contivity2. Contivity3 will assume the IP address 10.40.2.186 as well as the MAC address of Contivity2's interface. All IP traffic from Host2 to Host1 will now flow through Contivity3. Contivity3 will forward all of Contivity2's routed packets, but will drop all packets destined to Contivity2. For example, a data packet from Host2 to Host1 will be forwarded, but a ping request to 10.2.40.186 will be dropped. Host1 is *not* aware of such a change.

Routing configuration plays a vital role in this failover operation. Contivity2 and Contivity3 need to know that the path to Host1 is through Contivity1; and Contivity1 should know that there are two paths to Host2: one through Contivity2 and another through Contivity3. The routing information on the each Contivity gateway can be manually populated using static routes, but dynamic routing

protocols such as RIP, BGP, or OSPF provide more reliable route information in networks that are considered dynamic or volatile (route information changes often). In this case, OSPF, BGP, or RIP would update Contivity2 so that Contivity1 no longer has a route to Host2.

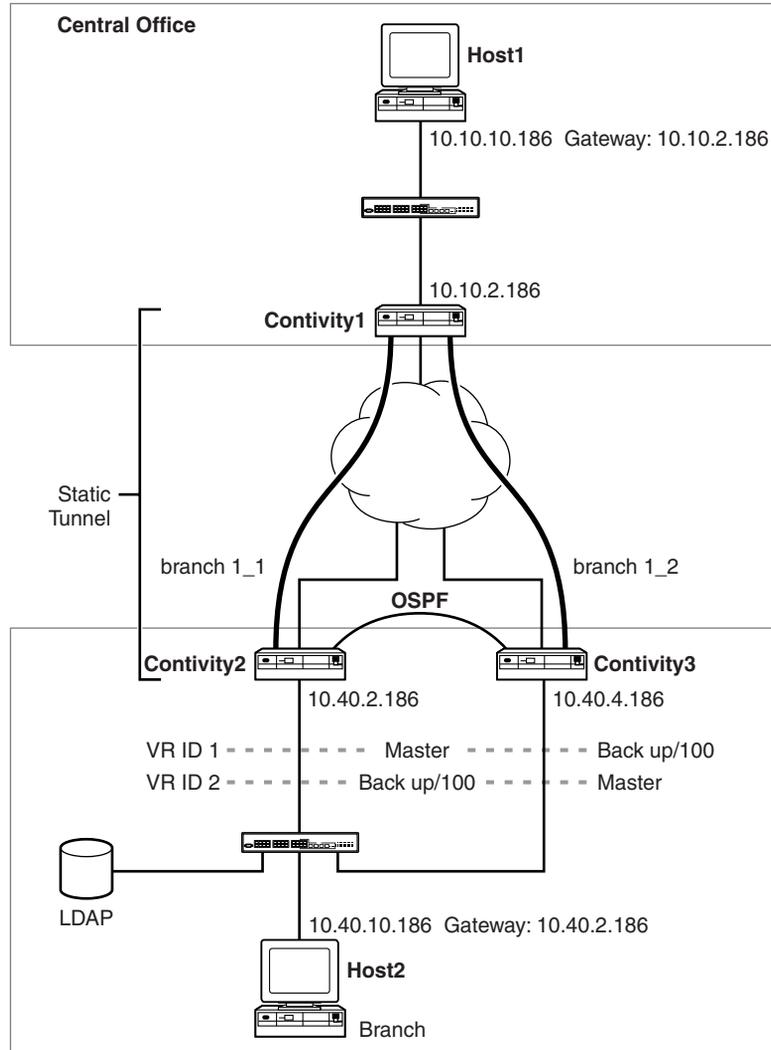
The VRRP failover occurs within 3 seconds based on the default configuration. Use of OSPF on the tunnels guarantees a maximum failover time of 40 seconds based on the default configuration. However, by setting the appropriate value for the OSPF hello interval, failover time can be drastically reduced.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

**Figure 12** Sample high-availability environment

In the previous example, if the branch office tunnels are static routes and Host2's default gateway Contivity2 encounters a public interface failure (private interface 10.40.2.186 remained active), VRRP would not failover. If Contivity2 is unaware that another route to Host1 exists through Contivity3, it will drop all traffic from Host2 destined to Host1. To correct this, another route to Host1 through Contivity3 must be added to Contivity2's route table. One way to add this route is shown in [Figure 13 on page 103](#).

In [Figure 13](#), an OSPF branch office tunnel is added between Contivity2 and Contivity3 to provide both with a secondary route to Host1. Because static routes are preferred over OSPF, both Contivity2 and Contivity3 will always use their static route to Host1 through Contivity1 if it is available. This inter-gateway branch office tunnel does not have to use OSPF. RIP or a static route of higher cost would work equally well.

**Figure 13** VRRP and static tunnels



## Configuring VRRP on the Contivity gateway

To configure VRRP for LAN and VLAN on the Contivity gateway:

- 1 Go to the Routing > VRRP window, and check Enabled.
- 2 Enter the IP address for the virtual router, and click Create.
- 3 In the VRRP > Edit VRRP IP Address window, enter a decimal value from 1 through 255 for the VRID. This number must be unique to the LAN or VLAN segment running VRRP and common to all Contivity gateways that participate within this virtual router group.
- 4 In the Advertise Interval box, enter the rate the virtual router advertises its hello messages. The range is 1 through 255 seconds and the default is 1 second.
- 5 Select None or Simple as the authentication type for this virtual router. None means that VRRP protocol exchanges are not authenticated; Simple means they are authenticated by a simple text password.
- 6 If you choose Simple authentication, enter up to 8 characters of text for the authentication string and confirm it.
- 7 The Master Delay mode controls when a Contivity gateway takes mastership of an address it owns. Normally, this occurs as soon as the interface is enabled. Master Delay mode makes it is possible to delay when the master's assertion happens. Master Delay mode operates in one of two possible ways: Delay or Time of Day. The default for a VR in Master Delay mode is disabled (None).



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

---

- 8 Click OK.
- 9 Go to the Routing > Interfaces window and click the Configure button next to VRRP for the appropriate interface. The LAN (with corresponding physical address on the box) and VLAN interfaces are automatically displayed in the

Master Status section and all others are displayed in the Backed up Addresses section.

- 10** Check or uncheck Enable to enable or disable VRRP for this interface.
- 11** In the Master Status section, enable all interfaces that you want to be master and click OK.

The Current Backed up Addresses section displays information about the currently configured backups. Displayed are the IP addresses this subinterface is backing up, the VRID it is using, its configured state (which can be Enabled or Disabled), and the current operational state and its priority.

- 12** In the New Backed up Address section, back up an IP address by selecting an IP address from the menu.
- 13** Enter a priority number in the Priority box.
- 14** Click Add.

## Configuring IP addresses for backups

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

For example, for Contivity2 to be the master of VRID 1 and Contivity3 to be its backup, configure the following:

- 1** On Contivity2, go to Routing > VRRP and add IP address 10.40.2.186 with VRID 1.
- 2** In the Routing > Interfaces window, select 10.40.2.186 and configure and check the Master Box.
- 3** On Contivity3, go to Routing > VRRP and add IP Address 10.40.4.186 with a VRID not equal to 1 (use 2) and add IP Address 10.40.2.186 with a VRID equal to 1.
- 4** In the Routing > Interface window, select 10.40.4.186 and configure. From the New Backed up Address, select 10.40.2.186, VRID 1 and click ADD.

To configure Contivity3 to be the master of VRID 2:

- 1 On Contivity2, go to Routing > VRRP and add IP address 10.40.2.186 with VRID 1.
- 2 In the Routing > VRRP window, add IP address 10.40.4.186 with VRID 2.
- 3 In the Routing > Interfaces window, select 10.40.2.186 and the Master Box next to 10.40.2.186. The Backed Up list contains 10.40.4.186 VRID 2.
- 4 On Contivity3, go to Routing > VRRP and add IP address 10.40.4.186 with VRID 2.
- 5 In the Routing > VRRP window, add IP address 10.40.2.186 with VRID 1.
- 6 In the Routing > Interfaces window, select 10.40.4.186 and the Master Box next to 10.40.4.186. Backed Up list contains 10.40.2.186 VRID 1.

For example, for a VLAN to be the master of VRID 1 and Contivity3 to be its backup, configure the following:

- 1 On VLAN, go to Routing > VRRP and add IP address 1.1.1.1 with VRID 1.
- 2 In the Routing > Interfaces window, select 1.1.1.1 and configure and check the Master Box.
- 3 On Contivity3, go to Routing > VRRP and add IP Address 10.40.4.186 with a VRID not equal to 1 (use 2) and add IP Address 1.1.1.1 with a VRID equal to 1.
- 4 In the Routing > Interface window, select 10.40.4.186 and configure. From the New Backed up Address, select 1.1.1.1, VRID 1 and click ADD.

## Interface groups and critical interface failover

Interface groups support the backup interface services (BIS), which is an automated mechanism to back up an interface when a designated primary connection fails, and VRRP critical interface failover. An interface group is a logical grouping of interfaces (physical or tunnel) defined in a Contivity gateway. The group may consist of a single physical interface, an IP address, a list of

physical interfaces, a tunnel, a list of tunnels or any combination of physical interfaces and tunnels defined on the gateway. Status of a critical interface is defined to be up when at least one of its members is up and is down when all of its members are down.



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

---

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

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



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

---

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

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

To configure interface groups:

- 1 Go to the Routing > Interface Group window.
- 2 Click Add to access the Interface Group > Add window.
- 3 Enter a name for the group in the name field.

- 4** Select and move the available interfaces that you want to include in the group into the Interfaces in Group list.
- 5** To find interface groups with a given interface, enter the IP address and click Search.
- 6** Click OK.
- 7** Go to the Routing > Interfaces > Configure VRRP window for the VRRP interface that will be associated with the critical interface group.
- 8** Under Master Status, select the interface group from the list.
- 9** Click Enabled, and then click OK to enable the VRRP critical interface.

---

# Index

---

## A

Accept policies 83  
advanced routing key 42, 54  
Announce policies 83

## B

BGP 55

## C

client address redistribution 89  
  configuring 92  
  sample 90  
  summarization 90

## D

default route 50

## E

equal cost multipath (ECMP) 49

## I

interface filter  
  Permit All 96

## L

load balancing 50

## M

multicast 96

multicast relay 95

## O

OSPF  
  configuration 44  
  overview 41

## P

Permit All 96  
ping  
  validating public default route 81  
poison reverse 34  
publications  
  hard copy 19

## R

RIP 33  
  using 33, 41  
route redistribution 85  
route selection 30  
route table 23, 27  
  lookup 29  
routes  
  default 50  
  dynamic 27  
  static 27  
routing  
  dynamic 22  
  enhanced 22  
  integrated firewall 22  
  loops 34  
  overview 21

- policy 87
- policy service 83
- route lookup 29
- route table types 29
- rules of redistribution 85
- table 27

## S

- split horizon 34
- static routes 23, 79
- status 24

## T

- technical publications 19
- triggered updates 35

## U

- U tunnel 89

## V

- virtual links 43
- VRRP
  - configuring 104
  - failover 101
  - high availability 100
  - master interface 102
- VRRP overview 99