

Part No. 314721-B Rev 00  
May 2003

4655 Great America Parkway  
Santa Clara, CA 95054

# Configuring BGP Services

**Passport 8000 Series  
Software Release 3.5**



**NORTEL**  
**NETWORKS™**

## Copyright © 2003 Nortel Networks

All rights reserved. May 2003.

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, the Globemark, Unified Networks, and PASSPORT are trademarks of Nortel Networks.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation.

Adobe and Acrobat Reader are trademarks of Adobe Systems Incorporated.

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

Juniper Networks and JUNOS are trademarks of Juniper Networks, Inc.

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>17</b>
Before you begin .....	18
Text conventions .....	18
Hard-copy technical manuals .....	19
How to get help .....	20
<b>Chapter 1</b>	
<b>BGP concepts</b> .....	<b>21</b>
Autonomous systems .....	22
Internal/external BGP routing .....	22
BGP speaker .....	23
Transit AS .....	25
Stub and multihomed autonomous systems .....	25
Peers .....	25
Supernet advertisement .....	26
Reducing bandwidth and maintenance .....	26
Consolidating routing information .....	26
CIDR and aggregate addresses .....	27
Creating a supernet address .....	28
Aggregate routes .....	31
Confederations .....	32
IBGP route reflection .....	34
BGP updates .....	38
Withdrawn Routes Length field .....	39
Withdrawn Routes field .....	39
Total Path Attributes Length field .....	40
Path Attributes field .....	40
Attribute Type field .....	41

Attribute Length field . . . . .	43
Attribute Value field . . . . .	43
Network Layer Reachability Information field . . . . .	43
Equal-cost multipath . . . . .	43
TCP MD5 message authentication . . . . .	44
Entering and Storing MD5 Authentication Keys . . . . .	45
Generation of MD5 signatures on ingress BGP TCP packets . . . . .	46
Verification of MD5 signatures on egress BGP TCP packets . . . . .	46
Circuitless IP . . . . .	47
<b>Chapter 2</b>	
<b>Configuration considerations and limitations. . . . .</b>	<b>49</b>
BGP implementation notes . . . . .	50
Configuration guidelines . . . . .	51
BGP Neighbor Maximum Prefix . . . . .	52
BGP-4 support . . . . .	53
RIP redistribution . . . . .	53
BGP/OSPF interaction . . . . .	53
MD5 authentication . . . . .	54
<b>Chapter 3</b>	
<b>Using Device Manager to configure BGP . . . . .</b>	<b>55</b>
Configuring general parameters . . . . .	56
Displaying global BGP statistics . . . . .	61
Configuring aggregate addresses . . . . .	65
Configuring allowed network addresses . . . . .	68
Configuring and displaying peer information . . . . .	69
Peers tab . . . . .	70
Configuring and editing peers . . . . .	70
Displaying peer statistics . . . . .	76
Peers Info tab . . . . .	85
Peer Groups tab . . . . .	88
Displaying route information . . . . .	93
BGP community attributes . . . . .	97
Displaying dampened routes information . . . . .	100

---

Configuring redistribute entries .....	101
Configuring Circuitless IP .....	103
Configuring a circuitless IP interface .....	104
Enabling OSPF on a circuitless IP interface .....	106
Deleting a circuitless IP interface .....	107
Copying a circuitless IP interface .....	108
Exporting circuitless IP interface data .....	108
Specifying global parameters .....	109
<b>Chapter 4</b>	
<b>Using the CLI to configure BGP .....</b>	<b>113</b>
Roadmap of BGP CLI commands .....	114
Configuring general BGP parameters .....	120
Configuring BGP confederations .....	128
Configuring BGP peers or peer groups .....	129
Configuring route redistribution parameters .....	136
Configuring BGP route redistribution parameters .....	136
Configuring OSPF route redistribution parameters .....	136
Configuring Direct route redistribution parameters .....	138
Configuring RIP route redistribution parameters .....	140
Configuring Static route redistribution parameters .....	141
Configuring AS-path lists .....	142
Configuring community lists .....	144
Showing BGP configurations .....	146
Showing BGP aggregates .....	146
Showing BGP CIDR routes .....	148
Showing the BGP configuration .....	149
Showing flap-dampened routes .....	151
Showing global flap-dampening configurations .....	152
Showing a BGP redistribution list .....	153
Showing BGP imported routes .....	154
Showing BGP network configurations .....	156
Showing BGP peer groups .....	156
Showing BGP routes .....	158
Showing BGP global statistics .....	159

Showing BGP summaries .....	164
Showing BGP peers .....	165
Showing BGP peer statistics .....	166
Showing BGP peer routes .....	171
Circuitless IP .....	173
Configuring Circuitless IP .....	174
Showing Circuitless IP output .....	176
Configuring BGP debug commands .....	177
Tips for using the debug commands .....	177
BGP global debug commands .....	178
config ip bgp global-debug mask command .....	178
config ip bgp neighbor-debug-all mask command .....	179
Global debug command output examples .....	179
All debug output .....	179
Error debug output .....	181
Packet debug output .....	181
Event debug output .....	182
Warning debug output .....	182
Init debug output .....	183
Filter debug output .....	183
Update debug output .....	184
BGP peer/peer group debug commands .....	185
Peer and peer group debug command output examples .....	186
All debug output .....	186
Error debug output .....	188
Packet debug output .....	188
Event debug output .....	190
Warning debug output .....	190
State debug output .....	191
Filter debug output .....	192
Update debug output .....	194
<b>Chapter 5</b>	
<b>Configuration examples .....</b>	<b>195</b>
Basic BGP example .....	196

---

Policies .....	202
Creating OSPF and BGP route distribution policies .....	202
Configuration commands .....	202
Redistribution considerations and tips .....	204
Creating Direct, RIP, and Static route distribution policies .....	204
Injecting a default route when using OSPF as an IGP .....	205
CIDR and aggregate addressing example .....	209
EBGP multihop and EBGP load balance .....	212
BGP synchronization and next-hop self .....	215
Example 1 — initial configuration .....	215
Example 2 — unexpected changes to the initial topology .....	222
Example 3 — how to correct the next hop problem .....	225
MD5 authentication .....	230
BGP peer groups .....	233
BGP path attributes .....	236
Origin attribute .....	237
AS path attribute .....	241
AS path filtering .....	243
Local preference attribute .....	246
Adding preferences to specific routes .....	250
Multi-exit discriminator (MED) attribute .....	253
Community attribute .....	258
IBGP scalability issues .....	263
BGP confederations .....	264
Route reflectors .....	269
Multiple route reflectors .....	274
Route flap dampening .....	281
<b>Appendix A</b>	
<b>Translating Cisco Systems-to-Nortel Networks command equivalents .</b>	<b>285</b>
Configuration Commands .....	285
Interpreting Cisco Systems-to-Nortel Networks command equivalents .....	295
Operational Commands .....	298
Interpreting Cisco Systems-to-Nortel Networks Operational equivalents .....	299
Route preferences .....	300

<b>Appendix B</b>	
<b>Translating Juniper Networks-to-Nortel Networks equivalents. . . . .</b>	<b>301</b>
Configuration Commands . . . . .	301
Interpreting Juniper Networks-to-Nortel Networks command equivalents . . . . .	313
Operational Commands . . . . .	316
Interpreting Juniper Networks-to-Nortel Networks BGP Operational equivalents	317
Route preferences . . . . .	318
<b>Glossary . . . . .</b>	<b>319</b>
<b>Index . . . . .</b>	<b>323</b>

---

## Figures

---

Figure 1	Internetwork segmented into three autonomous systems	23
Figure 2	Transit AS, stub AS, multihomed AS, and peer-to-peer sessions,	24
Figure 3	Converting to binary notation	29
Figure 4	Combining into two aggregate ranges (1 of 2)	30
Figure 5	Combining into two aggregate ranges (2 of 2)	31
Figure 6	Aggregating routes with CIDR	32
Figure 7	Confederations	33
Figure 8	Fully meshed AS with IBGP speakers	35
Figure 9	AS with route reflector	36
Figure 10	Route reflector with client and nonclient peers	37
Figure 11	Update message format	38
Figure 12	IP Prefix format	39
Figure 13	Path attribute encoding	40
Figure 14	Attribute Type fields	41
Figure 15	Routers with I-BGP connections	47
Figure 16	BGP dialog box—Generals tab	56
Figure 17	BGP dialog box—Global Stats tab	61
Figure 18	BGP dialog box—Aggregates tab	65
Figure 19	BGP Insert Aggregates dialog box	66
Figure 20	BGP dialog box—Network tab	68
Figure 21	BGP Insert Network dialog box	68
Figure 22	BGP dialog box—Peers tab	71
Figure 23	BGP Insert Peers dialog box	72
Figure 24	BGP dialog box—Peers tab	76
Figure 25	BGP Peer Stat dialog box—General Stats tab	77
Figure 26	BGP Peer Stat dialog box—Receive Stats tab	80
Figure 27	BGP Peer Stat dialog box—Transmit Stats tab	83
Figure 28	BGP dialog box—Peers Info tab	86
Figure 29	BGP dialog box—Peer Groups tab	89

Figure 30	BGP Insert Peer Groups dialog box . . . . .	90
Figure 31	BGP dialog box—BGP Routes tab (1 of 2) . . . . .	93
Figure 32	BGP Route Range dialog box . . . . .	94
Figure 33	BGP dialog box—BGP Routes tab (2 of 2) . . . . .	95
Figure 34	BGP dialog box—BGP Community Attributes tab (1 of 2) . . . . .	97
Figure 35	BGP Route Range dialog box . . . . .	97
Figure 36	BGP dialog box—BGP Community Attributes tab (2 of 2) . . . . .	99
Figure 37	BGP dialog box—Dampened Routes tab . . . . .	100
Figure 38	BGP dialog box—BGP Redistribute tab . . . . .	102
Figure 39	BGP, Insert BGP Redistribute dialog box . . . . .	102
Figure 40	IP dialog box — Globals tab . . . . .	104
Figure 41	IP dialog box — Circuitless IP tab . . . . .	105
Figure 42	IP, Insert Circuitless dialog box . . . . .	105
Figure 43	OspfCircuitless dialog box . . . . .	107
Figure 44	IP dialog box—Globals tab . . . . .	110
Figure 45	config ip bgp info command output . . . . .	127
Figure 46	config ip bgp confederation info command output . . . . .	128
Figure 47	config ip bgp neighbor info command output . . . . .	135
Figure 48	config ip bgp redistribute ospf info command output . . . . .	138
Figure 49	config ip bgp redistribute direct info command output . . . . .	139
Figure 50	config ip bgp redistribute rip info command output . . . . .	141
Figure 51	config ip bgp redistribute static info command output . . . . .	142
Figure 52	config ip as-list info command output . . . . .	143
Figure 53	config ip community-list info command output . . . . .	145
Figure 54	show ip bgp aggregates command output . . . . .	147
Figure 55	show ip bgp cidr-only command output . . . . .	148
Figure 56	show ip bgp conf command output . . . . .	150
Figure 57	show ip bgp dampened-paths command output . . . . .	151
Figure 58	show ip bgp flap-damp-config command output . . . . .	152
Figure 59	show ip bgp redistribution command output . . . . .	153
Figure 60	show ip bgp imported-routes command output . . . . .	155
Figure 61	show ip bgp networks command output . . . . .	156
Figure 62	show ip bgp peer-group command output . . . . .	157
Figure 63	show ip bgp route command output . . . . .	158
Figure 64	showing global BGP statistics . . . . .	160

---

Figure 65	show ip bgp summary command output . . . . .	164
Figure 66	showing BGP peers . . . . .	166
Figure 67	showing BGP peer statistics . . . . .	167
Figure 68	show ip bgp neighbor route command output . . . . .	172
Figure 69	show ip circuitless-ip-int info command output . . . . .	176
Figure 70	All debug output example . . . . .	180
Figure 71	Error debug output example . . . . .	181
Figure 72	Packet debug output example . . . . .	181
Figure 73	Event debug output example . . . . .	182
Figure 74	Warning debug output example . . . . .	182
Figure 75	Init debug output example . . . . .	183
Figure 76	Filter debug output example . . . . .	183
Figure 77	Update debug output example . . . . .	184
Figure 78	Peer/peer group all debug output example . . . . .	187
Figure 79	Peer/peer group error debug output example . . . . .	188
Figure 80	Peer/peer group packet debug output example . . . . .	189
Figure 81	Peer/peer group event debug output example . . . . .	190
Figure 82	Peer/peer group warning debug output example . . . . .	191
Figure 83	Peer/peer group state debug output example . . . . .	191
Figure 84	Peer/peer group filter debug output example . . . . .	193
Figure 85	Update debug output example . . . . .	194
Figure 86	Basic BGP configuration . . . . .	197
Figure 87	show vlan info brouter-port command output . . . . .	199
Figure 88	config ip bgp redistribute static info command output . . . . .	205
Figure 89	Injecting a default route example . . . . .	206
Figure 90	Aggregation example . . . . .	209
Figure 91	show ip bgp route command output . . . . .	211
Figure 92	EBGP load balance example . . . . .	212
Figure 93	BGP synchronization and next-hop self example . . . . .	216
Figure 94	BGP synchronization and next hop-self example 2 . . . . .	222
Figure 95	show ip bgp route command output (Example 1) from R2 . . . . .	223
Figure 96	show ip route info command output from R2 . . . . .	224
Figure 97	show ip bgp route command output (Example 3A) . . . . .	226
Figure 98	show ip route info command output (Example 3B) . . . . .	227
Figure 99	show ip route info command output (Example 3C) . . . . .	229

Figure 100 MD5 authentication example .....	230
Figure 101 Peer group example .....	233
Figure 102 Static route distribution .....	238
Figure 103 show ip bgp route command output .....	240
Figure 104 Using the AS Path attribute to influence inbound traffic flow .....	241
Figure 105 AS path filtering example .....	243
Figure 106 Using the local preference attribute to influence outbound traffic flow ..	246
Figure 107 Using the local preference attribute to influence specific routes .....	250
Figure 108 Using the MED attribute to influence specific routes .....	253
Figure 109 Using the MED attribute for load balancing routes .....	255
Figure 110 Using the community attribute to control routes .....	259
Figure 111 show ip bgp route community enable command output .....	262
Figure 112 Confederation example .....	264
Figure 113 Route reflector example .....	270
Figure 114 Multiple route reflector example .....	275
Figure 115 Route flap dampening example .....	281
Figure 116 show ip bgp flap-damp-config command output .....	282
Figure 117 show ip bgp dampened-paths 10.2.2.2 command output (Example 1) ..	282
Figure 118 show ip bgp dampened-paths 10.2.2.2 command output (Example 2) ..	283

---

## Tables

---

Table 1	CIDR Conversion	27
Table 2	BGP mandatory path attributes	42
Table 3	MD5 Signature Verification Rules on BGP TCP Packets	46
Table 4	BGP dialog box—Generals tab fields	57
Table 5	Global Stats tab descriptions	62
Table 6	BGP Insert Aggregates dialog box fields	66
Table 7	BGP Insert Network dialog box	69
Table 8	BGP Insert Peers dialog box	73
Table 9	General Stats tab fields	78
Table 10	Receive Stats tab fields	80
Table 11	Transmit Stats tab fields	84
Table 12	BGP Peers Info tab fields	86
Table 13	BGP Insert Peer Groups dialog box fields	91
Table 14	BGP Route Range dialog box options	94
Table 15	BGP Routes tab fields	96
Table 16	BGP Route Range dialog box options	98
Table 17	BGP Routes Info tab fields	99
Table 18	Dampened Routes tab dialog box	101
Table 19	BGP Insert BGP Redistribute dialog box fields	102
Table 20	IP dialog box—Circuitless IP tab fields	105
Table 21	IP, Insert Circuitless dialog box fields	106
Table 22	IP dialog box—Globals tab fields	110
Table 23	show ip bgp cidr-only command parameters	149
Table 24	show ip bgp dampened-paths command parameters	151
Table 25	show ip bgp dampened-paths command parameters	153
Table 26	show ip bgp redistribution command parameters	154
Table 27	show ip bgp imported-routes command parameters	155
Table 28	show ip bgp route command parameters	158
Table 29	BGP global statistics descriptions	161

Table 30	show ip bgp route command parameters . . . . .	164
Table 31	BGP peer statistics descriptions . . . . .	168
Table 32	show ip bgp neighbor route command parameters . . . . .	173
Table 33	Cisco Systems-to-Nortel Networks command equivalents . . . . .	285
Table 34	Cisco Systems-to-Nortel Networks operational commands . . . . .	298
Table 35	Cisco Systems-to-Nortel Networks route preference comparison . . . . .	300
Table 36	Juniper Networks-to-Nortel Networks command equivalents . . . . .	301
Table 37	Juniper Networks-to-Nortel Networks operational commands . . . . .	316
Table 38	Juniper Networks-to-Nortel Networks route preference comparison . . .	318

---

## Preface

---

This guide provides instructions for configuring Border Gateway Protocol (BGP) services for Passport\* 8000 Series switches. The instructions include information about using both the Device Manager graphical user interface (GUI) and the command line interface (CLI) to perform BGP configuration and management operations.

BGP routers form *peer* relationships with other *neighboring* BGP routers. When you use the CLI to configure your system, the neighbor command lets you configure peers and peer groups. When you use Device Manager to configure your system, you can use the Peer tabs to configure peers and peer groups.

This guide has five chapters, two appendixes, a glossary, and an index:

<b>If you want to:</b>	<b>Go to:</b>
Learn about BGP concepts and key features for this release	Chapter 1
Review important configuration limitations and interoperability issues that should be considered when configuring BGP	Chapter 2
Use Device Manager to configure and manage BGP	Chapter 3
Use the Command Line Interface (CLI) to configure and manage BGP	Chapter 4
See examples of common BGP configuration tasks, with included CLI commands used to create the configuration	Chapter 5
See Cisco Systems-to-Nortel Networks command equivalents	Appendix A
See Juniper Networks-to-Nortel Networks command equivalents	Appendix B
Review a glossary of terms used with BGP	Glossary
See an alphabetical listing of the topics and subtopics in this guide, with cross-references to relevant information	Index

## Before you begin

This guide is intended for network administrators with the following background:

- Basic knowledge of networks, Ethernet bridging, and IP routing
- Familiarity with networking concepts and terminology
- Basic knowledge of network topologies
- Experience with windowing systems or GUIs

## 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 <code>ping &lt;ip_address&gt;</code> , you enter <code>ping 192.32.10.12</code>
<b>bold Courier text</b>	Indicates command names and options and text that you need to enter. Example: Use the <b>dinfo</b> command. Example: Enter <b>show ip {alerts routes}</b> .
braces ({} )	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. Example: If the command syntax is <code>show ip {alerts routes}</code> , you must enter either <code>show ip alerts</code> or <code>show ip routes</code> , but not both.
brackets ([ ])	Indicate optional elements in syntax descriptions. Do not type the brackets when entering the command. Example: If the command syntax is <code>show ip interfaces [-alerts]</code> , you can enter either <code>show ip interfaces</code> or <code>show ip interfaces -alerts</code> .

ellipsis points ( . . . )	Indicate that you repeat the last element of the command as needed. Example: If the command syntax is <code>ethernet/2/1 [&lt;parameter&gt; &lt;value&gt;] . . .</code> , you enter <code>ethernet/2/1</code> and as many parameter-value pairs as needed.
<i>italic text</i>	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. Example: If the command syntax is <code>show at &lt;valid_route&gt;</code> , <code>valid_route</code> is one variable and you substitute one value for it.
plain Courier text	Indicates command syntax and system output, for example, prompts and system messages. Example: <code>Set Trap Monitor Filters</code>
separator ( > )	Shows menu paths. Example: <code>Protocols &gt; IP</code> identifies the IP option on the Protocols menu.
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 <code>show ip {alerts routes}</code> , you enter either <code>show ip alerts</code> or <code>show ip routes</code> , but not both.

## Hard-copy technical manuals

You can print selected technical manuals and release notes free, directly from the Internet. Go to the [www.nortelnetworks.com/documentation](http://www.nortelnetworks.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 at the [www.adobe.com](http://www.adobe.com) URL to download a free copy of the Adobe Acrobat Reader.

You can purchase printed books and documentation sets from Vervante. To order printed documentation, go to Vervante at the [www.vervante.com/nortel](http://www.vervante.com/nortel) URL.

## How to get help

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

If you purchased a Nortel Networks\* service program, contact one of the following Nortel Networks Technical Solutions Centers:

Technical Solutions Center	Telephone
Europe, Middle East, and Africa	(33) (4) 92-966-968
North America	(800) 4NORTEL or (800) 466-7835
Asia Pacific	(61) (2) 9927-8800
China	(800) 810-5000

Additional information about the Nortel Networks\* Technical Solutions Centers is available from the [www.nortelnetworks.com/help/contact/global](http://www.nortelnetworks.com/help/contact/global) URL.

An Express Routing Code (ERC) is available for many Nortel Networks\* products and services. When you use an ERC, your call is routed to a technical support person who specializes in supporting that product or service. To locate an ERC for your product or service, go to the <http://www.nortelnetworks.com/help/contact/erc/index.html> URL.

---

# Chapter 1

## BGP concepts

---

BGP is an inter-domain routing protocol that provides loop-free inter-domain routing between *autonomous systems* (ASs) or within an AS. BGP systems can exchange network layer reachability information (NLRI) with other BGP systems for the purpose of constructing a graph of AS connectivity. BGP uses this information to prune routing loops and enforce AS-level policy decisions. BGP provides features that allow you to consolidate routing information and to control the flow of BGP updates.

The following sections provide an overview of BGP and includes descriptions of features you can use to optimize your BGP system.



**Note:** See Chapter 5, “Configuration examples,” on page 195, for configuration examples, including commands, for most of the concepts described in this chapter.

---

This chapter includes the following topics:

Topic	Page
<a href="#">Autonomous systems</a>	22
<a href="#">Consolidating routing information</a>	26
<a href="#">BGP updates</a>	38
<a href="#">Equal-cost multipath</a>	43
<a href="#">TCP MD5 message authentication</a>	44
<a href="#">Circuitless IP</a>	47

## Autonomous systems

An AS is a group of routers and hosts run by a single technical administrator that has a single, clearly defined routing policy. Each autonomous system has its own unique *AS number* assigned by the appropriate Internet Registry entity.



**Note:** LANs and WANs interconnected by IP routers form a group of networks called an *internetwork*. For administrative purposes, internetworks are divided into boundaries known as autonomous systems.

---

Figure 1 shows a sample internetwork segmented into three autonomous systems.

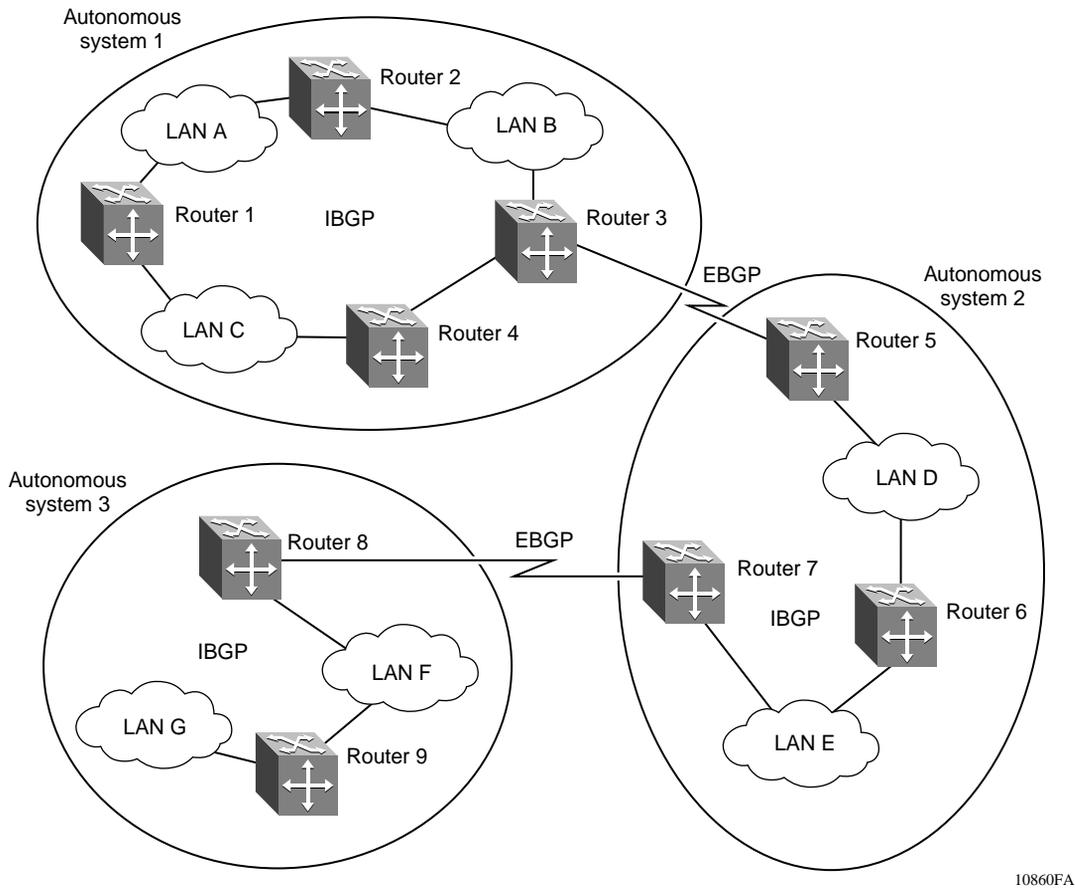
BGP exchanges information between ASs as well as between routers within the same AS. As shown in [Figure 1 on page 23](#), routers that are members of the same AS and exchange BGP updates run *internal BGP (IBGP)*, and routers that are members of different ASs and exchange BGP updates run *external BGP (EBGP)*.

After reviewing the AS concepts in this section, you can find more information about ASs in the section titled, “[Basic BGP example](#)” on page 196. That section provides a basic configuration example, including CLI commands used to create the topology.

## Internal/external BGP routing

Nortel Networks\* supports both Internal BGP (IBGP) intra-AS routing and External BGP (EBGP) external-AS routing. With IBGP, each router within an AS runs an interior gateway protocol (IGP), such as routing information protocol (RIP), and open shortest path first (OSPF). The IBGP information, along with the IGP route to the originating BGP border router, determines the next hop to use for exchanging information with an external AS. Each router uses IBGP exclusively to determine reachability to external ASs. When a router receives an IBGP update that is destined for an external AS, the update is passed to IP for inclusion in the routing table only if a viable IGP route to the correct border gateway is available.

EBGP is used to communicate routing information between BGP speakers that are in different ASs.

**Figure 1** Internetwork segmented into three autonomous systems

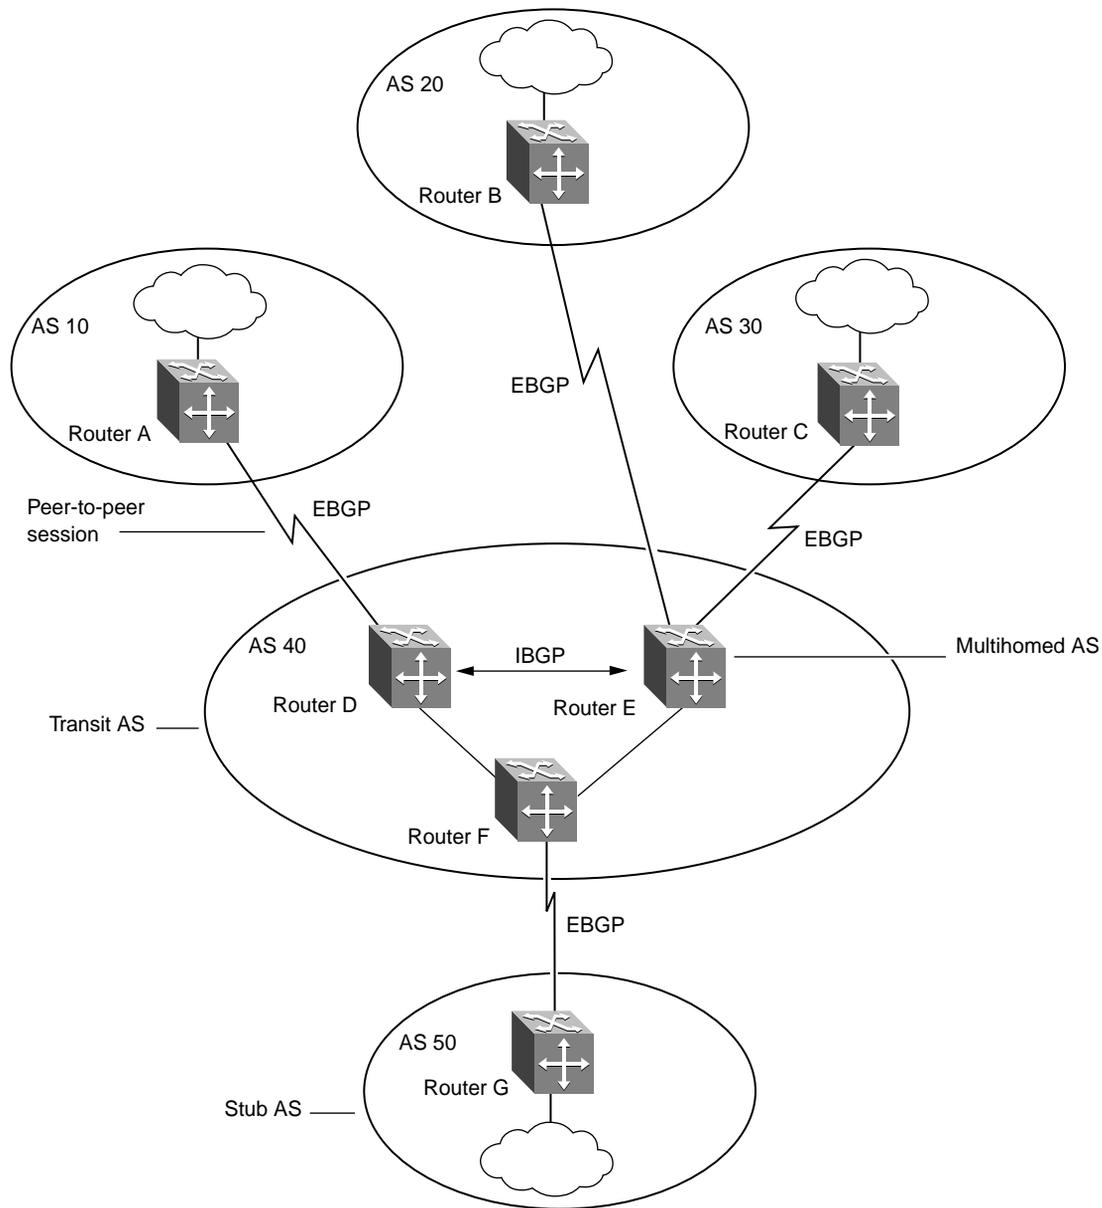
10860FA

## BGP speaker

BGP routers employ an entity within the router, referred to as a *BGP speaker*, that transmits and receives BGP messages and acts upon them. BGP speakers communicate with other BGP speakers by establishing a *peer-to-peer session*.

All BGP speakers within an AS must be fully meshed (see [Figure 2 on page 24](#)).

**Figure 2** Transit AS, stub AS, multihomed AS, and peer-to-peer sessions,



10861FA

## Transit AS

An AS with more than one BGP speaker can use IBGP to provide a transit service for networks located outside the AS. An AS that provides this service is called a *transit AS*. As shown in [Figure 2 on page 24](#), AS 40 is the transit AS. It provides information about its internal networks, as well as transit networks, to the remaining ASs. The IBGP connections between routers D, E, and F provide consistent routing information to the ASs.

## Stub and multihomed autonomous systems

As shown in [Figure 2 on page 24](#), an AS can include one or more BGP speakers that establish peer-to-peer sessions with BGP speakers in other ASs to provide external route information for the networks within the AS:

- A *stub* AS has a single BGP speaker that establishes a peer-to-peer session with one external BGP speaker. In this case, the BGP speaker provides external route information only for the networks contained within its own AS.
- A *multihomed* AS has multiple BGP speakers.

## Peers

The transport protocol used with BGP is Transmission Control Protocol (TCP). When any two routers open a TCP connection to each other for the purpose of exchanging routing information, they form a *peer-to-peer* relationship. In [Figure 2 on page 24](#), Routers A and D are BGP peers, as are Routers B and E, C and E, F and G, and Routers D, E, and F.

Note that although Routers A and D are running EBGP, Routers D, E, and F within AS 40 are running IBGP. The EBGP peers are directly connected, while the IBGP peers are not. As long as an IGP is running that allows any two neighbors to logically communicate, the IBGP peers do not require a direct connection.

Because all BGP speakers within an AS must be fully meshed logically, the IBGP mesh can grow to large proportions and become difficult to manage. You can reduce the number of peers within an AS by creating *confederations* and *route reflectors* (see [“Confederations” on page 32](#), and [“IBGP route reflection” on page 34](#)).

BGP peers exchange complete routing information only after the peer connection is established. Thereafter, BGP peers exchange routing updates. An update message consists of a network number, a list of autonomous systems that the routing information passed through (the AS path), and other path attributes that describe the route to a set of destination networks. When multiple paths are available, BGP compares the path attributes to choose the preferred path. For more information about update messages, see [“BGP updates” on page 38](#).

## Supernet advertisement

BGP has no concept of address classes. Each network listed in the network layer reachability information (NLRI) portion of an update message contains a prefix length field, which describes the length of the mask associated with the network. The prefix length field allows for both supernet and subnet advertisement. The supernet advertisement is what makes classless interdomain routing (CIDR) possible (see [“CIDR and aggregate addresses” on page 27](#)).

## Reducing bandwidth and maintenance

BGP also provides two features that reduce the high bandwidth and maintenance costs associated with a large full-mesh topology:

- Confederations
- Route reflectors

Confederations and route reflectors are discussed in [“Consolidating routing information,”](#) next.

## Consolidating routing information

This section describes BGP features that allow you to reduce the size of your routing tables.

The section includes the following topics:

- [“CIDR and aggregate addresses,”](#) next
- [“Aggregate routes” on page 31](#)

- “Confederations” on page 32
- “IBGP route reflection” on page 34

## CIDR and aggregate addresses

Classless interdomain routing (CIDR) is an addressing scheme (also known as supernetting) that eliminates the concept of classifying networks into class types. Earlier addressing schemes identified five classes of networks: Class A, Class B, Class C, Class D, and Class E. Classes D (used for multicast) and E (reserved and currently not used) are not discussed in this book.

For example, network 195.215.0.0, an illegal Class C network number, becomes a legal supernet when it is represented in CIDR notation as 195.215.0.0/16. The /16 is called the prefix length and becomes a way of expressing the explicit mask that CIDR requires. In this case, the addition of the prefix /16 indicates that the subnet mask consists of 16 bits (counting from the left).

Note that with this method, supernet 195.215.0.0/16 represents 195.215.0.0 255.255.0.0 (see [Table 1](#)).

**Table 1** CIDR Conversion

Prefix	Dotted-decimal	Binary	Network class
/1	128.0.0.0	1000 0000 0000 0000 0000 0000 0000 0000	128 Class As
/2	192.0.0.0	1100 0000 0000 0000 0000 0000 0000 0000	64 Class As
/3	224.0.0.0	1110 0000 0000 0000 0000 0000 0000 0000	32 Class As
/4	240.0.0.0	1111 0000 0000 0000 0000 0000 0000 0000	16 Class As
/5	248.0.0.0	1111 1000 0000 0000 0000 0000 0000 0000	8 Class As
/6	252.0.0.0	1111 1100 0000 0000 0000 0000 0000 0000	4 Class As
/7	254.0.0.0	1111 1110 0000 0000 0000 0000 0000 0000	2 Class As
/8	255.0.0.0	1111 1111 0000 0000 0000 0000 0000 0000	1 Class A or 256 Class Bs
/9	255.128.0.0	1111 1111 1000 0000 0000 0000 0000 0000	128 Class Bs
/10	255.192.0.0	1111 1111 1100 0000 0000 0000 0000 0000	64 Class Bs
/11	255.224.0.0	1111 1111 1110 0000 0000 0000 0000 0000	32 Class Bs
/12	255.240.0.0	1111 1111 1111 0000 0000 0000 0000 0000	16 Class Bs

**Table 1** CIDR Conversion (continued)

Prefix	Dotted-decimal	Binary	Network class
/13	255.248.0.0	1111 1111 1111 1000 0000 0000 0000 0000	8 Class Bs
/14	255.252.0.0	1111 1111 1111 1100 0000 0000 0000 0000	4 Class Bs
/15	255.254.0.0	1111 1111 1111 1110 0000 0000 0000 0000	2 Class Bs
/16	255.225.0.0	1111 1111 1111 1111 0000 0000 0000 0000	1 Class B or 256 Class Cs
/17	255.255.128.0	1111 1111 1111 1111 1000 0000 0000 0000	128 Class Cs
/18	255.255.192.0	1111 1111 1111 1111 1100 0000 0000 0000	64 Class Cs
/19	255.255.224.0	1111 1111 1111 1111 1110 0000 0000 0000	32 Class Cs
/20	255.255.240.0	1111 1111 1111 1111 1111 0000 0000 0000	16 Class Cs
/21	255.255.248.0	1111 1111 1111 1111 1111 1000 0000 0000	8 Class Cs
/22	255.255.252.0	1111 1111 1111 1111 1111 1100 0000 0000	4 Class Cs
/23	255.255.254.0	1111 1111 1111 1111 1111 1110 0000 0000	2 Class Cs
/24	255.255.225.0	1111 1111 1111 1111 1111 1111 0000 0000	1 Class C

CIDR also allows you to assign network prefixes of *arbitrary* lengths, as opposed to the now obsolete class system which assigned prefixes as even multiples of an octet.

For example, you can assign a single routing table supernet entry of 195.215.16/21 to represent 8 separate Class C network numbers: 195.215.16.0 through 195.215.23.0.

## Creating a supernet address

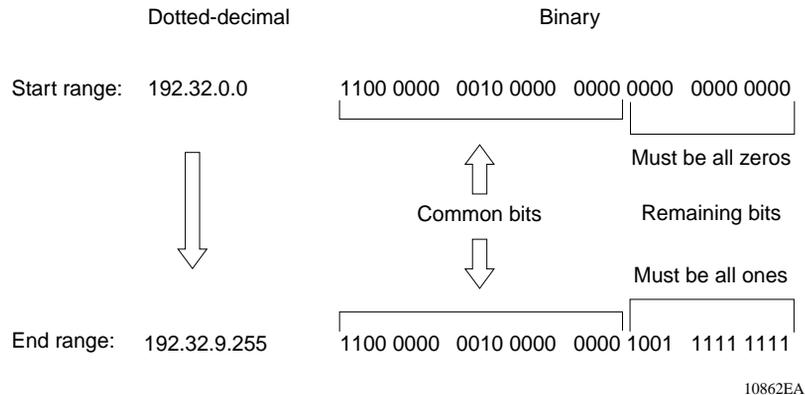
You can create a supernet address that covers any address range.

For example, to create a supernet address that covers an address range of 192.32.0.0 to 192.32.9.255, use the following procedure:

- 1 Convert the starting and ending address range from dotted-decimal notation to binary notation (Figure 3):

- a Locate the *common* bits in both ranges, and then ensure that the *remaining* bits in the start range are zeros, and the remaining bits in the end range are all ones.
- b If the remaining bits in the end range are *not* all ones (as shown in [Figure 3](#)), you must recalculate to find the IP prefix that has only ones in the remaining bits in the end range (Step 2).

**Figure 3** Converting to binary notation



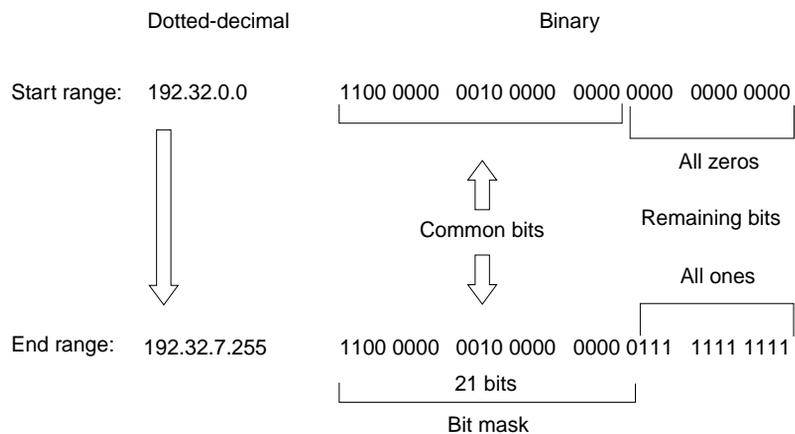
**2** Recalculate to find a network prefix that has all ones in the remaining end range bits ([Figure 4](#)):

- a In [Figure 4](#), 192.32.7.255 is the closest IP prefix that matches the start range's common bits.

The 21 bits that match the common bits form the *Prefix length*. Note that the Prefix length is the number of binary bits that form the explicit *mask* (in dotted-decimal notation) for this IP prefix.

- b As shown in [Figure 4](#), the resulting *first* aggregate 192.32.0.0/21 represents all of the IP prefixes from 192.32.0.0 to 192.32.7.255.
- c The remaining aggregate must be formed from 192.32.8.0 to the end range, 192.32.9.255 (see [Step 3](#)).

**Figure 4** Combining into two aggregate ranges (1 of 2)



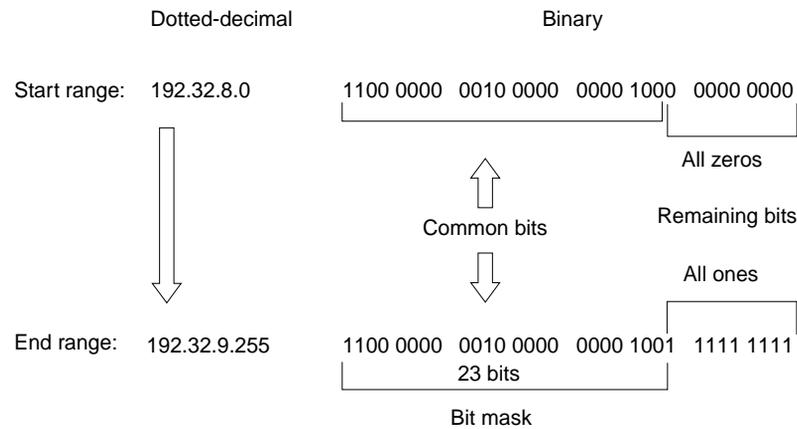
$$\text{Resulting First aggregate} = 192.32.0.0 + \frac{255.255.248.0}{\text{Explicit mask}} = 192.32.0.0/21$$

Prefix length

10863EA

- 3** [Figure 5](#) shows the results after forming the remaining aggregate from 192.32.9.0 to the end range, 192.32.9.255.

As shown in [Figure 5](#), the resulting aggregate 192.32.8.0/23 represents all of the IP prefixes from 192.32.8.0 to 192.32.9.255.

**Figure 5** Combining into two aggregate ranges (2 of 2)

$$\text{Resulting last aggregate} = 192.32.8.0 + \frac{255.255.254.0}{\text{Explicit mask}} = 192.32.8.0/23$$

Prefix length

10864EA

- 4** The final result of calculating the supernet address that ranges from 192.32.0.0 to 192.32.9.255 is as follows:
- a** 192.32.0.0 (with mask) 255.255.248.0 = 192.32.0.0/21
  - b** 192.32.8.0 (with mask) 255.255.254.0 = 192.32.8.0/23

For more information about CIDR, see [“CIDR and aggregate addressing example”](#) on page 209.

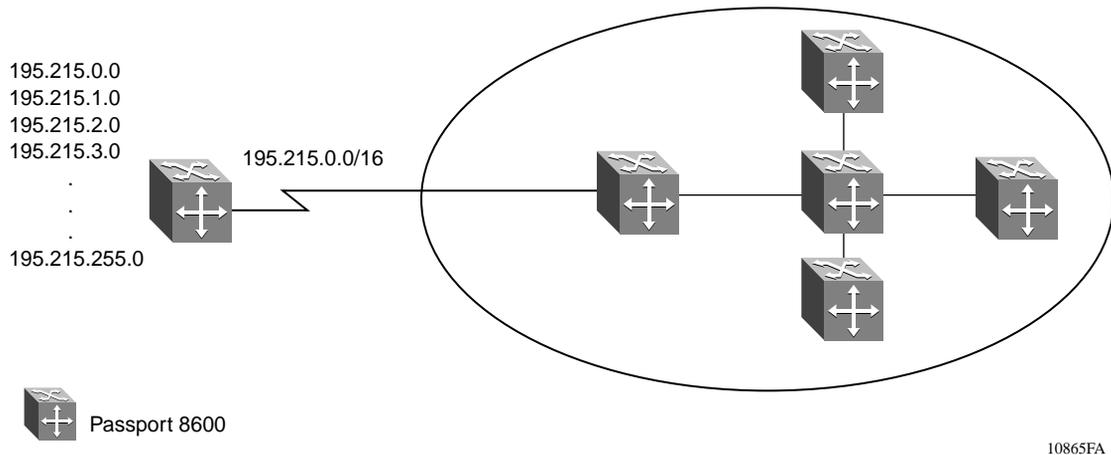
To see how CIDR lets you reduce the size of your routing tables by creating aggregate routes, see [“Aggregate routes,”](#) next.

## Aggregate routes

Eliminating the idea of network classes provides an easy method to aggregate routes. Rather than advertise a separate route for each destination network in a supernet, BGP uses a supernet address to advertise a single route (called an *aggregate* route) that represents all the destinations. CIDR also reduces the size of the routing tables used to store advertised IP routes.

Figure 6 shows an example of route aggregation using CIDR. In this example, a single supernet address 195.215.0.0/16 is used to advertise 256 separate Class C network numbers 195.215.0.0 through 195.215.255.0.

**Figure 6** Aggregating routes with CIDR



## Confederations

A BGP router configured for IBGP establishes a peer-to-peer session with every other IBGP speaker in the AS. In an AS with a large number of IBGP speakers, this *full-mesh topology* can result in high bandwidth and maintenance costs.

As shown in the following example, a full-mesh topology for an AS with 50 IBGP speakers requires 1225 internal peer-to-peer connections:

Example:

$$n \times (n-1)/2 = n \text{ IBGP sessions}$$

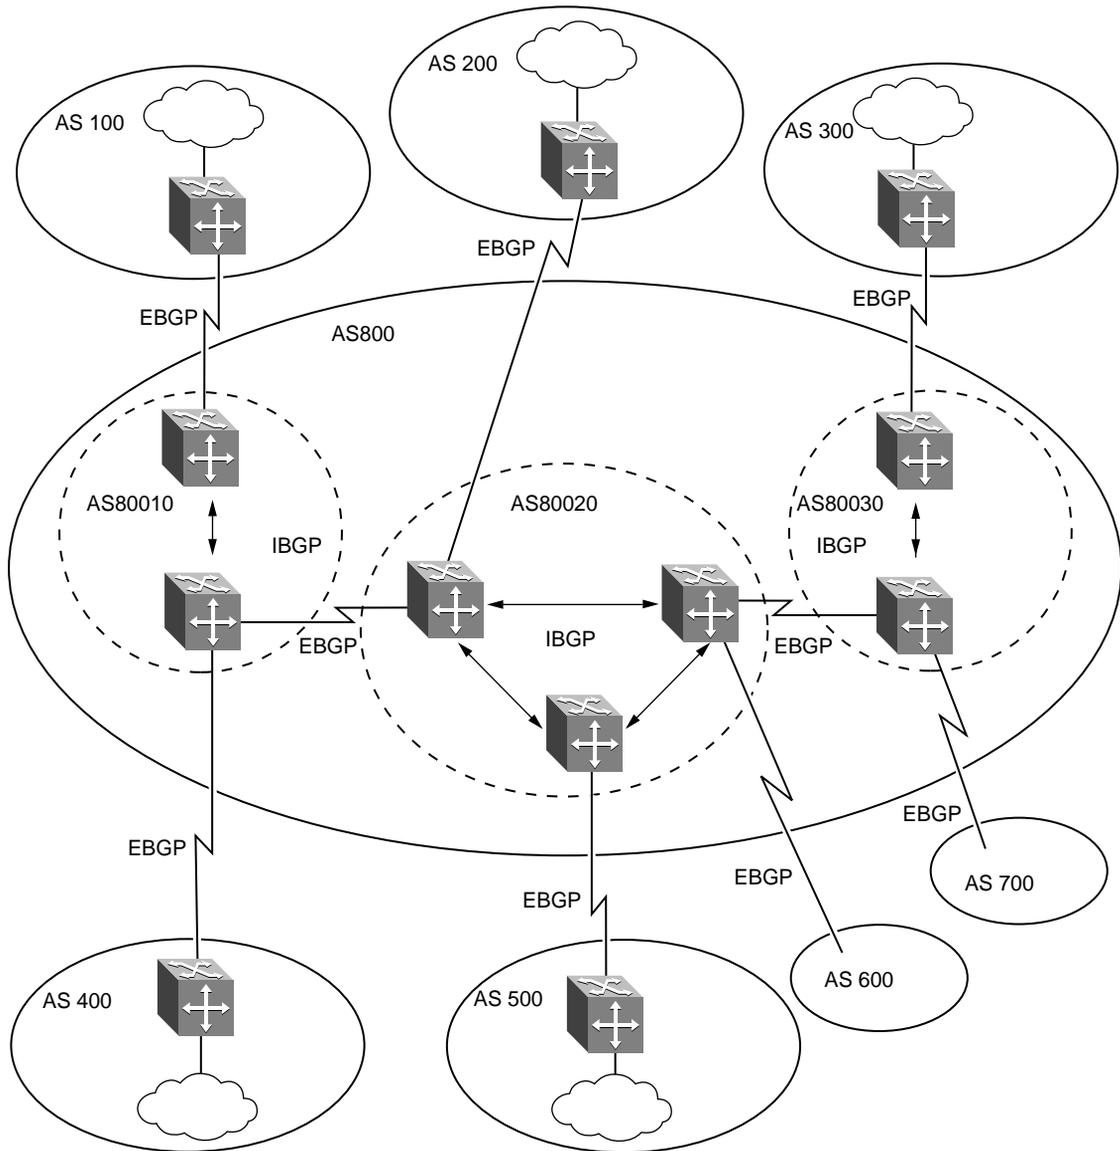
where:

$$50 \times (50-1)/2 = 1225 \text{ number of unique IBGP sessions}$$

You can reduce the high bandwidth and maintenance costs associated with a large full-mesh topology by dividing the AS into multiple smaller ASs (sub-ASs), and then grouping them into a single *confederation* [Figure 7 on page 33](#).

As shown in [Figure 7](#), each sub-AS is fully meshed within itself and has EBGP sessions with other sub-ASs that are in the same confederation.

**Figure 7** Confederations



10866FA

Although the peers that are located in different ASs have EBGP sessions with the various sub-AS peers, they preserve the next-hop, MED, and local preference information and exchange routing updates as if they were IBGP peers. This method allows all of the ASs to retain a single interior gateway protocol (IGP). When the confederation is assigned its own confederation identifier, the group of sub-ASs appear as a single AS (with the confederation identifier as the AS number).

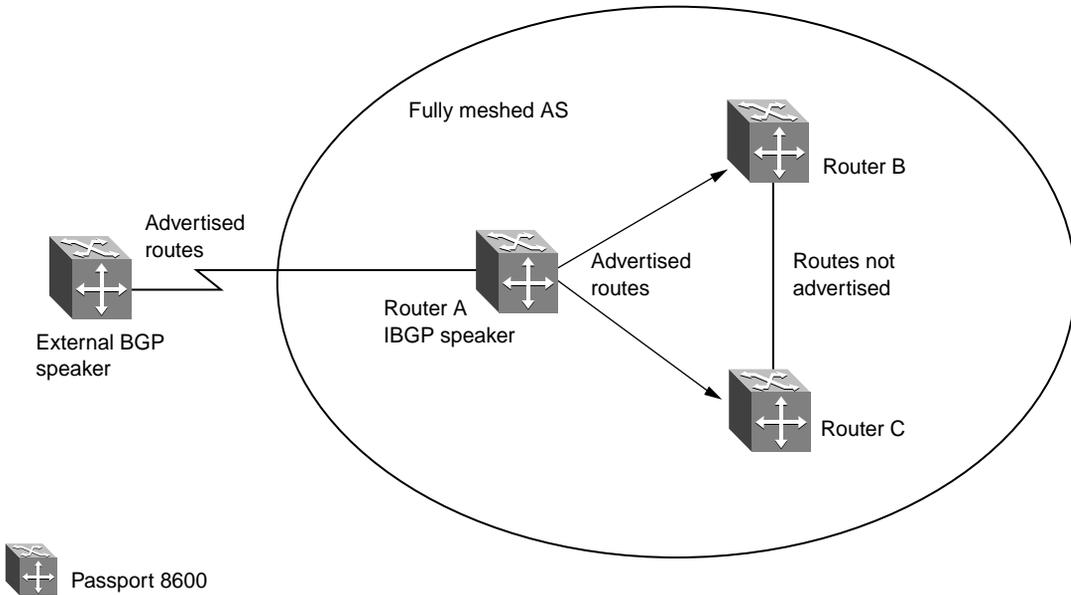
For more information about confederations, see [“BGP confederations” on page 264](#).

## IBGP route reflection

Another way to reduce the IBGP mesh that is inherent in an AS with a large number of IBGP speakers, is to configure a *route reflector (RR)*. Using this method, when an IBGP speaker needs to communicate with other BGP speakers in the AS, the speaker establishes a single peer-to-peer *RR client* session with the IBGP route reflector.

For example, [Figure 8](#), shows a simple IBGP configuration with three IBGP speakers (Routers A, B, and C). Without route reflectors configured, when Router A receives an advertised route from an external neighbor, it must advertise the route to Routers B and C.

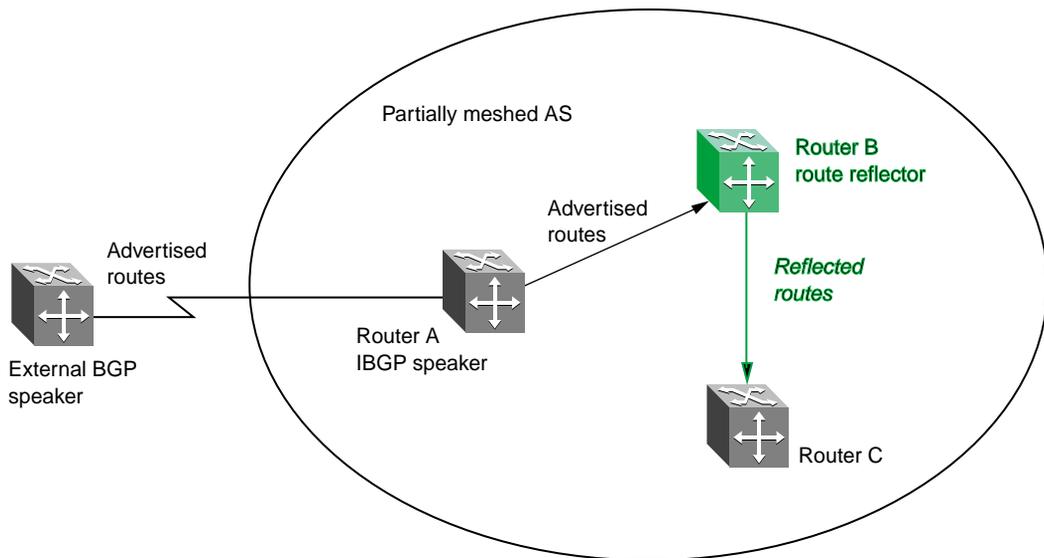
Routers B and C do not readvertise the IBGP learned routes to other IBGP speakers (BGP does not allow routers to pass routes learned from internal neighbors on to other internal neighbors, thus avoiding routing information loops).

**Figure 8** Fully meshed AS with IBGP speakers

As shown in [Figure 9](#), when you configure an internal BGP peer (Router B) as a route reflector, all of the IBGP speakers are not required to be fully meshed. In this case, the assigned route reflector assumes the responsibility for passing IBGP learned routes to a set of IBGP neighbors.

When Router B (the route reflector) in [Figure 9](#) receives routes advertised from Router A (the IBGP speaker) it advertises them to router C. Conversely, when the route reflector receives routes from internal peers, it advertises those routes to Router A. IBGP sessions are not required between Routers A and C.

**Figure 9** AS with route reflector

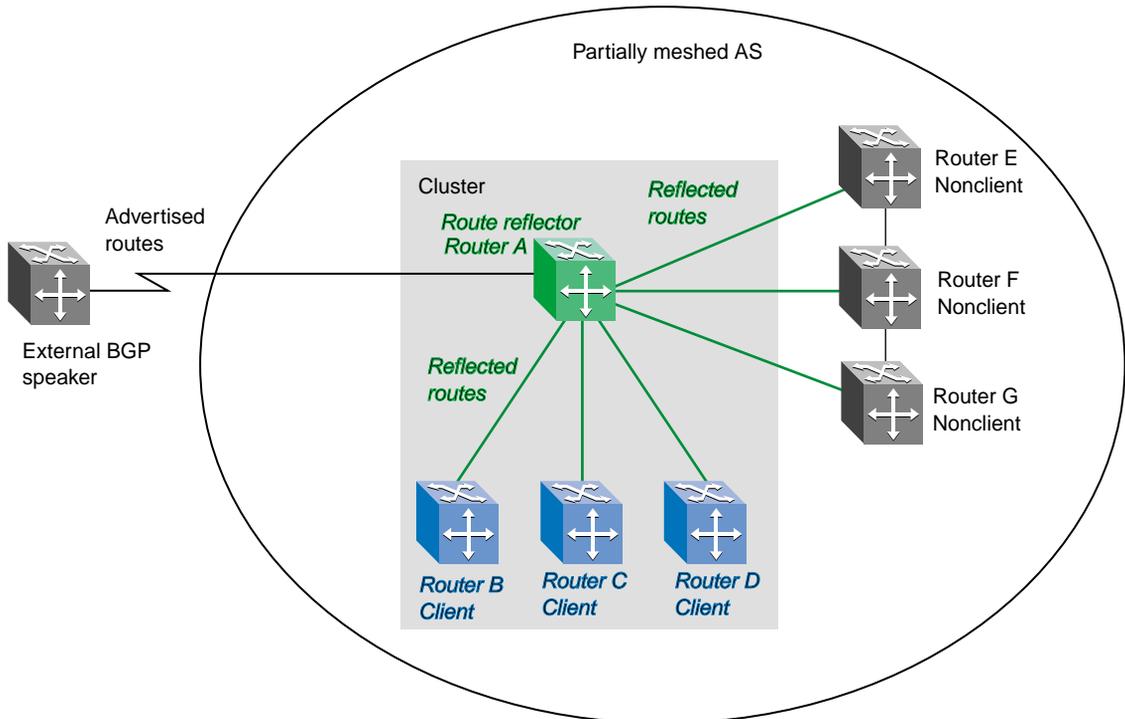


10868FA

Route reflectors separate internal peers into two groups: client peers and nonclient peers. The route reflector and its clients form a *cluster*. The client peers in the cluster are not required to be fully meshed, and do not communicate with IBGP speakers outside their cluster. Nonclient peers must be fully meshed with each other.

This concept is shown in [Figure 10](#), where Router A is shown as the route reflector in a cluster with client Routers B, C, and D. Routers E, F, and G are fully meshed, nonclient routers.

**Figure 10** Route reflector with client and nonclient peers



 Passport 8600

10869FA

For more information about route reflection, see [“Route reflectors”](#) on page 269.

## BGP updates

BGP uses update messages to communicate information between two BGP speakers. The update message can be used to advertise a single feasible route to a peer, or to withdraw multiple unfeasible routes from service.

Figure 11 shows the format of an update message.

**Figure 11** Update message format

Withdrawn Routes Length (2 octets)	
Withdrawn Routes (variable length)	
Total Path Attributes Length (2 octets)	
Path Attributes (variable length)	
Network Layer Reachability Information (variable length)	

This section describes how BGP uses the update message fields to communicate information between BGP speakers.

The following topics are included:

- [“Withdrawn Routes Length field,”](#) next
- [“Withdrawn Routes field”](#) on page 39
- [“Total Path Attributes Length field”](#) on page 40
- [“Path Attributes field”](#) on page 40
- [“Network Layer Reachability Information field”](#) on page 43

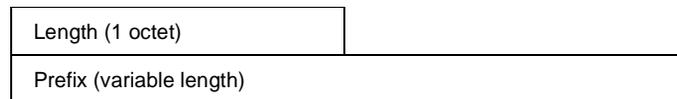
## Withdrawn Routes Length field

The Withdrawn Routes Length field (referred to in RFC 1771 as the Unfeasible Routes Length field) indicates the total length of the Withdrawn Routes field, in octets. The value of the Withdrawn Routes Length field is used to allow the length of the Network Layer Reachability Information field to be calculated. For example, a value of 0 indicates that no routes are being withdrawn from service, and that the Withdrawn Routes field is not present in this Update message. For more information about the Network Layer Reachability Information field, see [“Network Layer Reachability Information field” on page 43](#).

## Withdrawn Routes field

The Withdrawn Routes field is a variable length field that contains a list of IP prefixes for routes that are being withdrawn from service. [Figure 12](#) shows the format of an IP Prefix.

**Figure 12** IP Prefix format



### Length field

The Length field indicates the number of bits in the prefix (also called the network mask).

For example, 195.215.0.0/16 is equivalent to 195.215.0.0 255.255.0.0 (where: the network mask 255.255.0.0 is represented by the /16 which indicates the number of bits in the Length field).

### Prefix field

The Prefix field contains the IP address prefix itself, followed by enough trailing bits to make the length of the whole field an integer multiple of 8 bits (1 octet).

**Note:** The value of trailing bits is irrelevant.

For more information about IP Prefixes and how it is used in supernetting, see [“CIDR and aggregate addresses” on page 27](#).

## Total Path Attributes Length field

The Total Path Attributes Length field indicates the total length of the Path Attributes field in octets.

The value of the Total Path Attributes Length field is used to allow the length of the Network Layer Reachability Information field to be calculated. For example, a value of 0 indicates that no Network Layer Reachability Information field is present in this update message.

For more information about the Network Layer Reachability Information field, see [“Network Layer Reachability Information field” on page 43](#).

## Path Attributes field

The Path Attributes field is a variable length sequence of path attributes that is present in every BGP Update. The path attributes contain BGP attributes that are associated with the prefixes in the Network Layer Reachability Information field (see [“Network Layer Reachability Information field” on page 43](#)).

For example, the attribute values allow you to specify the prefixes that can be exchanged in the BGP session, which of the multiple paths of a specified prefix to use, and so on.

The attributes carry the following information about the associated prefixes:

- Path origin
- The AS paths through which the prefix has been advertised
- Metrics that display degrees of preference for this prefix

[Figure 13](#) shows the encoding used with the Path Attribute field. The fields are described in the sections that follow.

**Figure 13** Path attribute encoding

Attribute Type (2 octet)	
Attribute Length (1 or 2 octets)	
Attribute Value (variable length)	

## Attribute Type field

As shown in [Figure 14](#), The Attribute Type field is a two-octet field that comprises two sub-fields: the Attribute Flags field, and the Attribute Type Code field.

**Figure 14** Attribute Type fields

Attribute Flags (1 octet)	Attribute Type Code (1 octet)
---------------------------	-------------------------------

### *Attribute Flags*

The Attribute Flags field is a bit string that contains four binary values that describe the attribute, and four bits that are unused. The bit descriptions (from the high-order bit to the low-order bit) are:

- The high-order bit (bit 0) is the Optional bit. When set (1) the attribute is *optional*. When this bit is clear (0), the attribute is *well-known*. Well-known attributes must be recognized by all BGP implementations and, when appropriate, passed on to BGP peers. Optional attributes are not required to be present in all BGP implementations.
- The second high-order bit (bit 1) is the Transitive bit. For well-known attributes, this bit must be set to 1. For optional attributes, it defines whether the attribute is transitive (when set to 1) or non-transitive (when set to 0).
- The third high-order bit (bit 2) is the Partial bit. It defines whether the information contained in the optional transitive attribute is partial (when set to 1) or complete (when set to 0). For well-known attributes and for optional non-transitive attributes the Partial bit must be set to 0.
- The fourth high-order bit (bit 3) is the Extended Length bit. It defines whether the Attribute Length is one octet (when set to 0) or two octets (when set to 1). Extended Length may be used only if the length of the attribute value is greater than 255 octets.
  - If the Extended Length bit of the Attribute Flags octet is set to 0, the third octet of the Path Attribute contains the length of the attribute data in octets.
  - If the Extended Length bit of the Attribute Flags octet is set to 1, then the third and the fourth octets of the path attribute contain the length of the attribute data in octets.
- The lower-order four bits of the Attribute Flags octet are unused. They must be zero (and must be ignored when received).

## Attribute Type Code

The Attribute Type Code field contains the attribute type code, as defined by the Internet Assigned Numbers Authority (IANA). The Attribute Type Code field value is used to uniquely identify the attribute from all others. The remaining octets of the Path Attribute represent the attribute value and are interpreted according to the Attribute Flags and the Attribute Type Code fields (see “Attribute Flags” on page 41 and “Attribute Type Code” on page 42).

The supported Attribute Type Codes are shown in [Table 2](#):

**Table 2** BGP mandatory path attributes

Attribute	Type code	Description
Origin	1	Defines the origin of the path information: <ul style="list-style-type: none"> <li>Value = 0 --- IGP (the path is valid all the way to the IGP of the originating AS)</li> <li>Value = 1--- EGP (the path was advertised using an EGP by the last AS in the AS path)</li> <li>Value = 2--- Incomplete (the path is valid only to the last AS in the AS path)</li> </ul>
AS path	2	Contains a list of the ASs that must be traversed to reach the given destinations. Each AS path segment is represented as follows: <ul style="list-style-type: none"> <li>Path segment type</li> <li>Path segment length</li> <li>Path segment value</li> </ul>
Next hop	3	Specifies the IP address of the border router to use as a next hop for the advertised destinations (destinations listed in the NLRI field of the Update message).
Multiexit discriminator	4	This attribute is used on external (internal-AS) links to discriminate among multiple exit or entry points to the same neighboring AS.
Local preference	5	Allows AS border routers to indicate the preference they assigned to a chosen route when advertising it to IBGP peers
Atomic aggregate	6	Ensures that certain network layer reachability information (NLRI) is not deaggregated
Aggregator	7	Identifies which AS performed the most recent route aggregation. This attribute contains the last AS number that formed the aggregate route followed by the IP address of the BGP speaker that formed the aggregate route.

## Attribute Length field

The Attribute Length field can be one or two octets in length, depending on the value of the Extended Length field in the Attributes Flag field (see [“Attribute Flags” on page 41](#)).

This field is used to indicate the length of the Attribute Value field (see [“Attribute Value field,”](#) next).

## Attribute Value field

The Attribute Value field contains the actual value of the specific attribute and is implemented according to the values in the Attribute Flags and the Attribute Type Code fields (see [“Attribute Flags” on page 41](#) and [“Attribute Type Code” on page 42](#)).

## Network Layer Reachability Information field

The Network Layer Reachability Information field is a variable length field that contains a list of prefixes. The number of prefixes in the list is limited only by the packet size that can be sent between BGP speakers. For details about the format of an IP Prefix, see [Figure 12 on page 39](#)).

For more information about BGP updates, see [“BGP path attributes” on page 236](#).

## Equal-cost multipath

Equal-cost multipath (ECMP) support allows a BGP speaker to perform route or traffic balancing within an AS by using multiple equal-cost routes submitted to the routing table by OSPF, RIP, or static routes.

For more information about ECMP, see [“EBGP multihop and EBGP load balance” on page 212](#).

## TCP MD5 message authentication

You can authenticate BGP messages by TCP MD5 signatures, in compliance with RFC 2385, “Protection of BGP Sessions via the TCP MD5 Signature Option.” When you enable BGP authentication, the BGP speaker verifies that the BGP messages it receives from its peers are actually from a peer and not from a third party masquerading as a peer.

BGP-4 TCP MD5 message authentication provides the following features:

- A TCP MD5 signature can exist for BGP peers.
- You can configure authentication and secret keys on a per-peer basis. Peers configured with common secret keys can authenticate each other and exchange routing information.
- Your configurations can concurrently have some BGP peers configured with authentication enabled and other BGP peers with authentication disabled.
- The secret keys are always stored encrypted.

When you enable BGP-4 TCP MD5 authentication, the router computes a Message Digest 5 (MD5) signature for each TCP packet, based on the TCP packet and a per-peer secret key. The router adds this MD5 signature to the TCP packet containing a BGP message and sends it with the packet, but it does not send the secret key.

The receiver of the TCP packet also knows the secret key and can verify the MD5 signature. A third party trying to masquerade as the sender, however, cannot generate an authentic signature because it does not know the secret key.

The per-peer secret keys provide the security. If the keys are compromised, then the authentication itself is compromised. To prevent this, the secret keys are stored in encrypted form in the configuration files.



**Note:** In the CLI commands, the term “Password” refers to “Secret Key.”

---

## Entering and Storing MD5 Authentication Keys

By default, BGP TCP MD5 authentication is set to disabled.

To enable BGP TCP MD5 authentication:

### ***With Device Manager:***

This feature currently cannot be configured using Device Manager. You can configure this feature using the CLI (see the next section).

### ***With the CLI:***

Use the following command (see [“Configuring BGP peers or peer groups” on page 129](#)):

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name> password  
<password> MD5-authentication <enable|disable>
```

where:

- *nbr\_ipaddr|peer-group-name* is the peer’s IP address or the peer’s group name.
- *password* is the secret key, which is a string length with a range 0 to 1536 alphanumeric characters.

The password (secret key) is encrypted and stored in the BGP configuration files.

To communicate, both peers *must* be configured with the same password (peers that are configured with common passwords only can authenticate each other and exchange routing information).

When an MD5 password (or secret key) is configured for a BGP peer that has MD5 Authentication set to enabled, BGP reads the encrypted password from the BGP configuration files and passes the unencrypted authentication key to TCP.

## Generation of MD5 signatures on ingress BGP TCP packets

BGP peers calculate MD5 signatures in BGP messages on the following elements:

- TCP pseudo-header
- TCP header, excluding options
- TCP segment data
- TCP MD5 authentication key

If TCP receives an MD5 authentication key, it reduces its maximum segment size (MSS) by 18 octets, the length of the TCP MD5 option. It also adds an MD5 signature to each transmitted packet. The peer inserts the resulting 16-byte MD5 signature into the following TCP options: kind=19, length=18.

## Verification of MD5 signatures on egress BGP TCP packets

Upon receiving a packet, TCP performs three tests ([Table 3](#)).

- If a packet passes a test, it proceeds to the next test. When a packet has passed all three tests, TCP accepts the packet and sends it to BGP.
- If a packet fails a test, TCP logs an event, increments the count of TCP connection errors (*wfTcpConnMd5Errors*), and discards the packet. The TCP connection remains open.

[Table 3](#) lists the tests and the event message that TCP logs if a test fails.

**Table 3** MD5 Signature Verification Rules on BGP TCP Packets

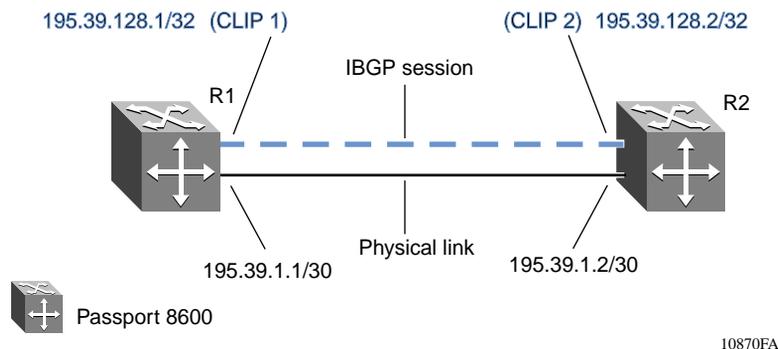
Condition Tested	Action on Success	Failure Event Message
Is the connection configured for MD5 authentication?	Verify that the packet contains a kind=19 option.	TCP MD5 No Signature
Is MD5 authentication enabled for this TCP connection?	TCP computes the expected MD5 signature.	TCP MD5 Authentication Disabled
Does the computed MD5 signature match the received MD5 signature?	TCP sends the packet to BGP.	TCP MD5 Invalid Signature

For more information about MD5, see [“MD5 authentication” on page 230](#).

## Circuitless IP

Circuitless IP (CLIP) is a virtual (or loop back) interface that is not associated with any *physical* port. You can use the CLIP interface to provide uninterrupted connectivity to your switch *as long as there is an actual path to reach the device*. For example, as shown in [Figure 15](#), a physical point-to-point link exists between R1 and R2 along with the associated addresses (195.39.1.1/30 and 195.39.1.2/30). Note also that an IBGP session exists between two additional addresses 195.39.128.1/32 (CLIP 1) and 195.39.128.2/32 (CLIP 2).

**Figure 15** Routers with I-BGP connections



The CLIP interface is treated as any other IP interface. The network associated with the CLIP is treated as a local network attached to the device. This route always exists and the circuit is always up because there is no physical attachment.

Routes are advertised to other routers in the domain either as external routes using the route-redistribution process or when you enable OSPF in a passive mode to advertise an OSPF internal route. You can configure the OSPF protocol only on the CLIP interface. When you create a CLIP interface, the system software programs a local route with the CPU as destID. All packets that are destined to the CLIP interface address are processed by the CPU. Any other packets with destination addresses associated with this network (but not to the interface address) are treated as if they are from any unknown host.

For more information about circuitless IP, see [“Basic BGP example”](#) on page 196.

For information about configuring CLIP using Device Manager, see [“Configuring Circuitless IP”](#) on page 103. For information about configuring CLIP using the CLI, see [“Circuitless IP”](#) on page 173.



---

## Chapter 2

# Configuration considerations and limitations

---

This chapter describes configuration limitations and interoperability issues that you should consider when configuring BGP.

For information about BGP concepts and terminology, see [Chapter 1, “BGP concepts,”](#) on page 21.

This chapter includes the following topics:

Topic	Page
<a href="#">BGP implementation notes</a>	50
<a href="#">Configuration guidelines</a>	51
<a href="#">BGP Neighbor Maximum Prefix</a>	52
<a href="#">BGP-4 support</a>	53
<a href="#">RIP redistribution</a>	53
<a href="#">BGP/OSPF interaction</a>	53
<a href="#">MD5 authentication</a>	54

## BGP implementation notes

The following guidelines are crucial to successful BGP configuration.



**Caution:** If you do not follow these guidelines, BGP either will not work efficiently or will become disabled on the interfaces involved.

---

- BGP will not operate with an IP router in nonforwarding (host-only) mode. Make sure that the routers you want BGP to operate with are in forwarding mode.
- If you are using BGP for a multihomed AS (one that contains more than a single exit point), Nortel Networks\* recommends that you use OSPF for your IGP and BGP for your sole exterior gateway protocol, or use intra-AS IBGP routing.
- If OSPF is the IGP, use the default OSPF tag construction. Using EGP or modifying the OSPF tags makes network administration and proper configuration of BGP path attributes difficult.
- For routers that support both BGP and OSPF, the OSPF router ID and the BGP identifier must be set to the same IP address. The BGP router ID automatically uses the OSPF router ID.
- In configurations where BGP speakers reside on routers that have multiple network connections over multiple IP interfaces (the typical case for IBGP speakers), consider using the address of the router's circuitless (virtual) IP interface as the local peer address. In this way, you ensure that BGP is reachable as long as there is an active circuit on the router.
- By default, BGP speakers do not advertise or inject routes into its IGP. You must configure route policies to enable route advertisement.
- Coordinate routing policies among all BGP speakers within an AS so that every BGP border router within an AS constructs the same path attributes for an external path.
- Configure accept and announce policies on all IBGP connections to accept and propagate all routes. Make consistent routing policy decisions on external BGP connections.

## Configuration guidelines

When configuring BGP on the Passport 8600, you must configure the following minimum parameters:

- Router ID
- Local AS Number
- Enable BGP Globally
- BGP Neighbor Peer Session: remote IP addresses
- Enable BGP peer
- When running both BGP and OSPF, the OSPF and BGP Router ID must be the same.
- The Router ID must be a valid IP address of an IP interface on the router or a circuitless IP address. This IP address is used in BGP Update messages
- By default, the BGP Router ID automatically uses the OSPF Router ID.

There is no option to configure the BGP Router ID. If you configure BGP before you have configured the OSPF Router ID, you must first disable BGP, and then enable BGP globally.

- BGP Policies can be added to the BGP peer configuration to influence route decisions (see [“Policies” on page 202](#)).
- Once the Passport 8600 is configured for BGP, some parameter changes may require to have either the BGP Global state or neighbor admin-state to be disabled/enabled. The CLI prompt will notify you if this is the case.
- BGP policies are dynamically modified. On the global level, the BGP redistribution command has an apply parameter that causes the policy to be applied, at that time.

You can use the following BGP neighbor CLI command to apply policies without bringing down the peer:

```
restart soft-reconfiguration <in> <out>
```

The following are some examples of these commands:

- To disable BGP global, enter:  

```
config ip bgp <enable/disable>
```

- To disable a BGP neighbor, enter:  

```
config ip bgp neighbor <ip address of neighbor>  
admin-state <enable/disable>
```
- To set BGP soft-reconfiguration:  

```
config ip bgp neighbor <ip address of neighbor>  
soft-reconfiguration-in <enable/disable>
```
- To set BGP policy redistribution:  

```
config ip bgp redistribute apply
```

When using Device Manager, use the following commands:

- To disable BGP globally, enter:  

```
IP_Routing>BGP>AdminStatus <enable/disable>
```
- To disable a BGP neighbor, enter:  

```
IP_Routing>BGP>AdminStatus>Peers <IP address of peer> <enable/  
disable>
```
- To set BGP soft-reconfiguration:  

```
IP_Routing>BGP>AdminStatus>Peers Info>SoftReconfiguration <in/out>  
IP_Routing>BGP>AdminStatus>Peers Info>SoftReconfigurationIn <enable/  
disable>
```

## BGP Neighbor Maximum Prefix

By default, the Maximum Prefix parameter value is set to limit 12,000 network layer reachability information (NLRI) messages per neighbor (the Maximum Prefix parameter limits the number of routes that the Passport 8600 can accept).

This value prevents large numbers of BGP routes from flooding the network in the event of a misconfiguration. You can configure the Maximum Prefix limit to any value, including 0, which allows unlimited routes. When you configure the Maximum Prefix value, consider the maximum number of active routes that your equipment configuration can support.

## BGP-4 support

Nortel Networks\* implementation of BGP supports BGP-4 as described in RFC 1771.

BGP-4 has no concept of address classes. Each network listed in the network layer reachability information (NLRI) portion of an update message contains a prefix length field, which describes the length of the mask associated with the network.

The prefix length field allows for both supernet and subnet advertisement. The supernet advertisement is what makes classless interdomain routing (CIDR) possible. See [“CIDR and aggregate addresses” on page 27](#).

In addition, BGP-4 supports BGP confederations and TCP MD5 message authentication.

## RIP redistribution

Nortel Networks\* implementation of BGP supports redistribution of RIP routes; however, for this current release, the RIP protocol does not have the ability to redistribute BGP routes.

## BGP/OSPF interaction

RFC 1745 defines the interaction between BGP and OSPF when OSPF is the IGP within an autonomous system. For routers running both protocols, the OSPF router ID and the BGP ID must be the same IP address. A BGP route policy must be configured to allow BGP advertisement of OSPF routes.

Interaction between BGP-4 and OSPF includes the ability to advertise supernets to support classless interdomain routing (CIDR). BGP-4 allows interdomain supernet advertisements; OSPF can carry supernet advertisements within a routing domain.

## **MD5 authentication**

The MD5 authentication feature is currently not configurable using Device Manager with release 3.3 software. You can configure this feature using the CLI commands.

---

## Chapter 3

# Using Device Manager to configure BGP

---

This chapter describes how to configure BGP using Device Manager.



**Note:** The screen examples in this chapter have hypertext links that lead you to relevant information about the screen. The links are accessible when viewed in online format (such as PDF). To access the links, drag your cursor across the screen and click on the item when the cursor changes to a pointing hand.

- For BGP concept details, see [Chapter 1, “BGP concepts,”](#) on page 21.
- For information about limitations and interoperability issues, see [Chapter 2, “Configuration considerations and limitations,”](#) on page 49.
- For configuration examples, including the required CLI commands, see [Chapter 5, “Configuration examples,”](#) on page 195.

This chapter includes the following topics:

Topic	Page
<a href="#">Configuring general parameters</a>	56
<a href="#">Displaying global BGP statistics</a>	61
<a href="#">Configuring aggregate addresses</a>	65
<a href="#">Configuring allowed network addresses</a>	68
<a href="#">Configuring and displaying peer information</a>	69
<a href="#">Displaying route information</a>	93
<a href="#">BGP community attributes</a>	97
<a href="#">Displaying dampened routes information</a>	100
<a href="#">Configuring redistribute entries</a>	101
<a href="#">Configuring Circuitless IP</a>	103

## Configuring general parameters

To configure general BGP parameters:

- From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (Figure 16).

The Generals tab information is grouped into three sections: General, Confederation, and Reflection.

**Figure 16** BGP dialog box—Generals tab

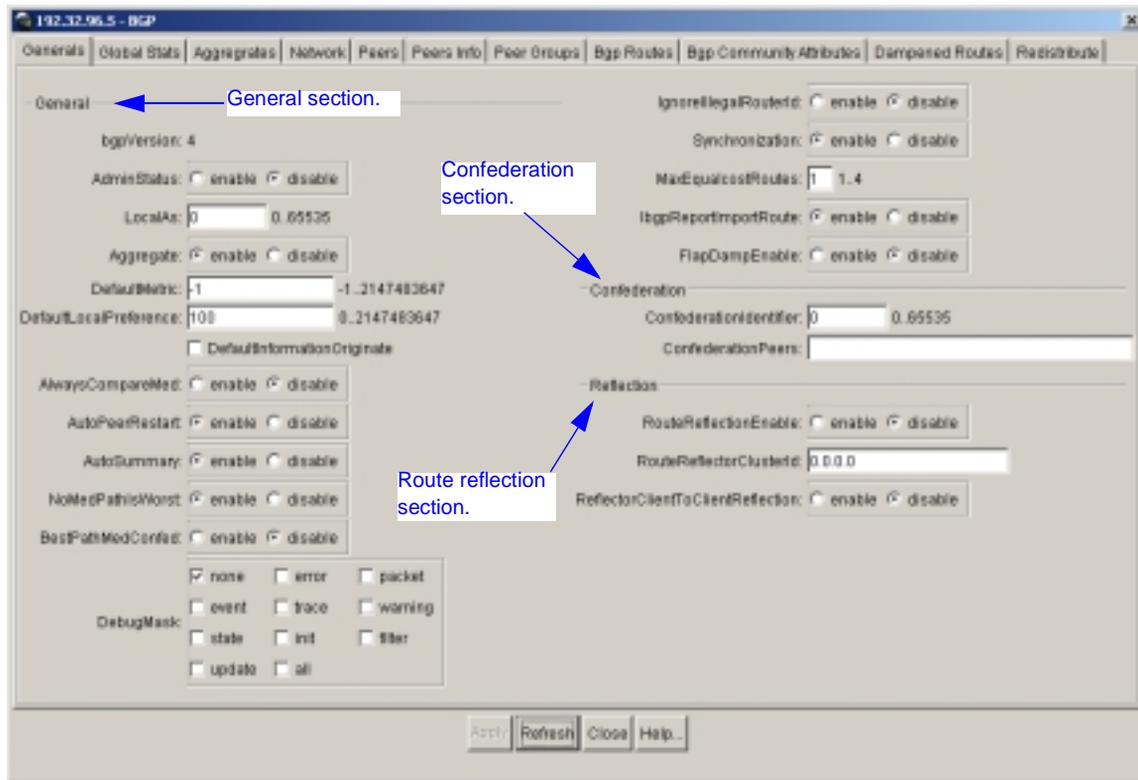


Table 4 describes the BGP dialog box—Generals tab fields.

**Table 4** BGP dialog box—Generals tab fields

Field	Description
AdminStatus	<p>Enables or disables BGP on the current system. The default value is disable.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul> <p><b>Note:</b> You cannot enable AdminStatus until you change the LocalAS value to any value other than 0.</p>
LocalAs	<p>Sets a local autonomous system number on the current system. The default value is 0.</p> <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 65535.</li> </ul> <p><b>Note:</b> You cannot enable AdminStatus until you change the LocalAS default value to any value other than 0. You cannot change configured values when AdminStatus is set to enable.</p>
Aggregate	<p>Enables or disables the aggregation feature on this interface. The default value is enable.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul> <p><b>Note:</b> You cannot change values when AdminStatus is set to enable.</p>
DefaultMetric	<p>Sets a value that is sent to a BGP neighbor to determine the cost of a route a neighbor is using. This option must be used in conjunction with the redistribute router configuration command to allow the current routing protocol to use the same metric value for all redistributed routes.</p> <p>The default value is -1.</p> <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 2147483647</li> </ul> <p><b>Note:</b> A default metric value helps solve the problems associated with redistributing routes that have incompatible metrics. For example, whenever metrics do not convert, using a default metric provides a reasonable substitute and allows the redistribution to proceed.</p>
DefaultLocalPreference	<p>Specifies the default value of the local preference attribute. The default value is 100.</p> <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 2147483647</li> </ul> <p><b>Note:</b> You cannot change the default value when AdminStatus is set to enable.</p>

**Table 4** BGP dialog box—Generals tab fields (continued)

Field	Description
DefaultInformationOriginate	Checkbox— When checked (enabled), allows the redistribution of network 0.0.0.0 into BGP. The default value is disable (not checked).
AlwaysCompareMed	<p>Enables or disables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different autonomous systems. A path with a lower MED is preferred over a path with a higher MED. The default value is disable.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul> <p><b>Note:</b> When this option is set to disable (the default value) during the best-path selection process, the MEDs are compared only among paths from the same autonomous system. If you enable this option, the MEDs are compared among paths received from any other autonomous systems.</p>
AutoPeerRestart	<p>Enables or disables the process that automatically restarts a connection to a BGP neighbor. The default value is enable.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>
AutoSummary	<p>When enabled, allows BGP to summarize networks based on class limits (For example, Class A, B, C networks). The default value is enable.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>
NoMedPathIsWorst	<p>When set to enable (the default value), BGP treats an update that is missing a multi-exit discriminator (MED) attribute, as the worst path.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>
BestPathMedConfed	<p>When enabled, allows you to compare multi-exit discriminator (MED) attributes within a confederation. The default value is disable.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>

**Table 4** BGP dialog box—Generals tab fields (continued)

Field	Description
DebugMask	<p>Checkboxes—When checked (enabled) allows you to display specified debug information for BGP global configuration. The default value is none (checked).</p> <p>Mask choices are:</p> <p><i>none</i>: disables all debug messages.</p> <p><i>all</i>: enables all debug messages.</p> <p><i>error</i>: enables display of debug error messages.</p> <p><i>packet</i>: enables display of debug packet messages.</p> <p><i>event</i>: enables display of debug event messages.</p> <p><i>trace</i>: enables display of debug trace messages.</p> <p><i>warning</i>: enables display of debug warning messages.</p> <p><i>state</i>: enables display of debug state transition messages.</p> <p><i>init</i>: enables display of debug initialization messages.</p> <p><i>filter</i>: enables display of debug messages related to filtering.</p> <p><i>update</i>: enables display of debug messages related to sending and receiving updates.</p>
IgnoreIllegalRouterid	<p>When enabled, allows BGP to overlook an illegal router ID. For example, you can set this command to enable or disable the acceptance of a connection from a peer that sends an open message using a router ID of 0 (zero). The default value is disable.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>
Synchronization	<p>Enables or disables the router from accepting routes from BGP peers without waiting for an update from the IGP. The default value is enable.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>
MaxEqualCostRoutes	<p>Sets the maximum number of equal-cost-paths that are available to a BGP router by limiting the number of equal-cost-paths that can be stored in the routing table. The default value is 1.</p> <ul style="list-style-type: none"> <li>Specify an integer value in the range 1 to 4.</li> </ul>

**Table 4** BGP dialog box—Generals tab fields (continued)

Field	Description
IbgpReportImportRoute	Configures BGP to report imported routes to an interior BGP (IBGP) peer. This command also enables or disables reporting of non-BGP imported routes to other IBGP neighbors. The default value is enable. <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>
FlapDampEnable	Enables or disables route suppression for routes that flap on and off. The default value is disable. <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>
ConfederationIdentifier	Specifies a BGP confederation identifier. <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 65535.</li> </ul> <b>Note:</b> You cannot configure this value when AdminStatus is set to enable.
ConfederationPeers	Lists adjoining ASs that are part of the confederation. <ul style="list-style-type: none"> <li>Specify a list of ASs separated by commas (5500,65535,0,10,.....).</li> </ul>
RouteReflectionEnable	Enables or disables the reflection of routes from IBGP neighbors. The default value is disable. <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>
RouteReflectorClusterId	Sets a cluster ID. This option is applicable only if the RouteReflectionEnable value is set to enable, and if multiple route reflectors are in a cluster. <ul style="list-style-type: none"> <li>Specify an IP address that is the cluster ID of the reflector router.</li> </ul>
ReflectorClienToClientReflection	Enables or disables route reflection between two route reflector clients. This option is applicable only if the RouteReflectionEnable value is set to enable. The default value is disable. <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul>

## Displaying global BGP statistics

To display global BGP statistics:

- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (see [Figure 16](#) on page 56).

- 2 Click the Global Stats tab.

The Global Stats tab opens and displays current BGP statistics ([Figure 17](#)).

**Figure 17** BGP dialog box—Global Stats tab

The screenshot shows the BGP dialog box with the Global Stats tab selected. The table displays various BGP statistics with columns for Absolute Value, Cumulative, Average/sec, Minimum/sec, Maximum/sec, and LastVal/sec. A vertical scrollbar on the right side of the table is highlighted with a blue arrow and a text box that says "Use the vertical scroll bar to display all rows." The table data is as follows:

	Absolute Value	Cumulative	Average/sec	Minimum/sec	Maximum/sec	LastVal/sec
Starts	7	0	0	0	0	0
Stops	1	0	0	0	0	0
Opens	6	0	0	0	0	0
Closes	3	0	0	0	0	0
Fails	6	0	0	0	0	0
Fatals	0	0	0	0	0	0
ConnExps	7	3	0.06	0	0.3	0.3
HoldExps	0	0	0	0	0	0
KeepExps	10	2	0.04	0	0.1	0.1
RxOpens	4	0	0	0	0	0
RxKeeps	7	3	0.06	0	0.18	0
RxUpdates	3,348	2,448	46.15	1.64	121.9	121.9
RxNotifs	0	0	0	0	0	0
TxOpens	6	0	0	0	0	0
TxKeeps	14	2	0.04	0	0.1	0.1
TxUpdates	4,187	4,187	79	0	237.4	1.4
TxNotifs	1	0	0	0	0	0
BadEvents	0	0	0	0	0	0
SynxFails	20	18	0.34	0	0.5	0.2
TxEvent	0	0	0	0	0	0
RxECdeHeader	0	0	0	0	0	0
RxECdeOpen	0	0	0	0	0	0
RxECdeUpdate	0	0	0	0	0	0
RxECdeHoldtimer	0	0	0	0	0	0
RxECdeFSM	0	0	0	0	0	0
RxECdeCease	0	0	0	0	0	0
RxHdCodeNoSyn	0	0	0	0	0	0
RxHdCodeInvalidMsgLen	0	0	0	0	0	0
RxHdCodeInvalidMsgType	0	0	0	0	0	0
RxOpCodeBadVar	0	0	0	0	0	0
RxOpCodeBadAs	0	0	0	0	0	0

[Table 5](#) describes the Global BGP statistics parameters.

**Table 5** Global Stats tab descriptions

<b>Field</b>	<b>Description</b>
Starts	Number of times BGP connection started
Stops	Number of times BGP connection stopped
Opens	Number of times BGP connection opened TCP
Closes	Number of times BGP connection closed TCP
Fails	Number of times a TCP attempt failed
Fatals	Number of times TCP crashed due to fatal error
ConnExps	Number of times the TCP retry timer expired
HoldExps	Number of times the hold timer expired
KeepExps	Number of times the keepalive timer expired
RxOpens	Number of Opens received by BGP
RxKeeps	Number of Keepalive messages received by BGP
RxUpdates	Number of Updates received by BGP
RxNotifys	Number of Notifications received by BGP
TxOpens	Number of transmitted by BGP
TxKeeps	Number of Keepalive messages transmitted by BGP
TxUpdates	Number of Updates transmitted by BGP
TxNotifys	Number of Notifications transmitted by BGP
BadEvents	Number of invalid events received by FSM
SyncFails	Number of times the FDB sync failed
TrEvent	Trace event
RxECodeHeader	Total number of Header errors received
RxECodeOpen	Total number of Open errors received
RxECodeUpdate	Total number of Update errors received
RxECodeHoldtimer	Total number of Hold Timer Expired errors received
RxECodeFSM	Total number of FSM errors received
RxECodeCease	Total number of Cease errors received
RxHdrCodeNoSync	Number of Header errors received as: Not Synchronized
RxHdrCodeInvalidMsgLength	Number of Header errors received as: Invalid Msg len

**Table 5** Global Stats tab descriptions (continued)

Field	Description
RxHdrCodeInvalidMsgType	Number of Header errors received as: Invalid Msg type
RxOpCodeBadVer	Number of Open errors received as: Bad version
RxOpCodeBadAs	Number of Open errors received as: Bad AS number
RxOpCodeBadRtID	Number of Open errors received as: Bad BGP Rtr ID
RxOpCodeUnsuppOption	Number of Open errors received as: Unsupported Option
RxOpCodeAuthFail	Number of Open errors received as: Auth Failure
RxOpCodeBadHold	Number of Open errors received as: Bad Hold Value
RxUpdCodeMalformedAttrList	Number of Update errors received as: Malformed Attr List
RxUpdCodeWelknownAttrUnrecog	Number of Update errors received as: Welknown Attr Unrecog
RxUpdCodeWelknownAttrMiss	Number of Update errors received as: Welknown Attr Missing
RxUpdCodeAttrFlagError	Number of Update errors received as: Attr Flag Error
RxUpdCodeAttrLenError	Number of Update errors received as: Attr Len Error
RxUpdCodeBadORIGINAttr	Number of Update errors received as: Bad ORIGIN Attr
RxUpdCodeASRoutingLoop	Number of Update errors received as: AS Routing Loop
RxUpdCodeBadNHAttr	Number of Update errors received as: Bad NEXT-HOP Attr
RxUpdCodeOptionalAttrError	Number of Update errors received as: Optional Attr Error
RxUpdCodeBadNetworkField	Number of Update errors received as: Bad Network Field
RxUpdCodeMalformedASPath	Number of Update errors received as: Malformed AS Path
TxECodeHeader	Total number of Header errors transmitted
TxECodeOpen	Total number of Open errors transmitted
TxECodeUpdate	Total number of Update errors transmitted
TxECodeHoldtimer	Total number of Hold Timer Expired errors transmitted
TxECodeFSM	Total number of FSM errors transmitted

**Table 5** Global Stats tab descriptions (continued)

<b>Field</b>	<b>Description</b>
TxECodeCease	Total number of Cease errors transmitted
TxHdrCodeNoSync	Number of Header errors transmitted as: Not Synchronized
TxHdrCodeInvalidMsgLen	Number of Header errors transmitted as: Invalid Msg len
TxHdrCodeInvalidMsgType	Number of Header errors transmitted as: Invalid Msg type
TxOpCodeBadVer	Number of Open errors transmitted as: Bad version
TxOpCodeBadAs	Number of Open errors transmitted as: Bad AS number
TxOpCodeBadRtrID	Number of Open errors transmitted as: Bad BGP Rtr ID
TxOpCodeUnsuppOption	Number of Open errors transmitted as: Unsupported Option
TxOpCodeAuthFail	Number of Open errors transmitted as: Auth Failure
TxOpCodeBadHold	Number of Open errors transmitted as: Bad Hold Value
TxUpdCodeMalformedAttrList	Number of Update errors transmitted as: Malformed Attr List
TxUpdCodeWelknownAttrUnrecog	Number of Update errors transmitted as: Welknown Attr Unrecog
TxUpdCodeWelknownAttrMiss	Number of Update errors transmitted as: Welknown Attr Missing
TxUpdCodeAttrFlagError	Number of Update errors transmitted as: Attr Flag Error
TxUpdCodeAttrLenError	Number of Update errors transmitted as: Attr Len Error
TxUpdCodeBadORIGINAttr	Number of Update errors transmitted as: Bad ORIGIN Attr
TxUpdCodeASRoutingLoop	Number of Update errors transmitted as: AS Routing Loop
TxUpdCodeBadNHAttr	Number of Update errors transmitted as: Bad NEXT-HOP Attr
TxUpdCodeOptionalAttrError	Number of Update errors transmitted as: Optional Attr Error

**Table 5** Global Stats tab descriptions (continued)

Field	Description
TxUpdCodeBadNetworkField	Number of Update errors transmitted as: Bad Network Field
TxUpdCodeMalformedASPath	Number of Update errors transmitted as: Malformed AS Path

## Configuring aggregate addresses

To configure aggregate addresses:

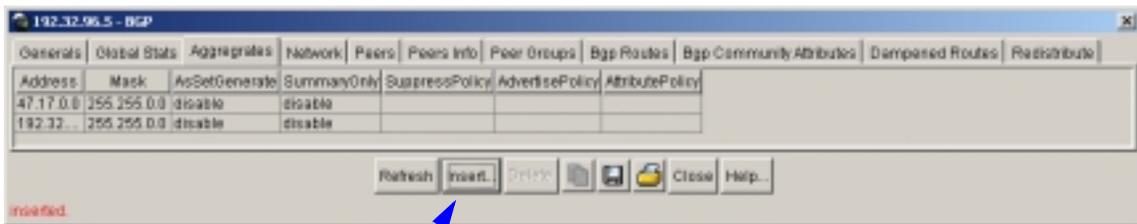
- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (see [Figure 16 on page 56](#)).

- 2 Click the Aggregates tab.

The Aggregates tab opens and displays current aggregate addresses ([Figure 18](#)).

**Figure 18** BGP dialog box—Aggregates tab



Click Insert to open the Aggregates dialog box.

- 3 Click Insert.

The BGP Insert Aggregates dialog box opens ([Figure 19](#)).

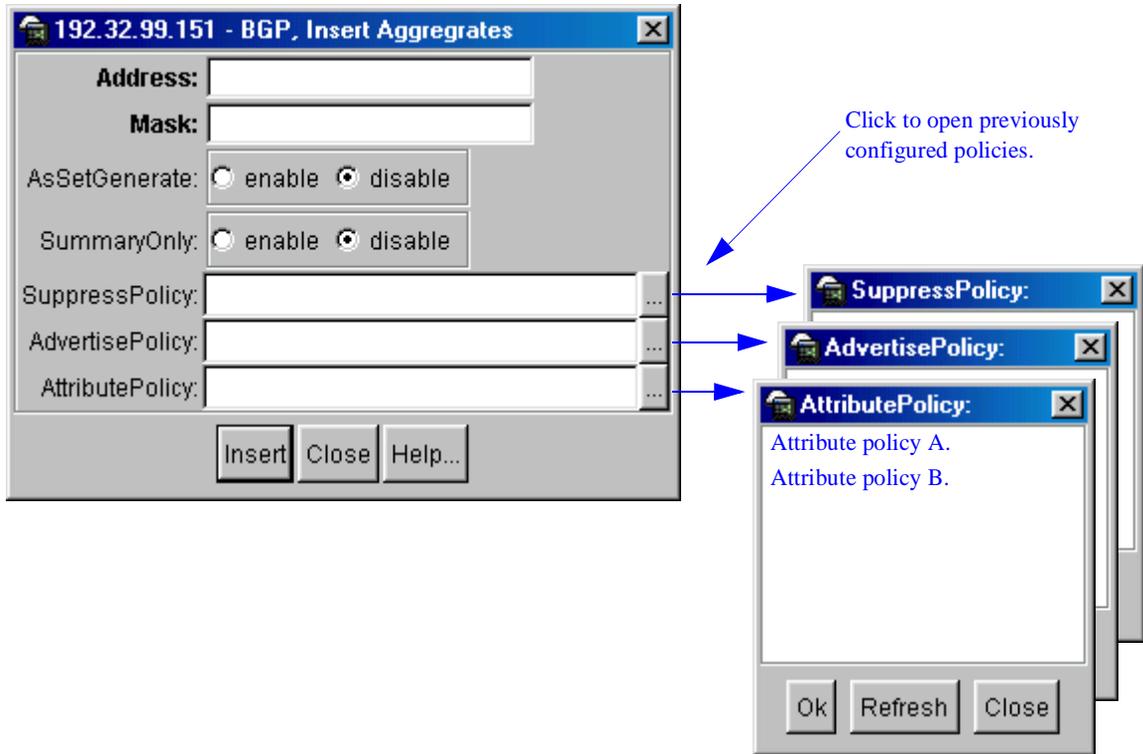
**Figure 19** BGP Insert Aggregates dialog box

Table 6 describes the BGP Insert Aggregates dialog box fields.

**Table 6** BGP Insert Aggregates dialog box fields

Field	Description
Address	The aggregate's IP address that you want to add or modify.
Mask	The aggregate's subnet mask.
AsSetGenerate	Enables or disables autonomous system information. The default value is disable. <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the feature.</li> </ul> You can also enable or disable this feature in the Aggregates tab by clicking in the current field and selecting enable or disable from the pull-down menu (see <a href="#">Figure 18 on page 65</a> ).

**Table 6** BGP Insert Aggregates dialog box fields (continued)

Field	Description
SummaryOnly	<p>Enables or disables the summarization of routes not included in routing updates. This parameter creates the aggregate route and suppresses advertisements of more specific routes to all neighbors. The default value is disable.</p> <ul style="list-style-type: none"> <li>Click the appropriate radio button to enable or disable the option.</li> </ul> <p>You can also enable or disable this feature in the Aggregates tab by clicking in the current field and selecting enable or disable from the pull-down menu (see <a href="#">Figure 18 on page 65</a>).</p>
SuppressPolicy	<p>Sets the route policy by name (any string length between 0 and 64 characters long) to be used for the suppressed route list. This parameter creates the aggregate route and suppresses advertisements of the specified routes.</p> <p>Click the ellipse button and choose a policy (if configured) from the list in the SuppressPolicy dialog box (see <a href="#">Figure 19 on page 66</a>). To deselect an entry, press [Ctrl] and click the left mouse button.</p>
AdvertisePolicy	<p>Sets the route policy by name (any string length between 0 and 64 characters long) to be used for route advertisements.</p> <p>Click the ellipse button and choose a policy (if configured) from the list in the AdvertisePolicy dialog box (see <a href="#">Figure 19 on page 66</a>). To deselect an entry, press [Ctrl] and click the left mouse button.</p>
AttributePolicy	<p>Sets the route policy by name (any string length between 0 and 64 characters long) to be used for setting aggregate route attributes.</p> <p>Click the ellipse button and choose a policy (if configured) from the list in the AttributePolicy dialog box (see <a href="#">Figure 19 on page 66</a>). To deselect an entry, press [Ctrl] and click the left mouse button.</p>

## Configuring allowed network addresses

To configure allowed network addresses:

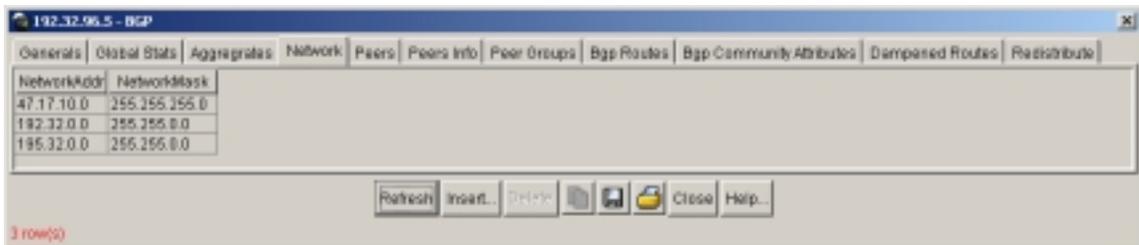
- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (see [Figure 16 on page 56](#)).

- 2 Click the Network tab.

The Network tab opens and displays current allowed network addresses ([Figure 20](#)).

**Figure 20** BGP dialog box—Network tab



- 3 Click Insert.

The BGP Insert Network dialog box opens ([Figure 21](#)).

**Figure 21** BGP Insert Network dialog box



[Table 7](#) describes the BGP Insert Network dialog box fields.

**Table 7** BGP Insert Network dialog box

Field	Description
NetworkAddr	The IP address that you want to add or modify.
NetworkMask	The network subnet mask.

## Configuring and displaying peer information

Three tabs are available for configuring and displaying information about peers and peer groups:

- Peers tab

The Peers tab (see [“Peers tab” on page 70](#)) allows you to display and edit current peer information, and to configure new peers.

You can also use the Peers tab to display the following peer-related statistical data:

- General (peer) statistics
- Receive statistics
- Transmit statistics

- Peers Info tab

The Peers Info tab allows you to display additional information about the current peers that is not available (due to space limitations) in the Peers tab. As in other screen displays, some fields are also editable (see [“Peers Info tab” on page 85](#)).

- Peer Groups tab

The Peer Groups tab allows you to display and edit current peer group information, and to configure new peer groups (see [“Peer Groups tab” on page 88](#)).

## Peers tab

This section includes the following topics:

- [“Configuring and editing peers,”](#) next
- [“Displaying peer statistics”](#) on page 76

### Configuring and editing peers

To configure or edit peers:

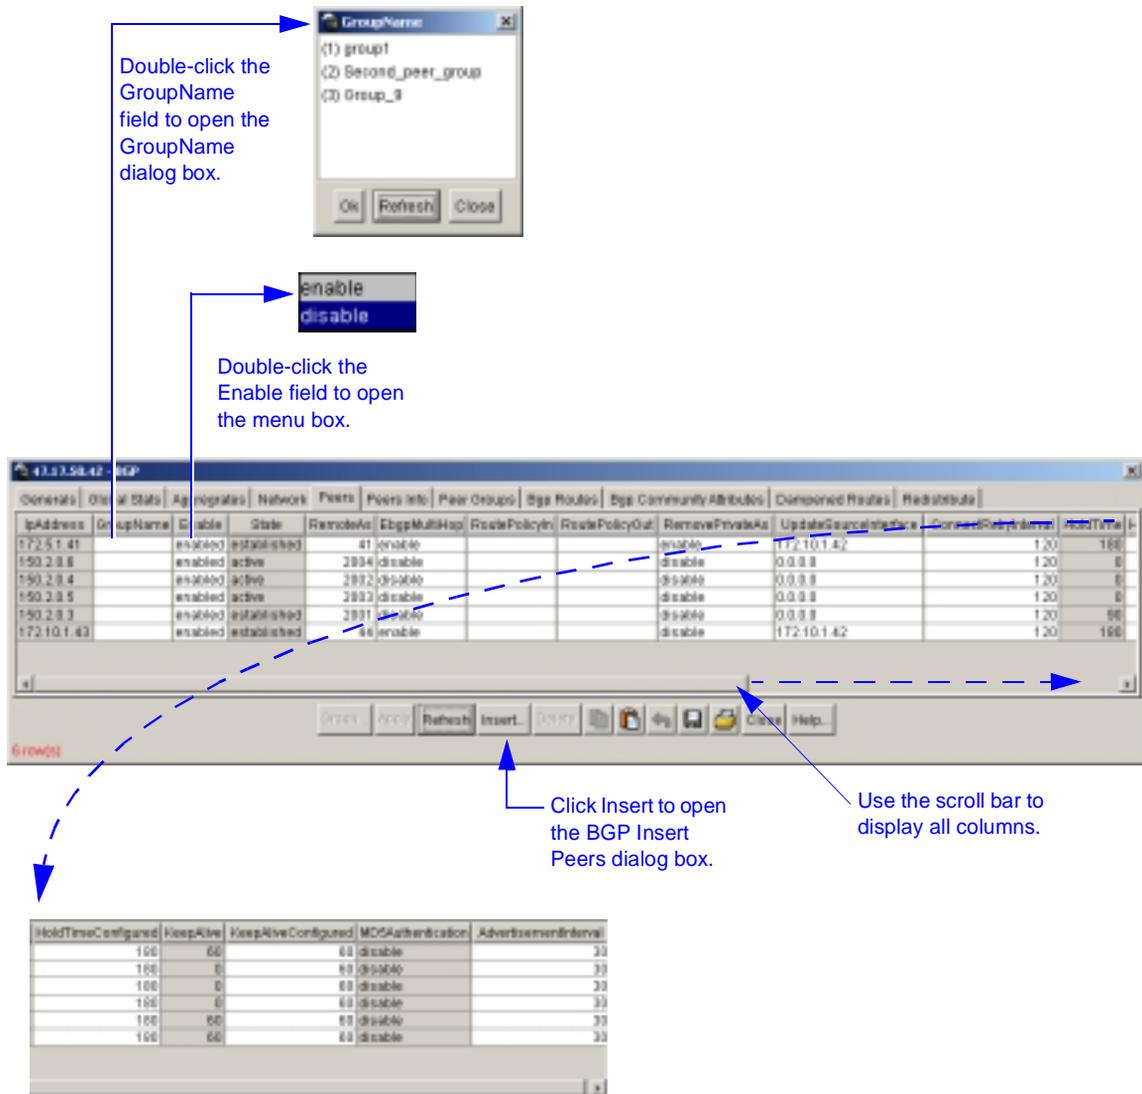
- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (see [Figure 16 on page 56](#)).

- 2 Click the Peers tab.

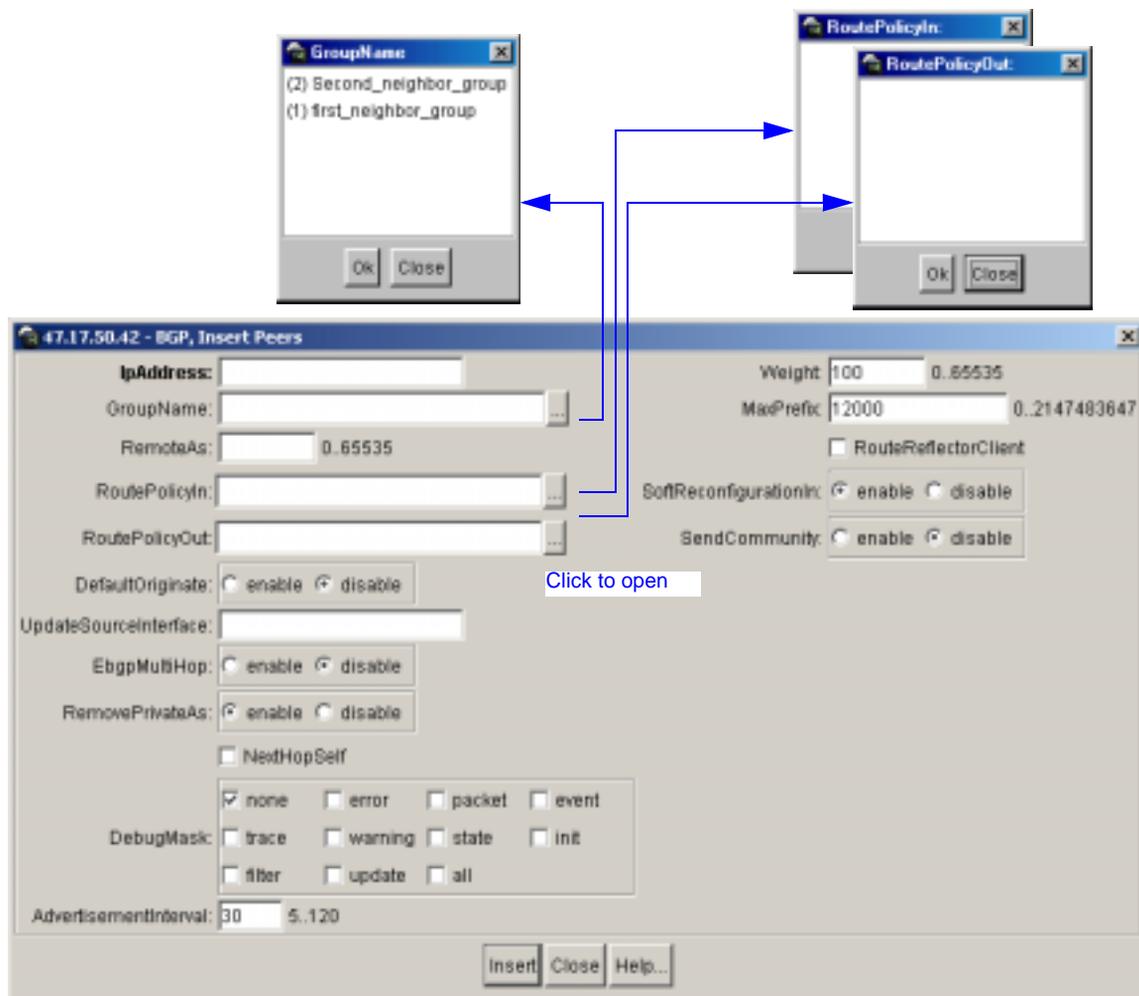
The Peers tab opens and displays current peer configurations ([Figure 22](#)).

Figure 22 BGP dialog box—Peers tab



### 3 Click Insert.

The BGP Insert Peers dialog box opens (Figure 23).

**Figure 23** BGP Insert Peers dialog box

After you enter your required configuration parameters in the Insert Peers dialog box (as described in [Table 8 on page 73](#)), and then click Insert, the Insert Peers dialog box closes and displays your new configuration parameters in the BGP dialog box—Peers tab (see [Figure 22 on page 71](#)).

By default, new peer configuration parameters are set to disable. To enable your new configuration, click the new peer's IP address in the BGP dialog box—Peers tab, and then select enable from the pull-down menu in the Enable column (see [Figure 22 on page 71](#)).

Note that, in [Figure 22 on page 71](#), the specified peer's connection *state* is shown in the (read-only) State field. The State field can display any of the following connection states:

- idle
- connect
- active
- opensent
- openconfirm
- established

[Table 8](#) describes the BGP Insert Peers dialog box fields.

**Table 8** BGP Insert Peers dialog box

Field	Description
IpAddress	The IP address of this peer.
GroupName	<p>Adds a BGP peer to the specified subscriber group.</p> <p>Click the ellipse button and choose a group (if configured) from the list in the GroupName dialog box (see <a href="#">Figure 22 on page 71</a>). To deselect an entry, press [Ctrl] and click the left mouse button.</p> <p>You can also choose a group (if configured) from the Peers tab by clicking in the current field and choosing from the list in the GroupName dialog box (see <a href="#">Figure 22 on page 71</a>).</p>
RemoteAs	<p>Configures a remote-as for the peer or peer-group.</p> <ul style="list-style-type: none"> <li>• Specify an integer value in the range 0 to 65535.</li> </ul>
RoutePolicyIn	<p>Applies an incoming route policy rule to all routes that are learned from, or sent to, the local BGP router's peers, or peer groups.</p> <p>Click the ellipse button and choose a route policy (if configured) from the list in the RoutePolicyIn dialog box (see <a href="#">Figure 23 on page 72</a>). To deselect an entry, press [Ctrl] and click the left mouse button.</p>

**Table 8** BGP Insert Peers dialog box (continued)

Field	Description
RoutePolicyOut	<p>Applies an outgoing route policy rule to all routes that are learned from, or sent to, the local BGP router's peers, or peer groups.</p> <p>Click the ellipse button and choose a route policy (if configured) from the list in the RoutePolicyOut dialog box (see <a href="#">Figure 23 on page 72</a>). To deselect an entry, press [Ctrl] and click the left mouse button.</p>
DefaultOriginate	<p>When enabled, specifies that the current route originated from the BGP peer. This field enables or disables sending the default route information to the specified neighbor or peer. The default value is disable.</p>
UpdateSourceInterface	<p>Specifies the source IP address to be used when sending EBGP packets to this peer or peer group.</p>
EbgpMultiHop	<p>Enables or disables a connection to a BGP peer that is more than one hop away from the local router. The default value is disable.</p>
RemovePrivateAs	<p>When enabled, strips private AS numbers when sending an update. This feature is especially useful within a confederation. The default value is enable.</p>
NextHopSelf	<p>Check box—When checked (enabled), specifies that the next-hop attribute in an IBGP update is the address of the local router or the router that is generating the IBGP update. The default value is not checked (disabled).</p>
DebugMask	<p>Allows you to display specified debug information for the specified BGP peer. The default value is none.</p> <p>Mask choices are:</p> <ul style="list-style-type: none"> <li><i>none</i>: disables all debug messages.</li> <li><i>error</i>: enables display of debug error messages.</li> <li><i>packet</i>: enables display of debug packet messages.</li> <li><i>event</i>: enables display of debug event messages.</li> <li><i>trace</i>: enables display of debug trace messages.</li> <li><i>warning</i>: enables display of debug warning messages.</li> <li><i>state</i>: enables display of debug state transition messages.</li> <li><i>init</i>: enables display of debug initialization messages.</li> <li><i>filter</i>: enables display of debug messages related to filtering.</li> <li><i>update</i>: enables display of debug messages related to sending and receiving updates.</li> <li><i>all</i>: enables all debug messages.</li> </ul>

**Table 8** BGP Insert Peers dialog box (continued)

Field	Description
AdvertisementInterval	Specifies the time interval (in seconds) that transpires between each transmission of an advertisement from a BGP neighbor. The default value is 5 seconds. <ul style="list-style-type: none"> <li>Specify an integer value in the range 5 to 120 seconds.</li> </ul>
ConnectRetryInterval	Sets the time interval (in seconds) for the ConnectRetry Timer. The default value is 120 seconds. <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 65535 seconds.</li> </ul>
Weight	Specifies this peer's or peer groups' weight, or the priority of updates that can be received from this BGP peer. The default value is 100. <p><b>Note:</b> A weight is a numerical value you assign to a path that allows you to control the path selection process. The administrative weight is local to the router. For example, if you have particular neighbors that you want to prefer for most of your traffic, you can assign a higher weight to all routes learned from that neighbor.</p> <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 65535 seconds.</li> </ul>
MaxPrefix	Sets a limit on the number of routes that can be accepted from a neighbor. The default value is 12,000 routes (see <a href="#">"Specifying number of routes learned" on page 201</a> ). <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 2147483647.</li> </ul> <p><b>Note:</b> A value of 0 (zero) indicates that there is no limit to the number of routes that can be accepted.</p>
RouteReflectorClient	Check box—When checked (enabled), specifies this peer or peer group as a route reflector client. The default value is disable. <p><b>Note:</b> All peers that are configured become members of the client group and the remaining IBGP peers become members of the nonclient group for the local route reflector.</p>
SoftReconfigurationIn	Allows the router to relearn routes from the specified neighbor or group of neighbors without resetting the connection when the policy changes in the inbound direction. The default value is disable.
SendCommunity	Enables or disables sending the update message's community attribute to the specified peer. The default value is disable.

## Displaying peer statistics

To display statistics for a specified peer:

- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (see [Figure 16](#) on page 56).

- 2 Click the Peers tab.

The Peers tab opens and displays current peer configurations ([Figure 24](#)).

**Figure 24** BGP dialog box—Peers tab

IpAddress	GroupName	Enable	State	RemoteAs	SigAuthHop	RoutePolicyIn	RoutePolicyOut	RemovePrivateAs	UpdateSourceInterface	ConnectRetryInterval	HoldTime
172.5.1.41		enabled	established	41	enable			enable	172.10.1.42	120	180
150.2.8.6		enabled	active	2004	disable			disable	8.8.8.0	120	0
150.2.8.4		enabled	active	2002	disable			disable	8.8.0.0	120	0
150.2.8.5		enabled	active	2003	disable			disable	8.8.0.0	120	0
150.2.8.3		enabled	established	2001	disable			disable	8.8.0.0	120	90
172.10.1.43		enabled	established	44	enable			enable	172.10.1.42	120	180

Graph button  
(shown in inactive state).

Use the scroll bar to  
display all columns.

- 3 Select the IP address of the peer you want to query.

The selected peer's IP address is highlighted and the Graph button is now active.

- 4 Click Graph.

The BGP Peer Stat dialog box opens with the General Stats tab displayed ([Figure 25](#)).

**Figure 25** BGP Peer Stat dialog box—General Stats tab

	AbsoluteValue	Cumulative	Average/sec	Minimum/sec	Maximum/sec	LastVal/sec
InUpdates	73	0	0	0	0	0
OutUpdates	3,424	0	0	0	0	0
InTotalMessages	185	1	0.03	0	0.05	0.05
OutTotalMessages	3,536	1	0.03	0	0.05	0.05
FsmEstablishedTransitions	1	0	0	0	0	0
FsmEstablishedTime	83	N/A	N/A	N/A	N/A	N/A
InUpdateElapsedTime	135	N/A	N/A	N/A	N/A	N/A
Starts	2	0	0	0	0	0
Stops	0	0	0	0	0	0
Opens	1	0	0	0	0	0
Closes	0	0	0	0	0	0
Fails	1	0	0	0	0	0
Fatalis	0	0	0	0	0	0
ConnExps	0	0	0	0	0	0
HoldExps	0	0	0	0	0	0
KeepExps	110	1	0.03	0	0.05	0.05
BadEvents	0	0	0	0	0	0
SyncFails	0	0	0	0	0	0
RcvdTooShort	0	0	0	0	0	0
NoMarker	0	0	0	0	0	0
Dropped	0	0	0	0	0	0
BadMsgTypes	0	0	0	0	0	0
TrEvent	0	0	0	0	0	0

As shown in [Figure 25](#), the BGP Peer Stat dialog box comprises three tabs:

- “General Stats tab,” next
- “Receive Stats tab” on page 79
- “Transmit Stats tab” on page 82

### *General Stats tab*

The General Stats tab (shown in [Figure 25](#)) displays general statistics for a specified peer.

[Table 9](#) describes the General Stats tab fields.

**Table 9** General Stats tab fields

Field	Description
InUpdates	Number of updates received by the peer.
OutUpdates	Number of updates transmitted by the peer.
InTotalMessages	Number of total messages received by the peer.
OutTotalMessages	Number of total messages transmitted by the peer.
FsmEstablishedTransitions	The total number of times the BGP FSM transitioned into the established state.
FsmEstablishedTime	Number of seconds this peer has been in the Established state or how long since this peer was last in the Established state. This value It is set to zero when a new peer is configured or the router is booted.
InUpdateElapsedTime	Elapsed time in seconds since the last BGP UPDATE message was received from the peer.
Starts	Number of times peer BGP connection started.
Stops	Number of times peer BGP connection stopped.
Opens	Number of times peer opened TCP.
Closes	Number of times peer closed TCP.
Fails	Number of times a peer TCP attempt failed.
Fatals	Number of times peer TCP crashed due to fatal error.
ConnExps	Number of times the peer TCP retry timer expired.
HoldExps	Number of times the peer hold timer expired.
KeepExps	Number of times the peer keepalive timer expired.
BadEvents	Number of invalid events received by the peer.
SyncFails	Number of times the peer FDB sync failed.
RcvdTooShort	Number of “too short” messages received by the peer.
NoMarker	Number of messages received by the peer with no marker.
Dropped	Number of messages dropped by the peer.
BadMsgTypes	Number of messages received by the peer as “invalid message type.”
TrEvent	Peer trace event.

➤ To display statistics *received* by the peer, see “[Receive Stats tab](#),” next.

- To display statistics *transmitted* by the peer, see “[Transmit Stats tab](#)” on [page 82](#).

### *Receive Stats tab*

To display statistical information received by the specified peer:

- From the The BGP Peer Stat dialog box, click the Receive Stats tab (see [Figure 25 on page 77](#)).

The Receive Stats tab opens and displays current statistical information received by the specified peer (see [Figure 26 on page 80](#)).

**Figure 26** BGP Peer Stat dialog box—Receive Stats tab

	Absolute Value	Cumulative	Average/sec	Minimum/sec	Maximum/sec	Last/Val/sec
RxMsgs	107	3	0.02	0	0.1	0.1
RxInCompPkts	0	0	0	0	0	0
RxOpens	1	0	0	0	0	0
RxKeeps	113	3	0.02	0	0.1	0.1
RxUpdates	73	0	0	0	0	0
RxNotifs	0	0	0	0	0	0
RxRoutesAdded	0	0	0	0	0	0
RxRoutesReplaced	0	0	0	0	0	0
RxNil	0	0	0	0	0	0
RxValidUpdates	73	0	0	0	0	0
RxECodeHeader	0	0	0	0	0	0
RxECodeOpen	0	0	0	0	0	0
RxECodeUpdate	0	0	0	0	0	0
RxECodeHoldTimer	0	0	0	0	0	0
RxECodeFSM	0	0	0	0	0	0
RxECodeCease	0	0	0	0	0	0
RxHdrCodeNoSync	0	0	0	0	0	0
RxHdrCodeInvalidMsgLen	0	0	0	0	0	0
RxHdrCodeInvalidMsgType	0	0	0	0	0	0
RxOpCodeBadVer	0	0	0	0	0	0
RxOpCodeBadAs	0	0	0	0	0	0
RxOpCodeBadRID	0	0	0	0	0	0
RxOpCodeUnsuppOption	0	0	0	0	0	0
RxOpCodeAuthFail	0	0	0	0	0	0
RxOpCodeBadHold	0	0	0	0	0	0
RxUpdCodeMalformedAttrList	0	0	0	0	0	0
RxUpdCodeWellknownAttrUnrecog	0	0	0	0	0	0
RxUpdCodeWellknownAttrMiss	0	0	0	0	0	0
RxUpdCodeAttrFlagError	0	0	0	0	0	0
RxUpdCodeAttrLenError	0	0	0	0	0	0
RxUpdCodeBadORIGINAttr	0	0	0	0	0	0
RxUpdCodeASRoutingLoop	0	0	0	0	0	0

Use the vertical scrollbar to display all rows.

Table 10 describes the Receive Stats tab fields.

**Table 10** Receive Stats tab fields

Field	Description
RxMsgs	Number of messages received by the peer.
RxInCompPkts	Number of Incomplete messages received by the peer.
RxOpens	Number of Opens received by the peer.

**Table 10** Receive Stats tab fields (continued)

Field	Description
RxKeeps	Number of Keepalive messages received by the peer.
RxUpdates	Number of Updates received by the peer.
RxNotifys	Number of Notifications received by the peer.
RxRoutesAdded	Number of routes added into loc_rib by this peer.
RxRoutesReplaced	Number of routes that were replaced by routes received by the peer.
RxNLri	Number of network layer reachability information (NLRI) messages received by the peer.
RxValidUpdates	Number of valid Updates received by the peer.
RxECodeHeader	Number of Header errors received by the peer.
RxECodeOpen	Number of Open errors received by the peer.
RxECodeUpdate	Number of Update errors received by the peer.
RxECodeHoldtimer	Number of Hold errors received by the peer.
RxECodeFSM	Number of FSM errors received by the peer.
RxECodeCease	Number of Cease errors received by the peer.
RxHdrCodeNoSync	Number of Header errors received by the peer as: Not Synchronized.
RxHdrCodeInvalidMsgLen	Number of Header errors received by the peer as: Invalid Message Length.
RxHdrCodeInvalidMsgType	Number of Header errors received by the peer as: Invalid Message Type.
RxOpCodeBadVer	Number of Open errors received by the peer as: Bad Version.
RxOpCodeBadAs	Number of Open errors received by the peer as: Bad AS.
RxOpCodeBadRtID	Number of Open errors received by the peer as: Bad BGP ID.
RxOpCodeUnsuppOption	Number of Open errors received by the peer as: Unsupported Options.
RxOpCodeAuthFail	Number of Open errors received by the peer as: Authorization Failures.
RxOpCodeBadHold	Number of Open errors received by the peer as: Bad Hold Value.
RxUpdCodeMalformedAttrList	Number of Update errors received by the peer as: Malformed Attr List.
RxUpdCodeWelknownAttrUnrecog	Number of Update errors received by the peer as: Wellknown Attr Unrecog.
RxUpdCodeWelknownAttrMiss	Number of Update errors received by the peer as: Wellknown Attr Missing.
RxUpdCodeAttrFlagError	Number of Update errors received by the peer as: Attr Flag Error.
RxUpdCodeAttrLenError	Number of Update errors received by the peer as: Attr Length Error.

**Table 10** Receive Stats tab fields (continued)

Field	Description
RxUpdCodeBadORIGINAttr	Number of Update errors received by the peer as: Attr Flag Error.
RxUpdCodeASRoutingLoop	Number of Update errors received by the peer as: AS Routing Loop.
RxUpdCodeBadNHAttr	Number of Update errors received by the peer as: Bad Next-Hop Attr.
RxUpdCodeOptionalAttrError	Number of Update errors received by the peer as: Optional Attr Error.
RxUpdCodeBadNetworkField	Number of Update errors received by the peer as: Bad Network Field.
RxUpdCodeMalformedASPath	Number of Update errors received by the peer as: Malformed AS PATH.

- To display statistics *transmitted* by the peer, see [“Transmit Stats tab” on page 82](#).

### *Transmit Stats tab*

To display statistical information transmitted by the specified peer:

- From the The BGP Peer Stat dialog box, click the Transmit Stats tab (see [Figure 25 on page 77](#)).

The Transmit Stats tab opens and displays current statistical information transmitted by the specified peer (see [Figure 27 on page 83](#)).

Figure 27 BGP Peer Stat dialog box—Transmit Stats tab

	Absolute Value	Cumulative	Average/sec	Minimum/sec	Maximum/sec	LastVal/sec
TxMsgs	3,535					
TxOpens	1					
TxKeeps	110					
TxUpdates	3,424					
TxNotfys	0					
TxRoutes	11,134					
TxECodeHeader	0					
TxECodeOpen	0					
TxECodeUpdate	0					
TxECodeHoldtimer	0					
TxECodeFSM	0					
TxECodeCease	0					
TxHdrCodeNoSync	0					
TxHdrCodeInvaldMsgLen	0					
TxHdrCodeInvaldMsgType	0					
TxOpCodeBadVer	0					
TxOpCodeBadAs	0					
TxOpCodeBadRTID	0					
TxOpCodeUnsuppOption	0					
TxOpCodeAuthFail	0					
TxOpCodeBadHold	0					
TxUpdCodeMalformedAttrList	0					
TxUpdCodeWelknownAttrUnrecog	0					
TxUpdCodeWelknownAttrMiss	0					
TxUpdCodeAttrFlagError	0					
TxUpdCodeAttrLenError	0					
TxUpdCodeBadORIGINAttr	0					
TxUpdCodeASRoutingLoop	0					
TxUpdCodeBadNHAttr	0					
TxUpdCodeOptionalAttrError	0					
TxUpdCodeBadNetworkField	0					
TxUpdCodeMalformedASPath	0					

Clear Counters Close Help... Poll Interval: 10s 00h:00m:00s

Table 11 describes the Transmit Stats tab fields.

**Table 11** Transmit Stats tab fields

Field	Description
TxMsgs	Number of messages transmitted by the peer.
TxOpens	Number of Opens transmitted by the peer.
TxKeeps	Number of Keepalive messages transmitted by the peer.
TxUpdates	Number of Updates transmitted by the peer.
TxNotifys	Number of Notifications transmitted by the peer.
TxRoutes	Number of network layer reachability information (NLRI) messages transmitted by the peer.
TxECodeHeader	Number of Header errors transmitted by the peer.
TxECodeOpen	Number of Open errors transmitted by the peer.
TxECodeUpdate	Number of Update errors transmitted by the peer.
TxECodeHoldtimer	Number of Hold errors transmitted by the peer.
TxECodeFSM	Number of FSM errors transmitted by the peer.
TxECodeCease	Number of Cease errors transmitted by the peer.
TxHdrCodeNoSync	Number of Header errors transmitted by the peer as: Not Synchronized.
TxHdrCodeInvalidMsgLen	Number of Header errors transmitted by the peer as: Invalid Message Length.
TxHdrCodeInvalidMsgType	Number of Header errors transmitted by the peer as: Invalid Message Type.
TxOpCodeBadVer	Number of Open errors transmitted by the peer as: Bad Version.
TxOpCodeBadAs	Number of Open errors transmitted by the peer as: Bad AS.
TxOpCodeBadRtID	Number of Open errors transmitted by the peer as: Bad BGP ID.
TxOpCodeUnsuppOption	Number of Open errors transmitted by the peer as: Unsupported Options.
TxOpCodeAuthFail	Number of Open errors transmitted by the peer as: Authorization Failures.
TxOpCodeBadHold	Number of Open errors transmitted by the peer as: Bad Hold Value.
TxUpdCodeMalformedAttrList	Number of Update errors transmitted by the peer as: Malformed Attr List.

**Table 11** Transmit Stats tab fields (continued)

Field	Description
TxUpdCodeWelknownAttrUnrecog	Number of Update errors transmitted by the peer as: Wellknown Attr Unrecog.
TxUpdCodeWelknownAttrMiss	Number of Update errors transmitted by the peer as: Wellknown Attr Missing.
TxUpdCodeAttrFlagError	Number of Update errors transmitted by the peer as: Attr Flag Error.
TxUpdCodeAttrLenError	Number of Update errors transmitted by the peer as: Attr Length Error.
TxUpdCodeBadORIGINAttr	Number of Update errors transmitted by the peer as: Attr Flag Error.
TxUpdCodeASRoutingLoop	Number of Update errors transmitted by the peer as: AS Routing Loop.
TxUpdCodeBadNHAttr	Number of Update errors transmitted by the peer as: Bad Next-Hop Attr.
TxUpdCodeOptionalAttrError	Number of Update errors transmitted by the peer as: Optional Attr Error.
TxUpdCodeBadNetworkField	Number of Update errors transmitted by the peer as: Bad Network Field.
TxUpdCodeMalformedASPath	Number of Update errors transmitted by the peer as: Malformed AS Path.

## Peers Info tab

To display or edit current peer information:

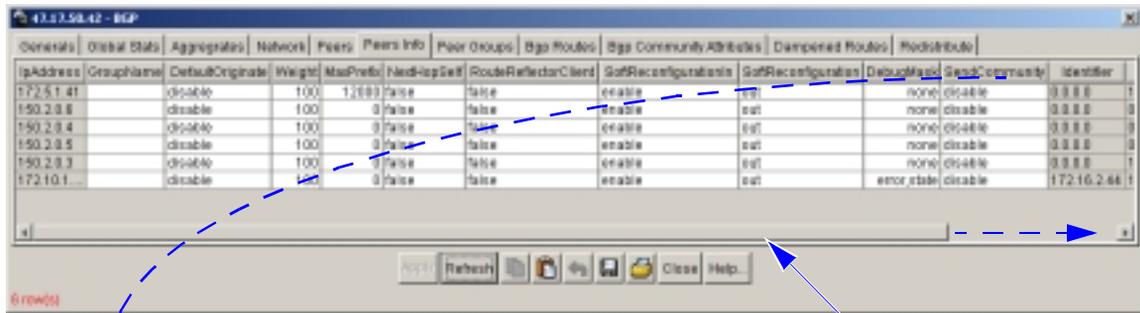
- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (see [Figure 16 on page 56](#)).

- 2 Click the Peers Info tab.

The Peers Info tab opens and displays current peer configurations ([Figure 28](#)).

**Figure 28** BGP dialog box—Peers Info tab



Use the scroll bar to display all columns.

LocalAddr	ConnectRetryInterval
172.18.1.42	120
0.0.0.0	120
0.0.0.0	120
0.0.0.0	120
168.2.0.1	120
172.18.1.42	120

Table 12 describes the BGP Peers Info tab fields.

**Table 12** BGP Peers Info tab fields

Field	Description
IpAddress	Read-only field — displays the IP address of configured peers.
GroupName	Displays the peer’s subscriber group name.
DefaultOriginate	When set to enable, indicates that this peer or peer group sends default route information to the specified neighbor or peer.
Weight	Specifies this peer’s or peer groups’ weight value, or the priority of updates that can be received from this BGP peer. <b>Note:</b> A weight is a numerical value you assign to a path that allows you to control the path selection process. The administrative weight is local to the router. For example, if you have particular neighbors that you want to prefer for most of your traffic, you can assign a higher weight to all routes learned from that neighbor.

**Table 12** BGP Peers Info tab fields (continued)

Field	Description
MaxPrefix	Displays the configured limit value on the number of routes that can be accepted from a neighbor. The default value is 12,000 routes (see <a href="#">“Specifying number of routes learned” on page 201</a> ). <b>Note:</b> A value of 0 (zero) indicates that there is no limit to the number of routes that can be accepted.
NextHopSelf	When set to true (true = enabled, false = disabled), specifies that the next-hop attribute in an IBGP update is the address of the local router or the router that is generating the IBGP update.
RouteReflectorClient	When set to true (true = enabled, false = disabled), specifies this peer or peer group as a route reflector client.
SoftReconfigurationIn	When set to enable, allows the router to relearn routes from the specified neighbor or group of neighbors without resetting the connection when the policy changes in the inbound direction.
SoftReconfiguration	Indicates that the specified neighbor sends out all updates to the remote neighbor without resetting the connection.
DebugMask	Indicates the specified debug setting for the specified BGP peer. The default value is none. Mask choices are: <i>none</i> : disables all debug messages. <i>error</i> : enables display of debug error messages. <i>packet</i> : enables display of debug packet messages. <i>event</i> : enables display of debug event messages. <i>trace</i> : enables display of debug trace messages. <i>warning</i> : enables display of debug warning messages. <i>state</i> : enables display of debug state transition messages. <i>init</i> : enables display of debug initialization messages. <i>filter</i> : enables display of debug messages related to filtering. <i>update</i> : enables display of debug messages related to sending and receiving updates. <i>all</i> : enables all debug messages.
SendCommunity	Enables or disables sending the update message's community attribute to the specified peer. The default value is disable.

**Table 12** BGP Peers Info tab fields (continued)

Field	Description
Identifier	Indicates the BGP identifier of this entry's BGP peer.
ConnectRetryInterval	Time interval (in seconds) for the ConnectRetry Timer. The default value is 120 seconds. <ul style="list-style-type: none"><li>Specify an integer value in the range 1 to 65535 seconds.</li></ul>

## Peer Groups tab

To configure or edit peer groups:

- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (see [Figure 16 on page 56](#)).

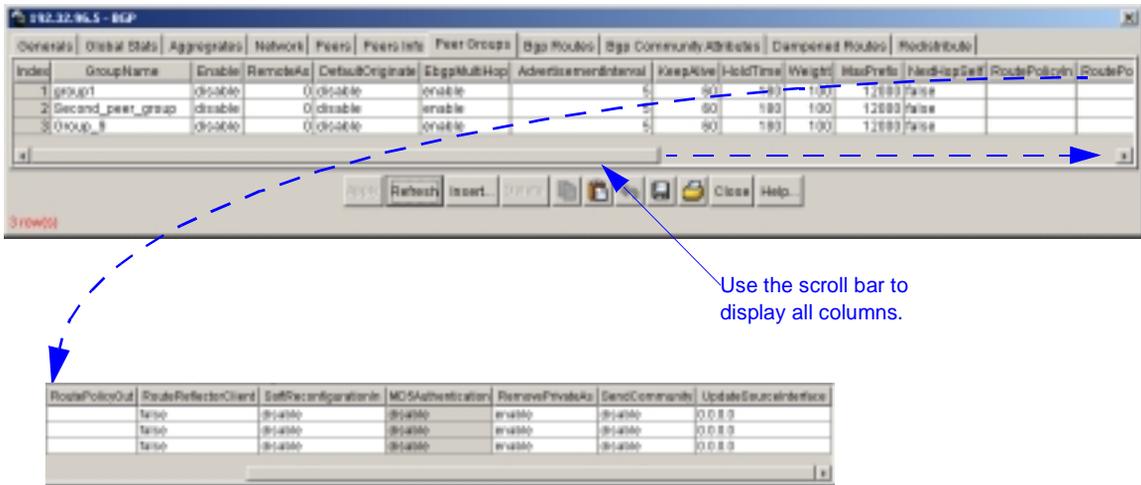
- 2 Click the Peer Groups tab.

The Peer Groups tab opens and displays current peer group configurations ([Figure 29](#)). You can edit any of the fields in this screen by double-clicking in the field. Your changes are applied after you click Apply.



**Note:** Configuration changes that you apply in this screen affect all members of the associated peer group.

---

**Figure 29** BGP dialog box—Peer Groups tab**3** Click Insert.

The BGP Insert Peer Groups dialog box opens (Figure 30).

**Figure 30** BGP Insert Peer Groups dialog box

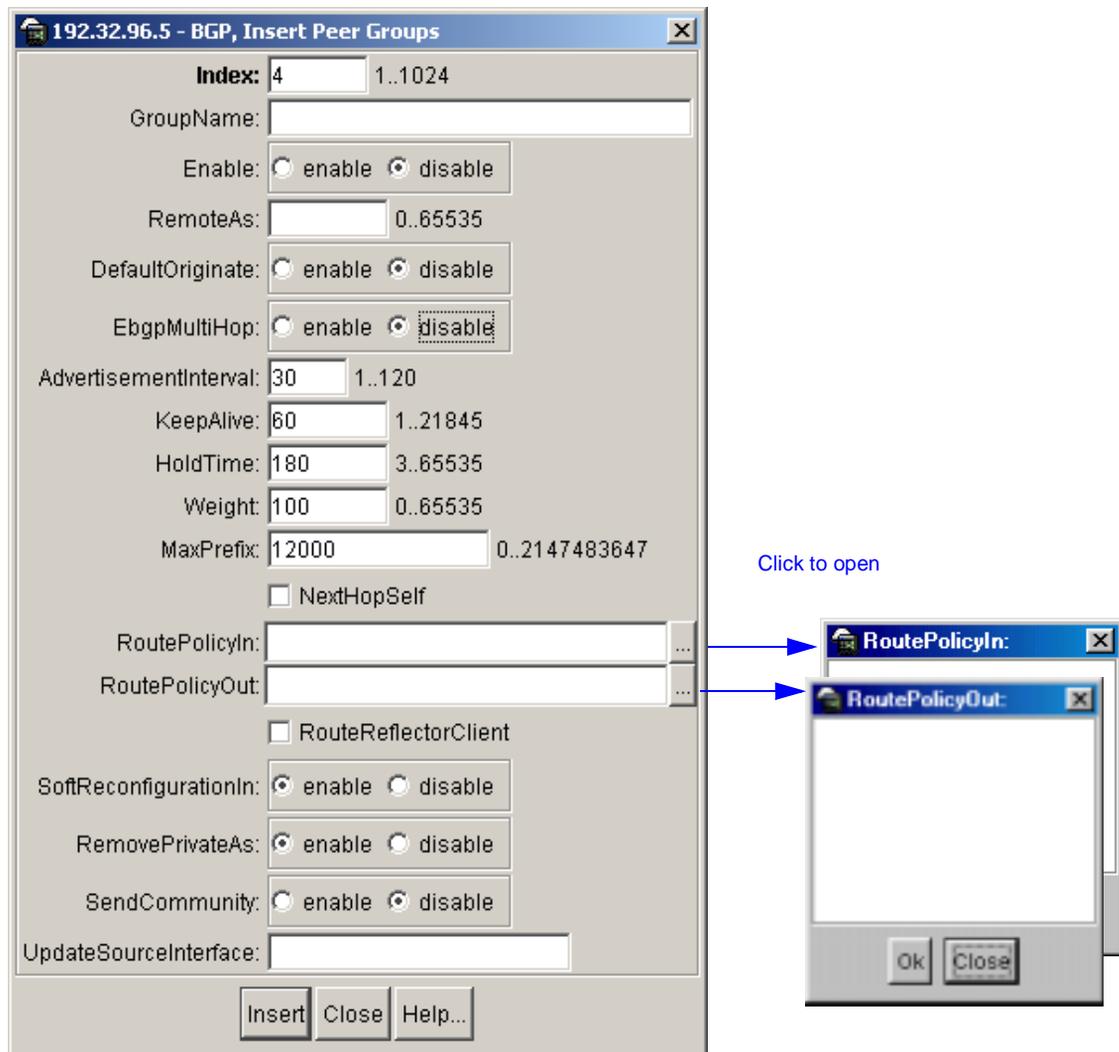


Table 13 describes the BGP Insert Peer Groups dialog box fields.

**Table 13** BGP Insert Peer Groups dialog box fields

Field	Description
Index	Sets a value which corresponds to a unique entry in the Peer Group Table. The default value is 1. <ul style="list-style-type: none"> <li>Specify an integer value in the range 1 to 1024.</li> </ul>
GroupName	Name of the group you are creating.
Enable	Enables or disable the group. The default value is disable.
RemoteAs	Configures a remote-as for all peer-group members. <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 65535.</li> </ul>
DefaultOriginate	When enabled, specifies that the current route originated from the BGP peer. This field enables or disables sending the default route information to the specified neighbor or peer. The default value is disable.
EbgpMultiHop	Enables or disables a connection to a BGP peer that is more than one hop away from the local router. The default value is disable.
AdvertisementInterval	Specifies the time interval (in seconds) that transpires between each transmission of an advertisement from a BGP neighbor. The default value is 30 seconds. <ul style="list-style-type: none"> <li>Specify an integer value in the range 1 to 120 seconds.</li> </ul>
KeepAlive	Sets the time interval (in seconds) that transpires between transmissions of the local BGP router's keep-alive packets. The keep-alive packets indicate the enabled status of the local BGP router to peers. The default value is 60 seconds. <ul style="list-style-type: none"> <li>Specify an integer value in the range 1 to 21845 seconds.</li> </ul>
HoldTime	Sets the BGP speaker's time interval (in seconds) for this peer. The default value is 180 seconds. <ul style="list-style-type: none"> <li>Specify an integer value in the range 3 to 21845 seconds (integer values 1 and 2 are not valid).</li> </ul>

**Table 13** BGP Insert Peer Groups dialog box fields (continued)

Field	Description
Weight	<p>Specifies this peer's or peer groups' weight, or the priority of updates that can be received from this BGP peer. The default value is 100.</p> <p><b>Note:</b> A weight is a numerical value you can assign a path so that you can control the path selection process. The administrative weight is local to the router. If you have particular neighbors that you want to prefer for most of your traffic, you can assign a higher weight to all routes learned from that neighbor.</p> <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 65535 seconds.</li> </ul>
MaxPrefix	<p>Sets a limit on the number of routes that can be accepted from a neighbor. The default value is 12,000 routes (see <a href="#">"Specifying number of routes learned" on page 201</a>).</p> <ul style="list-style-type: none"> <li>Specify an integer value in the range 0 to 2147483647.</li> </ul> <p><b>Note:</b> A value of 0 (zero) indicates that there is no limit to the number of routes that can be accepted.</p>
NextHopSelf	<p>Check box—When checked (enabled), specifies that the next-hop attribute in an IBGP update is the address of the local router or the router that is generating the IBGP update. The default value is not checked (disabled).</p>
RoutePolicyIn	<p>Applies an incoming route policy rule to all routes that are learned from, or sent to, the local BGP router's peers, or peer groups.</p> <p>Click the ellipse button and choose a route policy (if configured) from the list in the RoutePolicyIn dialog box (see <a href="#">Figure 30 on page 90</a>). To deselect an entry, press [Ctrl] and click the left mouse button.</p>
RoutePolicyOut	<p>Applies an outgoing route policy rule to all routes that are learned from, or sent to, the local BGP router's peers, or peer groups.</p> <p>Click the ellipse button and choose a route policy (if configured) from the list in the RoutePolicyOut dialog box. To deselect an entry, press [Ctrl] and click the left mouse button.</p>
RouteReflectorClient	<p>Check box—When checked (enabled), specifies this peer or peer group as a route reflector client. The default value is disable.</p> <p><b>Note:</b> All peers that are configured become members of the client group and the remaining IBGP peers become members of the nonclient group for the local route reflector.</p>

**Table 13** BGP Insert Peer Groups dialog box fields (continued)

Field	Description
SoftReconfigurationIn	Allows the router to relearn routes from the specified neighbor or group of neighbors without resetting the connection when the policy changes in the inbound direction. The default value is disable.
RemovePrivateAs	When enabled, strips private AS numbers when sending an update. This feature is especially useful within a confederation. The default value is enable.
SendCommunity	Enables or disables sending the update message's community attribute to the specified peer. The default value is disable.
UpdateSourceInterface	Specifies the source IP address to be used when sending BGP packets to this peer or peer group.

## Displaying route information

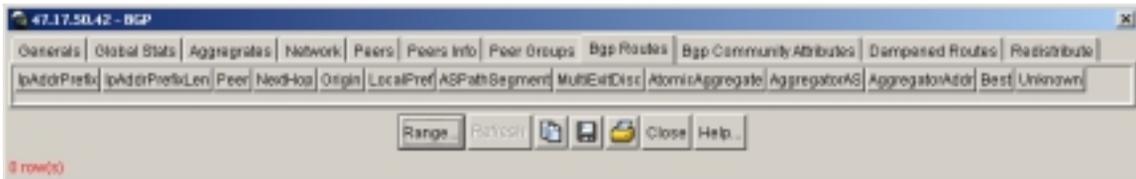
To display current BGP route information:

- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (see [Figure 16 on page 56](#)).

- 2 Click the BGP Routes tab.

The BGP Routes tab opens ([Figure 31](#)).

**Figure 31** BGP dialog box—BGP Routes tab (1 of 2)

- 3 Click Range.

The BGP Route Range dialog box opens ([Figure 32](#)).

**Figure 32** BGP Route Range dialog box

- 4 Choose the type of information you want to display by selecting the options in the Route Range dialog box.

[Table 14](#) describes the BGP Route Range dialog box options.

**Table 14** BGP Route Range dialog box options

Options	Description
Type	Specifies the route types to query: <ul style="list-style-type: none"> <li>• All = all routes in the route table</li> <li>• Exact Prefixes = only routes that exactly match the prefix/len</li> <li>• Longer Prefix = all routes that match the prefix with longer masks.</li> </ul>
IpAddrPrefix	The IP address of the route destination (for example 2.2.2.2).
IpAddrPrefixLen	The mask of the route destination. Must be an integer value between 0 and 32 (for example 16).

- 5 Click Set Range.

- 6 Click Close.

The Route Range dialog box closes.

- 7 In the BGP Routes tab, click Refresh

The BGP Routes tab refreshes and displays route information based on your selected options in the Route Range dialog box ([Figure 33](#)).



**Note:** Each time you choose new options from the Route Range dialog box, you must click Refresh in the BGP Routes tab to allow the screen to refresh and display your new configuration choices.

**Figure 33** BGP dialog box—BGP Routes tab (2 of 2)

IpAddr/Prefix	IpAddr/PrefixLen	Peer	NextHop	Origin	LocalPref	ASPathSegment	MultiExitDx
1.0.1.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
1.0.3.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
1.0.5.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
1.0.7.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
1.0.9.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
1.0.11.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
1.0.13.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
1.0.15.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
1.0.17.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
1.0.19.0	24	150.2.0.3	150.2.0.3	igp	100	SEQ 2001,1000	
12.0.0.0	8	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,7018	
12.0.48.0	20	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,208,1742	
12.0.262.0	23	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,701,16927	
12.1.83.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,7911,5696,14767	
12.1.86.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,701,19024,18024,14259,22306	
12.1.245.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,7018,11521	
12.2.6.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,7018,11101	
12.2.7.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,7018,11101	
12.2.41.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,208,13606	
12.2.86.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,208,22224	
12.2.88.0	22	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,1238,11101	
12.2.84.0	23	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,1239,11101	
12.2.87.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,1,7269,10856	
12.2.99.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,1,7269,10856	
12.2.109.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,1,7269,10856	
12.2.110.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,1,7269,10856	
12.2.142.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,7018,14381	
12.2.169.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,7018,11006	
12.2.192.0	24	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,701,14984,14984,14984,14984	
12.2.218.0	23	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,208,15052	
12.2.220.0	22	172.10.1.43	172.16.2.5	igp	100	SEQ 8600,762,11296,3561,208,15052	

stopped at 454 row(s)

Use the scroll bar to display all rows.

Use the scroll bar to display all columns.

Table 15 describes the BGP Routes tab fields.

**Table 15** BGP Routes tab fields

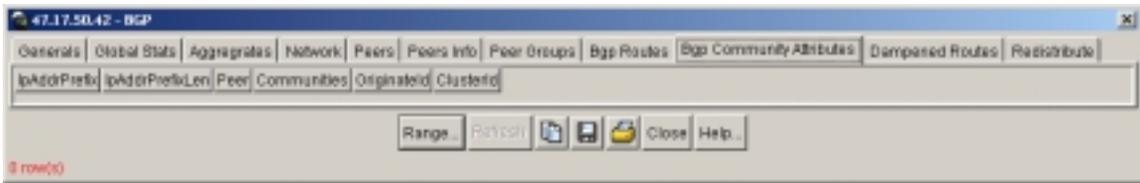
Field	Description
IpAddrPrefix	The IP address of the route destination (for example 2.2.2.2).
IpAddrPrefixLen	The mask of the route destination. Must be an integer value between 0 and 32 (for example /16).
Peer	The IP address of the peer where the path info was learned.
NextHop	The address of the border router that should be used for the destination network.
Origin	The ultimate origin of the path information: <ul style="list-style-type: none"> <li>• IGB = Networks are interior</li> <li>• EGB = Networks learned via EGP</li> <li>• Incomplete = Undetermined</li> </ul>
LocalPref	The originating BGP4 speaker's degree of preference for an advertised route. A value of -1 indicates the absence of this attribute.
ASPathSegment	Indicates the sequence of one or more AS path segments.
MultiExitDisc	This metric is used to discriminate between multiple exit points to an adjacent autonomous system. A value of -1 indicates the absence of this attribute.
AtomicAggregate	Specifies whether or not the local system has selected a less specific route, without selecting a more specific route.
AggregatorAS	The AS number of the last BGP4 speaker that performed route aggregation. A value of zero (0) indicates the absence of this attribute.
AggregatorAddr	The IP address of the last BGP4 speaker that performed route aggregation. A value of 0.0.0.0 indicates the absence of this attribute.
Best	Indicates whether or not this route was chosen as the best BGP4 route.
Unknown	One or more path attributes not understood by this BGP4 speaker. Size zero (0) indicates the absence of such attribute(s). Octets beyond the maximum size, if any, are not recorded by this object.

## BGP community attributes

To display information about BGP community attributes:

- 1 From the Device Manager menu bar, choose IP Routing > BGP.  
The BGP dialog box opens with the Generals tab displayed (see [Figure 16 on page 56](#)).
- 2 Click the BGP Community Attributes tab.  
The BGP Community Attributes tab opens ([Figure 34](#)).

**Figure 34** BGP dialog box—BGP Community Attributes tab (1 of 2)



- 3 Click Range.  
The BGP Route Range dialog box opens ([Figure 35](#)).

**Figure 35** BGP Route Range dialog box



- 4 Choose the type of information you want to display by selecting the options in the BGP Route Range dialog box.

[Table 16](#) describes the BGP Route Range dialog box options.

**Table 16** BGP Route Range dialog box options

Options	Description
Type	Specifies the route types to query: <ul style="list-style-type: none"><li>• All = all routes in the route table</li><li>• Exact Prefixes = only routes that exactly match the prefix/len</li><li>• Longer Prefix = all routes that match the prefix with longer masks.</li></ul>
IPAddrPrefix	The IP address of the route destination (for example 2.2.2.2).
IPAddrPrefixLen	The mask of the route destination. Must be an integer value between 0 and 32 (for example /16).

**5** Click Set Range.

**6** Click Close.

The Route Range dialog box closes.

**7** In the BGP Community Attributes tab, click Refresh.

The BGP Community Attributes tab refreshes and displays route information based on your selected options in the Route Range dialog box (Figure 36).



**Note:** Each time you choose new options from the Route Range dialog box, you must click Refresh in the BGP Community Attributes tab to allow the screen to refresh and display your new configuration choices.

---

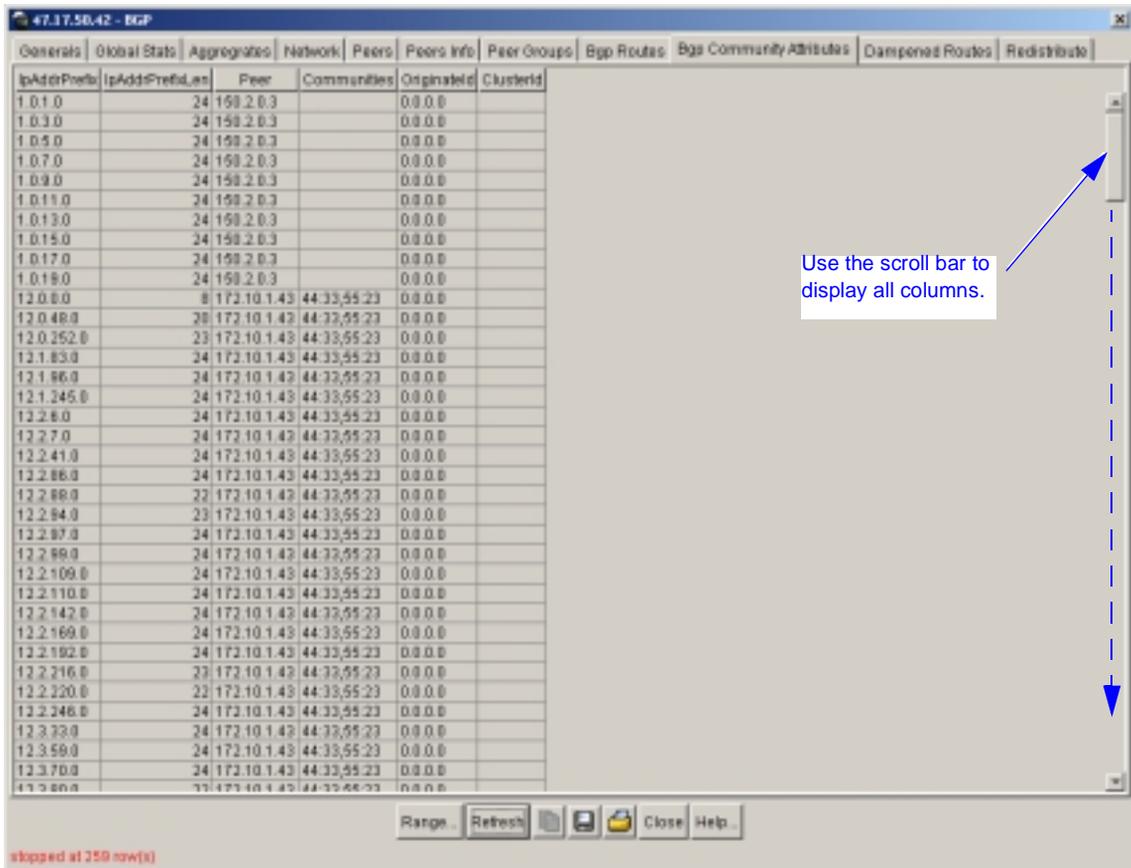
**Figure 36** BGP dialog box—BGP Community Attributes tab (2 of 2)

Table 17 describes the BGP Routes Info tab fields.

**Table 17** BGP Routes Info tab fields

Field	Description
IpAddrPrefix	The IP address of the route destination (for example 2.2.2.2).
IpAddrPrefixLen	The mask of the route destination. Must be an integer value between 0 and 32 (for example /16).
Peer	The IP address of the peer where the path info was learned.

**Table 17** BGP Routes Info tab fields (continued)

Field	Description
Communities	A string value that represents multiple community path attributes.
Originated	The IP address of the route's origin.
ClusterId	A string value that represents multiple cluster ID path attributes.

## Displaying dampened routes information

To display information about current receive paths:

- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed [Figure 16 on page 56](#).

- 2 Click the Dampened Routes tab.

The Dampened Routes tab opens and displays current dampened path information ([Figure 37](#)).

**Figure 37** BGP dialog box—Dampened Routes tab

IpAddr/Prefix	IpAddr/Prefix/Len	Peer	FlapPenalty	FlapCount	RouteDampened	ReuseTime
12.162.160.0	19	172.5.1.41	1024	1	announced	180
150.1.0.0	24	172.5.1.41	1024	1	announced	180
172.5.1.0	24	172.5.1.41	1024	1	announced	180

[Table 18](#) describes the Dampened Routes tab dialog box fields.

**Table 18** Dampened Routes tab dialog box

Field	Description
IpAddrPrefix	The IP address of the route destination (for example 2.2.2.2).
IpAddrPrefixLen	The mask of the route destination. Must be an integer value between 0 and 32 (for example /16).
Peer	The IP address of the peer where the path info was learned.
FlapPenalty	Penalty value based on number of route flaps.
FlapCount	Number of times a route flapped since the last time the penalty was reset to zero.
RouteDampened	Indicates whether this route is currently being suppressed or announced.
ReuseTime	A read-only value that allows a suppressed route to be added back to the routing table after the penalty value falls below this limit.

## Configuring redistribute entries

You can configure a redistribute entry for BGP to announce routes of a certain source type, for example, direct, static, RIP, and OSPF. If a route policy field is not configured for a redistribute entry, then the default action is taken on the basis of metric, metric-type, and subnet configured. This is called basic redistribution. Otherwise, you use the route policy specified to perform detailed redistribution. If no redistribution entry is configured, no external LSA is generated for non-BGP routes.

To configure a BGP redistribute entry:

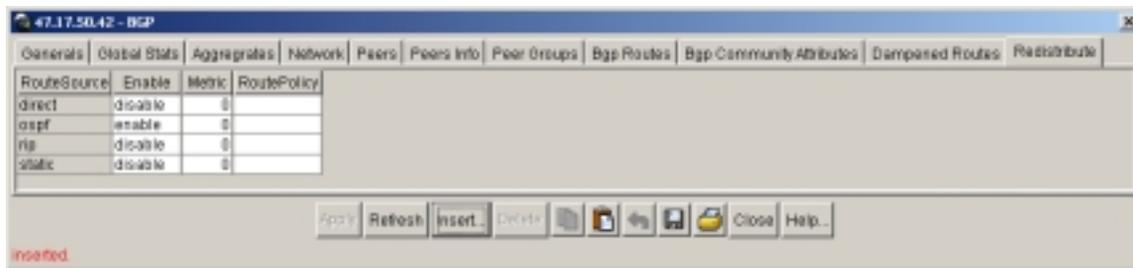
- 1 From the Device Manager menu bar, choose IP Routing > BGP.

The BGP dialog box opens with the Generals tab displayed (see [Figure 16 on page 56](#)).

- 2 Click the BGP Redistribute tab.

The BGP Redistribute tab opens ([Figure 38](#)).

**Figure 38** BGP dialog box—BGP Redistribute tab



3 Click Insert.

The BGP, Insert BGP Redistribute dialog box opens (Figure 39).

**Figure 39** BGP, Insert BGP Redistribute dialog box

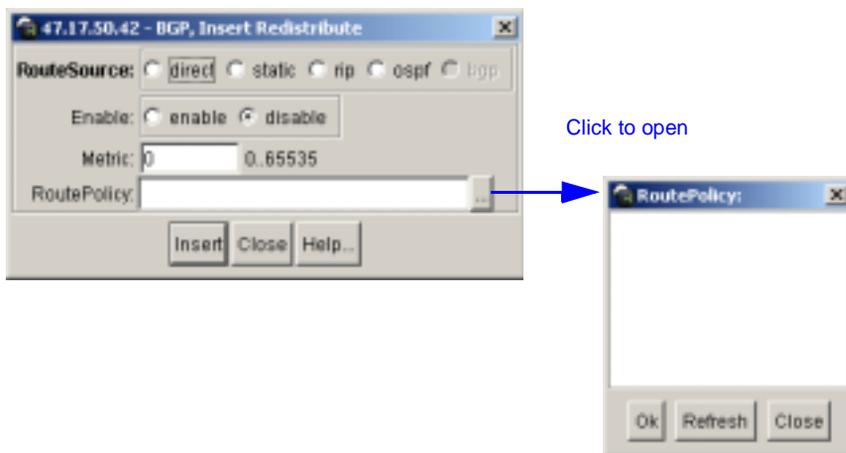


Table 19 describes the BGP, Insert BGP Redistribute dialog box fields.

**Table 19** BGP Insert BGP Redistribute dialog box fields

Field	Description
RouteSource	Select the route source protocol for the redistribution entry.
Enable	Enables (or disables) a BGP redistribute entry for a specified source type. You can also enable or disable this feature in the Redistribute tab of the BGP dialog box by clicking in the field and selecting enable or disable from the pull-down menu.

**Table 19** BGP Insert BGP Redistribute dialog box fields (continued) (continued)

Field	Description
Metric	<p>Sets a numerical metric value for the redistributed route. The value can be a range between 0 to 65535. The default value is 0.</p> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>If you do not specify a value for this option, and no value is specified in the <code>config ip bgp default-metric</code> command, the default metric value is 0. Nortel Networks* recommends that you use a value that is consistent with the destination protocol.</li> <li>When you apply the route-policy parameter (see next field description), the <i>metric-value</i> that is configured within that policy takes precedence over the <i>metric-value</i> setting for the metric parameter described here.</li> </ul> <p><i>metric-value</i> is an integer value between 0 and 65535.</p>
RoutePolicy	<p>Sets the route policy by name to be used for the detailed redistribution of external routes from a specified source into a BGP domain.</p> <p>Click the ellipse button and choose a policy (if configured) from the list in the Route Policy dialog box (see <a href="#">Figure 39 on page 102</a>). To deselect an entry, press [Ctrl] and click the left mouse button.</p>

## Configuring Circuitless IP

This section describes how to configure the circuitless IP feature and includes the following topics:



**Note:** You can configure a maximum of 32 circuitless IP interfaces on each device.

- [“Configuring a circuitless IP interface,”](#) next
- [“Enabling OSPF on a circuitless IP interface”](#) on page 106
- [“Deleting a circuitless IP interface”](#) on page 107
- [“Copying a circuitless IP interface”](#) on page 108
- [“Exporting circuitless IP interface data”](#) on page 108
- [“Specifying global parameters”](#) on page 109

For conceptual information about the circuitless IP feature, see [“Circuitless IP” on page 47](#).

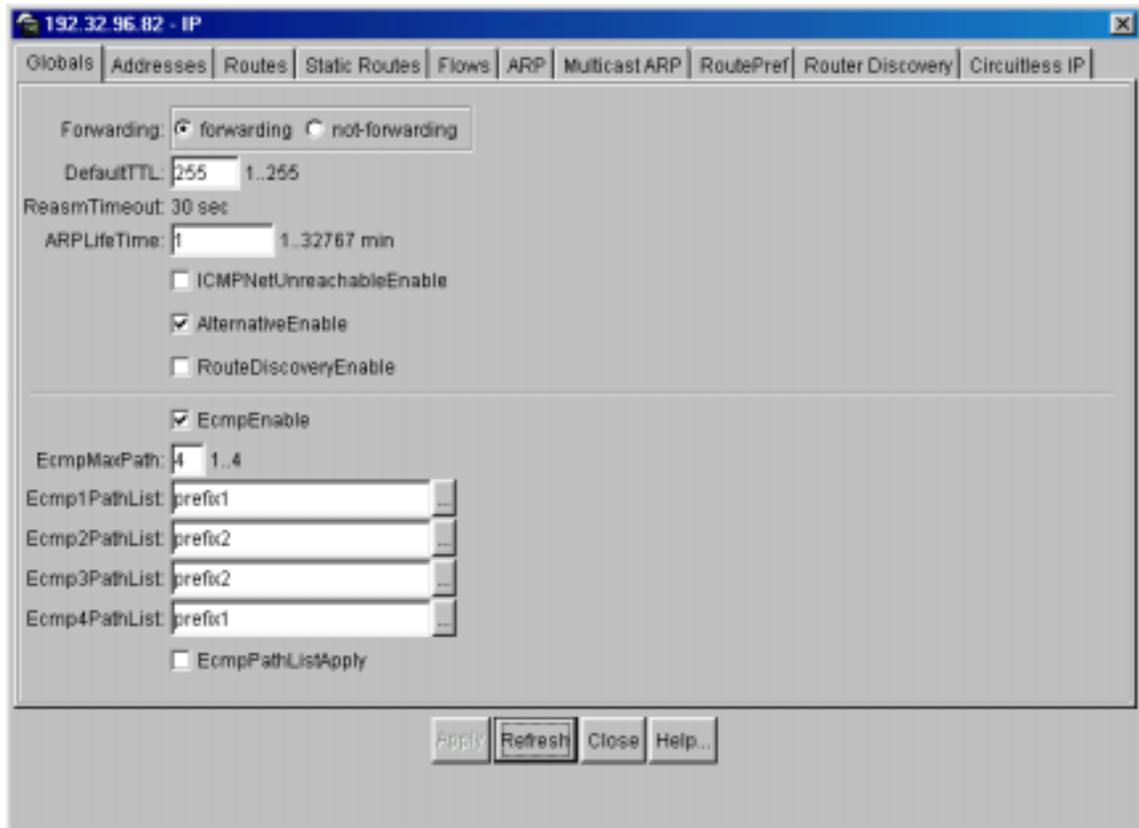
## Configuring a circuitless IP interface

To configure a circuitless IP interface:

- 1 From the Device Manager menu bar, choose IP routing > IP.

The IP dialog box opens with the Globals tab displayed ([Figure 40](#)).

**Figure 40** IP dialog box — Globals tab



- 2 Click the Circuitless IP tab.

The Circuitless IP tab opens ([Figure 41](#)).

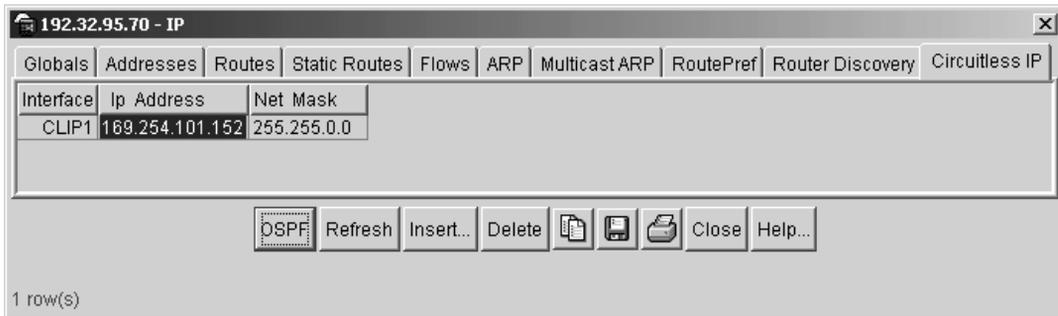
**Figure 41** IP dialog box — Circuitless IP tab

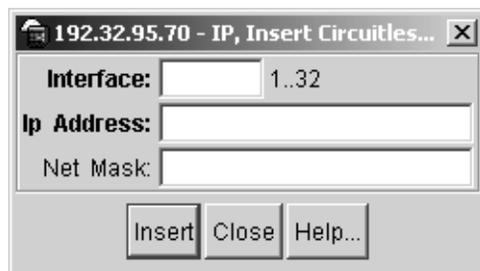
Table 20 describes the Circuitless IP tab fields.

**Table 20** IP dialog box—Circuitless IP tab fields

Field	Description
Interface	Displays the number assigned to the interface.
IP Address	Displays the IP address of the interface you are specifying as circuitless.
Net Mask	Displays the Net Mask address of the interface you are specifying as circuitless.

### 3 Click Insert.

The IP, Insert Circuitless dialog box opens (Figure 42).

**Figure 42** IP, Insert Circuitless dialog box

[Table 21](#) describes the IP, Insert Circuitless dialog box fields.

**Table 21** IP, Insert Circuitless dialog box fields

Field	Description
Interface	Number assigned to the interface. The range is 1...32.
IP Address	IP address of the interface you are specifying as circuitless.
Net Mask	Net Mask address of the interface you are specifying as circuitless.

- 4 Enter an integer value in the Interface field (in the range 1 and 32).
- 5 Enter the IP address.
- 6 Enter the network Mask.
- 7 Click Insert.

The new interface is created and appears in the Circuitless IP tab (see [Figure 41 on page 105](#)).



**Note:** Your interface selection appears in the Circuitless IP tab dialog box as CLIP1, CLIP2, etc.

---

## Enabling OSPF on a circuitless IP interface

To enable OSPF on an interface:

- 1 Select the interface (CLIP1, CLIP2, etc.) in the Circuitless IP tab dialog box.



**Note:** You must enable OSPF for circuitless IP to function.

---

- 2 Click OSPF.

The OspfCircuitless dialog box opens.

**Figure 43** OspfCircuitless dialog box

- 3 Click the Enable checkbox.
- 4 Click Apply to enable OSPF.



**Note:** When OSPF is enabled, the circuitless IP interface is configured to OSPF backbone AreaId (0.0.0.0) until you change the configuration.

- 5 To change the OSPF backbone AreaId:
  - a Choose IP Routing > OSPF from the Device Manager menu bar.
  - b Click the Interfaces tab.
  - c Click in the current AreaID field to make the change to the OSPF backbone area.
- 6 Close the dialog box.

## Deleting a circuitless IP interface

To delete a circuitless IP interface:

- 1 From the Device Manager menu bar, choose IP routing > IP.  
The IP dialog box opens with the Globals tab displayed (see [Figure 40 on page 104](#)).
- 2 Click the Circuitless IP tab.  
The Circuitless IP tab opens (see [Figure 41 on page 105](#)).

- 3 In the Interface column, select the CLIP number of the interface you want to delete.
- 4 Click Delete.  
The new interface is deleted from the list of interfaces.
- 5 Close the dialog box.

## Copying a circuitless IP interface

To copy a Circuitless IP interface:

- 1 From the Device Manager menu bar, choose IP routing > IP.
- 2 Click the Circuitless IP tab.  
The Circuitless IP tab opens (see [Figure 41 on page 105](#)).
- 3 In the Interface column, select the CLIP number of the interface you want to copy.
- 4 Click Copy.  
The interface is copied to the clipboard.
- 5 Create a text file.
- 6 Use Paste to paste the information in the text file.
- 7 Close the dialog box.

## Exporting circuitless IP interface data

To export circuitless IP interface data:

- 1 From the Device Manager menu bar, choose IP routing > IP.
- 2 Click the Circuitless IP tab.  
The Circuitless IP tab opens (see [Figure 41 on page 105](#)).
- 3 In the Interface column, select the appropriate interface CLIP number of the interface that contains the data you want to export.

- 4 Click Export.

The Export dialog box opens.

- 5 Specify a file name for the new export file data.

- 6 Specify a location where you want the new export file to be created.

- 7 Click OK.

The interface is deleted from the list of interfaces.

- 8 Click Close.

## Specifying global parameters

The Global tab displays default parameters that are in effect. You can change them, as desired.

To change default parameters that are currently in effect:

- 1 From the Device Manager menu bar, choose IP routing > IP.

The IP dialog box opens with the Globals tab displayed ([Figure 44](#)).

**Figure 44** IP dialog box—Globals tab

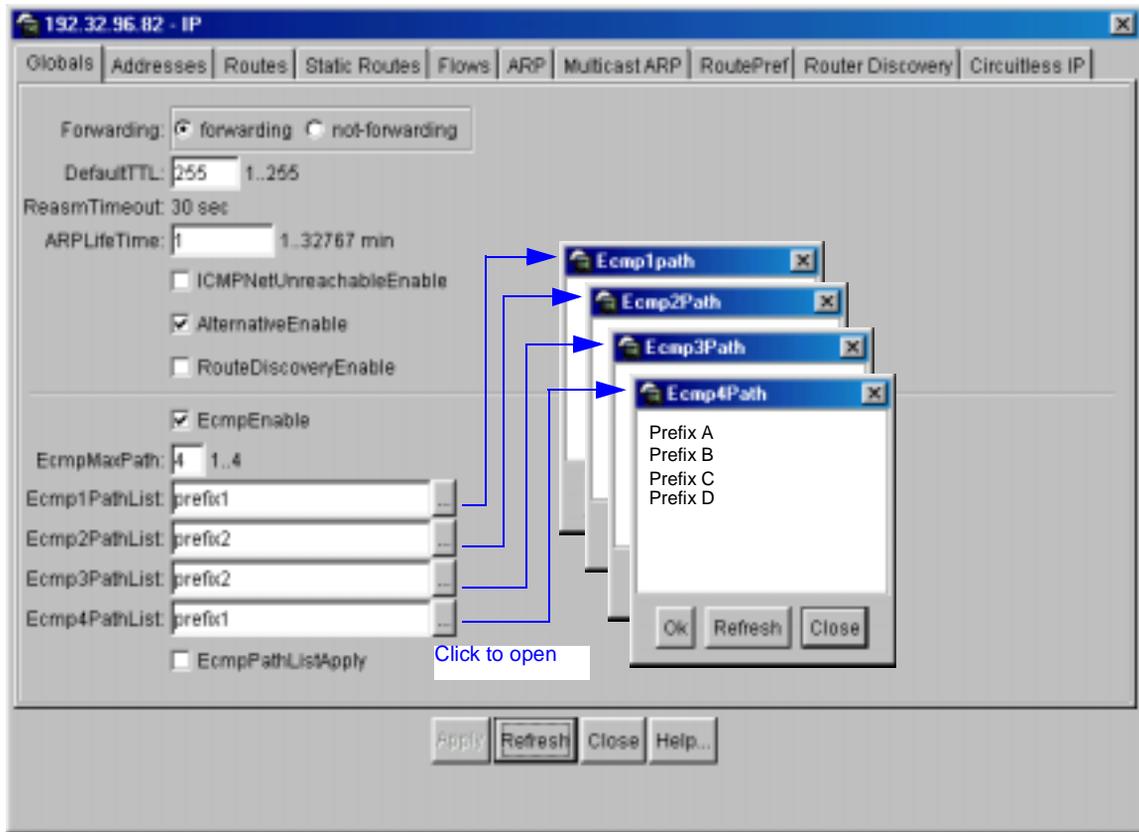


Table 22 describes the IP dialog box—Globals tab fields.

**Table 22** IP dialog box—Globals tab fields

Field	Description
Forwarding	Sets the switch for forwarding (routing) or non forwarding. The default value is forwarding.
DefaultTTL	Sets the default TTL value, the maximum number of seconds before a packet is discarded, for a routed packet. Enter an integer between 1 and 255. The default value (255) is inserted if one is not supplied in the datagram header.
ReasmTlmeout	Read-only field--The maximum number of seconds that received fragments are held while they are waiting for reassembly at this entity. The default value is 30 seconds.

**Table 22** IP dialog box—Globals tab fields (continued)

Field	Description
ARPLifeTime	The lifetime of an ARP entry within the system, global to the switch. The default value is 360 minutes. The range value is 1 through 32767 minutes.
ICMPNetUnreachableEnable	Enables or disables the generation of Internet Control Message Protocol (ICMP) net unreachable messages if the destination network is not reachable from this router. These messages assist in determining if the routing switch is reachable over the network. The default is value disable.
AlternateEnable	Enables or disables the alternative-route feature globally. For more information about alternative routes, see Networking Concepts for the Passport 8000 Series Switch. The default value is enable. Note: If the alternative-route parameter is disabled, all alternative routes are removed. When the parameter is enabled alternative routes are added back.
RouteDiscoveryEnable	Used to globally enables or disables router discovery. The default is disable.
EcmpEnable	Globally enables or disables the Equal Cost Multipath (ECMP) feature. The default is disabled. When ECMP is disabled, the EcmpMaxPath is reset to the default value of 1.
EcmpMaxPath	Globally configures the maximum number of ECMP paths. The range for this value is from 1 to 4. You cannot configure this field unless ECMP is enabled globally.
Ecmp1PathList	Specifies that the network routes associated with the configured prefix-list can have only one ECMP path, regardless of the EcmpMaxPath field value. Click the ellipse button and choose a prefix-list (if configured) from the list in the RoutePolicyIn dialog box (see <a href="#">Figure 44 on page 110</a> ).
Ecmp2PathList	Specifies that the network routes associated with the configured prefix-list can have up to two ECMP paths, regardless of the EcmpMaxPath field value. Click the ellipse button and choose a prefix-list (if configured) from the list in the RoutePolicyIn dialog box (see <a href="#">Figure 44 on page 110</a> ).
Ecmp3PathList	Specifies that the network routes associated with the configured prefix-list can have up to three ECMP paths, regardless of the EcmpMaxPath field value. Click the ellipse button and choose a prefix-list (if configured) from the list in the RoutePolicyIn dialog box (see <a href="#">Figure 44 on page 110</a> ).
Ecmp4PathList	Specifies that the network routes associated with the configured prefix-list can have up to four ECMP paths, regardless of the EcmpMaxPath field value. Click the ellipse button and choose a prefix-list (if configured) from the list in the RoutePolicyIn dialog box (see <a href="#">Figure 44 on page 110</a> ).
EcmpPathListApply	Applies any changes in the ECMP path list configuration or in any prefix-lists that are configured to be used as a path list.



---

## Chapter 4

# Using the CLI to configure BGP

---

This chapter describes how to configure BGP using the command line interface (CLI).

- For conceptual information about BGP, see [Chapter 1, “BGP concepts,”](#) on page 21.
- For information about limitations and interoperability issues, see [Chapter 2, “Configuration considerations and limitations,”](#) on page 49.
- For configuration examples, including the required CLI commands, see [Chapter 5, “Configuration examples,”](#) on page 195.

This chapter includes the following topics:

Topic	Page
<a href="#">Roadmap of BGP CLI commands</a>	114
<a href="#">Configuring general BGP parameters</a>	120
<a href="#">Configuring BGP confederations</a>	128
<a href="#">Configuring BGP peers or peer groups</a>	129
<a href="#">Configuring route redistribution parameters</a>	136
<a href="#">Configuring AS-path lists</a>	142
<a href="#">Configuring community lists</a>	144
<a href="#">Showing BGP configurations</a>	146
<a href="#">Circuitless IP</a>	173
<a href="#">Configuring BGP debug commands</a>	177

## Roadmap of BGP CLI commands

The following roadmap lists all the BGP commands and their parameters. Use this list as a quick reference or click on any entry for more information:

Command	Parameter
<code>config ip bgp</code>	<code>info</code> <code>aggregate-address &lt;prefix/len&gt;</code> <code>&lt;add del&gt; [as-set &lt;value&gt;]</code> <code>[summary-only &lt;value&gt;]</code> <code>[suppress-map &lt;value&gt;]</code> <code>[advertise-map &lt;value&gt;]</code> <code>[attribute-map &lt;value&gt;]</code> <code>aggregation &lt;enable disable&gt;</code> <code>always-cmp-med &lt;enable disable&gt;</code> <code>auto-peer-restart &lt;enable disable&gt;</code> <code>auto-summary &lt;enable disable&gt;</code> <code>cl-to-cl-reflection</code> <code>&lt;enable disable&gt;</code> <code>cluster-id &lt;ipaddr&gt; &lt;add del&gt;</code> <code>comp-bestpath-med-confed</code> <code>&lt;enable disable&gt;</code> <code>default-local-pref &lt;value&gt;</code> <code>&lt;add del&gt;</code> <code>default-metric &lt;value&gt; &lt;add del&gt;</code> <code>disable</code> <code>enable</code> <code>flap-dampening &lt;enable disable&gt;</code> <code>global-debug mask &lt;value&gt;</code> <code>ibgp-report-import-rt</code> <code>&lt;enable disable&gt;</code> <code>ignore-illegal-rtrid</code> <code>&lt;enable disable&gt;</code> <code>local-as &lt;asnum&gt;</code>

<b>Command</b>	<b>Parameter</b>
	max-equalcost-routes <value> <enable disable>
	neighbor-debug-all mask <value>
	network <prefix/len> <add del>
	no-med-path-is-worst <enable disable>
	orig-def-route <enable disable>
	restart
	route-reflection <enable disable>
	stats-clear
	synchronization <enable disable>
	debug-screen [<setting>]
config ip bgp confederation	info
	identifier <value> <add del>
	peers <as-list>
config ip bgp neighbor <nbr_ipaddr peer-group-name>	info
	create
	delete
	remote-as <asnum>
	connect-retry-interval <value> <add del>
	admin-state <enable disable>
	ebgp-multihop <enable disable>
	hold-time <value> <add del>
	keepalive-time <value> <add del>
	max-prefix <value> <add del>
	nexthop-self <enable disable>
	originate-def-route <enable disable>

Command	Parameter
	password <password> <add del>
	MD5-authentication <enable disable>
	neighbor-debug mask <value>
	peer-group <peer-group-name> <add del>
	remove-private-as <enable disable>
	restart [soft-reconfiguration <value>]
	route-advertisement- interval <value> <add del>
	route-policy <in out> <route-map name> <add del>
	route-reflector-client <enable disable>
	send-community <enable disable>
	soft-reconfiguration-in <enable disable>
	stats-clear
	update-source-interface <ipaddr> <add del>
	weight <value> <add del>
config ip bgp redistribute apply	
config ip bgp redistribute ospf	info
	apply
	create
	disable
	delete
	enable
	metric <metric-value>
	route-policy <policy name>

<b>Command</b>	<b>Parameter</b>
<code>config ip bgp redistribute direct</code>	<code>info</code> <code>apply</code> <code>create</code> <code>disable</code> <code>delete</code> <code>enable</code> <code>metric &lt;metric-value&gt;</code> <code>route-policy &lt;policy name&gt;</code>
<code>config ip bgp redistribute rip</code>	<code>info</code> <code>apply</code> <code>create</code> <code>disable</code> <code>delete</code> <code>enable</code> <code>metric &lt;metric-value&gt;</code> <code>route-policy &lt;policy name&gt;</code>
<code>config ip bgp redistribute static</code>	<code>info</code> <code>apply</code> <code>create</code> <code>disable</code> <code>delete</code> <code>enable</code> <code>metric &lt;metric-value&gt;</code> <code>route-policy &lt;policy name&gt;</code>

Command	Parameter
<code>config ip as-list &lt;listid&gt;</code>	<code>info</code> <code>create</code> <code>delete</code> <code>add-as-path &lt;memberId&gt;</code> <code>&lt;permit deny&gt; &lt;as path&gt;</code> <code>remove-as-path [memberId &lt;value&gt;]</code> <code>[as-path &lt;value&gt;]</code>
<code>config ip community-list</code> <code>&lt;listid&gt;</code>	<code>info</code> <code>create</code> <code>delete</code> <code>add-community memberId</code> <code>&lt;permit deny&gt; community-string</code> <code>remove-community [memberId &lt;value&gt;]</code> <code>[community-string &lt;value&gt;]</code>
<code>show ip bgp aggregates [&lt;prefix/</code> <code>len&gt;]</code>	
<code>show ip bgp cidr-only [exact</code> <code>&lt;value&gt;]</code>	
<code>show ip bgp conf</code>	
<code>show ip bgp dampened-paths</code> <code>&lt;ipaddr&gt; [&lt;prefix&gt;]</code> <code>[longer-prefixes]</code>	
<code>show ip bgp flap-damp-config</code>	
<code>show ip bgp redistribution</code>	
<code>show ip bgp imported-routes</code>	
<code>show ip bgp networks</code>	
<code>show ip bgp peer-group</code> <code>[&lt;peer-group name&gt;]</code>	
<code>show ip bgp route [&lt;prefix&gt;]</code> <code>[longer-prefixes] [community</code> <code>&lt;value&gt;]</code>	

**Command****Parameter**

```
show ip bgp stats
show ip bgp summary
show ip bgp neighbor info
[<ipaddr>]
show ip bgp neighbor stats
<ipaddr>
show ip bgp neighbor route
<ipaddr> [<prefix>]
[longer-prefixes] [community
<value>]
```

## Configuring general BGP parameters

To configure general BGP parameters, use the following command:

```
config ip bgp
```

This command includes the following options:

<b>config ip bgp</b> followed by:	
info	Displays the current BGP system configuration (Figure 45).
<pre>aggregate-address &lt;prefix/len&gt; &lt;add del&gt;  [as-set &lt;value&gt;] [summary-only &lt;value&gt;] [suppress-map &lt;value&gt;] [advertise-map &lt;value&gt;] [attribute-map &lt;value&gt;]</pre>	<p>Adds or deletes an aggregate address in a BGP routing table. Because the routes in the table are only aggregated to EBGP peers, routing traffic is minimized.</p> <p><b>Note:</b> If you add or modify the <i>suppress-map</i>, <i>advertise-map</i>, or <i>attribute-map</i> options, be sure that the route-policy for those options are already configured.</p> <ul style="list-style-type: none"> <li>• <i>prefix/len</i> is an ip address and an integer value (between 0 and 32)</li> <li>• <i>add del</i> adds or deletes the entry</li> </ul> <p><b>Optional values:</b></p> <ul style="list-style-type: none"> <li>• <i>as-set &lt;value&gt;</i> enables or disables autonomous system information. The default value is disable.</li> <li>• <i>summary-only &lt;value&gt;</i> enables or disables the summarization of routes not included in routing updates. This parameter creates the aggregate route and suppresses advertisements of more specific routes to all neighbors. The default value is disable.</li> <li>• <i>suppress-map &lt;value&gt;</i> is the route map name (string length between 0 and 64 characters long) for the suppressed route list.</li> <li>• <i>advertise-map &lt;value&gt;</i> is the route map name (any string length between 0 and 64 characters long) for route advertisements.</li> <li>• <i>attribute-map &lt;value&gt;</i> is the route map name (string length between 0 and 64 characters long).</li> </ul>

<b>config ip bgp</b> followed by:	
aggregation <enable   disable>	Enables or disables the aggregation feature on this interface. The default value is enable. <b>Note:</b> You cannot change the default value when BGP is enabled.
always-cmp-med <enable   disable>	Enables or disables the comparison of the multi-exit discriminator (MED) parameter for paths from neighbors in different autonomous systems. A path with a lower MED is preferred over a path with a higher MED. The default value is disable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the option.</li> </ul> <b>Note:</b> When this option is set to disable (the default value) during the best-path selection process, the MEDs are compared only among paths from the same autonomous system. If you enable this option, the MEDs are compared among paths received from any other autonomous systems.
auto-peer-restart <enable   disable>	Enables or disables the process that automatically restarts a connection to a BGP neighbor. The default value is enable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the option.</li> </ul>
auto-summary <enable   disable>	When enabled, allows BGP to summarize networks based on class limits (For example, Class A, B, C networks). The default value is enable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the option.</li> </ul>
cl-to-cl-reflection <enable   disable>	Enables or disables route reflection between two route reflector clients. This option is applicable only if the route reflection value is set to enable. The default value is enable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the option.</li> </ul> <b>Note:</b> Route reflection may be enabled even when clients are fully meshed. In this event, route reflection is not required.
cluster-id <ipaddr> <add   del>	Sets a cluster ID. This option is applicable only if the route reflection value is set to enable, and if multiple route reflectors are in a cluster. <ul style="list-style-type: none"> <li>• <i>ipaddr</i> is the cluster ID of the reflector router.</li> <li>• <i>add   del</i> adds or deletes the cluster ID.</li> </ul>

<b>config ip bgp</b> followed by:	
<code>comp-bestpath-med-confed</code> <code>&lt;enable   disable&gt;</code>	When enabled, allows you to compare multi-exit discriminator (MED) attributes within a confederation. The default value is <code>disable</code> . <ul style="list-style-type: none"> <li><code>enable   disable</code> enables or disables the option.</li> </ul>
<code>default-local-pref</code> <code>&lt;value&gt; &lt;add   del&gt;</code>	Specifies the default value of the local preference attribute. The default value is 0. <ul style="list-style-type: none"> <li><code>value</code> is an integer value between 0 and 2147483647.</li> <li><code>add   del</code> adds or deletes the configuration.</li> </ul> <b>Note:</b> You cannot change the default value when BGP is enabled.
<code>default-metric</code> <code>&lt;value&gt; &lt;add   del&gt;</code>	Sets a value that is sent to a BGP neighbor to determine the cost of a route a neighbor is using. This option must be used in conjunction with the redistribute router configuration command to allow the current routing protocol to use the same metric value for all redistributed routes. The default value is 0. <ul style="list-style-type: none"> <li><code>value</code> is an integer value between -1 and 2147483647</li> <li><code>add   del</code> adds or deletes the configuration.</li> </ul> <b>Note:</b> A default metric value helps solve the problems associated with redistributing routes that have incompatible metrics. For example, whenever metrics do not convert, using a default metric provides a reasonable substitute and allows the redistribution to proceed.
<code>disable</code>	Disables BGP.
<code>enable</code>	Enables BGP. <b>Note:</b> You cannot enable BGP until you change the LocalAS value to any value other than 0.
<code>flap-dampening</code> <code>&lt;enable   disable&gt;</code>	Enables or disables route suppression for routes that flap on and off. The default value is <code>disable</code> . <ul style="list-style-type: none"> <li><code>enable   disable</code> enables or disables the option.</li> </ul>

<b>config ip bgp</b> followed by:	
<pre>global-debug mask &lt;value&gt;</pre>	<p>Allows you to display specified debug information for BGP global configuration. The default value is none.</p> <ul style="list-style-type: none"> <li><i>value</i> is a list of mask choices that you enter (separated by comma's with no space between choices).</li> </ul> <p>For example:  <code>{&lt;mask&gt;, &lt;mask&gt;, &lt;mask&gt;...}</code> where  <i>mask</i>:  <code>{none, all, error, packet, event, trace, warning, state, init, filter, update}</code>.</p> <p>Mask choices are:</p> <p><i>none</i>: disables all debug messages.</p> <p><i>all</i>: enables all debug messages.</p> <p><i>error</i>: enables display of debug error messages.</p> <p><i>packet</i>: enables display of debug packet messages.</p> <p><i>event</i>: enables display of debug event messages.</p> <p><i>trace</i>: enables display of debug trace messages.</p> <p><i>warning</i>: enables display of debug warning messages.</p> <p><i>state</i>: enables display of debug state transition messages.</p> <p><i>init</i>: enables display of debug initialization messages.</p> <p><i>filter</i>: enables display of debug messages related to filtering.</p> <p><i>update</i>: enables display of debug messages related to sending and receiving updates.</p>
<pre>ibgp-report-import-rt &lt;enable   disable&gt;</pre>	<p>Configures BGP to report imported routes to an interior BGP (IBGP) peer. This command also enables or disables reporting of non-BGP imported routes to other IBGP neighbors. The default value is enable.</p> <ul style="list-style-type: none"> <li><i>enable   disable</i> enables or disables the option.</li> </ul>

<b>config ip bgp</b> followed by:	
<code>ignore-illegal-rtrid</code> <code>&lt;enable disable&gt;</code>	<p>When enabled, allows BGP to overlook an illegal router ID. For example, you can set this command to enable or disable the acceptance of a connection from a peer that sends an open message using a router ID of 0 (zero). The default value is enable.</p> <ul style="list-style-type: none"> <li>• <code>enable disable</code> enables or disables the option.</li> </ul>
<code>local-as &lt;asnum&gt;</code>	<p>Sets a local autonomous system number on the current system. The default value is 0.</p> <ul style="list-style-type: none"> <li>• <code>asnum</code> is an integer value between 0 and 65535.</li> </ul> <p><b>Note:</b> You cannot enable BGP until you change the LocalAS value to any value other than 0. You cannot change configured values when BGP is enabled.</p>
<code>max-equalcost-routes</code> <code>&lt;value&gt; &lt;enable disable&gt;</code>	<p>Sets the maximum number of equal-cost-paths that are available to a BGP router by limiting the number of equal-cost-paths that can be stored in the routing table. The default value is enable.</p> <ul style="list-style-type: none"> <li>• <code>value</code> is the max-equalcost-routes allowed. the range is 1 to 4 equal-cost-paths. The default value is 1.</li> <li>• <code>enable disable</code> enables or disables the option.</li> </ul>

<b>config ip bgp</b>	
followed by:	
neighbor-debug-all mask <value>	<p>Allows you to display specified debug information for BGP neighbors. The default value is none.</p> <ul style="list-style-type: none"> <li><i>value</i> is a list of mask choices that you enter (separated by comma's with no space between choices).</li> </ul> <p>For example:  {&lt;mask&gt;, &lt;mask&gt;, &lt;mask&gt;...} where  <i>mask</i>:  {none, all, error, packet, event, trace, warning, state, init, filter, update}.</p> <p>Mask choices are:</p> <p><i>none</i>: disables all debug messages.</p> <p><i>all</i>: enables all debug messages.</p> <p><i>error</i>: enables display of debug error messages.</p> <p><i>packet</i>: enables display of debug packet messages.</p> <p><i>event</i>: enables display of debug event messages.</p> <p><i>trace</i>: enables display of debug trace messages.</p> <p><i>warning</i>: enables display of debug warning messages.</p> <p><i>state</i>: enables display of debug state transition messages.</p> <p><i>init</i>: enables display of debug initialization messages.</p> <p><i>filter</i>: enables display of debug messages related to filtering.</p> <p><i>update</i>: enables display of debug messages related to sending and receiving updates.</p>
network <prefix/len> <add del>	<p>Specifies IGP network prefixes for BGP to advertise for redistribution.</p> <ul style="list-style-type: none"> <li><i>prefix/len</i> is the network address and mask.</li> <li><i>add del</i> adds or deletes the configuration.</li> </ul>
no-med-path-is-worst <enable disable>	<p>Enables or disables BGP from treating an update without a multi-exit discriminator (MED) attribute as the worst path. The default value is disable.</p> <ul style="list-style-type: none"> <li><i>enable disable</i> enables or disables the feature.</li> </ul>

<b>config ip bgp</b> followed by:	
<code>orig-def-route</code> <code>&lt;enable   disable&gt;</code>	Allows the redistribution of network 0.0.0.0 into BGP. The default value is disable. <ul style="list-style-type: none"> <li><code>enable   disable</code> enables or disables the feature.</li> </ul>
<code>restart</code>	Restarts BGP on the current system.
<code>route-reflection</code> <code>&lt;enable   disable&gt;</code>	Enables or disables the reflection of routes from IBGP neighbors. The default value is enable. <ul style="list-style-type: none"> <li><code>enable   disable</code> enables or disables the feature.</li> </ul>
<code>stats-clear</code>	Resets all displayed counters to 0 (zero).
<code>synchronization</code> <code>&lt;enable   disable&gt;</code>	Enables or disables the router from accepting routes from BGP peers without waiting for an update from the IGP. The default value is enable. <ul style="list-style-type: none"> <li><code>enable   disable</code> enables or disables the feature.</li> </ul>
<code>debug-screen [&lt;setting&gt;]</code>	Lets you display debug messages on your console, or to save them in a log file. <ul style="list-style-type: none"> <li><code>setting</code> is either of two values {off   on} that you enter to disable BGP screen logging (off), or to enable BGP screen logging (on).</li> </ul>

Figure 45 shows `config ip bgp info` command output.

**Figure 45** config ip bgp info command output

```
Passport-8610:5# config ip bgp info

BGP version - 4
local-as - 2
Identifier - 192.32.96.18

                                BGP on/off - ON
                                local-as - 2
                                aggregation - enable
                                always-cmp-med - disable
                                auto-peer-restart - enable
                                auto-summary - enable
                                comp-bestpath-med-confed - disable
                                default-local-preference - 100
                                default-metric - 0
                                flap-dampening - disable
                                global-debug - none
                                ibgp-report-import-rt - enable
                                ignore-illegal-rtrid - disable
                                max-equalcost-routes - 1
                                no-med-path-is-worst - enable
                                orig-def-route - disable
                                synchronization - enable
                                route-reflection - disable
                                cluster-id - 0.0.0.0
                                cl-to-cl-reflection - disable
                                decision state - Idle
                                confederation identifier - 0
```

## Configuring BGP confederations

To configure BGP confederations, use the following command:

```
config ip bgp confederation
```

This command includes the following options.

<b>config ip bgp confederation</b> followed by:	
<code>info</code>	Displays current config ip bgp confederation info command output ( <a href="#">Figure 46</a> ).
<code>identifier &lt;value&gt;</code> <code>&lt;add del&gt;</code>	Specifies a BGP confederation identifier. <ul style="list-style-type: none"> <li><code>value</code> is an integer value between 0 and 65535.</li> <li><code>add del</code> adds or deletes the configuration.</li> </ul> <p><b>Note:</b> You cannot configure this option when BGP is enabled.</p>
<code>peers &lt;as-list&gt;</code>	Adds or deletes other autonomous systems to a confederation. <ul style="list-style-type: none"> <li><code>as-list</code> is an integer string length value between 0 and 256 that represents the AS number.</li> </ul> <p>Enter one or more as numbers within quotes (for example, "20 30 40 50") to add AS systems. Enter another set of numbers within quotes, or a single number within quotes to replace the previous configuration. Enter only a set of quotes (or 0) to delete the configuration.</p>

[Figure 46](#) shows `config ip bgp confederation info` command output.

**Figure 46** config ip bgp confederation info command output

```
Passport-8610:5# config ip bgp confederation info
confederation identifier 35
confederation peer as 50 40 30 20
```

## Configuring BGP peers or peer groups

As is often the case with BGP speakers, many neighbors are configured with similar update policies (for example, many neighbors use the same distribute lists, filter lists, outbound route maps, update source, and so on). You can group the neighbors that have the same update policies into *peer groups* and *peer* associations. This association and grouping allows you to simplify your configurations and makes updates more efficient.

To configure BGP peers or peer groups, use the following neighbor command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>
```

where:

***nbr\_ipaddr|peer-group-name*** indicates that you enter the peer's IP address or the peer's group name.

This command includes the following options.

<b>config ip bgp neighbor &lt;nbr_ipaddr peer-group-name&gt;</b> followed by:	
info	Displays current configuration information of a BGP peer or peer-group ( <a href="#">Figure 47</a> ).
create	Creates a BGP peer or a peer-group.
delete	Deletes a BGP peer or a peer-group.
remote-as <asnum>	Configures a remote-as for a BGP peer or a peer-group. <ul style="list-style-type: none"> <li><i>asnum</i> is an integer value between 0 and 65535.</li> </ul> <p><b>Note:</b> You cannot configure this option when BGP peer is enabled.</p>
connect-retry-interval <value> <add del>	Sets the time interval (in seconds) for the ConnectRetry Timer. The default value is 120 seconds. <ul style="list-style-type: none"> <li><i>value</i> is an integer value between 1 and 65535 seconds.</li> <li><i>add del</i> adds or deletes the configuration.</li> </ul>

<b>config ip bgp neighbor &lt;nbr_ipaddr   peer-group-name&gt;</b> followed by:	
admin-state <enable   disable>	Enables or disables the administrative state of a BGP peer. The default value is disable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the administrative state.</li> </ul>
ebgp-multihop <enable   disable>	Allows you to enable or disable a connection to a BGP peer that is more than one hop away from the local router. The default value is disable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the connection.</li> </ul>
hold-time <value> <add   del>	Sets the time interval (in seconds) for the BGP speaker for this peer. The default value is 180 seconds. <ul style="list-style-type: none"> <li>• <i>value</i> is an integer value with a range 0, 3 to 180 seconds (the integer values 1 and 2 are not valid).</li> <li>• <i>add   del</i> adds or deletes the configuration.</li> </ul>
keepalive-time <value> <add   del>	Specifies the time interval (in seconds) that transpires between transmissions of the local BGP router's keep-alive packets. The keep-alive packets indicate the enabled status of the local BGP router to peers. The default value is 60 seconds. <ul style="list-style-type: none"> <li>• <i>value</i> is an integer value between 0 and 21845 seconds.</li> <li>• <i>add   del</i> adds or deletes the configuration.</li> </ul>
max-prefix <value> <add   del>	Sets a limit on the number of routes that can be accepted from a neighbor. The default value is 12,000 routes (see <a href="#">"Specifying number of routes learned" on page 201</a> ). <ul style="list-style-type: none"> <li>• <i>value</i> is an integer value between 0 and 2147483647 routes.</li> <li>• <i>add   del</i> adds or deletes the configuration.</li> </ul> <p><b>Note:</b> A value of 0 (zero) indicates that there is no limit to the number of routes that can be accepted.</p>
nexthop-self <enable   disable>	When enabled, specifies that the next-hop attribute in an IBGP update is the address of the local router or the router that is generating the IBGP update. The default value is disable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the feature.</li> </ul> <p><b>Note:</b> This feature can only be configured when the neighbor option is disabled.</p>

<b>config ip bgp neighbor &lt;nbr_ipaddr   peer-group-name&gt;</b> followed by:	
<code>originate-def-route</code> <code>&lt;enable   disable&gt;</code>	When enabled, specifies that the current route originated from the BGP peer. This command enables or disables sending the default route information to the specified neighbor or peer. The default value is <code>disable</code> . <ul style="list-style-type: none"> <li>• <code>enable   disable</code> enables or disables the feature.</li> </ul>
<code>password &lt;password&gt;</code> <code>&lt;add   del&gt;</code>	Allows you to configure a password for TCP MD5 authentication between two peers. <ul style="list-style-type: none"> <li>• <code>password</code> is an alphanumeric string length from 0 to 1536 characters.</li> <li>• <code>add   del</code> adds or deletes the password.</li> </ul>
<code>MD5-authentication</code> <code>&lt;enable   disable&gt;</code>	Enables or disables TCP MD5 authentication between two peers. The default value is <code>disable</code> . <ul style="list-style-type: none"> <li>• <code>enable   disable</code> enables or disables the feature.</li> </ul>

<b>config ip bgp neighbor &lt;nbr_ipaddr   peer-group-name&gt;</b> followed by:	
<pre>neighbor-debug mask &lt;value&gt;</pre>	<p>Allows you to display specified debug information for a BGP peer. The default value is none.</p> <ul style="list-style-type: none"> <li><i>value</i> is a list of mask choices that you enter (separated by comma's with no space between choices).</li> </ul> <p>For example:    {&lt;mask&gt; , &lt;mask&gt; , &lt;mask&gt; . . . } where  <i>mask</i>:    {none , all , error , packet , event ,    trace , warning , state , init ,    filter , update }.</p> <p>Mask choices are:</p> <p><i>none</i>: disables all debug messages.</p> <p><i>all</i>: enables all debug messages.</p> <p><i>error</i>: enables display of debug error messages.</p> <p><i>packet</i>: enables display of debug packet messages.</p> <p><i>event</i>: enables display of debug event messages.</p> <p><i>trace</i>: enables display of debug trace messages.</p> <p><i>warning</i>: enables display of debug warning messages.</p> <p><i>state</i>: enables display of debug state transition messages.</p> <p><i>init</i>: enables display of debug initialization messages.</p> <p><i>filter</i>: enables display of debug messages related to filtering.</p> <p><i>update</i>: enables display of debug messages related to sending and receiving updates.</p>
<pre>peer-group &lt;peer-group-name&gt; &lt;add   del&gt;</pre>	<p>Adds a BGP peer to the specified subscriber group.</p> <ul style="list-style-type: none"> <li><i>peer-group-name</i> is a string length from 0 to 1536 characters.</li> <li><i>add   del</i> adds or deletes the configuration.</li> </ul> <p><b>Note:</b> You must create the specified subscriber group before you issue this command.</p>

<b>config ip bgp neighbor &lt;nbr_ipaddr   peer-group-name&gt;</b> followed by:	
remove-private-as <enable   disable>	When enabled, allows you to strip private AS numbers when sending an update. This feature is especially useful within a confederation. The default value is enable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the feature.</li> </ul>
restart [soft-reconfiguration <value>]	Allows you to reset a BGP inbound or outbound (or both inbound and outbound) connection using BGP soft reconfiguration. <ul style="list-style-type: none"> <li>• <i>soft-reconfiguration value</i> allows you to choose either inbound or outbound soft configuration {<i>in   out</i>}.</li> </ul> <p><b>Note:</b> If no {<i>in   out</i>} choice is made, both inbound and outbound soft configurations are triggered.</p>
route-advertisement-interval <value> <add   del>	Specifies the time interval (in seconds) that transpires between each transmission of an advertisement from a BGP neighbor. The default value is 5 seconds. <ul style="list-style-type: none"> <li>• <i>value</i> is an integer value between 5 and 120 seconds.</li> <li>• <i>add   del</i> adds or deletes the configuration.</li> </ul>
route-policy <in   out> <route-map name> <add   del>	Applies a route policy rule to all routes that are learned from, or sent to, the local BGP router's peers, or peer groups. <p><b>Note:</b> The local BGP router is the BGP router that allows or disallows routes and sets attributes in incoming or outgoing updates.</p> <ul style="list-style-type: none"> <li>• <i>in   out</i> indicates whether the route is incoming or outgoing.</li> <li>• <i>route-map name</i> is an alphanumeric string length (0 to 256 characters) that indicates the name of the route map.</li> <li>• <i>add   del</i> adds or deletes the configuration.</li> </ul>
route-reflector-client <enable   disable>	Configures the specified neighbor or group of neighbors as its route reflector client. The default value is disable. <p><b>Note:</b> All neighbors that are configured become members of the client group and the remaining IBGP peers become members of the nonclient group for the local route reflector.</p> <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the feature.</li> </ul>

<b>config ip bgp neighbor &lt;nbr_ipaddr   peer-group-name&gt;</b> followed by:	
send-community <enable   disable>	Enables or disables sending the update message's community attribute to the specified peer. The default value is disable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the feature.</li> </ul>
soft-reconfiguration-in <enable   disable>	Allows the router to relearn routes from the specified neighbor or group of neighbors without resetting the connection when the policy changes in the inbound direction. The default value is disable. <ul style="list-style-type: none"> <li>• <i>enable   disable</i> enables or disables the feature.</li> </ul>
stats-clear	Resets all displayed counters to 0 (zero).
update-source-interface <ipaddr> <add   del>	Specifies the source IP address when sending BGP packets to this peer or peer group. <p><b>Note:</b> You cannot configure this option when BGP peer is enabled.</p> <ul style="list-style-type: none"> <li>• <i>ipaddr</i> is the specified source IP address.</li> <li>• <i>add   del</i> adds or deletes the configuration.</li> </ul>
weight <value> <add   del>	Specifies the weight of a BGP peer or peer groups, or the priority of updates that can be received from that BGP peer. The default value is 0. <p><b>Note:</b> A weight is a numerical value you can assign a path so that you can control the path selection process. The administrative weight is local to the router. If you have particular neighbors that you want to prefer for most of your traffic, you can assign a higher weight to all routes learned from that neighbor.</p> <ul style="list-style-type: none"> <li>• <i>value</i> is an integer value between 0 and 65535.</li> <li>• <i>add   del</i> adds or deletes the configuration.</li> </ul>

Figure 47 shows **config ip bgp neighbor <nbr\_ipaddr | peer-group-name>** command output.

**Figure 47** config ip bgp neighbor info command output

```
Passport-8606:5/config/ip/bgp/neighbor/128.1.1.2# info

BGP neighbor is 128.1.1.2 remote AS 0, Internal Peer,
BGP state [Idle]
remote router ID 0.0.0.0

          admin-state - BGP OFF
connect-retry-interval - 120
          ebgp-multihop - disable
              hold-time - 0
          keepalive-time - 0
          hold-time-configured - 180
keepalive-time-configured - 60
          max-prefix - 12000
          nexthop-self - disable
originate-def-route - disable
          MD5-authentication - disable
          neighbor-debug - none
          remove-private-as - enable
route-advertisement-interval - 30
          route-reflector-client - disable
          send-community - disable
soft-reconfigurationin-in - enable
          updt-source-interface - 0.0.0.0
              weight - 100
          Route Policy In -
          Route Policy Out -
```

## Configuring route redistribution parameters

BGP uses IGP within the AS to distribute BGP update information between BGP speakers (the Passport 8600 supports either RIP or OSPF for IGP). The IGP itself carries no BGP information. Each BGP speaker in the AS uses IBGP exclusively to determine reachability to external networks.

This section describes the commands used to configure route distribution parameters, and includes the following topics:

- [“Configuring BGP route redistribution parameters,” next](#)
- [“Configuring OSPF route redistribution parameters” on page 136](#)
- [“Configuring Direct route redistribution parameters” on page 138](#)
- [“Configuring RIP route redistribution parameters” on page 140](#)
- [“Configuring Static route redistribution parameters” on page 141](#)

### Configuring BGP route redistribution parameters

To allow policy updates to take effect for BGP route redistribution context, use the following command:

```
config ip bgp redistribute apply
```

### Configuring OSPF route redistribution parameters

To configure and apply OSPF route redistribution parameters, use the following command:

```
config ip bgp redistribute ospf
```

This command includes the following options.

<b>config ip bgp redistribute ospf</b> followed by:	
info	Displays current config ip bgp redistribute ospf info command output (Figure 48).
apply	Applies current OSPF route redistribution parameters to BGP.
create	Creates an OSPF route redistribution entry.
disable	Disables route redistribution from OSPF.
delete	Deletes the current OSPF route redistribution entry.
enable	Enables the current OSPF route redistribution entry.
metric <metric-value>	<p>Sets a numerical metric value for the redistributed route. The value can be a range between 0 to 65535. The default value is 0.</p> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>If you do not specify a value for this option, and no value is specified in the <code>config ip bgp default-metric</code> command, the default metric value is 0. Nortel Networks* recommends that you use a value that is consistent with the destination protocol.</li> <li>When you apply the route-policy parameter (see next field description), the <i>metric-value</i> that is configured within that policy takes precedence over the <i>metric-value</i> setting for the metric parameter described here.</li> <li><i>metric-value</i> is an integer value between 0 and 65535.</li> </ul>
route-policy <policy name>	<p>Specifies the route policy filter that is used to import OSPF routes to BGP.</p> <ul style="list-style-type: none"> <li><i>policy name</i> is an alphanumeric string length (0 to 64 characters) that indicates the name of the route policy.</li> </ul>

Figure 48 shows `config ip bgp redistribute ospf info` command output.

**Figure 48** config ip bgp redistribute ospf info command output

```

Passport-8610:5# config ip bgp redistribute ospf info
      create:
      delete: N/A
      enable: FALSE
      metric: 1
      route-policy:

```

## Configuring Direct route redistribution parameters

To configure and apply directly-connected IP interface route redistribution parameters, use the following command:

```
config ip bgp redistribute direct
```

This command includes the following options.

<b>config ip bgp redistribute direct</b> followed by:	
info	Displays current config ip bgp redistribute direct info command output ( <a href="#">Figure 49</a> ).
apply	Applies the current directly-connected IP interface route redistribution parameters to BGP.
create	Creates a directly-connected IP interface route redistribution entry.
disable	Disables route redistribution from directly-connected IP interfaces.
delete	Deletes the current directly-connected IP interface route redistribution entry.
enable	Enables the current directly-connected IP interface route redistribution entry.

<b>config ip bgp redistribute direct</b> followed by:	
<code>metric &lt;metric-value&gt;</code>	Sets a numerical metric value for the redistributed route. The default value is 0. <b>Note:</b> If you do not specify a value for this option, and no value is specified in the <code>config ip bgp default-metric</code> command, the default metric value is 0. Nortel Networks* recommends that you use a value that is consistent with the destination protocol. <ul style="list-style-type: none"> <li>• <i>metric-value</i> is an integer value between 0 and 65535.</li> </ul>
<code>route-policy &lt;policy name&gt;</code>	Specifies the route policy filter that is used to import directly-connected IP interface routes to BGP. <ul style="list-style-type: none"> <li>• <i>policy name</i> is an alphanumeric string length (0 to 64 characters) that indicates the name of the route policy.</li> </ul>

Figure 49 shows `config ip bgp redistribute direct info` command output.

**Figure 49** config ip bgp redistribute direct info command output

```
Passport-8610:5# config ip bgp redistribute direct info
      create:
      delete: N/A
      enable: FALSE
      metric: 1
      route-policy:
```

## Configuring RIP route redistribution parameters

To configure and apply RIP route redistribution parameters, use the following command:

```
config ip bgp redistribute rip
```

This command includes the following options.

<b>config ip bgp redistribute rip</b> followed by:	
info	Displays current config ip bgp redistribute rip info command output (Figure 50).
apply	Applies current RIP route redistribution parameters to BGP.
create	Creates a RIP route redistribution entry.
disable	Disables route redistribution from RIP.
delete	Deletes the current RIP route redistribution entry.
enable	Enables the current RIP route redistribution entry.
metric <metric-value>	<p>Sets a numerical metric value for the redistributed route. The default value is 0.</p> <p><b>Note:</b> If you do not specify a value for this option, and no value is specified in the <b>config ip bgp default-metric</b> command, the default metric value is 0. Nortel Networks* recommends that you use a value that is consistent with the destination protocol.</p> <ul style="list-style-type: none"> <li><i>metric-value</i> is an integer value between 0 and 65535.</li> </ul>
route-policy <policy name>	<p>Specifies the route policy filter that is used to import RIP routes to BGP.</p> <ul style="list-style-type: none"> <li><i>policy name</i> is an alphanumeric string length (0 to 64 characters) that indicates the name of the route policy.</li> </ul>

Figure 50 shows `config ip bgp redistribute rip info` command output.

**Figure 50** config ip bgp redistribute rip info command output

```
Passport-8610:5# config ip bgp redistribute rip info
      create:
      delete: N/A
      enable: FALSE
      metric: 1
      route-policy:
```

## Configuring Static route redistribution parameters

To configure and apply IP Static route redistribution parameters, use the following command:

```
config ip bgp redistribute static
```

This command includes the following options.

<b>config ip bgp redistribute static</b> followed by:	
info	Displays current config ip bgp redistribute static info command output (Figure 51).
apply	Applies current IP Static route redistribution parameters to BGP.
create	Creates a IP Static route redistribution entry.
disable	Disables route redistribution from IP Static routes.
delete	Deletes the current IP Static route redistribution entry.
enable	Enables the current IP Static route redistribution entry.

<b>config ip bgp redistribute static</b> followed by:	
<code>metric &lt;metric-value&gt;</code>	Sets a numerical metric value for the redistributed route. The default value is 0. <b>Note:</b> If you do not specify a value for this option, and no value is specified in the <code>config ip bgp default-metric</code> command, the default metric value is 0. Nortel Networks* recommends that you use a value that is consistent with the destination protocol. <ul style="list-style-type: none"> <li><code>metric-value</code> is an integer value between 0 and 65535.</li> </ul>
<code>route-policy &lt;policy name&gt;</code>	Specifies the route policy filter that is used to import IP Static routes to BGP. <ul style="list-style-type: none"> <li><code>policy name</code> is an alphanumeric string length (0 to 64 characters) that indicates the name of the route policy.</li> </ul>

Figure 51 shows `config ip bgp redistribute static info` command output.

**Figure 51** `config ip bgp redistribute static info` command output

```
Passport-8610:5# config ip bgp redistribute static info
      create:
      delete: N/A
      enable: FALSE
      metric: 1
      route-policy:
```

## Configuring AS-path lists

To create and modify an AS-path list for autonomous systems, use the following command:

```
config ip as-list <listid>
```

where:

*listid* is an integer value between 1 and 1024 that represents the AS-path list ID you want to create or modify.

This command includes the following options.

<b>config ip as-list &lt;listid&gt;</b> followed by:	
info	Displays the current content of the specified AS-path list ID ( <a href="#">Figure 52</a> ).
create	Creates the specified AS-path list entry.
delete	Deletes the specified AS-path list entry.
add-as-path <memberId> <permit deny> <as path>	<p>Adds a regular expression entry to the specified AS-path list.</p> <ul style="list-style-type: none"> <li>• <i>memberId</i> is an integer value between 0 and 65535 that represents the regular expression entry in the AS-path list.</li> <li>• <i>permit deny</i> permits or denies access for matching conditions.</li> <li>• <i>as path</i> is an integer value between 0 and 1536.</li> </ul>
remove-as-path [memberId <value>] [as-path <value>]	<p>Removes a regular expression entry from the specified AS-path list.</p> <ul style="list-style-type: none"> <li>• <i>memberId value</i> is an integer value between 0 and 65535 that represents the regular expression entry in the AS-path list.</li> <li>• <i>as-path value</i> is an integer value between 0 and 1536.</li> </ul>

[Figure 52](#) shows **config ip as-list info** command output.

**Figure 52** config ip as-list info command output

```
Passport-8610:5# config ip as-list 53 info
as-list id: 53
      memberId:      57 permit  127_224
```

## Configuring community lists

To create and modify community lists, use the following command:

```
config ip community-list <listid>
```

where:

*listid* is an integer value between 1 and 1024 that represents the community list ID you want to create or modify.

This command includes the following options.

<b>config ip community-list &lt;listid&gt;</b> followed by:	
info	Displays the current content of the specified community list ID ( <a href="#">Figure 53</a> ).
create	Creates the specified community list entry.
delete	Deletes the specified community list entry.

<b>config ip community-list &lt;listid&gt;</b> followed by:	
<pre>add-community memberId &lt;permit deny&gt; community-string</pre>	<p>Adds an entry to the community list.</p> <ul style="list-style-type: none"> <li><i>permit deny</i> allows you to set the access mode, which permits or denies access for matching conditions.</li> </ul> <p><b>Note:</b> The <i>memberId</i> is an integer value between 0 and 65535, that represents the member id in the community list.</p> <p>The <i>community-string</i> is an alphanumeric string value with a string length between 0 and 1536 characters (asnum:community-value) or (well-known community string).</p> <p><b>well-known community:</b></p> <ul style="list-style-type: none"> <li>internet</li> <li>no-export</li> <li>no-advertise</li> <li>local-as (known as NO_EXPORT_SUBCONFED)</li> </ul>
<pre>remove-community [memberId &lt;value&gt;] [community-string &lt;value&gt;]</pre>	<p>Removes an entry from the community list.</p> <ul style="list-style-type: none"> <li><i>memberId</i> is an integer value between 0 and 65535.</li> <li><i>community-string</i> is an alphanumeric string value with a string length between 0 and 1536 characters.</li> </ul>

Figure 53 shows `config ip community-list info` command output.

**Figure 53** config ip community-list info command output

```
Passport-8610:5# config ip community-list 29 info
community-list id: 29
  memberId:      37 permit  0:345
```

## Showing BGP configurations

This section describes show commands you can use to display current BGP configurations:

This section includes the following topics:

- [“Showing BGP aggregates” on page 146](#)
- [“Showing BGP CIDR routes” on page 148](#)
- [“Showing the BGP configuration” on page 149](#)
- [“Showing flap-dampened routes” on page 151](#)
- [“Showing global flap-dampening configurations” on page 152](#)
- [“Showing a BGP redistribution list” on page 153](#)
- [“Showing BGP imported routes” on page 154](#)
- [“Showing BGP network configurations” on page 156](#)
- [“Showing BGP peer groups” on page 156](#)
- [“Showing BGP routes” on page 158](#)
- [“Showing BGP global statistics” on page 159](#)
- [“Showing BGP summaries” on page 164](#)
- [“Showing BGP peers” on page 165](#)
- [“Showing BGP peer statistics” on page 166](#)
- [“Showing BGP peer routes” on page 171](#)

### Showing BGP aggregates

To display information about current aggregate addresses, use the following command:

```
show ip bgp aggregates [<prefix/len>]
```

where:

*prefix/len* is an IP address and an integer value (between 0 and 32).

[Figure 54](#) shows sample output for this command.

**Figure 54** show ip bgp aggregates command output

```
Passport-8610co:5# show ip bgp aggregates
aggregate address 10.10.10.0/255.255.255.0
as_set is not enabled
summary only is enabled
Attribute Map:
Advertise Map:
Suppression Map:
aggregate address 11.11.0.0/255.255.255.0
as_set is not enabled
summary only is enabled
Attribute Map:
Advertise Map:
Suppression Map:
```

## Showing BGP CIDR routes

To display information about current CIDR routes, use the following command:

```
show ip bgp cidr-only [exact <value>]
```

where:

exact *value* is an exact match of the prefix value (an IP address and an integer value between 0 and 32).

Figure 55 shows sample output for this command.

**Figure 55** show ip bgp cidr-only command output

```
Passport-8610:5/show ip bgp cidr-only
Network/Mask      Peer Rem Addr      NextHop Address Org      Loc Pref
-----
199.199.0.0/16    192.32.96.18      192.32.96.1      IGP 100  AS_PATH: (2 65123
762 11296 3561 701 5006)

200.200.0.0/16    192.32.96.18      192.32.96.1      IGP 100  AS_PATH: (2 65123
762 11296 3561 701 4230)

198.134.64.0/19   192.32.96.18      192.32.96.1      IGP 100  AS_PATH: (2 65123
762 11296 3561 1 4355)

200.24.208.0/20   192.32.96.18      192.32.96.1      IGP 100  AS_PATH: (2 65123
762 11296 3561 10910 10910 10910 11908 11908 11908 10530 19114)

198.12.112.0/22   192.32.96.18      192.32.96.1      IGP 100  AS_PATH: (2 65123
762 11296 3561 3908 6373 6373 )

200.56.240.0/22   192.32.96.18      192.32.96.1      IGP 100  AS_PATH: (2 65123
762 11296 3561 6453 17250)

The total number of CIDR routes is 7
```

Table 23 describes the cidr-only command parameters.

**Table 23** show ip bgp cidr-only command parameters

Field	Description
Network/Mask	The network IP address and exact mask length (must be an integer value between 0 and 32).
Peer Rem Addr	The IP address of the remote peer.
NextHop Address	The IP address of the next hop.
Org	The ultimate origin of the path information: <ul style="list-style-type: none"><li>• IGB = Networks are interior</li><li>• EGB = Networks learned via EGP</li><li>• INC (Incomplete) = Undetermined</li></ul>
Loc Pref	Local preference value.

## Showing the BGP configuration

To display information about your BGP configuration, use the following command:

```
show ip bgp conf
```

Figure 56 shows sample output for this command.

**Figure 56** show ip bgp conf command output

```
Passport-8606:5/show# ip bgp conf

BGP version - 4
local-as - 0
Identifier - 0.0.0.0

                                BGP on/off - OFF
                                local-as - 0
                                aggregation - enable
                                always-cmp-med - disable
                                auto-peer-restart - enable
                                auto-summary - enable
                                comp-bestpath-med-confed - disable
                                default-local-preference - 100
                                default-metric - -1
                                flap-dampening - disable
                                debug-screen - Off
                                global-debug - none
                                ibgp-report-import-rt - enable
                                ignore-illegal-rtrid - disable
                                max-equalcost-routes - 1
                                no-med-path-is-worst - enable
                                orig-def-route - disable
                                synchronization - enable
                                route-reflection - disable
                                cluster-id - 0.0.0.0
                                cl-to-cl-reflection - disable
                                decision state - Idle
                                confederation identifier - 0
```

## Showing flap-dampened routes

To display information about flap-dampened routes, use the following command:

```
show ip bgp dampened-paths <ipaddr> [<prefix>] [longer-prefixes]
```

where:

- *ipaddr* is the IP address
- *prefix* is the IP address and exact mask length (must be an integer value between 0 and 32).
- *longer-prefixes* indicates the mask length from any specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
- *community value* {disable|enable} allows you to enable or disable showing community attributes.

Figure 57 shows sample output for this command.

**Figure 57** show ip bgp dampened-paths command output

```
Passport-8610:5/config/ip/bgp/neighbor/192.32.96.18# show ip bgp
dampened-paths 192.32.96.18
Network/Mask      Peer Rem Addr  NextHop Address Org      Loc Pref
-----
10.0.0.0/8        192.32.96.18  192.32.96.18  INC 100 AS_PATH: (2)
24.154.160.0/20   192.32.96.18  192.32.96.1   IGP 100 AS_PATH: (2 65123
762 11296 3561 701 7046)
192.32.96.0/24    192.32.96.18  192.32.96.18  INC 100 AS_PATH: (2)
```

Table 24 describes the show ip bgp dampened-paths command parameters.

**Table 24** show ip bgp dampened-paths command parameters

Field	Description
Network/Mask	The network IP address and exact mask length (must be an integer value between 0 and 32).
Peer Rem Addr	The IP address of the remote peer.

**Table 24** show ip bgp dampened-paths command parameters (continued)

Field	Description
NextHop Address	The IP address of the next hop.
Org	The ultimate origin of the path information: <ul style="list-style-type: none"> <li>• IGB = Networks are interior</li> <li>• EGB = Networks learned via EGP</li> <li>• INC (Incomplete) = Undetermined</li> </ul>
Loc Pref	Local preference value.

## Showing global flap-dampening configurations

To display global information about flap-dampening, use the following command:

```
show ip bgp flap-damp-config
```

Figure 58 shows sample output for this command.

**Figure 58** show ip bgp flap-damp-config command output

```

Passport-8610co:5# show ip bgp flap-damp-config

*****
Global Flap Dampening Configuration
*****
                        Status - disable
                        PolicyName - N/A
CutoffThreshold - 1536
ReuseThreshold - 512
                        Decay - 2
                        MaxHoldDown - 180

```

Table 25 describes the show ip bgp dampened-paths command parameters.

**Table 25** show ip bgp dampened-paths command parameters

Field	Description
Status	Indicates the global state of the route flap dampening feature. Valid values are enable or disable.
PolicyName	Not applicable for this release.
CutoffThreshold	Indicates the penalty level that causes a route to be suppressed.
Decay	Indicates the decay rate, based on the decay algorithm.
MaxHoldDown	Indicates the maximum length of time (in seconds) that the route will be suppressed.

## Showing a BGP redistribution list

To display the current BGP redistribution list configuration, use the following command:

```
show ip bgp redistribution
```

Figure 59 shows sample output for this command.

**Figure 59** show ip bgp redistribution command output

```

Passport-8610co:5# show ip bgp redistribution
=====
                                BGP Redistribute List
=====
SRC  COMM  LV  LPRF  MET  MTYP  NHOP              ORIGIN  SRCLVL  EN  RMAP
-----
RIP  0     0   0     1    0    0.0.0.0           0       0       T  STAT

```

[Table 26](#) describes the show ip bgp redistribution command parameters.

**Table 26** show ip bgp redistribution command parameters

Field	Description
SRC	Indicates the redistribution source: RIP, Local, Static, or OSPF.
COMM	Not applicable.
LV	Not applicable.
LPRF	Not applicable.
MET	Indicates the metric value
MTYP	Not applicable.
NHOP	Not applicable.
ORIGIN	Indicates the Redistribution origin: IBGP or EBGP.
SRCLVL	Not applicable.
EN	Indicates whether the redistribution policy is enabled (T) true or disabled (F) false.
RMAP	The route policy that is currently assigned to the redistribution policy.

## Showing BGP imported routes

To display information about BGP imported routes, use the following command:

```
show ip bgp imported-routes
```

[Figure 60](#) shows sample output for this command.

**Figure 60** show ip bgp imported-routes command output

```

Passport-8610co:5# show ip bgp imported-routes

-----Imported Routes-----
route                community localPref nexthop
10.0.0.0/255.0.0.0   0          100      0.0.0.0
10.20.0.0/255.255.0.0 0          100      0.0.0.0
20.20.0.0/255.255.0.0 0          100      0.0.0.0
47.0.0.0/255.0.0.0   0          100      0.0.0.0
47.17.20.8/255.255.255.252 0        100      0.0.0.0
47.17.20.12/255.255.255.252 0        100      0.0.0.0
47.17.20.16/255.255.255.252 0        100      0.0.0.0
...
...
...
...
192.168.154.128/255.255.255.224 0        100      0.0.0.0
192.168.154.160/255.255.255.224 0        100      0.0.0.0

Total number of imported routes is 345

```

[Table 27](#) describes the show ip bgp imported-routes command parameters.

**Table 27** show ip bgp imported-routes command parameters

Field	Description
route	The path's prefix address.
community	Not applicable
localpref	Local preference value
nexthop	Not applicable

## Showing BGP network configurations

To display information about BGP network configurations, use the following command:

```
show ip bgp networks
```

[Figure 61](#) shows sample output for this command.

**Figure 61** show ip bgp networks command output

```
Passport-8610:5 show ip bgp networks
132.245.0.0 mask 255.255.0.0
192.32.0.0 mask 255.255.254.0
```

## Showing BGP peer groups

To display information about BGP peer groups, use the following command:

```
show ip bgp peer-group [<peer-group name>]
```

where:

peer-group name is the name of the peer group.

[Figure 62](#) shows sample output for this command.

**Figure 62** show ip bgp peer-group command output

```

Passport-8610:5 show ip bgp peer-group

*****peer group info*****
BGP peer group name: group1
BGP peer group index: 1
remote AS 100

                admin-state - BGP OFF
                ebgp-multihop - enable
                   hold-time - 180
                keepalive-time - 60
                   max-prefix - 0
                nexthop-self - disable
                originate-def-route - disable
                MD5-authentication - disable
                remove-private-as - enable
route-advertisement-interval - 30
                route-reflector-client - disable
                   send-community - disable
                soft-reconfiguration-in - disable
                   weight - 100

                Route Policy In -
                Route Policy Out -

-----neighbor info-----
BGP neighbor is 192.32.96.18 remote AS 2, External Peer,
BGP state [Idle]
remote router ID 0.0.0.0

                admin-state - BGP ON
connect-retry-interval - 120
                ebgp-multihop - enable
                   hold-time - 180
                keepalive-time - 60
                hold-time-configured - 180
keepalive-time-configured - 60
                max-prefix - 0
                nexthop-self - disable
                originate-def-route - disable
                MD5-authentication - disable
                neighbor-debug - none
                remove-private-as - enable
route-advertisement-interval - 30
                route-reflector-client - disable
                   send-community - disable
                soft-reconfigurationin-in - enable
                updt-source-interface - 0.0.0.0
                   weight - 100
                Route Policy In -
                Route Policy Out -

```

## Showing BGP routes

To display information about BGP routes, use the following command:

```
show ip bgp route [<prefix>] [longer-prefixes] [community <value>]
```

where:

- *prefix* is the IP address and exact mask length (must be an integer value between 0 and 32).
- *longer-prefixes* indicates the mask length from any specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
- *community value* {disable|enable} allows you to enable or disable showing community attributes.

Figure 63 shows sample output for this command.

**Figure 63** show ip bgp route command output

```
Passport-8610:5# show ip bgp route
The total number of routes is 4

Network/Mask      Peer Rem Addr  NextHop Address Org      Loc Pref
-----
10.0.0.0/8        192.32.96.18  192.32.96.18   INC 100 AS_PATH: (2)
24.154.160.0/20   192.32.96.18  192.32.96.1    IGP 100 AS_PATH: (2 65123
762 11296 3561 701 7046)

192.32.96.0/24    192.32.96.18  192.32.96.18   INC 100 AS_PATH: (2)
192.32.97.0/24    192.32.97.177 192.32.97.177  INC 100 AS_PATH: (4)

MED:0
```

Table 28 describes the show ip bgp route command parameters.

**Table 28** show ip bgp route command parameters

Field	Description
Network/Mask	The path's prefix address.
Peer Rem Addr	The remote peer address.

**Table 28** show ip bgp route command parameters (continued)

Field	Description
NextHop Addr	The BGP next hop address.
Org	The route's origin value: INC (incomplete), IGP, EGP.
Local Pref	The local preference value.

## Showing BGP global statistics

To display global BGP statistics, use the following command:

```
show ip bgp stats
```

[Figure 64](#) shows sample output for this command.

**Figure 64** showing global BGP statistics

```

Passport-8610co:5# show ip bgp stats
BGP Protocol: Global Statistics for BGP
=====
Conn-Starts..... : 0                Conn-Stops..... : 0
Conn-Opens.....   : 0                Conn-Closes..... : 0
Conn-Failures.... : 0                Conn-TCP Crashes. : 0
Conn-Expired..... : 0                Hold-Expired..... : 0
Bad-Events.....   : 0                Sync-Fails.....  : 0

Keepalive-Expired : 0
Recv-Opens.....   : 0                Recv-Keepalives.. : 0
Recv-Updates.... : 0                Recv-Notifys..... : 0
Sent-Opens.....   : 0                Sent-Keepalives.. : 0
Sent-Updates.... : 0                Sent-Notifys..... : 0
Notify Error Statistics:
  Hdr Errors..... : Rcvd 0                Sent 0
    Not Synchronized..... : Rcvd 0                Sent 0
    Invalid msg len.....  : Rcvd 0                Sent 0
    Invalid msg type..... : Rcvd 0                Sent 0
  Open Errors..... : Rcvd 0                Sent 0
    Bad Version.....      : Rcvd 0                Sent 0
    Bad AS Number.....    : Rcvd 0                Sent 0
    Bad BGP Rtr ID.....   : Rcvd 0                Sent 0
    Unsupported Option... : Rcvd 0                Sent 0
    Auth Failures.....    : Rcvd 0                Sent 0
    Bad Hold Value.....   : Rcvd 0                Sent 0
  Update Errors.... : Rcvd 0                Sent 0
    Malformed Attr List... : Rcvd 0                Sent 0
    Welknown Attr Unrecog. : Rcvd 0                Sent 0
    Welknown Attr Missing. : Rcvd 0                Sent 0
    Attr Flag Error.....  : Rcvd 0                Sent 0
    Attr Len Error.....   : Rcvd 0                Sent 0
    Bad ORIGIN Attr.....  : Rcvd 0                Sent 0
    AS Routing Loop.....  : Rcvd 0                Sent 0
    Bad NEXT-HOP Attr.... : Rcvd 0                Sent 0
    Optional Attr Error... : Rcvd 0                Sent 0
    Bad Network Field.... : Rcvd 0                Sent 0
    Malformed AS Path.... : Rcvd 0                Sent 0
  Hold Timer Exp...  : Rcvd 0                Sent 0
  FSM Error.....     : Rcvd 0                Sent 0
  Cease.....         : Rcvd 0                Sent 0

```

Table 29 describes the BGP statistics parameters.

**Table 29** BGP global statistics descriptions

Field	Description
Conn-Starts	Number of times BGP connection started
Conn-Stops	Number of times BGP connection stopped
Conn-Opens	Number of times BGP connection opened TCP
Conn-Closes	Number of times BGP connection closed TCP
Conn-Failures	Number of times a TCP attempt failed
Conn-TCP Crashes	Number of times TCP crashed due to fatal error
Conn-Expired	Number of times the TCP retry timer expired
Hold-Expired	Number of times the hold timer expired
Bad-Events	Number of invalid events received by FSM
Sync-Fails	Number of times the FDB sync failed
Keepalive-Expired	Number of times the keepalive timer expired
Recv-Opens	Number of Opens received by BGP
Recv-Keepalives	Number of Keepalive messages received by BGP
Recv-Updates	Number of Updates received by BGP
Recv-Notifys	Number of Notifications received by BGP
Sent-Opens	Number of Opens transmitted by BGP
Sent-Keepalives	Number of Keepalive messages transmitted by BGP
Sent-Updates	Number of Updates transmitted by BGP
Sent-Notifys	Number of Notifications transmitted by BGP
Hdr Errors (Rcvd/Sent)	Total number of Header errors received/transmitted
Not Synchronized (Rcvd/Sent)	Number of Header errors received/transmitted as: Not Synchronized
Invalid Msg len (Rcvd/Sent)	Number of Header errors received/transmitted as: Invalid Msg len
Invalid Msg type (Rcvd/Sent)	Number of Header errors received/transmitted as: Invalid Msg type
Open Errors (Rcvd/Sent)	Total number of Open errors received/transmitted
Bad Version (Rcvd/Sent)	Number of Open errors received/transmitted as: Bad version

**Table 29** BGP global statistics descriptions (continued)

Field	Description
Bad AS number (Rcvd/Sent)	Number of Open errors received/transmitted as: Bad AS number
Bad BGP Rtr ID (Rcvd/Sent)	Number of Open errors received/transmitted as: Bad BGP Rtr ID
Unsupported Option (Rcvd/Sent)	Number of Open errors received/transmitted as: Unsupported Option
Auth Failure (Rcvd/Sent)	Number of Open errors received/transmitted as: Auth Failure
Bad Hold Value (Rcvd/Sent)	Number of Open errors received/transmitted as: Bad Hold Value
Update Errors (Rcvd/Sent)	Total number of Update errors received/transmitted
Malformed Attr List (Rcvd/Sent)	Number of Update errors received/transmitted as: Malformed Attr List
Welknown Attr Unrecog (Rcvd/Sent)	Number of Update errors received/transmitted as: Welknown Attr Unrecog
Welknown Attr Missing (Rcvd/Sent)	Number of Update errors received/transmitted as: Welknown Attr Missing
Attr Flag Error (Rcvd/Sent)	Number of Update errors received/transmitted as: Attr Flag Error
Attr Len Error (Rcvd/Sent)	Number of Update errors received/transmitted as: Attr Len Error
Bad ORIGIN Attr (Rcvd/Sent)	Number of Update errors received/transmitted as: Bad ORIGIN Attr
AS Routing Loop (Rcvd/Sent)	Number of Update errors received/transmitted as: AS Routing Loop
Bad NEXT-HOP Attr (Rcvd/Sent)	Number of Update errors received/transmitted as: Bad NEXT-HOP Attr
Optional Attr Error (Rcvd/Sent)	Number of Update errors received/transmitted as: Optional Attr Error
Bad Network Field (Rcvd/Sent)	Number of Update errors received/transmitted as: Bad Network Field
Malformed AS Path (Rcvd/Sent)	Number of Update errors received/transmitted as: Malformed AS Path
Hold Timer Exp (Rcvd/Sent)	Total number of Hold Timer Expired errors received/transmitted

**Table 29** BGP global statistics descriptions (continued)

<b>Field</b>	<b>Description</b>
FSM Error (Rcvd/Sent)	Total number of FSM errors received/transmitted
Cease (Rcvd/Sent)	Total number of Cease errors received/transmitted

## Showing BGP summaries

To display information about BGP summaries, use the following command:

```
show ip bgp summary
```

Figure 65 shows sample output for this command.

**Figure 65** show ip bgp summary command output

```

Passport-8610co:5# show ip bgp summary

BGP version - 4
local-as - 0
Identifier - 0.0.0.0
Decision state - Idle
The total number of routes is 0

BGP NEIGHBOR INFO :
Neighbor RemoteAS BGP state HoldTime KeepAlive HoldTimeConfig KeepAliveConfig Weight ConnRetryInterval
=====
2.2.2.2      0    Idle      0      0      180      60    100      120
2.3.2.1      0    Idle      0      0      180      60    100      120
2.3.2.2      0    Idle      0      0      90       30     0       3
2.3.4.2      0    Idle      0      0      90       30     0      120

BGP CONFEDERATION PEER INFO :
confederation identifier 22
confederation peer as

BGP NETWORK INFO :
```

Table 30 describes the show ip bgp route command parameters.

**Table 30** show ip bgp route command parameters

Field	Description
Neighbor	The IP address of the remote peer.
RemoteAS	The AS number of the remote peer.
BGP state	The peers operating state: Idle, Accept, Connect, Open, Open-sent, and Established.
HoldTime	The negotiated holdtime value.
KeepAlive	The keepalive value.

**Table 30** show ip bgp route command parameters (continued)

Field	Description
HoldTimeConfig	The configured holdtime value.
KeepAliveConfig	The configured keepalive value.
Weight	The weight value assigned to the peer.
ConnRetryInterval	The reestablished timer value.

## Showing BGP peers

To display information about BGP peers, use the following neighbor command:

```
show ip bgp neighbor info [<ipaddr>]
```

where:

*ipaddr* is an optional parameter you provide to display information about a specific peer.

[Figure 66](#) shows sample output for this command.

**Figure 66** showing BGP peers

```
Passport-8610:5# show ip bgp neighbor info

BGP neighbor is 47.17.50.41 remote AS 100, External Peer,
BGP state [Active]
remote router ID 0.0.0.0

                                admin-state - BGP ON
connect-retry-interval - 120
                                ebgp-multihop - enable
                                hold-time - 0
                                keepalive-time - 0
                                hold-time-configured - 180
keepalive-time-configured - 60
                                max-prefix - 0
                                nexthop-self - disable
originate-def-route - disable
                                MD5-authentication - disable
                                neighbor-debug - none
                                remove-private-as - enable
route-advertisement-interval - 30
                                route-reflector-client - disable
                                send-community - disable
                                soft-reconfigurationin-in - enable

--More-- (q = quit)
```

## Showing BGP peer statistics

To display statistical information about a BGP peer, use the following neighbor command:

```
show ip bgp neighbor stats <ipaddr>
```

where:

*ipaddr* is the specific peer's IP address.

[Figure 67](#) shows sample output for this command.

Figure 67 showing BGP peer statistics

```

Passport-8610:5# show ip bgp neighbor stats 192.92.32.18
BGP Protocol: Peer 192.92.32.18 Statistics
=====
Conn-Starts..... : 0                Conn-Stops..... : 0
Conn-Opens.....   : 0                Conn-Closes..... : 0
Conn-Failures.... : 0                Conn-TCP Crashes. : 0
Conn-Expired..... : 0                Hold-Expired..... : 0
Keepalive-Expired : 0

Sent-Messages.... : 0
Sent-Opens.....   : 0                Sent-Keepalives.. : 0
Sent-Updates..... : 0                Sent-Notifys..... : 0
Rcv-Messages....  : 0                Rcv-Incomp-pkts.  : 0
Rcv-No-Markers..  : 0                Rcv-Opens.....    : 0
Rcv-Keepalives..  : 0                Rcv-Updates.....  : 0
Rcv-Notifys.....  : 0
Sent-NLRI-Count.. : 0                Rcv-Valid-update  : 0
Rcv-NLRI-Count..  : 0                Rcv-NLRI-Added..  : 0
Rcv-NLRI-Replace  : 0

Notify Error Statistics:
  Hdr Errors..... : Rcvd 0                Sent 0
    Not Synchronized..... : Rcvd 0                Sent 0
    Invalid msg len.....  : Rcvd 0                Sent 0
    Invalid msg type..... : Rcvd 0                Sent 0
  Open Errors.... : Rcvd 0                Sent 0
    Bad Version.....      : Rcvd 0                Sent 0
    Bad AS Number.....    : Rcvd 0                Sent 0
    Bad BGP Rtr ID.....   : Rcvd 0                Sent 0
    Unsupported Option... : Rcvd 0                Sent 0
    Auth Failures.....    : Rcvd 0                Sent 0
    Bad Hold Value.....   : Rcvd 0                Sent 0
  Update Errors... : Rcvd 0                Sent 0
    Malformed Attr List... : Rcvd 0                Sent 0
    Welknown Attr Unrecog. : Rcvd 0                Sent 0
    Welknown Attr Missing. : Rcvd 0                Sent 0
    Attr Flag Error.....   : Rcvd 0                Sent 0
    Attr Len Error.....    : Rcvd 0                Sent 0
    Bad ORIGIN Attr.....   : Rcvd 0                Sent 0
    AS Routing Loop.....   : Rcvd 0                Sent 0
    Bad NEXT-HOP Attr.... : Rcvd 0                Sent 0
    Optional Attr Error... : Rcvd 0                Sent 0
    Bad Network Field.... : Rcvd 0                Sent 0
    Malformed AS Path.... : Rcvd 0                Sent 0
  Hold Timer Exp... : Rcvd 0                Sent 0
  FSM Error.....    : Rcvd 0                Sent 0
  Cease.....        : Rcvd 0                Sent 0

```

Table 31 describes the BGP peer statistics parameters.

**Table 31** BGP peer statistics descriptions

Field	Description
Conn-Starts	Number of times peer BGP connection started
Conn-Stops	Number of times peer BGP connection stopped
Conn-Opens	Number of times peer opened TCP
Conn-Closes	Number of times peer closed TCP
Conn-Failures	Number of times peer TCP attempt failed
Conn-TCP Crashes	Number of times peer TCP crashed due to fatal error
Conn-Expired	Number of times the peer TCP retry timer expired
Hold-Expired	Number of times the peer hold timer expired
Keepalive-Expired	Number of times the peer keepalive timer expired
Sent-Opens	Number of Opens transmitted by the peer
Sent-Keepalives	Number of Keepalive messages transmitted by the peer
Sent-Updates	Number of Updates transmitted by the peer
Sent-Notifys	Number of Notifications transmitted by the peer
Recv-Messages	Total number of messages received by the peer
Recv-Incomp-pkts	Number of incomplete messages received by the peer
Recv-No-Markers	Number of messages without markers received by the peer
Recv-Opens	Number of Opens received by the peer
Recv-Keepalives	Number of Keepalive messages received by the peer
Recv-Updates	Number of Updates received by the peer
Recv-Notifys	Number of Notifications received by the peer
Sent-NLRI-Count	Number of network layer reachability information (NLRI) messages transmitted by the peer
Recv-Valid-Update	Number of valid Updates received by the peer
Recv-NLRI-Count	Number of network layer reachability information (NLRI) messages received by the peer
Recv-NLRI-Added	Number of routes added to the loc_rib by this peer

**Table 31** BGP peer statistics descriptions (continued)

Field	Description
Recv-NLRI-Replace	Number of routes that were replaced by routes received by the peer
Hdr Errors (Rcvd/Sent)	Total number of Header errors received/transmitted by the peer
Not Synchronized (Rcvd/Sent)	Number of Header errors received/transmitted by the peer as: Not Synchronized
Invalid Msg len (Rcvd/Sent)	Number of Header errors received/transmitted by the peer as: Invalid Msg len
Invalid Msg type (Rcvd/Sent)	Number of Header errors received/transmitted by the peer as: Invalid Msg type
Open Errors (Rcvd/Sent)	Total number of Open errors received/transmitted by the peer
Bad Version (Rcvd/Sent)	Number of Open errors received/transmitted by the peer as: Bad version
Bad AS number (Rcvd/Sent)	Number of Open errors received/transmitted by the peer as: Bad AS number
Bad BGP Rtr ID (Rcvd/Sent)	Number of Open errors received/transmitted by the peer as: Bad BGP Rtr ID
Unsupported Option (Rcvd/Sent)	Number of Open errors received/transmitted by the peer as: Unsupported Option
Auth Failure (Rcvd/Sent)	Number of Open errors received/transmitted by the peer as: Auth Failure
Bad Hold Value (Rcvd/Sent)	Number of Open errors received/transmitted by the peer as: Bad Hold Value
Update Errors (Rcvd/Sent)	Total number of Update errors received/transmitted by the peer
Malformed Attr List (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Malformed Attr List
Welknown Attr Unrecog (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Welknown Attr Unrecog
Welknown Attr Missing (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Welknown Attr Missing
Attr Flag Error (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Attr Flag Error
Attr Len Error (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Attr Len Error
Bad ORIGIN Attr (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Bad ORIGIN Attr

**Table 31** BGP peer statistics descriptions (continued)

<b>Field</b>	<b>Description</b>
AS Routing Loop (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: AS Routing Loop
Bad NEXT-HOP Attr (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Bad NEXT-HOP Attr
Optional Attr Error (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Optional Attr Error
Bad Network Field (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Bad Network Field
Malformed AS Path (Rcvd/Sent)	Number of Update errors received/transmitted by the peer as: Malformed AS Path
Hold Timer Exp (Rcvd/Sent)	Total number of Hold Timer Expired errors received/transmitted by the peer
FSM Error (Rcvd/Sent)	Total number of FSM errors received/transmitted by the peer
Cease (Rcvd/Sent)	Total number of Cease errors received/transmitted by the peer

## Showing BGP peer routes

To display information about BGP peer routes, use the following neighbor command:

```
show ip bgp neighbor route <ipaddr> [<prefix>]  
[longer-prefixes] [community <value>]
```

where:

- *ipaddr* is the specific peer's IP address.
- *prefix* is the IP address and exact mask length (must be an integer value between 0 and 32).
- *longer-prefixes* indicates the mask length from any specified prefix to 32 (for example, show from prefix a.b.c.d/len to a.b.c./32).
- *community value* {disable|enable} allows you to enable or disable showing community attributes.

Figure 68 shows sample output for this command.

**Figure 68** show ip bgp neighbor route command output

```

Passport-8610co:5# show ip bgp neighbor route 200.1.1.1

The total number of accepted routes from the neighbor is 1

Network/Mask      Peer Rem Addr   NextHop Address Org Loc Pref   Status
-----
11.11.1.0/30     200.1.1.1     10.1.1.13      IGP 10      Accepted
  AS_PATH: (40)
  MED:10

10.1.1.12/30     200.1.1.1     200.1.1.1      IGP 10      Accepted
  AS_PATH: path-is-empty
  MED:10

10.1.1.40/30     200.1.1.1     10.1.1.13      IGP 10      Best
  AS_PATH: (40)
  MED:10

10.1.1.60/30     200.1.1.1     10.1.1.13      IGP 10      Accepted
  AS_PATH: (40)
  MED:10

172.1.1.0/24     200.1.1.1     10.1.1.13      IGP 10      Used
  AS_PATH: (40 80)
  MED:10

172.1.2.0/30     200.1.1.1     10.1.1.13      IGP 10      Used
  AS_PATH: (40 80)
  MED:10

192.1.1.0/24     200.1.1.1     10.1.1.13      IGP 10      Accepted
  AS_PATH: (40 200)
  MED:10

200.1.1.0/30     200.1.1.1     200.1.1.1      IGP 10      Best
  AS_PATH: path-is-empty
  MED:10

200.1.30.0/30    200.1.1.1     200.1.1.1      IGP 10      Best
  AS_PATH: path-is-empty
  MED:10

200.1.20.0/30    200.1.1.1     200.1.1.1      IGP 10      Best
  AS_PATH: path-is-empty
  MED:10

200.1.40.0/24    200.1.1.1     200.1.1.1      IGP 10      Best
  AS_PATH: path-is-empty
  MED:10

```

[Table 32](#) describes the `show ip bgp neighbor route` command parameters.

**Table 32** show ip bgp neighbor route command parameters

Field	Description
Network/Mask	The IP mask of the direct route.
Peer Rem Addr	The IP address of the remote peer.
NextHop Address	The IP address of the next hop.
Org	Well-known mandatory attribute that specifies the source of a route: <ul style="list-style-type: none"> <li>• IGP — the route is interior to the originating AS that inserts this route into the BGP table (0 = IGP).</li> <li>• EGP — the route is learned via the Exterior Gateway Protocol (EGP) prior to being inserted into the BGP table (1 = BGP).</li> <li>• Incomplete — the origin of the route is unknown or learned by some other means. For example, these routes could be learned through RIP, OSPF, or static routes (2 = Incomplete).</li> </ul>
Local Pref	The value of the local preference attribute. The default value is 100 (any integer value in the range 0 to 2147483647).
Status	The route status: Accepted, Best, Used, Rejected

## Circuitless IP

This section describes how to configure the circuitless IP feature.



**Note:** You can configure a maximum of 32 circuitless IP interfaces on each device.

This section includes the following topics:

- [“Configuring Circuitless IP,”](#) next
- [“Showing Circuitless IP output”](#) on page 176

For conceptual information about the circuitless IP feature, see [“Circuitless IP”](#) on page 47.

## Configuring Circuitless IP

To configure circuitless IP, use the following command:

```
config ip circuitless-ip-int <id>
```

where:

<id> is an integer value in the range 1 to 32 that indicates the identification number for the specific circuitless IP interface.

This command includes the following options:

<b>config ip circuitless-ip-int</b> followed by:	
<code>info</code>	Displays the configured parameters for the circuitless IP interface
<code>area &lt;ipaddr&gt;</code>	Designates an area for the circuitless IP interface <ul style="list-style-type: none"> <li>• &lt;ipaddr&gt; is the IP address of the OSPF area that is associated with the circuitless IP interface (CLIP).</li> </ul>
<code>create &lt;ipaddr/mask&gt;</code>	Creates a circuitless IP interface <ul style="list-style-type: none"> <li>• &lt;ipaddr/mask&gt; is the IP address and Net Mask of the circuitless-IP interface.</li> </ul>
<code>delete &lt;ipaddr&gt;</code>	Deletes the specified circuitless IP interface <ul style="list-style-type: none"> <li>• &lt;ipaddr&gt; is the IP address of the Circuitless IP interface to be deleted.</li> </ul>
<code>ospf &lt;enable/disable&gt;</code>	Configures OSPF in passive mode for the circuitless IP interface. <ul style="list-style-type: none"> <li>• &lt;enable/disable&gt; enables or disables the option.</li> </ul>

### Configuration example

The following configuration example uses the above commands to configure circuitless IP, assign an interface number to the circuitless IP interface, and enable OSPF support.

The example also uses the **config ip circuitless-ip-int info** command to display information about the Circuitless IP setup

```
Passport-8010:5/config/ip/circuitless-ip-int/1# create 11.126.205.1/255.0.0.0
Passport-8010:5/config/ip/circuitless-ip-int/1# area 134.177.1.0
Passport-8010:5/config/ip/circuitless-ip-int/1# ospf enable
Passport-8010:5/config/ip/circuitless-ip-int/1# info
Sub-Context:
Current Context:
  Clip 1 :
    area : 134.177.1.0
    create : 11.126.205.1/255.0.0.0
    delete : N/A
    ospf : enabled
```

## Showing Circuitless IP output

To display information about the current circuitless IP configuration, use the following command:

```
show ip circuitless-ip-int info
```

Figure 69 shows sample output for this command.

**Figure 69** show ip circuitless-ip-int info command output

```
Passport-8610:5# show ip circuitless-ip-int info
=====
                        Circuitless Ip Interface
=====
INTERFACE   IP_ADDRESS      NET_MASK        OSPF_STATUS     AREA_ID
ID
-----
1           198.1.16.0      255.255.255.255 enable          0.0.0.0
2           200.4.0.0       255.255.255.255 enable          0.0.0.1
Passport-8610:5#
```

## Configuring BGP debug commands

This section describes BGP debug commands. You can use these commands to troubleshoot your BGP configuration.

This section includes the following topics:

- [“Tips for using the debug commands,”](#) next
- [“BGP global debug commands” on page 178](#)
- [“BGP peer/peer group debug commands” on page 185](#)

### Tips for using the debug commands

Debug command values allow you to control debug messages for global BGP message types, and for message types associated with a specified BGP peer or peer group.

The following tips can help you use the debug commands:

- You can display debug commands for multiple mask choices by entering the mask choices separated by commas, with no space between choice.

For example, to display the global debug command for mask choices *error* and *packet*, use the following command:

```
config ip bgp global-debug mask error,packet
```

- To end (disable) the display of debug messages, use the *none* mask choice. For example, to end the display of global debug messages, use the following command:

```
config ip bgp global-debug mask none
```

- You can save debug messages in a log file, or you can display the messages on your console. For example, to display (and log) a debug message, use the following command:

```
config ip bgp debug-screen [<setting>]
```

where:

*setting* is either of two values {off | on} that you enter to disable BGP screen logging (off), or to enable BGP screen logging (on).

## BGP global debug commands

This section describes global debug commands that allow you to display specific debug messages for your global BGP configuration, including the BGP neighbors.

You can display global debug messages for the following mask categories:

- none - disables the display of *all* debug messages
- all- sets the switch to display *all* categories of debug messages
- error - sets the switch to display *error* debug messages
- packet- sets the switch to display *packet* debug messages
- event- sets the switch to display *event* debug messages
- warning- sets the switch to display *warning* debug messages
- init- sets the switch to display *initialization* debug messages
- filter- sets the switch to display *filter -related* debug messages
- update- sets the switch to display *update -related* debug messages

This section includes the following topics:

- [“config ip bgp global-debug mask command,”](#) next
- [“config ip bgp neighbor-debug-all mask command”](#) on page 179
- [“Global debug command output examples”](#) on page 179

### config ip bgp global-debug mask command

To set the switch to display specific debug messages for your global BGP configuration, use the following command:

```
config ip bgp global-debug mask <value>
```

where:

*value* is one or more mask choice that you enter, separated by comma's with no space between choices. For example: [<mask>,<mask>,<mask>...].

For examples of global debug command output, see [“Global debug command output examples”](#) on page 179.

## config ip bgp neighbor-debug-all mask command

To set the switch to display specific debug messages for your global BGP neighbors, use the following command:

```
config ip bgp neighbor-debug-all mask <value>
```

where:

*value* is one or more mask choice that you enter, separated by comma's with no space between choices. For example: [<mask>,<mask>,<mask>...].

For examples of global debug command output, see [“Global debug command output examples,”](#) next.

## Global debug command output examples

This section includes the following global debug command output examples:

- [“All debug output,”](#) next
- [“Error debug output”](#) on page 181
- [“Packet debug output”](#) on page 181
- [“Event debug output”](#) on page 182
- [“Warning debug output”](#) on page 182
- [“Init debug output”](#) on page 183
- [“Filter debug output”](#) on page 183
- [“Update debug output”](#) on page 184

### All debug output

To set the switch to display *all* debug messages, use the following command:

```
config ip bgp global-debug mask all
```

[Figure 70](#) shows the all debug output.

**Figure 70** All debug output example

```
Passport-8610:5# config ip bgp global-debug mask all
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_EVENT:Closing socket 1
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_INIT:Closing connection cb
0x3d71094, connid 2
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_EVENT:Closing socket 2
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_WARNING: No duplicate peer for
192.32.97.175
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_INIT:Closing connection cb
0x3d714a8, connid 4
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_EVENT:Closing socket 4
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_WARNING: No duplicate peer for
192.32.96.3
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_EVENT:Created socket 1
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_EVENT:Bind socket 1, port 179:
myaddr 0.0.0.0
Passport-8610:5# [03/27/02 19:56:24] GLOBAL_EVENT:Listen on socket 1
Passport-8610:5# [03/27/02 19:56:25] GLOBAL_EVENT:Created socket 2
Passport-8610:5# [03/27/02 19:56:25] GLOBAL_EVENT:Bind socket 2, port 0:
myaddr 192.32.96.18
Passport-8610:5# [03/27/02 19:56:25]GLOBAL_EVENT:connect in progress:sockid 2,
errno -25
Passport-8610:5# [03/27/02 19:56:25] GLOBAL_EVENT:Created socket 3
Passport-8610:5# [03/27/02 19:56:25] GLOBAL_EVENT:Bind socket 3, port 0:
myaddr 192.32.96.18
Passport-8610:5# [03/27/02 19:56:25] GLOBAL_EVENT:connect in progress: sockid
3, errno -25
Passport-8610:5# [03/27/02 19:56:25] GLOBAL_EVENT:sockid 3: Client connected:
loc 192.32.96.18: rem 192.32.96.3
Passport-8610:5# [03/27/02 19:56:25] GLOBAL_EVENT:Closing socket 3
Passport-8610:5# [03/27/02 19:56:25] GLOBAL_EVENT:Created socket 3
Passport-8610:5# [03/27/02 19:56:26] GLOBAL_EVENT:Process incoming
Conn: sockid 1
Passport-8610:5# [03/27/02 19:56:26] ===== Incoming Connection =====
Passport-8610:5# [03/27/02 19:56:26] GLOBAL_EVENT:sockid 1: Accepted
conn: loc 192.32.96.18: rem 192.32.97.175
Passport-8610:5# [03/27/02 19:56:26] GLOBAL_WARNING: No peer for bgpid
192.32.97.175 (src 192.32.97.175)
Passport-8610:5# [03/27/02 19:56:26] GLOBAL_INIT:Closing connection cb
0x3d71094, connid 2
Passport-8610:5# [03/27/02 19:56:26] GLOBAL_EVENT:Closing socket 2
```

## Error debug output

To set the switch to display *error* debug messages, use the following command:

```
config ip bgp global-debug mask error
```

Figure 71 shows the error debug output.

**Figure 71** Error debug output example

```
Passport-8610:5# config ip bgp global-debug mask error
Passport-8610:5/config/ip/bgp# [03/27/02 19:58:48] GLOBAL_ERROR EVENT
(47.17.50.5): accept conn: no peer for src

Passport-8610:5/config/ip/bgp# [03/27/02 19:58:48] GLOBAL_ERROR:
Incoming conn rejected: sockid 1, src 47.17.50.5
```

## Packet debug output

To set the switch to display *packet* debug messages, use the following command:

```
config ip bgp global-debug mask packet
```

Figure 72 shows the packet debug output.

**Figure 72** Packet debug output example

```
Passport-8610:5# config ip bgp global-debug mask packet
Passport-8610:5# [03/27/02 20:05:16] GLOBAL_WARNING PKT: Send stalled
error -12, sockid 2
Passport-8610:5# [03/27/02 20:05:48] GLOBAL_WARNING PKT: Send stalled
error -12, sockid 4
Passport-8610:5# [03/27/02 20:05:49] GLOBAL_WARNING PKT: Send stalled
error -12, sockid 4
Passport-8610:5# [03/27/02 20:05:51] GLOBAL_WARNING PKT: Send stalled
error -12, sockid 4
Passport-8610:5# [03/27/02 20:05:52] GLOBAL_WARNING PKT: Send stalled
error -12, sockid 4
```

## Event debug output

To set the switch to display *event* debug messages, use the following command:

```
config ip bgp global-debug mask event
```

Figure 73 shows the event debug output.

**Figure 73** Event debug output example

```
Passport-8610:5# config ip bgp global-debug mask event
Passport-8610:5# [03/27/02 20:06:50] GLOBAL_EVENT:Process incoming Conn:
sockid 1
Passport-8610:5# [03/27/02 20:06:50] ===== Incoming Connection =====
Passport-8610:5# [03/27/02 20:06:50] GLOBAL_EVENT:sockid 1: Accepted conn: loc
192.32.96.18: rem 47.17.50.5
Passport-8610:5# [03/27/02 20:06:50] GLOBAL_ERROR EVENT (47.17.50.5): accept
conn: no peer for src
```

## Warning debug output

To set the switch to display *warning* debug messages, use the following command:

```
config ip bgp global-debug mask warning
```

Figure 74 shows the warning debug output.

**Figure 74** Warning debug output example

```
Passport-8610:5# config ip bgp global-debug mask warning
Passport-8610:5# [03/27/02 20:08:05] GLOBAL_WARNING: No peer for bgpid
192.32.97.175 (src 192.32.97.175)
Passport-8610:5# [03/27/02 20:08:05] GLOBAL_WARNING: No duplicate peer for
192.32.97.175
Passport-8610:5# [03/27/02 20:08:07] GLOBAL_WARNING: No peer for bgpid
192.32.96.3 (src 192.32.96.3)
Passport-8610:5# [03/27/02 20:08:08] GLOBAL_WARNING: No duplicate peer for
192.32.96.3
Passport-8610:5# [03/27/02 20:08:14] GLOBAL_WARNING PKT: Send stalled error
-12, sockid 4
```

## Init debug output

To set the switch to display *init* debug messages, use the following command:

```
config ip bgp global-debug mask init
```

Figure 75 shows the init debug output.

**Figure 75** Init debug output example

```
Passport-8610:5# config ip bgp global-debug mask init
Passport-8610:5# [03/27/02 20:09:27] GLOBAL_INIT:Closing connection cb
0x3d71094, connid 2
Passport-8610:5# [03/27/02 20:09:27] GLOBAL_INIT:Closing connection cb
0x3d714a8, connid 4
```

## Filter debug output

To set the switch to display *filter-related* debug messages, use the following command:

```
config ip bgp global-debug mask filter
```

Figure 76 shows the filter debug output.

**Figure 76** Filter debug output example

```
Passport-8610:5# config ip bgp global-debug mask filter
Passport-8610:5# [03/27/02 20:24:32] GLOBAL_FILTER: bgpCheckAdvMap No-Match
found
Passport-8610:5# [03/27/02 20:24:32] GLOBAL_FILTER: bgpCheckAdvMap match
permit
Passport-8610:5# [03/27/02 20:24:32] GLOBAL_FILTER: bgpCheckAdvMap No-Match
found
Passport-8610:5# [03/27/02 20:24:32] GLOBAL_FILTER: bgpCheckAdvMap No-Match
found
```

## Update debug output

To set the switch to display *update-related* debug messages, use the following command:

```
config ip bgp global-debug mask update
```

[Figure 77](#) shows the update debug output.

### Figure 77 Update debug output example

```
Passport-8610:5# config ip bgp global-debug mask update

Passport-8610:5# [03/27/02 20:14:57] GLOBAL_ERROR UPDATE: Nbr 192.32.97.175:
Reject NLRI 200.0.84.0/24, Exceeded limit 3700

Passport-8610:5# [03/27/02 20:14:57] GLOBAL_ERROR UPDATE: Nbr 192.32.97.175:
Reject NLRI 200.3.8.0/24, Exceeded limit 3700

Passport-8610:5# [03/27/02 20:14:57] GLOBAL_ERROR UPDATE: Nbr 192.32.97.175:
Reject NLRI 200.14.106.0/24, Exceeded limit 3700
```

## BGP peer/peer group debug commands

This section describes neighbor debug commands that allow you to display specific debug messages for BGP peers or peer groups. You can display neighbor debug messages for the following BGP peers or peer groups mask categories:

- none - disables the display of *all* debug messages
- all- sets the switch to display *all* categories of debug messages
- error - sets the switch to display *error* debug messages
- packet- sets the switch to display *packet* debug messages
- event- sets the switch to display *event* debug messages
- warning- sets the switch to display *warning* debug messages
- filter- sets the switch to display *filter -related* debug messages
- update- sets the switch to display *update -related* debug messages

To set the switch to display specific debug messages for BGP peers or peer groups, use the following neighbor command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>  
neighbor-debug mask <value>
```

where:

- *nbr\_ipaddr|peer-group-name* indicates that you enter the peer's IP address or the peer's group name.
- *value* is one or more mask choice that you enter, separated by comma's with no space between choices. For example: [<mask>,<mask>,<mask>...].

For examples of peer and peer group debug command output, see [“Peer and peer group debug command output examples,”](#) next.

## Peer and peer group debug command output examples

The following debug command output examples are included in this section:

- [“All debug output” on page 186](#)
- [“Error debug output” on page 188](#)
- [“Packet debug output” on page 188](#)
- [“Event debug output” on page 190](#)
- [“Warning debug output” on page 190](#)
- [“State debug output” on page 191](#)
- [“Filter debug output” on page 192](#)
- [“Update debug output” on page 194](#)

### All debug output

To set the switch to display *all* debug messages, use the following command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>  
neighbor-debug mask all
```

where:

*nbr\_ipaddr/peer-group-name* indicates that you enter the peer’s IP address or the peer group’s name.

[Figure 78](#) shows the peer and peer group all debug output.

**Figure 78** Peer/peer group all debug output example

```

Passport-8610:5# config ip bgp neighbor 192.32.97.175 neighbor-debug mask all
Passport-8610:5# [03/27/02 20:28:31] PEER_EVENT: 192.32.97.175: skt 2: event
STOP state ESTABLISHED
Passport-8610:5# [03/27/02 20:28:31] PEER_ERROR PKT EVENT:(192.32.97.175) Snd
notify: Cease Error, subcode 0
Passport-8610:5# [03/27/02 20:28:31]
===== SENT PKT - 21 bytes =====
Passport-8610:5# [03/27/02 20:28:31]
ff ff
00 15 03 06 00
Passport-8610:5# [03/27/02 20:28:31] PEER_EVENT:192.32.97.175: Del from init
update send list
Passport-8610:5# [03/27/02 20:28:31] State 192.32.97.175:
ESTABLISHED --> IDLE
Passport-8610:5# [03/27/02 20:28:31] PEER_UPDATE: nbr 192.32.97.175: Purge
RIBIN
Passport-8610:5# config ip bgp en
Passport-8610:5# [03/27/02 20:28:34] PEER_EVENT: 192.32.97.175: skt 0: event
START state IDLE
Passport-8610:5# [03/27/02 20:28:34] State 192.32.97.175: IDLE --> CONNECT
Passport-8610:5# [03/27/02 20:28:34] PEER_EVENT:Response to connect req: Peer
socket 2 Nbr: 192.32.97.175
Passport-8610:5# [03/27/02 20:28:34] PEER_EVENT: 192.32.97.175: skt 2: event
TCP_OPEN state CONNECT
Passport-8610:5# [03/27/02 20:28:34]
===== SENT PKT - 29 bytes =====
Passport-8610:5# [03/27/02 20:28:34]
ff ff
00 1d 01 04 00 02 00 b4 c0 20 60 12 00
Passport-8610:5# [03/27/02 20:28:34] PEER_EVENT:192.32.97.175: Setting state
to BGPST_OPENSENT
Passport-8610:5# [03/27/02 20:28:34] State 192.32.97.175: CONNECT --> OPENSENT
Passport-8610:5# [03/27/02 20:28:34] PEER_EVENT PKT      nbr 192.32.97.175,
rcvd full PDU: rlen 29
Passport-8610:5# [03/27/02 20:28:34]
===== RECEIVED PKT - 29 bytes =====
Passport-8610:5# [03/27/02 20:28:34]
ff ff
00 1d 01 04 fe 63 00 5a c0 20 61 af 00
Passport-8610:5# [03/27/02 20:28:34] PEER_EVENT: 192.32.97.175: skt 2: event
RX_OPEN state OPENSENT
Passport-8610:5# [03/27/02 20:28:34]
===== SENT PKT - 19 bytes =====

```

## Error debug output

To set the switch to display *error* debug messages for a specified peer or peer group, use the following command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>  
neighbor-debug mask error
```

where:

*nbr\_ipaddr/peer-group-name* indicates that you enter the peer's IP address or the peer group's name.

[Figure 79](#) shows the peer and peer group error debug output.

**Figure 79** Peer/peer group error debug output example

```
Passport-8610:5# config ip bgp neighbor 192.32.97.175 neighbor-debug mask  
error  
Passport-8610:5# [03/27/02 20:29:56] PEER_ERROR PKT EVENT:(192.32.97.175) Snd  
notify: Cease Error, subcode 0
```

## Packet debug output

To set the switch to display *packet* debug messages for a specified peer or peer group, use the following command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>  
neighbor-debug mask packet
```

where:

*nbr\_ipaddr/peer-group-name* indicates that you enter the peer's IP address or the peer group's name.

[Figure 80](#) shows the peer and peer group packet debug output.

**Figure 80** Peer/peer group packet debug output example

```
Passport-8610:5# config ip bgp neighbor 192.32.97.175 neighbor-debug mask
packet

Passport-8610:5# [03/27/02 20:30:56] PEER_EVENT PKT      nbr 192.32.97.175,
rcvd full PDU: rlen 29

Passport-8610:5# [03/27/02 20:30:56]
===== RECEIVED PKT - 29 bytes =====

Passport-8610:5# [03/27/02 20:30:56]
ff ff
00 1d 01 04 fe 63 00 5a c0 20 61 af 00

Passport-8610:5# [03/27/02 20:30:56]
===== SENT PKT - 19 bytes =====

Passport-8610:5# [03/27/02 20:30:56]
ff ff
00 13 04

Passport-8610:5# [03/27/02 20:30:56] PEER_EVENT PKT      nbr 192.32.97.175,
rcvd full PDU: rlen 19

Passport-8610:5# [03/27/02 20:30:56]
===== RECEIVED PKT - 19 bytes =====

Passport-8610:5# [03/27/02 20:30:56]
ff ff
00 13 04

Passport-8610:5# [03/27/02 20:30:57]
===== SENT PKT - 19 bytes =====

Passport-8610:5# [03/27/02 20:30:57]
ff ff
00 13 04
```

## Event debug output

To set the switch to display *event* debug messages for a specified peer or peer group, use the following command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>  
neighbor-debug mask event
```

where:

*nbr\_ipaddr/peer-group-name* indicates that you enter the peer's IP address or the peer group's name.

[Figure 81](#) shows the peer and peer group event debug output.

**Figure 81** Peer/peer group event debug output example

```
Passport-8610:5# config ip bgp neighbor 192.32.97.175 neighbor-debug mask  
event  
  
Passport-8610:5# [03/27/02 20:32:11] PEER_EVENT: 192.32.97.175: skt 2: event  
STOP state ESTABLISHED  
  
Passport-8610:5# [03/27/02 20:32:11] PEER_ERROR PKT EVENT:(192.32.97.175) Snd  
notify: Cease Error, subcode 0  
  
Passport-8610:5# [03/27/02 20:32:11] PEER_EVENT:192.32.97.175: Del from init  
update send list  
  
Passport-8610:5# [03/27/02 20:32:11] State 192.32.97.175: ESTABLISHED --> IDLE
```

## Warning debug output

To set the switch to display *warning* debug messages for a specified peer or peer group, use the following command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>  
neighbor-debug mask warning
```

where:

*nbr\_ipaddr/peer-group-name* indicates that you enter the peer's IP address or the peer group's name.

Figure 82 shows the peer and peer group warning debug output.

**Figure 82** Peer/peer group warning debug output example

```
Passport-8610:5# config ip bgp neighbor 192.32.97.175 neighbor-debug mask
warning

eventssport-8610:5# [03/27/02 20:36:09] GLOBAL_ERROR EVENT: (192.32.97.175)
Peer disable: already OFF
Passport-8610:5# [03/27/02 20:36:14] PEER_WARNING EVENT: (192.32.97.175) Peer
enable: Already enabled
```

### State debug output

To set the switch to display *state* debug messages for a specified peer or peer group, use the following command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>
neighbor-debug mask state
```

where:

*nbr\_ipaddr/peer-group-name* indicates that you enter the peer's IP address or the peer group's name.

Figure 83 shows the peer and peer group state debug output.

**Figure 83** Peer/peer group state debug output example

```
Passport-8610:5# config ip bgp neighbor 192.32.97.175 neighbor-debug mask
state
Passport-8610:5# [03/27/02 20:37:17] State 192.32.97.175: IDLE --> CONNECT

Passport-8610:5# [03/27/02 20:37:17] State 192.32.97.175: CONNECT --> OPENSENT

Passport-8610:5# [03/27/02 20:37:17] State 192.32.97.175: OPENSENT -->
OPENCONFIRM

Passport-8610:5# [03/27/02 20:37:17] State 192.32.97.175: OPENCONFIRM -->
ESTABLISHED
```

## Filter debug output

To set the switch to display *filter-related* debug messages for a specified peer or peer group, use the following command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>  
neighbor-debug mask filter
```

where:

*nbr\_ipaddr/peer-group-name* indicates that you enter the peer's IP address or the peer group's name.

[Figure 84](#) shows the peer and peer group filter debug output.

**Figure 84** Peer/peer group filter debug output example

```
Passport-8610:5# config ip bgp neighbor 192.32.97.175 neighbor-debug mask
filter
Passport-8610:5# [03/28/02 11:54:31] PEER_FILTER: prefix 200.1.193.0/24
Passport-8610:5# [03/28/02 11:54:31] PEER_FILTER:nbr 192.32.97.175: Inbound
Policy match access list
Passport-8610:5# [03/28/02 11:54:31] PEER_FILTER:nbr 192.32.97.175: Match
found in the current entry
Passport-8610:5# [03/28/02 11:54:31] PEER_FILTER:nbr 192.32.97.175: Match
found in the current entry
Passport-8610:5# [03/28/02 11:54:31] PEER_FILTER:nbr 192.32.97.175: match
permit
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175:
bgpPolicyUsesNlri: TRUE
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER: prefix 12.1.83.0/24
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: Inbound
Policy match access list
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: No-Match
found
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: keep rej
rt entry: 12.1.83.0/24
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175:
bgpPolicyUsesNlri: TRUE
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER: prefix 12.1.248.0/24
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: Inbound
Policy match access list
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: No-Match
found
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: keep rej
rt entry: 12.1.248.0/24
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175:
bgpPolicyUsesNlri: TRUE
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER: prefix 12.1.245.0/24
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: Inbound
Policy match access list
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: No-Match
found
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: keep rej
rt entry: 12.1.245.0/24
Passport-8610:5# [03/28/02 11:54:32] PEER_FILTER:nbr 192.32.97.175: Outbound
plcy is NULL
```

## Update debug output

To set the switch to display *update-related* debug messages for a specified peer or peer group, use the following command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>  
neighbor-debug mask update
```

where:

*nbr\_ipaddr/peer-group-name* indicates that you enter the peer's IP address or the peer group's name.

[Figure 85](#) shows the update debug output.

**Figure 85** Update debug output example

```
Passport-8610:5# config ip bgp neighbor 192.32.97.175 neighbor-debug mask  
update  
Passport-8610:5# [03/27/02 20:39:08] PEER_UPDATE: nbr 192.32.97.175: Purge  
RIBIN
```

---

## Chapter 5

# Configuration examples

---

The Border Gateway Protocol (BGP) is an exterior gateway protocol used by border routers to exchange network reachability information with other BGP systems. BGP routers form peer relationships with neighboring BGP routers. Using an entity called a BGP speaker, the BGP peers transmit and receive current routing information over a reliable transport layer connection, making periodic updates unnecessary. BGP peers exchange complete routing information only when they establish a peer connection. Thereafter, BGP peers exchange routing information in the form of *routing updates*.

A routing update includes a network number, a list of autonomous systems that the routing information has passed through (the AS path), and other path attributes that describe the route to a set of destination networks. When multiple paths are available, BGP compares the path attributes to choose the preferred path.

In addition to exchanging BGP information between autonomous systems, BGP exchanges information between routers in the same AS. To differentiate between these uses, the latter is called *interior BGP* (IBGP).

This chapter provides examples of common BGP configuration tasks and includes examples of the CLI commands used to create the configuration. For a complete description of all of the available CLI commands you can use to configure BGP, including those shown in this chapter, refer to [Chapter 4, “Using the CLI to configure BGP,” on page 113](#).

This chapter includes the following topics:

- [“Basic BGP example,” next](#)
- [“Policies” on page 202](#)
- [“CIDR and aggregate addressing example” on page 209](#)
- [“EBGP multihop and EBGP load balance” on page 212](#)
- [“BGP synchronization and next-hop self” on page 215](#)

- “MD5 authentication” on page 230
- “BGP peer groups” on page 233
- “BGP path attributes” on page 236
- “Community attribute” on page 258
- “IBGP scalability issues” on page 263
- “BGP confederations” on page 264
- “Route reflectors” on page 269
- “Route flap dampening” on page 281

## Basic BGP example

BGP uses two basic connections types: *IBGP* and *EBGP*.

- IBGP

Routers that belong to the same autonomous system (AS) and exchange BGP updates run *internal* BGP (IBGP).

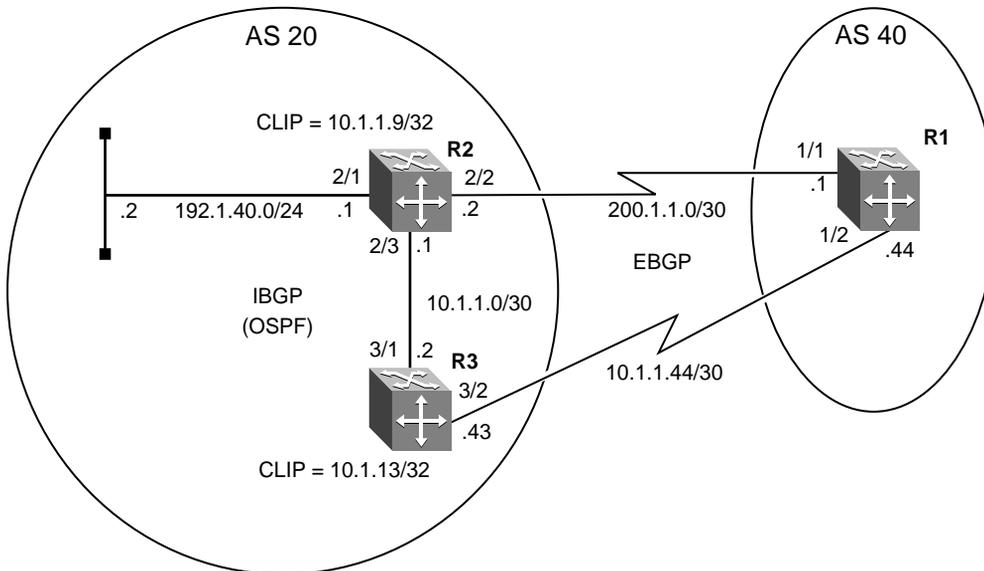
- EBG

Routers that belong to a different AS and exchange BGP updates run *external* BGP (EBGP).

Within any AS, routers run an interior gateway protocol, such as OSPF.

As shown in [Figure 86 on page 197](#), R1 in AS 40 is running EBGP to connect to R2 and R3 in AS 20. Note that, within AS 20, R2 and R3 run IBGP.

Following [Figure 86](#), is a step-by-step procedure that shows how to configure R2 for this example. The procedure does not show the configuration steps for R1 and R3, which must also be configured to complete the configuration example.

**Figure 86** Basic BGP configuration

10871FA

The following configuration tasks are required to configure R2:

- [Configure circuitless IP on R2](#)
- [Configure R2 ports for EBGP and IBGP interfaces](#)
- [Configure OSPF on R2](#)
- [Configure BGP on R2](#)
- [Configure BGP peer interfaces for R2](#)
- [Configure IGP network prefixes](#)
- [BGP route auto-summary](#)
- [Specifying number of routes learned](#)

To configure R2 in [Figure 86](#), complete the following steps:

### *Configure circuitless IP on R2*



**Note:** Circuitless IP (CLIP) ensures that, if one or more of the device's interfaces becomes disabled, the device is always reachable as long as a viable path to the device exists.

---

- 1 Define a circuitless IP (CLIP) address on R2:

```
Passport-8610:5# config ip circuitless-ip-int 1 create
10.1.1.9/32
```

- 2 Enable OSPF on the CLIP:

```
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

### *Configure R2 ports for EBGp and IBGP interfaces*



**Note:** In the example shown in [Figure 86](#), R2 router ports (2/2 and 2/3) are used as the BGP EBGp and IBGP interfaces. For this reason, VLAN ID's of 2090 and 2091 are used in the command line. Either a router port or a VLAN can be configured as the BGP interface.

---

- 1 Configure port 2/1 interface:

```
Passport-8610:5# config ethernet 2/1 ip create
192.1.40.1/24 2092
```

- 2 Configure port 2/2 interface:

```
Passport-8610:5# config ethernet 2/2 ip create 200.1.1.2/
30 2090
```

- 3 Configure port 2/3 interface:

```
Passport-8610:5# config ethernet 2/3 ip create 10.1.1.1/
30 2091
```

To display the router port VLAN ID's, use the following command:

```
Passport-8610:5# show vlan info router-port
```

Figure 87 shows the `show vlan info brouter-port` command output.

**Figure 87** show vlan info brouter-port command output

```

Passport-8610:5# show vlan info brouter-port

      Vlan Id          Port
      =====          ====
      2090             2/2
      2091             2/3
      2092             2/1

Passport-8610:5#

```

### Configure OSPF on R2



**Note:** Be sure that the router address (the router ID) is the same as the circuitless IP address.

- 1 Assign R2 an AS boundary router (ASBR):

This command allows R2 to accept external routes.

```

Passport-8610:5# config ip ospf admin-state enable
Passport-8610:5# config ip ospf as-boundary-router enable

```

- 2 Configure the OSPF router with the same address as the CLIP:

```

Passport-8610:5# config ip ospf router-id 10.1.1.9

```

- 3 Enable OSPF:

```

Passport-8610:5# config ip ospf enable

```

### Configure BGP on R2

- 1 Assign R2 to AS 20:

```

Passport-8610:5# config ip bgp local-as 20

```

**2** Disable synchronization on R2:

The following command disables R2 from accepting routes from BGP peers without waiting for an update from the IGP.

```
Passport-8610:5# config ip bgp synchronization disable
```

**3** Enable BGP on R2:

```
Passport-8610:5# config ip bgp enable
```

*Configure BGP peer interfaces for R2*

Because many neighbors use similar update policies, you can group the neighbors that have the same update policies into peer groups and peer associations. This association and grouping allows you to simplify your configurations and makes updates more efficient. You can configure peers and peer groups using the neighbor commands, as shown in the following steps:

**1** Assign neighboring router (R1) as an R2 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.1.1 create
Passport-8610:5# config ip bgp neighbor 200.1.1.1
remote-as 40
Passport-8610:5# config ip bgp neighbor 200.1.1.1
admin-state enable
```

**2** Assign neighboring router (R3) as an R2 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.2 create
Passport-8610:5# config ip bgp neighbor 10.1.1.2
remote-as 20
Passport-8610:5# config ip bgp neighbor 10.1.1.2
admin-state enable
```

*Configure IGP network prefixes*

Configure BGP networks that you want R2 to advertise for redistribution.



**Note:** The networks must be present in the routing table before BGP can advertise them.

---

```
Passport-8610:5# config ip bgp network 192.1.40.0/24 add
```

## *BGP route auto-summary*

By default, the Passport 8600 summarizes network routes based on class limits (for example, Class A, B, C networks). To disable this feature, use the following command.

```
Passport-8610:5# config ip bgp auto-summary disable
```

## *Specifying number of routes learned*

The BGP implementation for the 3.3 release currently has a default number of routes that can be accepted (the default value is 12,000 routes). To accept more than 12,000 routes, you must change the max-prefix parameter value.



**Note:** The max-prefix parameter controls the maximum number of routes that a peer can accept. The purpose is to prevent non M mode configurations from accepting more routes than it can forward. Use a setting of 0 to accept an unlimited number of prefixes. For more information about the max-prefix parameter, see [“Configuring BGP peers or peer groups” on page 129](#).

To modify the max-prefix parameter value, use the following command:

```
config ip bgp neighbor <nbr_ipaddr|peer-group-name>
max-prefix 0 add
```

Example:

```
Passport-8610:5# config ip bgp neighbor 150.1.0.3 max-prefix
0 add
```

For more information about...	See...
Showing BGP routes...	<a href="#">“Showing BGP routes” on page 158.</a>
Showing BGP routes advertised by neighbors...	<a href="#">“Showing BGP peer routes” on page 171.</a>
Showing BGP peers operational state...	<a href="#">“Showing BGP summaries” on page 164</a>
Showing BGP imported routes...	<a href="#">“Showing BGP imported routes” on page 154</a>

## Policies

BGP uses IGP in the AS to distribute BGP update information between BGP speakers (the Passport 8600 supports either RIP or OSPF for IGP). The IGP itself carries no BGP information. Each BGP speaker in the AS uses IBGP exclusively to determine reachability to external networks.

This section provides examples of the commands you use to create redistribution policies that can inject external routes within an AS. For a complete description of the route distribution commands, see [“Configuring route redistribution parameters” on page 136](#).



**Note:** If the AS is running OSPF, the border router must be configured as an AS boundary router (ASBR) in order to accept external routes (see [“Basic BGP example” on page 196](#)).

---

This section includes the following topics:

- [“Creating OSPF and BGP route distribution policies,” next](#)
- [“Creating Direct, RIP, and Static route distribution policies” on page 204](#)
- [“Injecting a default route when using OSPF as an IGP” on page 205](#)

## Creating OSPF and BGP route distribution policies

This section describes commands you use to create OSPF and BGP route distribution policies.

### Configuration commands

To create OSPF and BGP route distribution policies, complete the following steps:

#### *Configure Circuitless IP*

Enter the following commands to configure a CLIP interface:

- 1 Define the circuitless IP (CLIP) address:

```
Passport-8610:5# config ip circuitless-ip-int 1 create
10.1.1.9/32
```

**2** Enable OSPF on the CLIP:

```
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

*Configure OSPF***1** Assign an AS boundary router (ASBR):

This command allows the border router to accept external routes.

```
Passport-8610:5# config ip ospf admin-state enable
Passport-8610:5# config ip ospf as-boundary-router enable
```

**2** Configure the OSPF router with the same address as the CLIP:

```
Passport-8610:5# config ip ospf router-id 10.1.1.9
```

**Note:** Configuring the router-id with the same address as the CLIP address ensures that, if one or more of the device's interfaces becomes disabled, the device is always reachable, as long as a viable path to the device exists.

**3** Enable OSPF:

```
Passport-8610:5# config ip ospf enable
```

*Configure OSPF to BGP redistribution*

Enter the following commands to configure a route policy for OSPF to BGP redistribution (if required):

```
Passport-8610:5# config ip ospf redistribute bgp create
Passport-8610:5# config ip ospf redistribute bgp enable
Passport-8610:5# config ip ospf redistribute bgp apply
```

*Configure BGP to OSPF redistribution*

Enter the following commands to configure a route policy for BGP to OSPF redistribution (if required):

```
Passport-8610:5# config ip bgp redistribute ospf create
Passport-8610:5# config ip bgp redistribute ospf enable
Passport-8610:5# config ip bgp redistribute ospf apply
```

## Redistribution considerations and tips

Consider the following tips when creating OSPF and BGP route distribution policies:

- The BGP route redistribution feature  
Use caution when configuring this feature. An improperly configured parameter could cause learned EBGP routes to be advertised out of your local AS. If this happens, other networks could be routed through your local AS.  
You should not enable this feature if you are peering to an Internet Service Provider (ISP) and do not wish to have traffic transit your local AS.
- Redistributing OSPF into BGP  
When you redistribute OSPF into BGP, route priority is in effect and can create routing loops. Because BGP has a higher route preference than OSPF External 1 & 2, if you redistribute OSPF external 1 & 2 routes into BGP, the BGP routes will be used and could cause a routing loop.
- Changing the BGP Router ID  
By default, the BGP Router ID automatically uses the OSPF Router ID.  
If you change the OSPF Router ID, you must restart BGP to use the new value. Note that OSPF uses a random Router ID, by default.

## Creating Direct, RIP, and Static route distribution policies

You can use the following command examples to configure the Passport 8600 for distributing Static, OSPF, Direct and RIP routes. Route policies can also be used with a distribution policy.

The following example shows how to enter the appropriate commands for distributing Static routes to BGP:

- 1 To view the range of redistribute commands, enter:

```
Passport-8610:5# config ip bgp ?  
  
Sub-Context: ospf direct rip static  
Current Context:
```

- 2 Configure and apply IP Static route redistribution parameters:

```
Passport-8610:5# config ip bgp redistribute static create
Passport-8610:5# config ip bgp redistribute static enable
Passport-8610:5# config ip bgp redistribute static apply
```

- 3 To display the command output, enter:

```
Passport-8610:5# config ip bgp redistribute static info
```

[Figure 88](#) shows `config ip bgp redistribute static info` command output.

**Figure 88** config ip bgp redistribute static info command output

```
Passport-8610:5# config ip bgp redistribute static info
      create:
      delete: N/A
      enable: FALSE
      metric: 1
      route-policy:
```

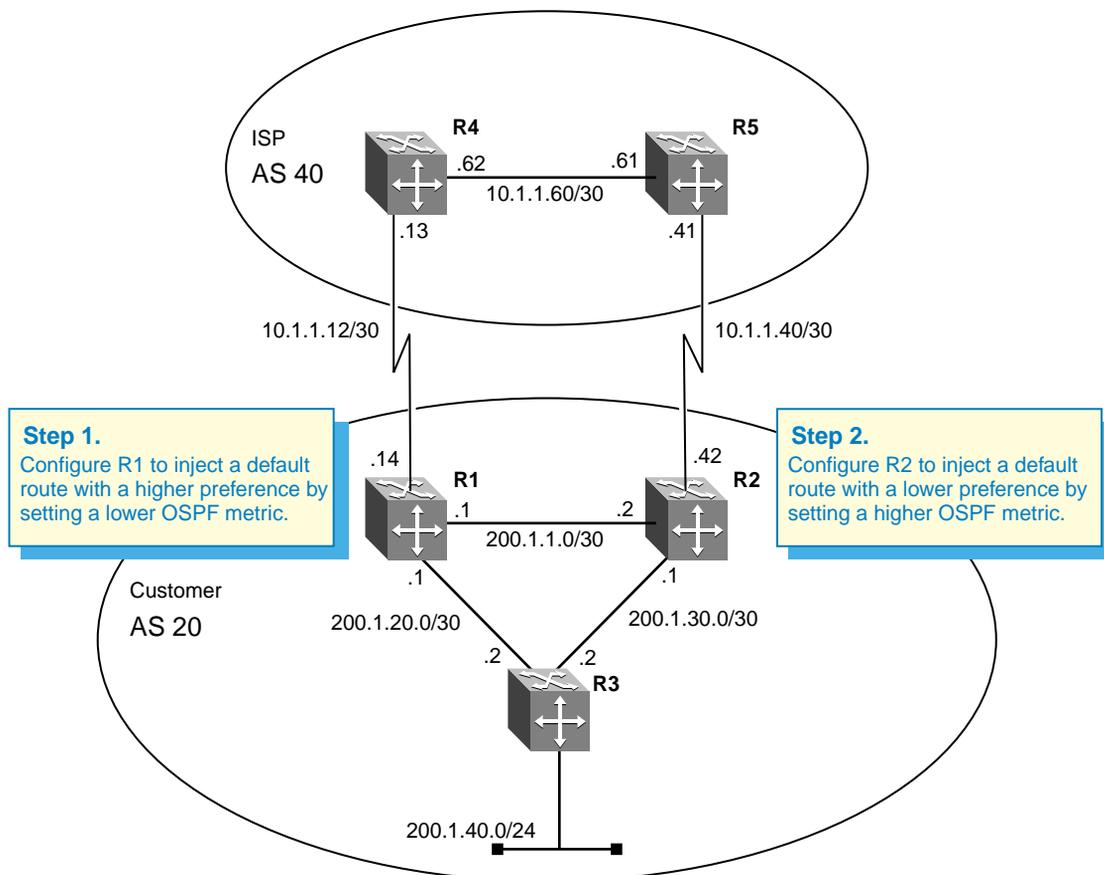
Notice the route-policy and metric settings. Both can be used influence the advertised route(s).

## Injecting a default route when using OSPF as an IGP

The example shown in [Figure 89](#) shows Customer AS 20 connecting to ISP AS 40 through R1 and R2. For this example, let's suppose that you want to configure both R1 and R2 to inject a default route into the IGB in AS 20. In addition, you also want to influence the path of the default route with route metrics that will allow R3 to use either interface 200.1.20.1 or interface 200.1.30.1 as the next hop.

In [Figure 89](#), OSPF is the IGP protocol. In the step-by-step procedure that follows [Figure 89](#), both R1 and R2 are configured to inject a default route into OSPF. The metrics for R1 are configured to allow R1 to be used as the default next-hop for router R3, as long as R1 remains operational

**Figure 89** Injecting a default route example



10872FA

## R1 Configuration

- 1 Configure the IP prefix list:

```
Passport-8610:5# config ip prefix-list DR add-prefix  
0.0.0.0/0
```

**Note:** For this example, DR represents the name used for this policy.

- 2 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy Default_OSPF  
Passport-8610:5# config ip route-policy seq 1 create  
Passport-8610:5# config ip route-policy seq 1 enable  
Passport-8610:5# config ip route-policy seq 1 action  
permit  
Passport-8610:5# config ip route-policy seq 1  
set-injectlist DR  
Passport-8610:5# config ip route-policy seq 1 set-metric  
100
```

**Note:** The set-metric value directly influences the OSPF route decision. For this example, R1 is set to a lower metric value than R2, which results in a higher preference value.

- 3 Configure route redistribution:

```
Passport-8610:5# config ip ospf redistribute bgp create  
Passport-8610:5# config ip ospf redistribute bgp  
route-policy Default_OSPF enable  
Passport-8610:5# config ip ospf redistribute bgp  
route-policy Default_OSPF apply
```

## *R2 Configuration*

- 1 Configure the IP prefix list:

```
Passport-8610:5# config ip prefix-list DR add-prefix
0.0.0.0/0
```

**Note:** For this example, DR represents the name used for this policy.

- 2 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy Default_OSPF
Passport-8610:5# config ip route-policy seq 1 create
Passport-8610:5# config ip route-policy seq 1 enable
Passport-8610:5# config ip route-policy seq 1 action
permit
Passport-8610:5# config ip route-policy seq 1
set-injectlist DR
Passport-8610:5# config ip route-policy seq 1 set-metric
300
```

**Note:** The set-metric value directly influences the OSPF route decision. For this example, R2 is set to a higher metric value than R1, which results in a lower preference value.

- 3 Configure route redistribution:

```
Passport-8610:5# config ip ospf redistribute bgp create
Passport-8610:5# config ip ospf redistribute bgp
route-policy Default_OSPF enable
Passport-8610:5# config ip ospf redistribute bgp
route-policy Default_OSPF apply
```

The end result of this configuration is that R3 will use the next hop to R1 for access the Internet.

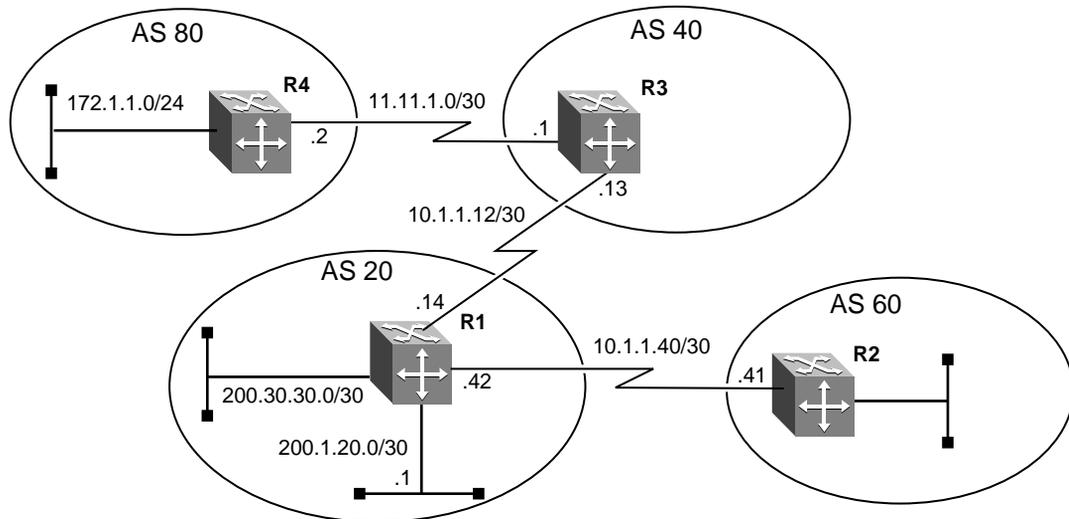
## CIDR and aggregate addressing example

BGP4 supports Classless Interdomain Routing (CIDR), an addressing scheme (also known as supernetting) that eliminates the concept of classifying networks into class types. The IP address 192.3.0.0/16, which would normally be an illegal Class C address, provides a good example of CIDR. With this method, supernet 192.3.0.0/16 represents 192.3.0.0 255.255.0.0 (for more information, see [“CIDR and aggregate addresses”](#) on page 27).

CIDR make it easy to aggregate multiple routes into a single route and considerably reduces the size of routing tables. In the example shown in [Figure 90](#), R1 in AS 20 is configured to advertise an aggregate route of 200.0.0.0 to AS 40.

Following [Figure 90](#), is a step-by-step procedure that shows how to configure R1 for this example.

**Figure 90** Aggregation example



10873FA

### *R1 Configuration*

- 1 Configure BGP on R1 as follows:

```
Passport-8610:5# config ip bgp network 200.30.30.0/30 add
Passport-8610:5# config ip bgp network 200.1.20.0/30 add
Passport-8610:5# config ip bgp aggregate-address
200.0.0.0/8 add summary-only enable
```

- 2 Assign R3 in AS 40 as a BGP peer to R1:

Use the neighbor command to assign peers and peer groups.

```
Passport-8610:5# config ip bgp neighbor 10.1.1.13 create
Passport-8610:5# config ip bgp neighbor 10.1.1.13
remote-as 40
```

- 3 Enable the peer's administrative state:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.13
admin-state enable
```

- 4 Enter the following show command to display the route table for the R3 peer:

```
Passport-8610:5# show ip bgp neighbor route 10.1.1.14
```

Figure 91 shows sample output for this command.



**Note:** The route table for the R3 peer now displays the aggregate address of 200.0.0.0/8.

---

**Figure 91** show ip bgp route command output

```

Passport-8610:5# show ip bgp neighbor route 10.1.1.14
The total number of accepted routes from the neighbor is 6

Network/Mask      Peer Rem Addr  NextHop Address Org Loc Pref  Status
-----
16.16.16.16/30   10.1.1.14     10.1.1.14     IGP 100    Accepted
    AS_PATH: (20 60)
10.1.1.12/30     10.1.1.14     10.1.1.14     IGP 100    Best
    AS_PATH: (20)
14.14.14.0/24    10.1.1.14     10.1.1.14     IGP 100    Accepted
    AS_PATH: (20 60)
10.1.1.40/30     10.1.1.14     10.1.1.14     IGP 100    Accepted
    AS_PATH: (20)
192.1.1.0/24     10.1.1.14     10.1.1.14     IGP 100    Accepted
    AS_PATH: (20 60 200)
200.0.0.0/8    10.1.1.14     10.1.1.14     IGP 100    Used
    AS_PATH: (20){25}
    AGGR-AS:20 AGGR-ADDR:10.1.1.14

```

<b>For more information about:</b>	<b>See:</b>
CIDR and aggregate addresses	<a href="#">"Consolidating routing information" on page 26.</a>

## EBGP multihop and EBGP load balance

By default, BGP enforces the one-hop rule for BGP peers whenever two EBGP speakers are directly connected by setting the TTL field in the update message to 1. This ensures that the remote peer is located on a directly attached network. However, there may be situations where you may not be able to use the next-hop address due to an indirect connection, such as a circuitless IP (CLIP) address. In this case, you can use BGP multihop to set the TTL field to a different value.



**Note:** BGP multihop is only used for EBGP connections, not for IBGP connections.

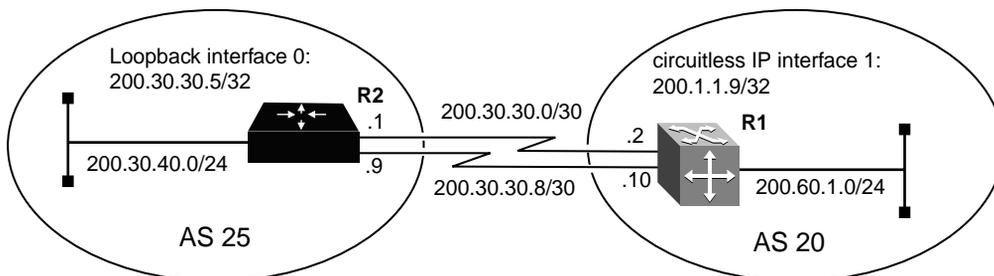
Figure 92 shows an example of when you should use BGP multihop. As shown in this example, when two or more links are used to connect two EBGP speakers, the neighbor remote-as configuration specifies the CLIP address (also referred to as a loopback address).



**Note:** Because the BGP peer is not directly connected when using BGP multihop, you must also configure static routes.

As shown in Figure 92, R1 in AS 20 is configured to perform EBGP load balance with R2, which is a Cisco Systems\* router. Following Figure 92, is a step-by-step procedure that shows how to configure R1 for this example.

**Figure 92** EBGP load balance example



Passport 8600



Cisco router

10874FA

---

## R1 Configuration

### *Configure Circuitless IP on R1*

Enter the following commands to configure a CLIP interface for R1:

- 1 Define the circuitless IP (CLIP) address:

```
Passport-8610:5# config ip circuitless-ip-int 1 create
200.1.1.9/255.255.255.255
```

### *Configure BGP on R1*

- 1 Assign R1 to AS 20:

```
Passport-8610:5# config ip bgp local-as 20
```

- 2 Set the equal-cost-path value for R1:

```
Passport-8610:5# config ip bgp max-equalcost-route 2
enable
```

- 3 Enable the equal cost multipath on R1:

```
Passport-8610:5# config ip ecmp enable
```

- 4 Enable BGP on R1:

```
Passport-8610:5# config ip bgp enable
```

### *Configure IGP network prefixes*

Configure BGP networks that you want R1 to advertise for redistribution.

```
Passport-8610:5# config ip bgp network 200.60.1.0/24 add
```

### *Configure BGP peer interface*

Assign neighboring router (R2) as an R1 peer:

```
Passport-8610:5# config ip bgp neighbor 200.30.30.5
create
Passport-8610:5# config ip bgp neighbor 200.30.30.5
remote-as 25
```

```
Passport-8610:5# config ip bgp neighbor 200.30.30.5
ebgp-multihop enable
Passport-8610:5# config ip bgp neighbor 200.30.30.5
update-source-interface 200.1.1.9 add
Passport-8610:5# config ip bgp neighbor 200.30.30.5
admin-state enable
```

### *Configure R1 static routes*

```
Passport-8610:5# config ip static-route create
200.30.30.5/255.255.255.255 next-hop 200.30.30.1
Passport-8610:5# config ip static-route create
200.30.30.5/255.255.255.255 next-hop 200.30.30.9
```

## **R2 Configuration (Cisco router)**

The following Cisco commands configure R2:

```
!Router 2
interface loopback 0
ip address 200.30.30.5 255.255.255.255
!
router bgp 25
network 200.30.40.0
neighbor 200.1.1.9 remote-as 20
neighbor 200.1.1.9 ebgp-multihop 2
neighbor 200.1.1.9 update-source loopback 0
!
ip route 200.1.1.9 255.255.255.255 200.30.30.2
ip route 200.1.1.9 255.255.255.255 200.30.30.10
```

<b>For more information about:</b>	<b>See:</b>
EBGP multihop and load balancing	<a href="#">"Configuring general BGP parameters" on page 120.</a>

---

## BGP synchronization and next-hop self

The BGP synchronization feature allows you to enable or disable the router from accepting routes from BGP peers without waiting for an update from the IGP. When you enable the synchronization parameter, the router does not advertise a route to EBGP peers until all the routers in the AS have learned the route via IGP.

Normally, synchronization should be enabled for most BGP configurations (the default value is enable). An exception is when the AS is a stub AS, which does not pass traffic from one AS to another, or if all the routers in the AS run BGP.

The BGP synchronization feature is predictable and, depending on your configuration settings, reacts to topology changes in predictable ways. This section describes some of the ways the BGP synchronization feature reacts to an unexpected change to the topology.

This section includes the following topics:

- [“Example 1 — initial configuration,”](#) next
- [“Example 2 — unexpected changes to the initial topology”](#) on page 222
- [“Example 3 — how to correct the next hop problem”](#) on page 225

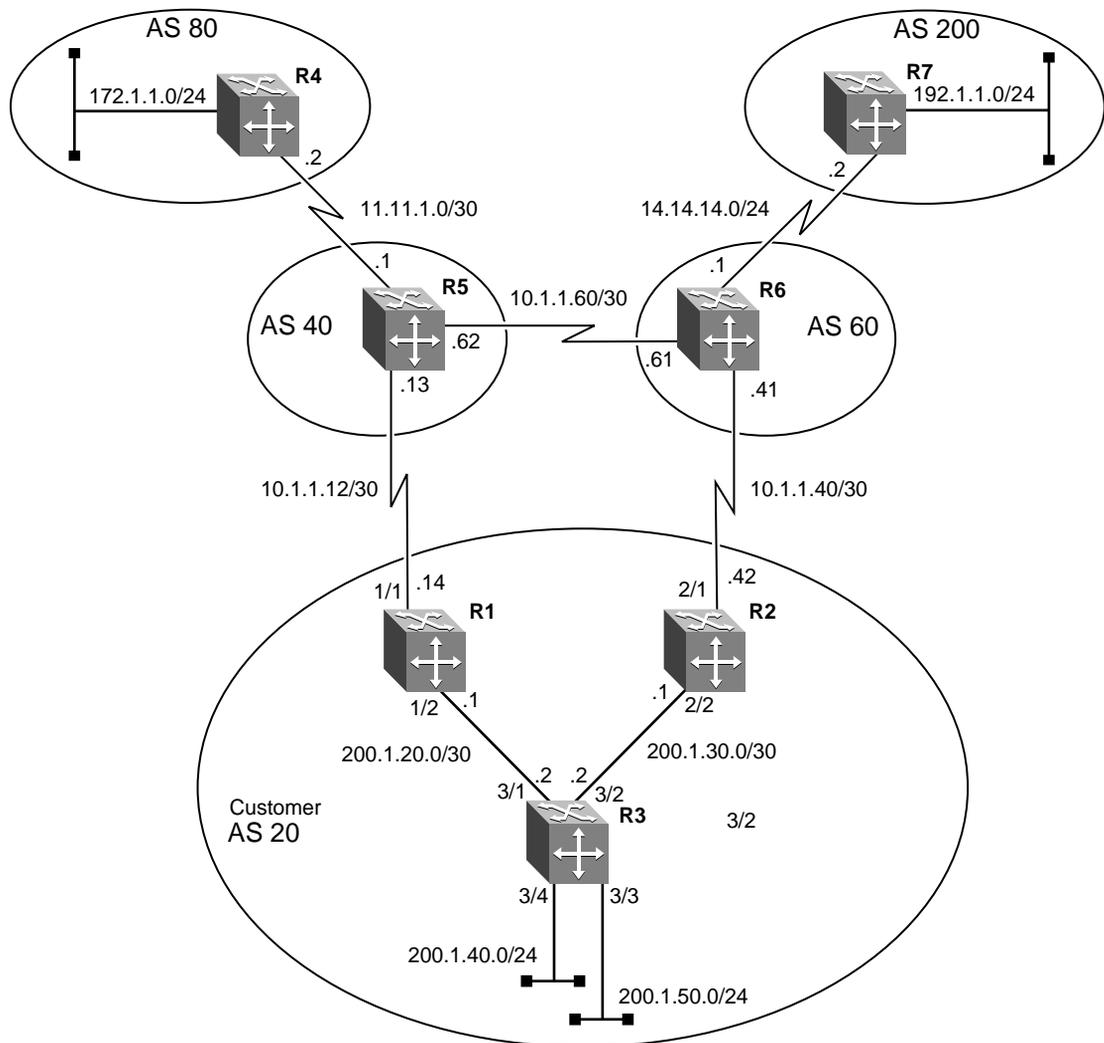
### Example 1 — initial configuration

For the example shown in [Figure 93 on page 216](#), synchronization is set to enabled. In this scenario, if the link between R2 to R6 goes down, R2 will not learn the external BGP routes from R1.



**Note:** This may be the effect you wish to have, however if your AS is a transit AS you may still want R2 to be able to reach the external BGP routes.

---

**Figure 93** BGP synchronization and next-hop self example

10875FA

Unless R2 learns about the 10.1.1.12 network, it will not be able to reach the external BGP routes. To resolve this issue, you can either enable *passive* OSPF on the 10.1.1.14 interface (if the IGP is OSPF), or you can enable BGP next hop-self on R1.

Another issue that must be addressed before R2 can reach the external routes depends on whether R3 is configured for BGP. If it is, there is no issue and R2 will place the external BGP routes into its routing table. However, if R3 is running an IGP only, then R1 and R2 should be configured to distribute BGP into IGP.

Following [Figure 93](#), is a step-by-step procedure that shows how to configure R1 and R2 for this example.

## R1 Configuration

### *Configure an OSPF interface*

- 1 Configure the OSPF interface on R1, port 1/2:

```
Passport-8610:5# config ethernet 1/2 ip create
200.1.20.1/255.255.255.252 2065
```

- 2 Enable OSPF:

```
Passport-8610:5# config ethernet 1/2 ip ospf enable
```

### *Configure circuitless IP*

- 1 Configure CLIP on R1:

```
Passport-8610:5# config ip circuitless-ip-int 1 create
200.1.1.9/255.255.255.255
```

- 2 Enable OSPF on the CLIP:

```
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

### *Configure OSPF*

- 1 Assign an AS boundary router (ASBR):

This command allows the border router to accept external routes.

```
Passport-8610:5# config ip ospf admin-state enable
Passport-8610:5# config ip ospf as-boundary-router enable
```

- 2 Configure the OSPF router with the same address as the CLIP:

```
Passport-8610:5# config ip ospf router-id 200.1.1.9
```

**Note:** Configuring the router-id with the same address as the CLIP address ensures that, if one or more of the device's interfaces becomes disabled, the device is always reachable, as long as a viable path to the device exists.

- 3 Enable OSPF:

```
Passport-8610:5# config ip ospf enable
```

### *Configure BGP on R1*

- 1 Assign R1 to AS 20:

```
Passport-8610:5# config ip bgp local-as 20
```

- 2 Enable synchronization on R2:

The following command prevents R1 from advertising a route until all routers in the AS have learned the route through the IGP.

```
Passport-8610:5# config ip bgp synchronization enable
```

- 3 Enable BGP on R1:

```
Passport-8610:5# config ip bgp enable
```

### *Configure IGP network prefixes*

Configure BGP networks that you want R1 to advertise for redistribution.

```
Passport-8610:5# config ip bgp network 10.1.1.12/30 add  
Passport-8610:5# config ip bgp network 200.1.20.0/30 add
```

### *Configure BGP peer interfaces for R1*

Because many neighbors use similar update policies, you can group the neighbors that have the same update policies into peer groups and peer associations. This association and grouping allows you to simplify your configurations and makes updates more efficient. You can configure peers and peer groups using the neighbor commands, as shown in the following steps:

- 1 Assign neighboring router (R2) as an R1 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.30.1 create
Passport-8610:5# config ip bgp neighbor 200.1.30.1
remote-as 20
Passport-8610:5# config ip bgp neighbor 200.1.30.1
admin-state enable
```

- 2 Assign neighboring router (R5) as an R1 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.13 create
Passport-8610:5# config ip bgp neighbor 10.1.1.13
remote-as 40
Passport-8610:5# config ip bgp neighbor 10.1.1.13
admin-state enable
```

## R2 Configuration

### *Configure an OSPF interface*

- 1 Configure the OSPF interface on R2, port 2/2:

```
Passport-8610:5# config ethernet 2/2 ip create
200.1.30.1/255.255.255.252 2065
```

- 2 Enable OSPF:

```
Passport-8610:5# config ethernet 2/2 ip ospf enable
```

### *Configure circuitless IP*

- 1 Configure CLIP on R2:

```
Passport-8610:5# config ip circuitless-ip-int 1 create
200.1.1.5/255.255.255.255
```

- 2 Enable OSPF on the CLIP:

```
Passport-8610:5# config ip circuitless-ip-int 1 ospf
enable
```

### *Configure OSPF*

- 1 Assign an AS boundary router (ASBR):

This command allows the border router to accept external routes.

```
Passport-8610:5# config ip ospf admin-state enable
Passport-8610:5# config ip ospf as-boundary-router enable
```

- 2 Configure the OSPF router with the same address as the CLIP:

```
Passport-8610:5# config ip ospf router-id 200.1.1.5
```

**Note:** Configuring the router-id with the same address as the CLIP address ensures that, if one or more of the device's interfaces becomes disabled, the device is always reachable, as long as a viable path to the device exists.

- 3 Enable OSPF:

```
Passport-8610:5# config ip ospf enable
```

### *Configure BGP on R2*

- 1 Assign R2 to AS 20:

```
Passport-8610:5# config ip bgp local-as 20
```

- 2 Enable synchronization on R2:

The following command prevents R2 from advertising a route until all routers in the AS have learned the route through the IGP.

```
Passport-8610:5# config ip bgp synchronization enable
```

- 3 Enable BGP on R2:

```
Passport-8610:5# config ip bgp enable
```

### *Configure IGP network prefixes*

Configure BGP networks that you want R2 to advertise for redistribution.

```
Passport-8610:5# config ip bgp network 10.1.1.40/30 add
Passport-8610:5# config ip bgp network 200.1.30.0/30 add
```

## *Configure BGP peer interfaces for R2*

Because many neighbors use similar update policies, you can group the neighbors that have the same update policies into peer groups and peer associations. This association and grouping allows you to simplify your configurations and makes updates more efficient. You can configure peers and peer groups using the neighbor commands, as shown in the following steps:

- 1 Assign neighboring router (R1) as an R2 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.20.1 create  
Passport-8610:5# config ip bgp neighbor 200.1.20.1  
remote-as 20  
Passport-8610:5# config ip bgp neighbor 200.1.20.1  
admin-state enable
```

- 2 Assign neighboring router (R6) as an R2 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.41 create  
Passport-8610:5# config ip bgp neighbor 10.1.1.41  
remote-as 40  
Passport-8610:5# config ip bgp neighbor 10.1.1.41  
admin-state enable
```

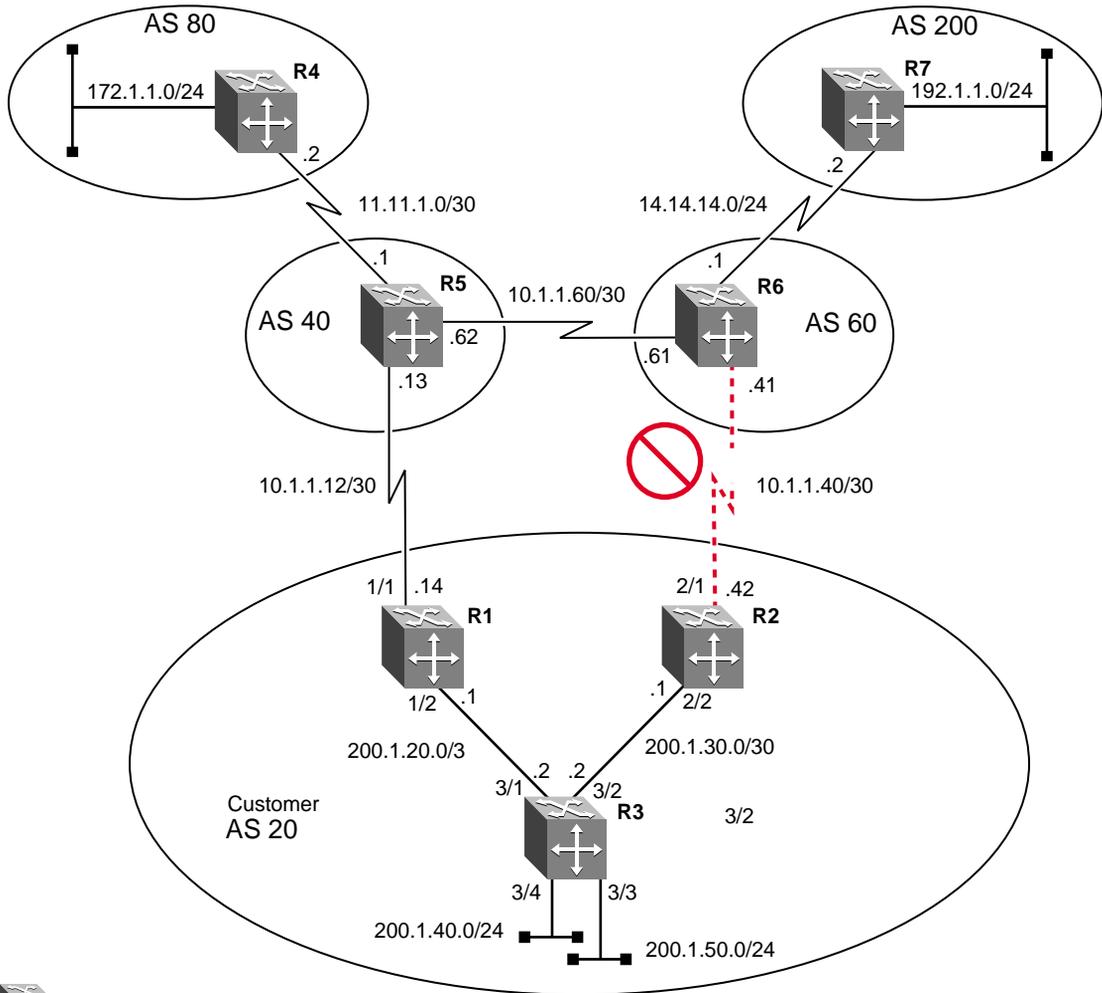
To understand how the BGP synchronization feature reacts to an unexpected change in the topology, see [“Example 2 — unexpected changes to the initial topology,”](#) next.

## Example 2 — unexpected changes to the initial topology

This section shows how the BGP synchronization feature reacts to an unexpected change in the topology (for example, if a critical link is down). For this example, R3 is running OSPF as an IGP and is not running BGP (Figure 94).

For demonstration purposes, we will assume that the connection between R2 and R6 is removed, simulating a link problem.

**Figure 94** BGP synchronization and next hop-self example 2



To view the changes to the topology, complete the following steps:

- 1 Enter the following show command:

```
Passport-8610:5# show ip bgp route
```

Figure 95 displays the output of the show ip bgp route command.

**Figure 95** show ip bgp route command output (Example 1) from R2

```
Passport-8610:5# show ip bgp route

The total number of routes is 11

Network/Mask      Peer Rem Addr  NextHop Address  Org  Loc  Pref
-----
11.11.1.0/30     200.1.20.1    10.1.1.13       IGP  100
  AS_PATH: (40)
10.1.1.12/30     200.1.20.1    200.1.20.1      IGP  100
  AS_PATH: path-is-empty
14.14.14.0/24    200.1.20.1    10.1.1.13       IGP  100
  AS_PATH: (40 60)
10.1.1.40/30     200.1.20.1    10.1.1.13       IGP  100
  AS_PATH: (40 60)
10.1.1.60/30     200.1.20.1    10.1.1.13       IGP  100
  AS_PATH: (40)
172.1.1.0/24     200.1.20.1    10.1.1.13       IGP  100
  AS_PATH: (40 80)
172.1.2.0/30     200.1.20.1    10.1.1.13       IGP  100
  AS_PATH: (40 80)
192.1.1.0/24     200.1.20.1    10.1.1.13       IGP  100
  AS_PATH: (40 60 200)
200.1.1.8/30     200.1.20.1    200.1.20.1      IGP  100
  AS_PATH: path-is-empty
200.20.20.0/24   200.1.20.1    200.1.20.1      IGP  100
  AS_PATH: path-is-empty
200.1.20.0/30    200.1.20.1    200.1.20.1      IGP  100
  AS_PATH: path-is-empty
```

Notice that because the connection to R6 is down, the only route available is the next hop attribute of 10.1.1.13 for all external routes.

2 Enter the following show command:

```
Passport-8610:5# show ip route info
```

Figure 96 displays the output of the show ip route info command.

Notice that there is no external route information in the standard show ip route info table shown in Figure 96. None of the BGP entries are in the route table because the next hop (10.1.1.13) for these entries is unreachable and is not learned through OSPF.

**Figure 96** show ip route info command output from R2

```
Passport-8610:5# show ip route info
=====
                                Ip Route
=====
DST                MASK                NEXT                COST  VLAN  PORT  PROT  AGE  TYPE  PRF
-----
200.1.1.4          255.255.255.252    200.1.1.5          1     0     -/-   LOC   0    DB    0
200.1.1.8          255.255.255.252    200.1.30.2         111   2065  2/2   OSPF  0    IB    20
200.1.20.0         255.255.255.252    200.1.30.2         101   2065  2/2   OSPF  0    IB    20
200.1.20.4         255.255.255.252    200.1.30.2         101   2065  2/2   OSPF  0    IB    20
200.1.30.0         255.255.255.252    200.1.30.1         0     0     2/2   LOC   0    DB    0

5 out of 5 Total Num of Dest Networks,5 Total Num of Route Entries displayed.
-----
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route,
E=Ecmp Route, U=Unresolved Route, N=Not in HW
```

## Example 3 — how to correct the next hop problem

This section describes how to resolve the next hop problem pointed out in the previous section.

To resolve the problem, complete the following steps:

### *Configure R1 for nexthop-self*

- 1 Disable the administration state for the R1 peer:

You must first disable the peer's administration state before you change next-hop parameter value.

```
Passport-8610:5# config ip bgp neighbor 200.1.30.1  
admin-state disable
```

- 2 Enable the R1 peer's nexthop-self parameter:

```
Passport-8610:5# config ip bgp neighbor 200.1.30.1  
nexthop-self enable
```

- 3 Enable the administration state for the R1 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.30.1  
admin-state enable
```

### *Viewing the BGP route table for R2*

To view the change to the topology, complete the following steps:

- 1 Enter the following show command:

```
Passport-8610:5# show ip bgp route
```

[Figure 97](#) displays the output of the `show ip bgp route` command.

**Figure 97** show ip bgp route command output (Example 3A)

```

Passport-8610:5# show ip bgp route

The total number of routes is 10

Network/Mask      Peer Rem Addr  NextHop Address  Org  Loc  Pref
-----
11.11.1.0/30      200.1.20.1    200.1.20.1      IGP  100
AS_PATH: (40)
10.1.1.12/30      200.1.20.1    200.1.20.1      IGP  100
AS_PATH: path-is-empty
14.14.14.0/24     200.1.20.1    200.1.20.1      IGP  100
AS_PATH: (40 60)
10.1.1.60/30      200.1.20.1    200.1.20.1      IGP  100
AS_PATH: (40)
172.1.1.0/24      200.1.20.1    200.1.20.1      IGP  100
AS_PATH: (40 80)
172.1.2.0/30      200.1.20.1    200.1.20.1      IGP  100
AS_PATH: (40 80)
192.1.1.0/24      200.1.20.1    200.1.20.1      IGP  100
AS_PATH: (40 60 200)
200.1.1.8/30      200.1.20.1    200.1.20.1      IGP  100
AS_PATH: path-is-empty
200.20.20.0/24    200.1.20.1    200.1.20.1      IGP  100
AS_PATH: path-is-empty
200.1.20.0/30     200.1.20.1    200.1.20.1      IGP  100
AS_PATH: path-is-empty

```

As shown in [Figure 97](#), the NextHop Address of 200.1.20.1 is displayed for all routes.

**2** Enter the following show command:

```
Passport-8610:5# show ip route info
```

[Figure 98](#) displays the output of the show ip route info command.

**Figure 98** show ip route info command output (Example 3B)

```

Passport-8610:5# show ip route info

=====
                        Ip Route
=====
DST                MASK                NEXT                COST  VLAN  PORT  PROT  AGE  TYPE  PRF
-----
200.1.1.4          255.255.255.252    200.1.1.5          1     0     -/-   LOC   0    DB   0
200.1.1.8          255.255.255.252    200.1.30.2         111   2065  2/2   OSPF  0    IB   20
200.1.20.0         255.255.255.252    200.1.30.2         101   2065  2/2   OSPF  0    IB   20
200.1.20.4         255.255.255.252    200.1.30.2         101   2065  2/2   OSPF  0    IB   20
200.1.30.0         255.255.255.252    200.1.30.1         1     2/2   LOC   0    DB   0

5 out of 5 Total Num of Dest Networks,5 Total Num of Route Entries displayed.
-----
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route,
E=Ecmp Route, U=Unresolved Route, N=Not in HW

```

Notice that the information in the `ip route info` table (Figure 98) remains the same, with no changes to the route information (for example, the BGP entries still do not appear in the IP routing table).

To display BGP routes in the IP routing table, you must disable BGP synchronization on R1 and R2. Because IGP is not synchronized with BGP, BGP entries are not entered into the IP forwarding table.

To disable BGP synchronization on R1 and R2, complete the following steps:

### *Disabling synchronization on R1*

#### **1** Disable BGP on R1:

You must first disable BGP administration state before you change the synchronization parameter value.

```
Passport-8610:5# config ip bgp disable
```

#### **2** Disable synchronization on R1:

```
Passport-8610:5# config ip bgp synchronization disable
```

**3** Enable BGP on R1:

```
Passport-8610:5# config ip bgp enable
```

*Disabling synchronization on R2*

**1** Disable BGP on R2:

You must first disable BGP administration state before you change the synchronization parameter value.

```
Passport-8610:5# config ip bgp disable
```

**2** Disable synchronization on R2:

```
Passport-8610:5# config ip bgp synchronization disable
```

**3** Enable BGP on R2:

```
Passport-8610:5# config ip bgp enable
```

*Viewing the BGP route table for R2*

To view the change to the topology, enter the following show command:

```
Passport-8610:5# show ip route info
```

As shown in [Figure 99](#), the output of the `show ip route info` command now displays the BGP routes.

**Figure 99** show ip route info command output (Example 3C)

```

Passport-8610:5# show ip route info
=====
                          Ip Route
=====
DST                MASK                NEXT                COST  VLAN  PORT  PROT  AGE  TYPE  PRF
-----
10.1.1.12          255.255.255.252     200.1.30.2         0     2065  2/2   BGP   0    IB   45
10.1.1.60          255.255.255.252     200.1.30.2         1     2065  2/2   BGP   0    IB   45
11.11.1.0          255.255.255.252     200.1.30.2         1     2065  2/2   BGP   0    IB   45
14.14.14.0         255.255.255.0       200.1.30.2         2     2065  2/2   BGP   0    IB   45
172.1.1.0          255.255.255.0       200.1.30.2         2     2065  2/2   BGP   0    IB   45
172.1.2.0          255.255.255.252     200.1.30.2         2     2065  2/2   BGP   0    IB   45
192.1.1.0          255.255.255.0       200.1.30.2         3     2065  2/2   BGP   0    IB   45
200.1.1.4          255.255.255.252     200.1.1.5          1      0     -/-   LOC   0    DB   0
200.1.1.8          255.255.255.252     200.1.30.2        111    2065  2/2   OSPF  0    IB   20
200.1.20.0         255.255.255.252     200.1.30.2        101    2065  2/2   OSPF  0    IB   20
200.1.20.4         255.255.255.252     200.1.30.2        101    2065  2/2   OSPF  0    IB   20
200.1.30.0         255.255.255.252     200.1.30.1         1      -     2/2   LOC   0    DB   0
200.20.20.0        255.255.255.0       200.1.30.2         0     2065  2/2   BGP   0    IB   45

13 out of 13 Total Num of Dest Networks,13 Total Num Route Entries displayed.
-----
TYPE Legend:
I=Indirect Route, D=Direct Route, A=Alternative Route, B=Best Route,
E=Ecmp Route, U=Unresolved Route, N=Not in HW

```

Although the route table problem in R2 is resolved, R2 may still not be able to communicate with any of the external networks. This is because R3 has no knowledge of these external routes.

One method to correct this problem is to redistribute BGP routes into OSPF (see [“Creating OSPF and BGP route distribution policies” on page 202](#)).

Another method is to enable BGP on R3.

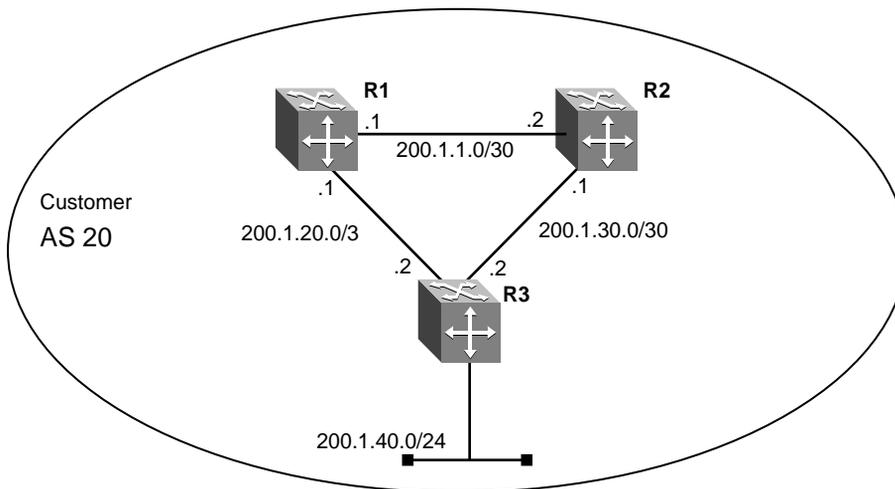
## MD5 authentication

The TCP MD5 authentication feature allows you to secure TCP connections that support BGP sessions. The TCP MD5 feature assigns an authentication key (also referred to as the *Password* or *Secret key*) to each BGP router, which then attaches a computed Message Digest 5 (MD5) signature to each TCP packet.

Figure 100 shows an example of an MD5 authentication configuration between R1 and R2. In this example, R1 and R2 are configured as peers. To communicate, both peers must be configured with the same password.

Following Figure 100, is a step-by-step procedure that shows how to configure R1 and R2 for this example.

**Figure 100** MD5 authentication example



Passport 8600

10877FA

To configure TCP MD5 authentication between R1 and R2, complete the following steps:



**Note:** This feature currently cannot be configured using Device Manager. You can configure this feature only using the CLI for this release.

### *Configure R1 for MD5 authentication*

- 1 Disable the administration state for the R1 peer:

MD5 authentication is set to disabled, by default. You must first disable the peer's administration state before you change the MD5 authentication parameter value.

```
Passport-8610:5# config ip bgp neighbor 200.1.1.2  
admin-state disable
```

- 2 Enable MD5 authentication on the R1 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.1.2  
MD5-authentication enable
```

- 3 Enter a password (the secret key) for the R1 peer:

The password value is a string length with a range 0 to 1536 alphanumeric characters.

```
Passport-8610:5# config ip bgp neighbor 200.1.1.2  
Password <password> add
```

- 4 Enable the administration state for the R1 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.1.2  
admin-state enable
```

### *Configure R2 for MD5 authentication*

- 1 Disable the administration state for the R1 peer:

MD5 authentication is set to disabled, by default. You must first disable the peer's administration state before you change the MD5 authentication parameter value.

```
Passport-8610:5# config ip bgp neighbor 200.1.1.1  
admin-state disable
```

- 2 Enable MD5 authentication on the R2 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.1.1  
MD5-authentication enable
```

**3** Enter a password (the secret key) for the R2 peer:

The password value is a string length with a range 0 to 1536 alphanumeric characters.

```
Passport-8610:5# config ip bgp neighbor 200.1.1.1  
Password <password> add
```

**4** Enable the administration state for the R2 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.1.1  
admin-state enable
```

<b>For more information about:</b>	<b>See:</b>
TCP MD5 authentication	<a href="#">“TCP MD5 message authentication” on page 44.</a>
Configuring TCP MD5 authentication using Device Manager	<a href="#">“Configuring and displaying peer information” on page 69.</a>
Configuring TCP MD5 authentication using the CLI	<a href="#">“Configuring BGP peers or peer groups” on page 129.</a>

## BGP peer groups

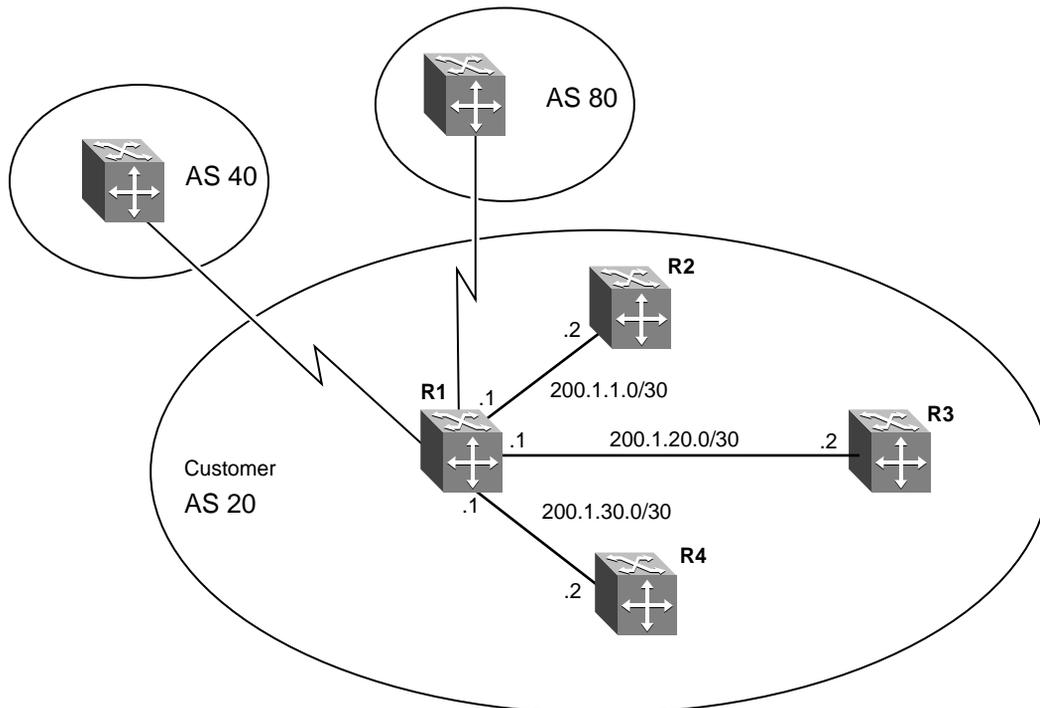
When you group existing BGP neighbors into peer groups, the changes that you make to the peer group are then applied to all current group members. This lets you simplify your configuration and make updates more efficient.



**Note:** Changes to a peer group only apply to the current members. This means that you have to add the neighbors before making peer group changes.

In the example shown in [Figure 101](#), the configuration policy for R1 is applied to all routers (R2, R3, and R4) within AS 20. Following the figure, is a step-by-step procedure that shows how to configure R1 for this example.

**Figure 101** Peer group example



Passport 8600

10878FA

## Configure a peer group on R1

For this example, R2, R3, and R4 are configured to use a BGP keepalive-time value of 60 and hold-time value of 180. When you configure the peer group on R1, you can change the keepalive-time and the hold-time values once, and then apply this peer group configuration to the peer group members (BGP neighbors R2, R3, and R4). If required, you can also add route policies to the peer group configuration.



**Note:** Changes to a peer group only apply to the current members. This means that you have to add the neighbors before making peer group changes.

---

To configure R1 for the example shown in [Figure 101](#), complete the following steps:

**1** Create the peer group (Group\_1):

Note that the assigned peer group name string is context-sensitive. For example, the name string “Group\_1,” is Not the same as “group\_1.”

```
Passport-8610:5# config ip bgp neighbor Group_1 create
```

**2** Create BGP peers:

```
Passport-8610:5# config ip bgp neighbor 200.1.1.2
create
```

```
Passport-8610:5# config ip bgp neighbor 200.1.20.2
create
```

```
Passport-8610:5# config ip bgp neighbor 200.1.30.2
create
```

**3** Add peers as members of Group\_1:

```
Passport-8610:5# config ip bgp neighbor 200.1.1.2
peer-group Group_1 add
```

```
Passport-8610:5# config ip bgp neighbor 200.1.20.2
peer-group Group_1 add
```

```
Passport-8610:5# config ip bgp neighbor 200.1.30.2
peer-group Group_1 add
```

**4** Assign Group\_1 to remote-as 20:

```
Passport-8610:5# config ip bgp neighbor Group_1
remote-as 20
```

**5** Assign the following parameter values to Group\_1:

```

Passport-8610:5# config ip bgp neighbor Group_1
keepalive-time 60 add
Passport-8610:5# config ip bgp neighbor Group_1
hold-time 180 add

```

**6** If required, add an existing policy (Pref\_AS20) to Group\_1:

```

Passport-8610:5# config ip bgp neighbor Group_1
route-policy in Pref_AS20 add

```

In this example, R2, R3, and R4 all use a BGP keepalive-time value of 60 and hold-time value of 180. By configuring a peer group on R1, you can change the keepalive-time and the hold-time values once, and then apply this peer group configuration to the peer group members (BGP neighbors R2, R3, and R4). If required, you can also add route policies to the peer group configuration.



**Note:** Changes to a peer group only apply to the current members. This means that you have to add the neighbors before making peer group changes.

For more information about:	See:
Peer concepts	<a href="#">“Autonomous systems” on page 22.</a>
Configuring peers and peer groups using Device Manager	<a href="#">“Configuring and displaying peer information” on page 69.</a>
Configuring peers and peer groups using the CLI	<a href="#">“Configuring BGP peers or peer groups” on page 129.</a>

## BGP path attributes

The Passport 8600 allows you to create route policies that control traffic flow. You can use the policies to control traffic over multiple connections, for inbound traffic from other ASs, and for outbound traffic that comes from outside a particular AS. You can create policies that control routes, work with default routing, control specific and aggregated routes, and manipulate BGP attributes.

There are four categories of BGP path attributes:

**1** Well-known mandatory

Mandatory attributes *must* be included in every BGP update message.

**2** Well-known discretionary

Discretionary attributes *may or may not* be sent in a particular BGP update message.

**3** Optional transitive

Transitive attributes are *accepted* and passed to other BGP peers.

**4** Optional non-transitive

Non-transitive attributes must be either accepted or ignored, *but must not be passed* along to other BGP peers.

Path attributes are used by border routers that utilize built-in algorithms or manually configured policies to select paths. BGP uses the following path attributes to control the path a BGP router chooses:

- Origin (well-known mandatory)
- AS\_path (well-known mandatory)
- Next Hop (well-known mandatory)
- Multi-Exit Discriminator Attribute (optional non-transitive)
- Local Preference (well-known discretionary)
- Atomic Aggregate (well-known discretionary)
- Aggregator (optional transitive)
- Community Attribute (optional transitive)

This section shows how you use BGP path attributes to create route policies. For more information about BGP path attributes, see [“BGP updates” on page 38](#).

This section includes the following topics:

- [“Origin attribute,” next](#)
- [“AS path attribute” on page 241](#)
- [“AS path filtering” on page 243](#)
- [“Local preference attribute” on page 246](#)
- [“Adding preferences to specific routes” on page 250](#)
- [“Multi-exit discriminator \(MED\) attribute” on page 253](#)
- [“Community attribute” on page 258](#)

## Origin attribute

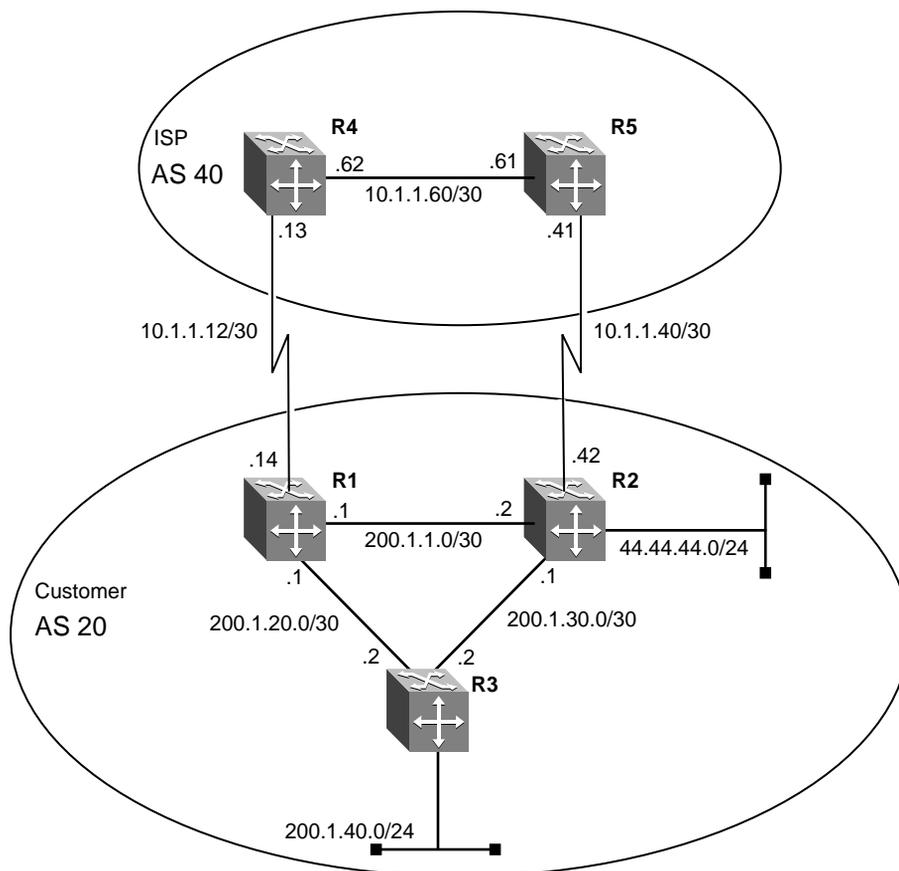
The Origin attribute is a well-known mandatory attribute that specifies the source of a route. The Origin is created by the AS that originates the route and includes the following possible values:

- IGP — the route is interior to the originating AS that inserts this route into the BGP table (0 = IGP).
- EGP — the route is learned via the Exterior Gateway Protocol (EGP) prior to being inserted into the BGP table (1 = BGP).
- Incomplete — the origin of the route is unknown or learned by some other means. For example, these routes could be learned through RIP, OSPF, or static routes (2 = Incomplete).

BGP uses the Origin attribute in its decision making process. BGP prefers the path with the lowest origin type; IGP is the lowest origin type, followed by EGP origin type, and Incomplete origin type.

In the example shown in [Figure 102](#), R1 is configured to distribute static routes. Following [Figure 102](#), is a step-by-step procedure that shows how to configure R1 for this example.

**Figure 102** Static route distribution



10879FA

## Configure R1

To configure R1, complete the following steps:

- 1 Create an IP static route policy on R1:

```
Passport-8610:5# config ip static-route create
44.44.44.0/255.255.255.0 next-hop 200.1.1.2 cost 1
preference 5
```

- 2 Create a static redistribution policy on R1:

```
Passport-8610:5# config ip bgp redistribute static create
Passport-8610:5# config ip bgp redistribute static enable
Passport-8610:5# config ip bgp redistribute static apply
```

**Note:** When you enable static route redistribution, the EBGP peer learns static routes with a BGP Origin attribute of Incomplete (INC).

- 3 Add network 200.1.1.0/30

```
Passport-8610:5# config ip bgp network 200.1.1.0/30 add
```

**Note:** When you inject networks with the BGP add network command, EBGP peers learn the network routes with a BGP Origin attribute of IGP.

## Displaying the static route from EPGP peer R4

To display the BGP route table, enter the following show command on R4:

```
Passport-8610:5# show ip bgp route
```

As shown in [Figure 103](#), the output of the `show ip bgp route` command now displays the BGP routes.

The **bolded** entry in the display shows the static route advertised from R1 with a BGP attribute of INC (incomplete). Because the source of the route is external to BGP, it is imported to BGP as INC.

Note also that when a route is imported using the network command (see Step 3, above), the route is considered an internal (IGP) route to BGP (see the second **bolded** entry in Figure 103).

**Figure 103** show ip bgp route command output

```

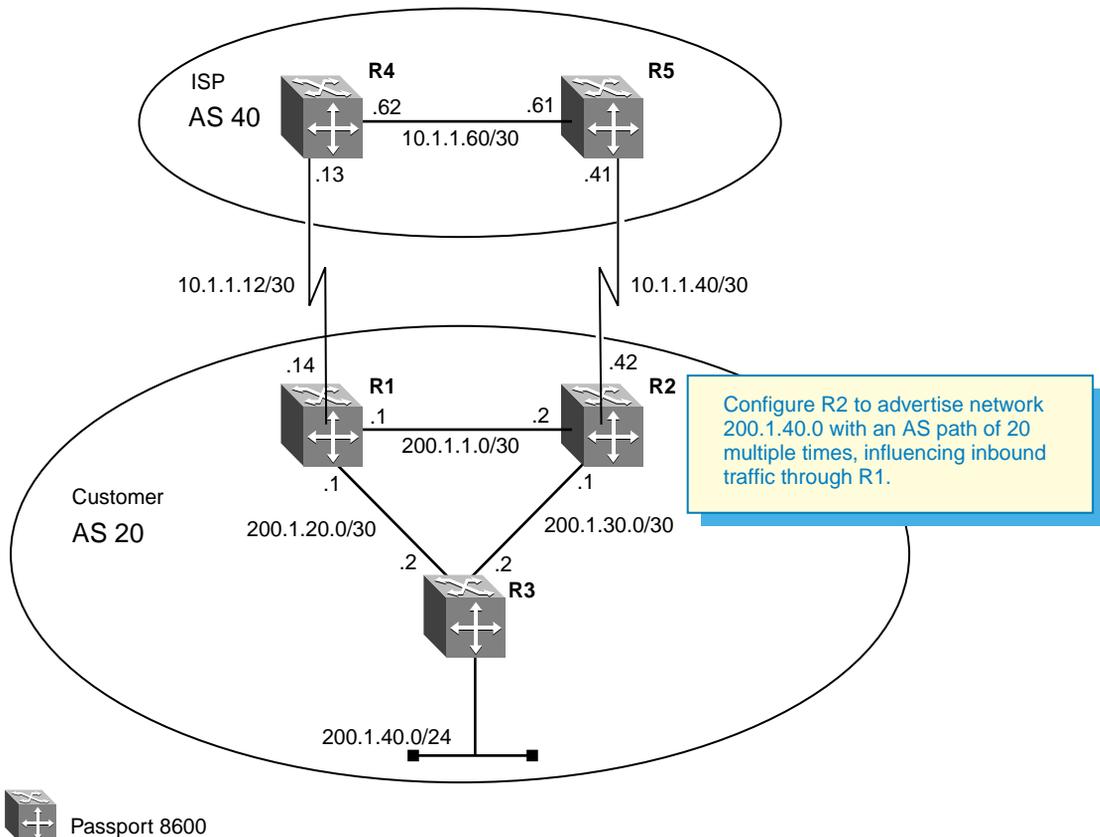
Passport-8610:5# show ip bgp route
Network/Mask      Peer Rem Addr    NextHop Address  Org  Loc  Pref  B/U  S1
-----
10.1.1.12/30     10.1.1.62      10.1.1.62       IGP  10   0     B/U  4
  As Path:
10.1.1.40/30     10.1.1.42      10.1.1.42       IGP  100  0     B    4
  As Path: <20>
10.1.1.40/30     10.1.1.62      10.1.1.62       IGP  0    0     B    4
  As Path:
10.1.1.60/30     10.1.1.62      10.1.1.62       IGP  0    0     B    4
  As Path:
11.11.1.0/30     10.1.1.62      10.1.1.62       IGP  0    0     B/U  4
  As Path:
44.44.44.0/24  10.1.1.42      10.1.1.42       INC  0    0     B    4
  As Path: <20>
172.1.1.0/24     10.1.1.62      11.11.1.2       IGP  10   0     B    4
  As Path: <80>
172.1.2.0/30     10.1.1.62      11.11.1.2       IGP  10   0     B    4
  As Path: <80>
192.1.1.0/24     14.14.14.2     14.14.14.2     IGP  100  0     B/U  4
  As Path: <200>
200.1.1.0/30  10.1.1.42      10.1.1.42       IGP  100  0     B/U  4
  As Path: <20>
200.1.1.0/30     10.1.1.62      10.1.1.14       IGP  10   0     B    4
  As Path: <20>
200.1.20.0/30    10.1.1.62      10.1.1.14       IGP  10   0     B    4
  As Path: <20>
200.1.30.0/30    10.1.1.42      10.1.1.42       IGP  100  0     B/U  4
  As Path: <20>
200.1.30.0/30    10.1.1.62      10.1.1.14       IGP  10   0     B    4
  As Path: <20>

```

## AS path attribute

Whenever a route passes from one AS to another, the new AS appends its AS number to the BGP update. This process creates an ordered list, referred to the *AS Sequence*. The AS Path attribute helps to ensure a loop-free topology. IBGP connections do not change the AS Path because the connections reside within a specific AS. Because BGP always prefers the shortest path, you can influence the incoming route selection when there are multiple paths to the local AS (referred to as AS-Path Prepending). You do this by manipulating the AS Path to a remote EBGP peer (see [Figure 104](#)).

**Figure 104** Using the AS Path attribute to influence inbound traffic flow



10880FA

For the configuration example shown in [Figure 104](#), R1 advertises the 200.1.40.0 network unchanged. R2 is configured to have its internal AS number inserted into the AS Path multiple times. This process influences all inbound traffic that is destined for the 200.1.40.0 network through R1.

The following section provides a step-by-step procedure that shows how to configure R2 for this example.

### *Configure R2*

- 1 Configure R2 to advertise network 200.1.40.0 for redistribution.

```
Passport-8610:5# config ip prefix-list 200.1.40.0
add-prefix 200.1.40.0/24
```

- 2 Configure a (multiple) IP AS list on R2:

```
Passport-8610:5# config ip as-list 1 create
Passport-8610:5# config ip as-list 1 add-as-path 1
permit "20 20 20"
```

- 3 Configure an IP route policy on R2:

```
Passport-8610:5# config ip route-policy AS_Prepend seq
1 create
Passport-8610:5# config ip route-policy AS_Prepend seq
1 enable
Passport-8610:5# config ip route-policy AS_Prepend seq
1 action permit
Passport-8610:5# config ip route-policy AS_Prepend seq
1 match-network 200.1.40.0
Passport-8610:5# config ip route-policy AS_Prepend seq
1 set-as-path 1
```

- 4 Configure a route policy that includes R5 as a peer:

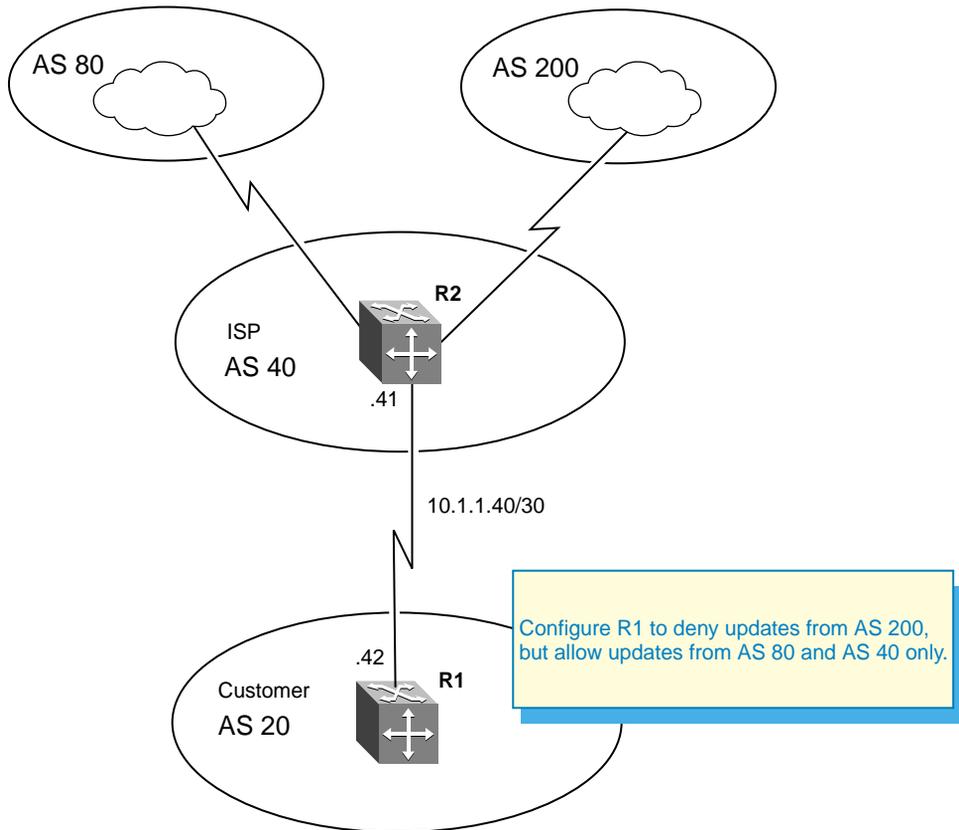
```
Passport-8610:5# config ip bgp neighbor 10.1.1.41
route-policy out AS_Prepend add
```

## AS path filtering

This section shows how you can use AS path filtering to set up specific access lists. In the example shown in [Figure 105](#), R1 is configured to deny any updates from AS 200, but allows updates from AS 40 and AS 80 only.

The following section provides a step-by-step procedure that shows how to configure R1 for this example.

**Figure 105** AS path filtering example



10879FA

## Configure R1

To configure R1, complete the following steps:

### 1 Configure an IP AS list on R1:

The following commands sets up an access list that denies updates from AS 200 (but allows updates from AS 40 and AS 80).

```
Passport-8610:5# config ip as-list 2 create
Passport-8610:5# config ip as-list 2 add-as-path 1
deny 200
```

```
Passport-8610:5# config ip as-list 3 create
Passport-8610:5# config ip as-list 3 add-as-path 1
permit 40
```

```
Passport-8610:5# config ip as-list 4 create
Passport-8610:5# config ip as-list 4 add-as-path 1
permit 80
```

### 2 Configure an IP route policy on R1:

```
Passport-8610:5# config ip route-policy AS seq 1 create
Passport-8610:5# config ip route-policy AS seq 1
enable
Passport-8610:5# config ip route-policy AS seq 1 action
deny
Passport-8610:5# config ip route-policy AS seq 1
match-as-path 2
```

```
Passport-8610:5# config ip route-policy AS seq 2 create
Passport-8610:5# config ip route-policy AS seq 2
enable
Passport-8610:5# config ip route-policy AS seq 2 action
permit
Passport-8610:5# config ip route-policy AS seq 2
match-as-path 3
Passport-8610:5# config ip route-policy AS seq 3 create
Passport-8610:5# config ip route-policy AS seq 3
enable
Passport-8610:5# config ip route-policy AS seq 3 action
permit
Passport-8610:5# config ip route-policy AS seq 3
match-as-path 4
```

- 3 Configure a route policy that includes R2 as a peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.41
route-policy in AS add
```

### *Alternate configuration method for R1*

The above configuration example is just one method of AS Path configuration. The same configuration can also be accomplished by using the following commands:

- 1 Configure an IP AS list on R1:

The following commands sets up an access list that denies updates from AS 200 (but allows updates from AS 40 and AS 80).

```
Passport-8610:5# config ip as-list 2 create
Passport-8610:5# config ip as-list 2 add-as-path 1
deny 200
Passport-8610:5# config ip as-list 2 add-as-path 2
permit 40
Passport-8610:5# config ip as-list 2 add-as-path 3
permit 80
```

- 2 Configure an IP route policy on R1:

```
Passport-8610:5# config ip route-policy AS seq 1 create
Passport-8610:5# config ip route-policy AS seq 1
enable
Passport-8610:5# config ip route-policy AS seq 1 action
permit
Passport-8610:5# config ip route-policy AS seq 1
match-as-path 2
```

- 3 Configure a route policy that includes R2 as a peer:

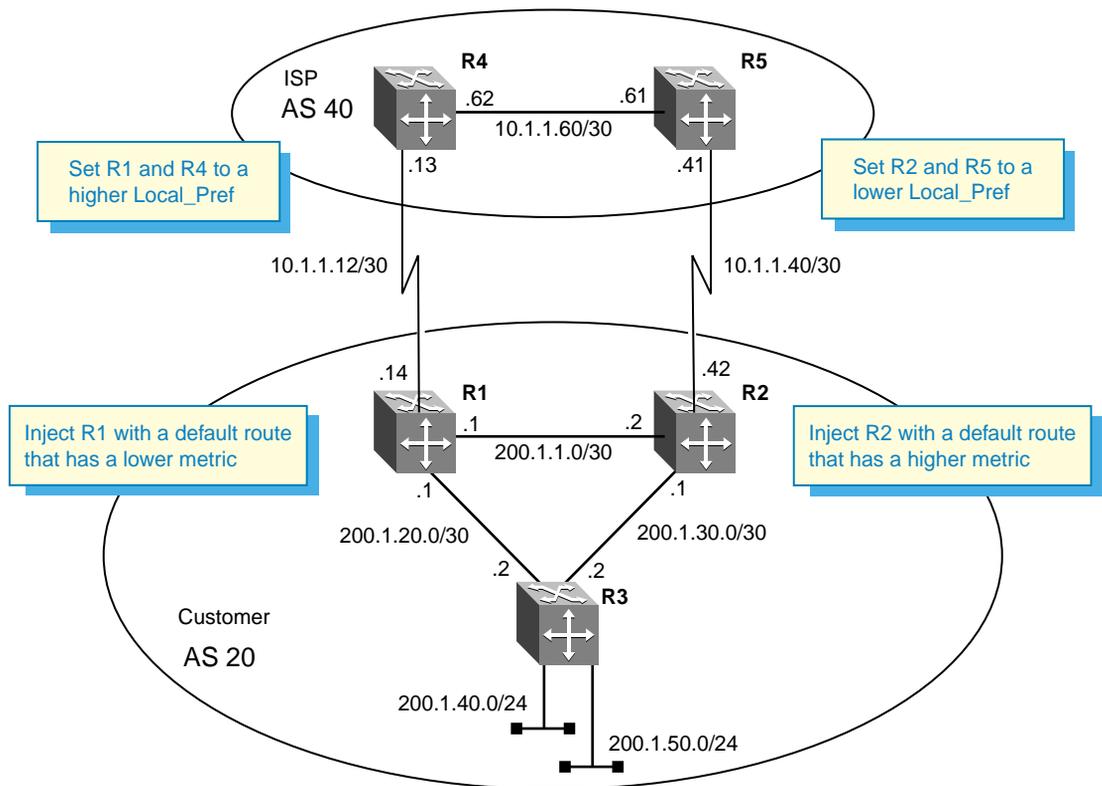
```
Passport-8610:5# config ip bgp neighbor 10.1.1.41
route-policy in AS add
```

## Local preference attribute

The local preference attribute is a well-known non-transitive attribute that influences the flow of outbound traffic by setting the exit point of an AS. Border routers within an AS calculate a local preference if the attribute is not configured in a BGP accept policy. The Local Preference attribute is local to ASs and is exchanged between IBGP peers only (it does not have any effect on the internal IGP protocol in use).

When BGP calculates a best route, and there are multiple paths to the same destination, the path with the higher preference is always chosen. For example, in the configuration shown in [Figure 106](#), you can influence the traffic to use the R4 to R1 path as the preferred path and use the R5 to R2 path for back up only.

**Figure 106** Using the local preference attribute to influence outbound traffic flow



In [Figure 106](#), R1 and R4 are configured with high local preference values, while R2 and R5 are configured with lower local preference values.

R1 is also configured to inject a default route with a lower OSPF metric than R2, which results in R1 having a higher preference. With this configuration, all traffic leaving AS 40 exits to the customer AS 20 through R4.

The local preference attribute can also be used to load-balance outbound traffic based on CIDR or network address groups.

The following sections provide step-by-step procedures that show how to configure the routers for the example shown in [Figure 106](#).

### *Configure R4*

- Configure a local preference value for R4:

Be sure to set the local preference value for R4 to a value that is higher than the local preference value you set for R5.

```
Passport-8610:5# config ip bgp default-local-pref 100 add
```

### *Configure R5*

- Configure a local preference value for R5:

Be sure to set the local preference value for R4 to a value that is lower than the local preference value you set for R4.

```
Passport-8610:5# config ip bgp default-local-pref 10 add
```

### *Configure R1*

- 1 Configure a local preference value for R1:

Be sure to set the local preference value for R1 to a value that is higher than the local preference value you set for R2.

```
Passport-8610:5# config ip bgp default-local-pref 100 add
```

**2** Configure the IP prefix list:

```
Passport-8610:5# config ip prefix-list DR add-prefix  
0.0.0.0/0
```

**Note:** For this example, DR represents the name used for this default route policy.

**3** Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy Default_OSPF  
Passport-8610:5# config ip route-policy seq 1 create  
Passport-8610:5# config ip route-policy seq 1 enable  
Passport-8610:5# config ip route-policy seq 1 action  
permit  
Passport-8610:5# config ip route-policy seq 1  
set-injectlist DR  
Passport-8610:5# config ip route-policy seq 1 set-metric  
100
```

**Note:** The set-metric value directly influences the OSPF route decision. For this example, R1 is set to a lower metric value than R2, which results in a higher preference value.

**4** Configure route redistribution:

```
Passport-8610:5# config ip ospf redistribute bgp create  
Passport-8610:5# config ip ospf redistribute bgp  
route-policy Default_OSPF enable  
Passport-8610:5# config ip ospf redistribute bgp  
route-policy Default_OSPF apply
```

## Configure R2

- 1 Configure a local preference value for R2:

Be sure to set the local preference value for R2 to a value that is lower than the local preference value you set for R1.

```
Passport-8610:5# config ip bgp default-local-pref 10 add
```

- 2 Configure the IP prefix list:

```
Passport-8610:5# config ip prefix-list DR add-prefix  
0.0.0.0/0
```

**Note:** For this example, DR represents the name used for this default route policy.

- 3 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy Default_OSPF  
Passport-8610:5# config ip route-policy seq 1 create  
Passport-8610:5# config ip route-policy seq 1 enable  
Passport-8610:5# config ip route-policy seq 1 action  
permit  
Passport-8610:5# config ip route-policy seq 1  
set-injectlist DR  
Passport-8610:5# config ip route-policy seq 1 set-metric  
300
```

**Note:** The set-metric value directly influences the OSPF route decision. For this example, R2 is set to a higher metric value than R1, which results in a lower preference value.

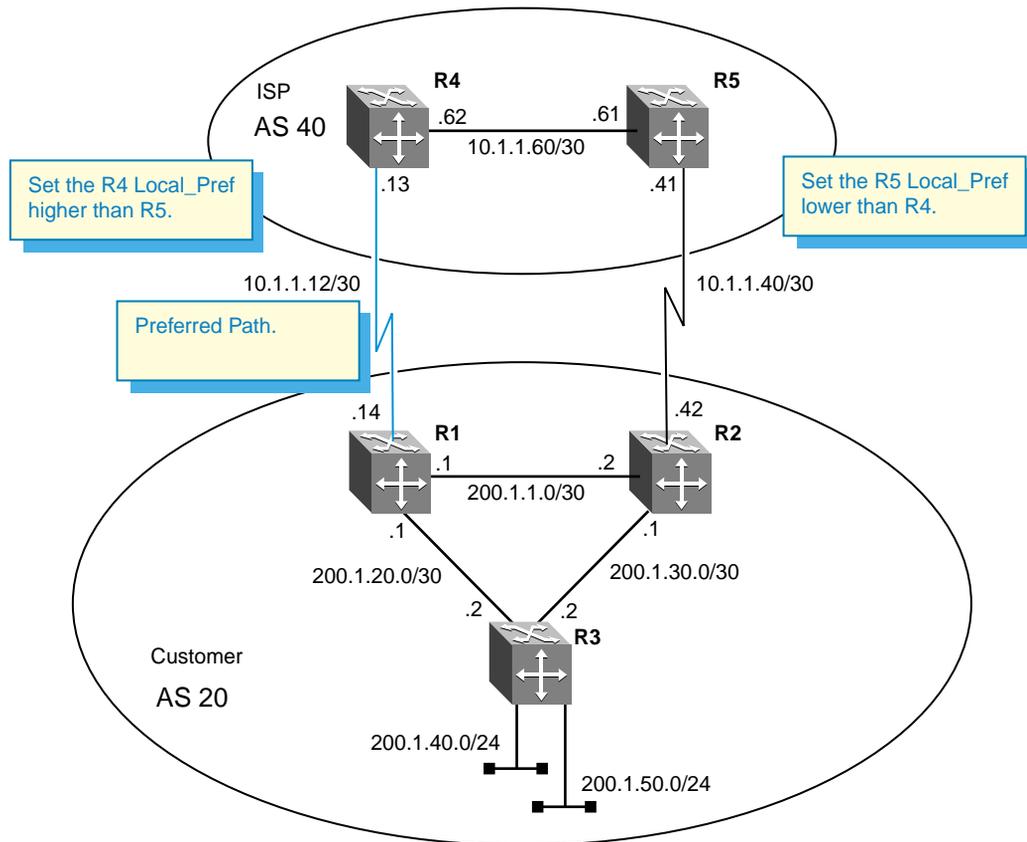
- 4 Configure route redistribution:

```
Passport-8610:5# config ip ospf redistribute bgp create  
Passport-8610:5# config ip ospf redistribute bgp  
route-policy Default_OSPF enable  
Passport-8610:5# config ip ospf redistribute bgp  
route-policy Default_OSPF apply
```

## Adding preferences to specific routes

The previous example shows how to configure the default local preference on the Passport 8600 to influence *all* networks. Alternatively, you can configure the Passport 8600 with a route policy that influences *only specified* networks. For example, to influence the traffic for network 200.1.40.0 to take the path between R4 and R1, you can configure a policy on R4 to have a higher Local Preference than R5 (Figure 107).

**Figure 107** Using the local preference attribute to influence specific routes



10883FA

The following sections provide step-by-step procedures that show how to configure R4 and R5 for the example shown in Figure 107.

## Configure R4

- 1 Configure the IP prefix list:

```
Passport-8610:5# config ip prefix-list 200.1.40.0/24
add-prefix 200.1.40.0/24
```

**Note:** For this example, `200.1.40.0/24` represents the name used for this prefix list. You can use any name.

- 2 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy Policy
```

**Note:** For this example, `Policy` represents the name used for this policy. You can use any name. (note that policy names are case-sensitive).

```
Passport-8610:5# config ip route-policy seq 1 create
Passport-8610:5# config ip route-policy seq 1 enable
Passport-8610:5# config ip route-policy seq 1 action
permit
Passport-8610:5# config ip route-policy seq 1
match-network 200.1.40.0
```

**Note:** For this example, `200.1.40.0` is the name of the IP Prefix List Configuration.

```
Passport-8610:5# config ip route-policy seq 1
set-local-pref 900
```

**Note:** For this example, the `set-local-pref` value must be set to a higher value than R5's `set-local-pref` value.

- 3 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy Policy
Passport-8610:5# config ip route-policy seq 2 create
```

**Note:** For this example, sequence 2 is required for the Route Policy above. If sequence 2 is not configured, then only the `200.1.40.0/24` network will be learned via BGP on R4.

```
Passport-8610:5# config ip route-policy seq 2 enable
Passport-8610:5# config ip route-policy seq 2 action
permit
```

- 4 Assign R1 as a peer to R4:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.14
route-policy in Policy add
```

### Configure R5

- 1 Configure the IP prefix list:

```
Passport-8610:5# config ip prefix-list 200.1.40.0/24
add-prefix 200.1.40.0/24
```

**Note:** For this example, `200.1.40.0/24` represents the name used for this prefix list. You can use any name.

- 2 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy Policy
```

**Note:** For this example, `Policy` represents the name used for this policy. You can use any name. (note that policy names are case-sensitive).

```
Passport-8610:5# config ip route-policy seq 1 create
Passport-8610:5# config ip route-policy seq 1 enable
Passport-8610:5# config ip route-policy seq 1 action
permit
Passport-8610:5# config ip route-policy seq 1
match-network 200.1.40.0
```

**Note:** For this example, `200.1.40.0` is the name of the IP Prefix List Configuration.

```
Passport-8610:5# config ip route-policy seq 1
set-local-pref 800
```

**Note:** For this example, the `set-local-pref` value must be set to a lower value than R4's `set-local-pref` value.

- 3 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy Policy
Passport-8610:5# config ip route-policy seq 2 create
```

**Note:** For this example, sequence 2 is required for the Route Policy above. If sequence 2 is not configured, then only the `200.1.40.0/24` network will be learned via BGP on R4.

```
Passport-8610:5# config ip route-policy seq 2 enable
Passport-8610:5# config ip route-policy seq 2 action
permit
```

- 4 Assign R5 as a peer to R2:

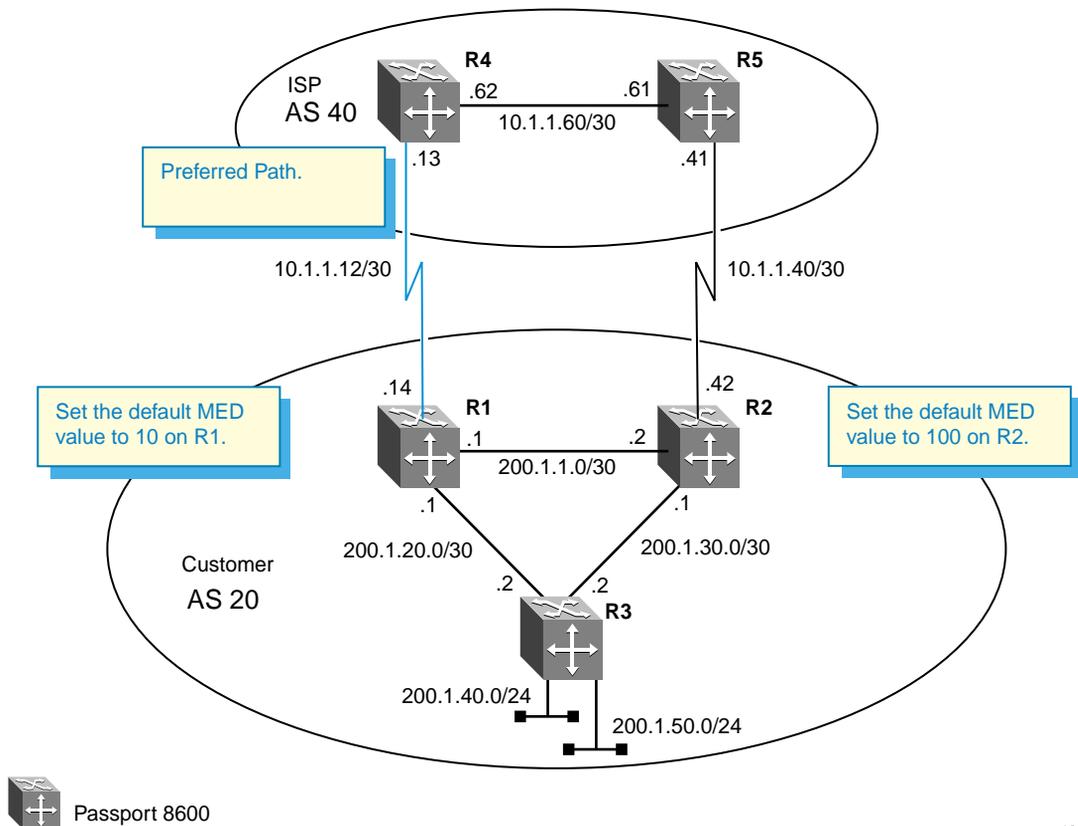
```
Passport-8610:5# config ip bgp neighbor 10.1.1.42
route-policy in Policy add
```

## Multi-exit discriminator (MED) attribute

The MED attribute is an optional non-transitive attribute that hints at preferred paths for routes that come from neighbors. The MEDs are only used with multiple connections to a neighboring AS in order to select a path for the return traffic. A lower MED value indicates a stronger path preference than a higher MED value. When an UPDATE message enters an AS with a certain MED value, that value is used to help the AS make the routing decision.

In the example shown in [Figure 108](#), R1's MED value is 10 and R2's MED value is 100. All traffic destined for AS 20 transverse over the 10.1.1.12 network.

**Figure 108** Using the MED attribute to influence specific routes



10884FA

The following sections provide step-by-step procedures that show how to configure R1 and R2 to complete the following tasks:

- “Influencing return traffic,” next
- “Load balancing traffic” on page 255

## Influencing return traffic

The following commands cause R1 to advertise all routes with a MED value of 10 and R2 to advertise all routes with a MED value of 100. As shown in [Figure 108](#), all traffic from AS 40 uses the preferred path over the 10.1.1.12 network.

### *Configure R1*

- Configure a MED value for R1:

Be sure to set the MED value for R1 to a value that is lower than the MED value you set for R2.

**Note:** A lower MED value indicates a stronger path preference than a higher MED value.

```
Passport-8610:5# config ip bgp default-metric 10 add
```

### *Configure R2*

- Configure a MED value for R2:

Be sure to set the MED value for R2 to a value that is higher than the MED value you set for R1.

**Note:** A higher MED value indicates a weaker path preference than a lower MED value.

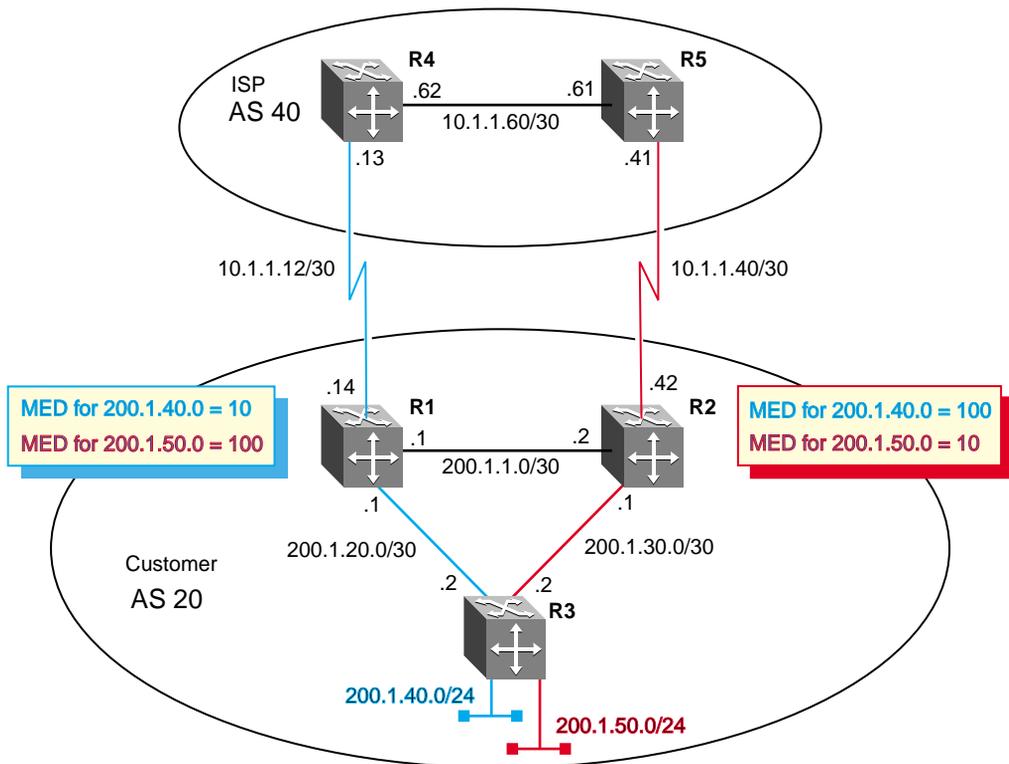
```
Passport-8610:5# config ip bgp default-metric 100 add
```

## Load balancing traffic

MEDs can also be used to load balance inbound traffic. For example, different MED values can be used to control different CIDR blocks.

This example shown in [Figure 109](#), shows how to influence all traffic destined for network 200.1.40.0/24 to transverse over R1, and for all traffic destined for 200.1.50.0/24 to transverse over R2.

**Figure 109** Using the MED attribute for load balancing routes



10885FA

For this example, R1 is configured to advertise the 200.1.40.0/24 network with a MED setting of 10 and the 200.1.50.0/24 network with a MED setting of 100.

R2 is configured with a policy that advertises the 200.1.40.0/24 network with a MED setting of 100 and the 200.1.50.0/24 network with a MED setting of 10.

This configuration results in all traffic destined for network 200.1.40.0/24 to transverse over R1, and traffic destined for network 200.1.50.0/24 to transverse over R2.

The following sections provide step-by-step procedures that show how to configure R1 and R2 for the example shown in [Figure 109](#).

### *Configure R1*

- 1 Configure the IP prefix list:

```
Passport-8610:5# config ip prefix-list 200.1.40.0
add-prefix 200.1.40.0/24
Passport-8610:5# config ip prefix-list 200.1.50.0
add-prefix 200.1.50.0/24
```

- 2 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy MED
Passport-8610:5# config ip route-policy seq 1 create
Passport-8610:5# config ip route-policy seq 1 enable
Passport-8610:5# config ip route-policy seq 1 action
permit
Passport-8610:5# config ip route-policy seq 1
match-network 200.1.40.0
Passport-8610:5# config ip route-policy seq 1 set-metric
10
```

```
Passport-8610:5# config ip route-policy MED
Passport-8610:5# config ip route-policy seq 2 create
Passport-8610:5# config ip route-policy seq 2 enable
Passport-8610:5# config ip route-policy seq 2 action
permit
Passport-8610:5# config ip route-policy seq 2
match-network 200.1.50.0
Passport-8610:5# config ip route-policy seq 1 set-metric
100
```

- 3 Assign R4 as a peer to R1:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.13
route-policy out MED add
```

---

## Configure R2

### 1 Configure the IP prefix list:

```
Passport-8610:5# config ip prefix-list 200.1.40.0
add-prefix 200.1.40.0/24
Passport-8610:5# config ip prefix-list 200.1.50.0
add-prefix 200.1.50.0/24
```

### 2 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy MED
Passport-8610:5# config ip route-policy seq 1 create
Passport-8610:5# config ip route-policy seq 1 enable
Passport-8610:5# config ip route-policy seq 1 action
permit
Passport-8610:5# config ip route-policy seq 1
match-network 200.1.40.0
Passport-8610:5# config ip route-policy seq 1 set-metric
100
```

```
Passport-8610:5# config ip route-policy MED
Passport-8610:5# config ip route-policy seq 2 create
Passport-8610:5# config ip route-policy seq 2 enable
Passport-8610:5# config ip route-policy seq 2 action
permit
Passport-8610:5# config ip route-policy seq 2
match-network 200.1.50.0
Passport-8610:5# config ip route-policy seq 2 set-metric
10
```

### 3 Assign R5 as a peer to R2:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.41
route-policy out MED add
```

## Community attribute

Community is an optional transitive attribute that groups destinations into communities to simplify policy administration in a BGP network. A community is a group of destinations that share a common administrative property.

The community attribute allows you to control your routing policies, with respect to destinations. It is common practice to create communities when you have more than one destination and want to share a common attribute.

The following are specific community types:

- internet — advertise this route to the Internet community
- No-export — do not advertise any destinations outside of a BGP confederation
- No Advertise — do not advertise to any BGP peer including IBGP peers
- No Export Subconfed — do not advertise to external BGP peers, even within the same confederation.



**Note:** For the community type “no export subconfed,” the Passport 8600 uses a setting of “local-as.”

---

You can use the community attribute to control which routing information to accept, prefer, or distribute to other BGP neighbors. If you specify the append option in the route policy, the specified community value is added to the existing value of the community attribute. Otherwise, the specified community value replaces any community value that was set previously.

[Figure 110](#) shows an example of how you can use the community attributes to control routing information that will be distributed to neighbors.

For this example, suppose you want to advertise network 200.1.20.0 beyond AS 40, however you do not want to advertise network 200.30.30.0. You can use the community attributes to configure R1, as follows:



**Note:** The Passport 8600 uses an IP Community List policy to specify the community.

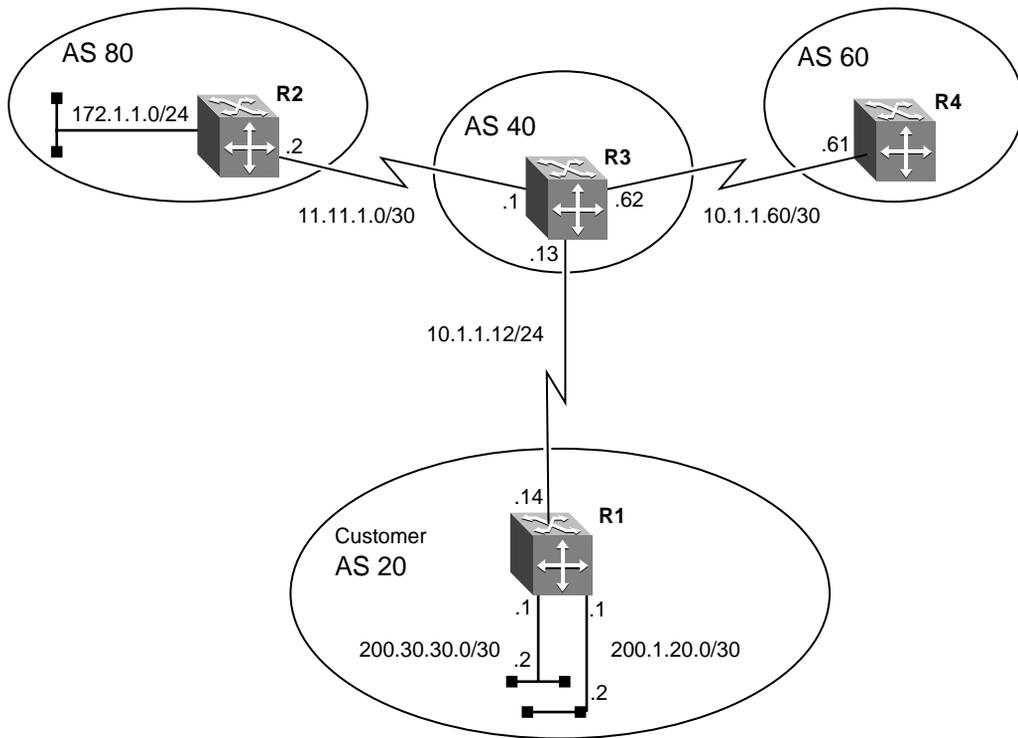
---

- Configure R1 with a community attribute of “no-export” to R3 in AS 40 for network 200.30.30.0.
- Configure R1 with a community attribute of “internet” for all other networks (network 200.1.20.0).

This informs R3 to not propagate the 200.30.30.0 network, but advertise all other routes learned from R1.

The following section provides a step-by-step procedure that show how to configure R1 for the example shown in [Figure 110](#).

**Figure 110** Using the community attribute to control routes



10886FA

## Configure R1

To enable the BGP peer to send community attributes, complete the following steps:

- 1 Configure the IP prefix list:

```
Passport-8610:5# config ip prefix-list 200.30.30.0
add-prefix 200.30.30.0/30
```

- 2 Configure the IP community list policy:

```
Passport-8610:5# config ip community-list 1 create
Passport-8610:5# config ip community-list 1 add-community
memberId 1 permit community-string 55:55
Passport-8610:5# config ip community-list 1 add-community
memberId 2 permit community-string no-export
```

```
Passport-8610:5# config ip community-list 2 create
Passport-8610:5# config ip community-list 2 add-community
memberId 1 permit community-string 55:55
Passport-8610:5# config ip community-list 2 add-community
memberId 2 permit community-string internet
```

- 3 Configure the IP Route Policy:

```
Passport-8610:5# config ip route-policy community
Passport-8610:5# config ip route-policy seq 1 create
Passport-8610:5# config ip route-policy seq 1 enable
Passport-8610:5# config ip route-policy seq 1 action
permit
Passport-8610:5# config ip route-policy seq 1
match-network 200.30.30.0
Passport-8610:5# config ip route-policy seq 1
set-community 1
Passport-8610:5# config ip route-policy seq 1
set-community-mode additive
```

**Note:** The following options are available for the set-community-mode:

```
set-community-mode <unchanged|additive|none>
```

- *unchanged* – do not change an exiting community
- *additive* – append the community to the exiting community
- *none* – remove the community

```

Passport-8610:5# config ip route-policy community
Passport-8610:5# config ip route-policy seq 2 create
Passport-8610:5# config ip route-policy seq 2 enable
Passport-8610:5# config ip route-policy seq 2 action
permit
Passport-8610:5# config ip route-policy seq 2
set-community 2
Passport-8610:5# config ip route-policy seq 2
set-community-mode additive

```

**Note:** The following options are available for the set-community-mode:

```
set-community-mode <unchanged|additive|none>
```

- *unchanged* – do not change an exiting community
- *additive* – append the community to the exiting community
- *none* – remove the community

#### 4 Assign R3 as a peer to R1:

```

Passport-8610:5# config ip bgp neighbor 10.1.1.13
admin-state disable
Passport-8610:5# config ip bgp neighbor 10.1.1.13
send-community enable
Passport-8610:5# config ip bgp neighbor 10.1.1.13
route-policy out community add
Passport-8610:5# config ip bgp neighbor 10.1.1.13
admin-state enable

```

To display the route table, enter the following show command on R3:

```
Passport-8610:5# show ip bgp route community enable
```

As shown in [Figure 111](#), the output of the `show ip bgp route community enable` command now displays the BGP routes. The output indicates that network 200.30.30.0 will not be advertised outside AS 40 while the 200.1.20.0 will be advertised outside AS 40.

**Figure 111** show ip bgp route community enable command output

```

Passport-8610:5# show ip bgp route community enable
The total number of routes is 11

Network/Mask      Peer Rem Addr   NextHop Address Org LocPref
-----
16.16.16.16/30    10.1.1.61      10.1.1.61      IGP 100
  AS_PATH: (60)
  COMMUNITY: no-community-attr
10.1.1.12/30      10.1.1.14      10.1.1.14      IGP 100
  AS_PATH: (20)
  COMMUNITY: 55:55 internet
14.14.14.0/24     10.1.1.61      10.1.1.61      IGP 100
  AS_PATH: (60)
  COMMUNITY: no-community-attr
10.1.1.40/30      10.1.1.61      10.1.1.61      IGP 100
  AS_PATH: (60)
  COMMUNITY: no-community-attr
172.1.1.0/24      11.11.1.2      11.11.1.2      IGP 100
  AS_PATH: (80)
  COMMUNITY: no-community-attr
172.1.2.0/30      11.11.1.2      11.11.1.2      IGP 100
  AS_PATH: (80)
  COMMUNITY: no-community-attr
192.1.1.0/24      10.1.1.61      10.1.1.61      IGP 100
  AS_PATH: (60 200)
  COMMUNITY: no-community-attr
200.30.30.0/30    10.1.1.14      10.1.1.14      IGP 100
  AS_PATH: (20)
  COMMUNITY: 55:55 no-export
200.1.1.4/30      10.1.1.61      10.1.1.61      IGP 100
  AS_PATH: (60 20)
  COMMUNITY: no-community-attr
200.1.30.0/30     10.1.1.61      10.1.1.61      IGP 100
  AS_PATH: (60 20)
  COMMUNITY: no-community-attr
200.1.20.0/30     10.1.1.14      10.1.1.14      IGP 100
  AS_PATH: (20)
  COMMUNITY: 55:55 internet

```

<b>For more information about:</b>	<b>See:</b>
BGP path attributes	<a href="#">“BGP updates” on page 38.</a>

## IBGP scalability issues

In an AS, an IGP protocol such as OSPF or RIP is used to provide routes to destinations. BGP can be configured to inject the routes it learns into the IGP to make these destinations known. However, IGPs do not understand or recognize BGP attributes such as AS path information.

In order to preserve and update BGP attributes, IBGP connections between border routers must be “fully-meshed.” Any external routing information must be re-distributed to all other routers with the AS. As the number of IBGP speakers increases, this full mesh requirement does not scale very well. With many border routers with many routes, IBGP peering can become an issue for resources such as CPU, bandwidth, and configuration management.

Because of scalability, BGP speakers within an AS must maintain  $n*(n-1)/2$  unique IBGP sessions.

BGP confederations and route reflectors can be used to eliminate the full-mesh scaling problem by minimizing the number of necessary peer sessions.

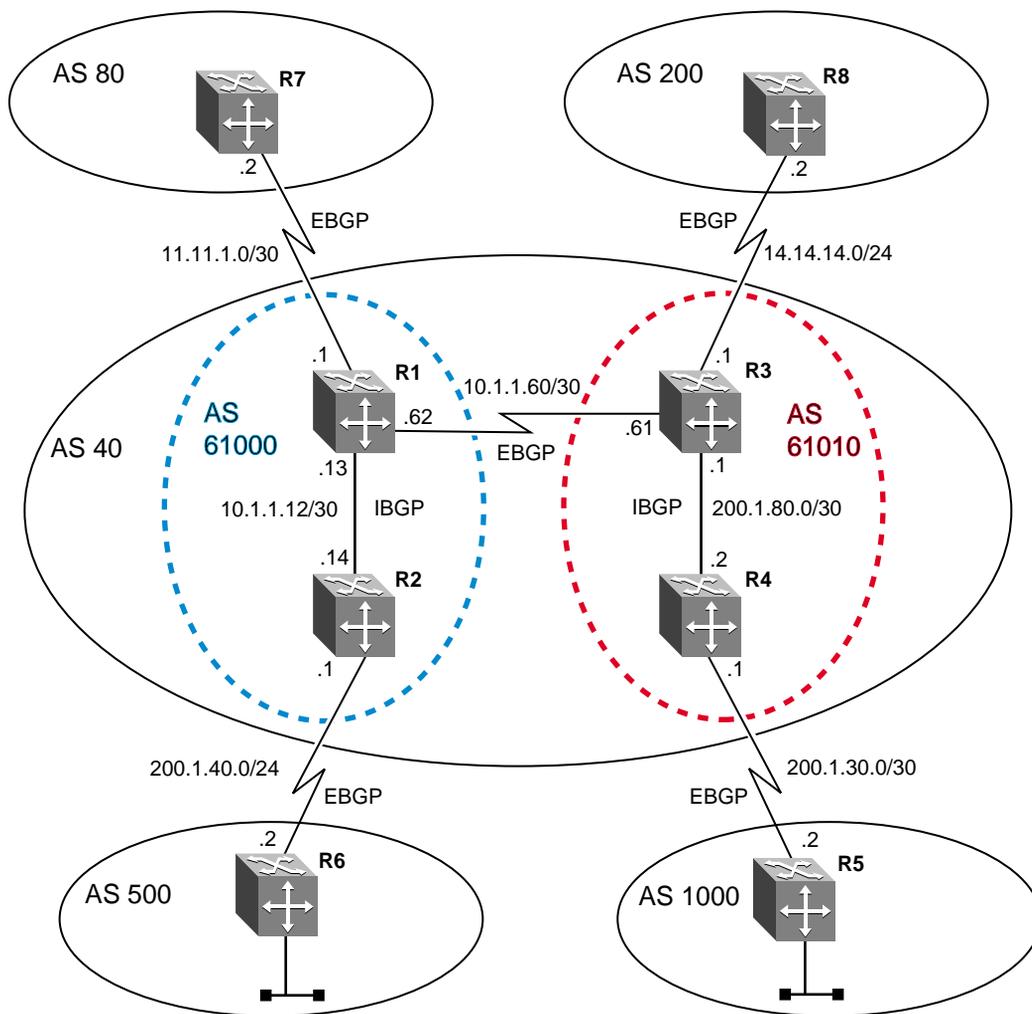
This section includes the following topics:

- [“BGP confederations,”](#) next
- [“Route reflectors”](#) on page 269
- [“Multiple route reflectors”](#) on page 274

## BGP confederations

You can reduce the high bandwidth and maintenance costs associated with a large full-mesh topology by dividing the AS into multiple smaller ASs (sub-ASs), and then grouping them into a single confederation (see [Figure 112](#)). The confederations reduce the total number of required peers within the AS.

**Figure 112** Confederation example



10887FA

As shown in [Figure 112 on page 264](#), IBGP speakers within each sub-system AS establish peer sessions only with other speakers in their own sub-system. One speaker from each sub-system establishes EBGP peer sessions with a single speaker from each of the other sub-systems. Although there can be multiple smaller sub-system ASs within the confederation, to the outside world, the confederation appears to be a single AS. Without confederations, all the routers in AS 40 must be fully meshed.

All routers within the confederation AS are fully meshed. Each confederation AS has a connection to the other confederation ASs and use EBGP to exchange routing updates. Even though EBGP is used between confederation ASs, the routing information exchanged is treated as if they are using IBGP. This preserves all the various IBGP information such as local preference and MED.

The following sections provides a step-by-step procedure that shows how to configure R1, R2, R3, and R4 for the example shown in [Figure 112 on page 264](#).

## R1 Configuration

### *Configure BGP on R1*

- 1 Assign R1 to AS 61000:

```
Passport-8610:5# config ip bgp local-as 61000
Passport-8610:5# config ip bgp confederation identifier
40 add
Passport-8610:5# config ip bgp confederation peers 61010
```

- 2 Enable BGP on R1:

```
Passport-8610:5# config ip bgp enable
```

### *Configure BGP peer interfaces for R1*

- 1 Assign neighboring router (R7) as an R1 peer:

```
Passport-8610:5# config ip bgp neighbor 11.11.1.2 create
Passport-8610:5# config ip bgp neighbor 11.11.1.2
remote-as 80
Passport-8610:5# config ip bgp neighbor 11.11.1.2
admin-state enable
```

- 2 Assign neighboring router (R3) as an R1 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.61 create
```

```
Passport-8610:5# config ip bgp neighbor 10.1.1.61
remote-as 61010
Passport-8610:5# config ip bgp neighbor 110.1.1.61
admin-state enable
```

- 3 Assign neighboring router (R2) as an R1 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.14 create
Passport-8610:5# config ip bgp neighbor 10.1.1.14
remote-as 61000
Passport-8610:5# config ip bgp neighbor 10.1.1.14
admin-state enable
```

## R2 Configuration

### *Configure BGP on R2*

- 1 Assign R2 to AS 61000:

```
Passport-8610:5# config ip bgp local-as 61000
Passport-8610:5# config ip bgp confederation identifier
40 add
```

- 2 Enable BGP on R2:

```
Passport-8610:5# config ip bgp enable
```

### *Configure BGP peer interfaces for R2*

- 1 Assign neighboring router (R1) as an R2 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.13 create
Passport-8610:5# config ip bgp neighbor 10.1.1.13
remote-as 61000
Passport-8610:5# config ip bgp neighbor 10.1.1.13
admin-state enable
```

- 2 Assign neighboring router (R6) as an R2 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.40.2 create
Passport-8610:5# config ip bgp neighbor 200.1.40.2
remote-as 500
Passport-8610:5# config ip bgp neighbor 200.1.40.2
admin-state enable
```

## R3 Configuration

### *Configure BGP on R3*

- 1 Assign R3 to AS 61000:

```
Passport-8610:5# config ip bgp local-as 61010
Passport-8610:5# config ip bgp confederation identifier
40 add
Passport-8610:5# config ip bgp confederation peers 61000
```

- 2 Enable BGP on R3:

```
Passport-8610:5# config ip bgp enable
```

### *Configure BGP peer interfaces for R3*

- 1 Assign neighboring router (R8) as an R3 peer:

```
Passport-8610:5# config ip bgp neighbor 14.14.14.2 create
Passport-8610:5# config ip bgp neighbor 14.14.14.2
remote-as 200
Passport-8610:5# config ip bgp neighbor 14.14.14.2
admin-state enable
```

- 2 Assign neighboring router (R1) as an R3 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.62 create
Passport-8610:5# config ip bgp neighbor 10.1.1.62
remote-as 61000
Passport-8610:5# config ip bgp neighbor 10.1.1.62
admin-state enable
```

- 3 Assign neighboring router (R4) as an R3 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.80.1 create
Passport-8610:5# config ip bgp neighbor 200.1.80.1
remote-as 61010
Passport-8610:5# config ip bgp neighbor 200.1.80.1
admin-state enable
```

## R4 Configuration

### *Configure BGP on R4*

- 1 Assign R4 to AS 61010:

```
Passport-8610:5# config ip bgp local-as 61010
Passport-8610:5# config ip bgp confederation identifier
40 add
```

- 2 Enable BGP on R4:

```
Passport-8610:5# config ip bgp enable
```

### *Configure BGP peer interfaces for R4*

- 1 Assign neighboring router (R3) as an R4 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.80.2 create
Passport-8610:5# config ip bgp neighbor 200.1.80.2
remote-as 61010
Passport-8610:5# config ip bgp neighbor 200.1.80.2
admin-state enable
```

- 2 Assign neighboring router (R5) as an R4 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.30.2 create
Passport-8610:5# config ip bgp neighbor 200.1.30.2
remote-as 1000
Passport-8610:5# config ip bgp neighbor 200.1.30.2
admin-state enable
```

<b>For more information about:</b>	<b>See:</b>
BGP confederation concepts	<a href="#">“Confederations” on page 32.</a>
Configuring confederations using Device Manager	<a href="#">“Configuring general parameters” on page 56.</a>
Configuring confederations using the CLI	<a href="#">“Configuring BGP confederations” on page 128.</a>

## Route reflectors

Route reflectors (RRs) provide another alternative for reducing the number of IBGP peers within an AS. Route reflectors allow routers to advertise or reflect IBGP routes to other IBGP speakers (for more information about route reflectors, see [“IBGP route reflection” on page 34](#)).

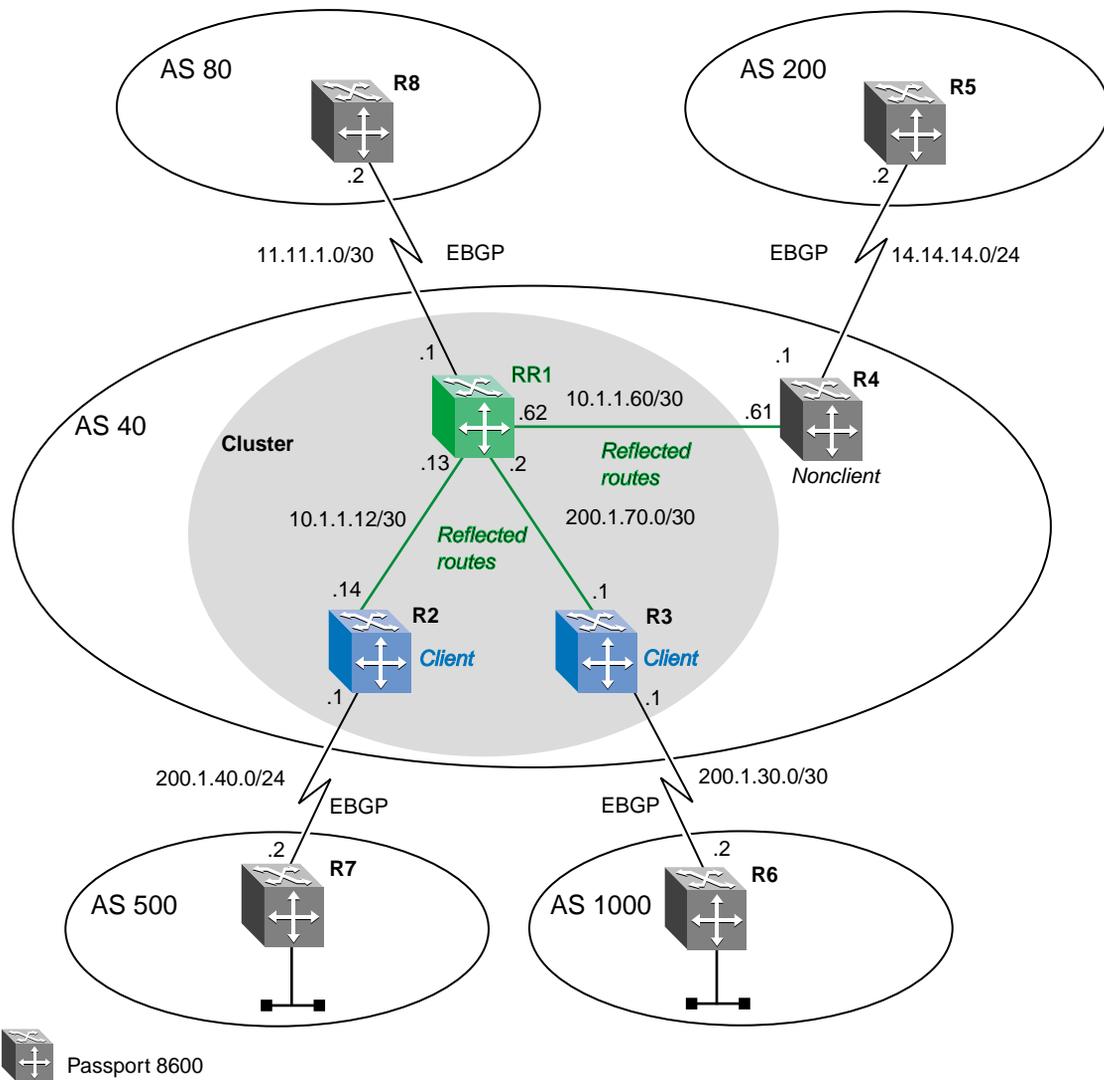
As shown in [Figure 113](#), the internal peers of route reflectors are divided into two groups: client peers and nonclient peers. A route reflector reflects routes between these two groups. The nonclient peers must be fully meshed, while the client peers do not need to be fully meshed.

Without a route reflector (RR1 in [Figure 113](#)), all routers in AS 40 would require a full IBGP mesh. For example, R3 would require IBGP peering with RR1, R2, and R4. With route reflection configured on RR1, IBGP peering on R3 is no longer required to R2 and R4.

When a router is configured with a route reflector, the configuration also includes the route reflector client configuration. The route reflector can also be configured to control whether the routes learned by a client are to be forwarded to other clients.

A route reflector and all its clients are called a cluster. Other IBGP peers of the route reflector, that are not route reflector clients, are called nonclients. In this example, RR1 is the route reflector, R2 and R3 are route reflector clients, and R4 is a nonclient.

In an AS, there can be more than one route reflector cluster. There can also be more than one route reflector in a cluster. When there is more than one reflector in a cluster, special care must be taken to prevent route loops.

**Figure 113** Route reflector example

The following sections provides a step-by-step procedure that shows how to configure RR1, R2, and R3 for the example shown in [Figure 113](#).

## RR1 Configuration

### *Disable synchronization on RR1*

#### 1 Disable BGP on RR1:

You must first disable BGP administration state before you change the synchronization parameter value.

```
Passport-8610:5# config ip bgp disable
```

#### 2 Disable synchronization on RR1:

```
Passport-8610:5# config ip bgp synchronization disable
```

#### 3 Assign RR1 to AS 40:

```
Passport-8610:5# config ip bgp local-as 40
```

#### 4 Enable route reflection on RR1:

```
Passport-8610:5# config ip bgp route-reflection enable
Passport-8610:5# config ip bgp cl-to-cl-reflection enable
```

**Note:** The `cl-to-cl-reflection` parameter lets you enable or disable the route reflector clients from distributing their respective BGP routes.

#### 5 Enable BGP on RR1:

```
Passport-8610:5# config ip bgp enable
```

### *Configure IGP network prefixes on RR1*

➤ Configure BGP networks that you want RR1 to advertise for redistribution.

```
Passport-8610:5# config ip bgp network 11.11.1.0/30 add
Passport-8610:5# config ip bgp network 10.1.1.60/30 add
```

### *Configure BGP peer interfaces for RR1*

#### 1 Assign neighboring router (R2) as an RR1 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.14 create
Passport-8610:5# config ip bgp neighbor 10.1.1.14
remote-as 40
Passport-8610:5# config ip bgp neighbor 10.1.1.14
route-reflector-client enable
Passport-8610:5# config ip bgp neighbor 10.1.1.14
admin-state enable
```

- 2 Assign neighboring router (R3) as an RR1 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.70.1 create
Passport-8610:5# config ip bgp neighbor 200.1.70.1
remote-as 40
Passport-8610:5# config ip bgp neighbor 200.1.70.1
route-reflector-client enable
Passport-8610:5# config ip bgp neighbor 200.1.70.1
admin-state enable
```

- 3 Assign neighboring router (R4) as an RR1 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.61 create
Passport-8610:5# config ip bgp neighbor 10.1.1.61
remote-as 40
Passport-8610:5# config ip bgp neighbor 10.1.1.61
admin-state enable
```

- 4 Assign neighboring router (R8) as an RR1 peer:

```
Passport-8610:5# config ip bgp neighbor 11.11.1.2 create
Passport-8610:5# config ip bgp neighbor 11.11.1.2
remote-as 80
Passport-8610:5# config ip bgp neighbor 11.11.1.2
admin-state enable
```

## R2 Configuration

### *Disable synchronization on R2*

- 1 Disable BGP on R2:

You must first disable BGP administration state before you change the synchronization parameter value.

```
Passport-8610:5# config ip bgp disable
```

- 2 Disable synchronization on R1:

```
Passport-8610:5# config ip bgp synchronization disable
```

- 3 Assign R2 to AS 40:

```
Passport-8610:5# config ip bgp local-as 40
```

- 4 Enable BGP on R1:

```
Passport-8610:5# config ip bgp enable
```

## Configure BGP peer interface for R2

- Assign neighbor (RR1) as an R1 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.13 create
Passport-8610:5# config ip bgp neighbor 10.1.1.13
remote-as 40
Passport-8610:5# config ip bgp neighbor 10.1.1.13
admin-state enable
```

- 5 Assign neighbor (R7) as an R2 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.40.2 create
Passport-8610:5# config ip bgp neighbor 200.1.40.2
remote-as 1000
Passport-8610:5# config ip bgp neighbor 200.1.40.2
admin-state enable
```

## R3 Configuration

### Disable synchronization on R3

- 1 Disable BGP on R3:

You must first disable BGP administration state before you change the synchronization parameter value.

```
Passport-8610:5# config ip bgp disable
```

- 2 Disable synchronization on R3:

```
Passport-8610:5# config ip bgp synchronization disable
```

- 3 Assign R3 to AS 40:

```
Passport-8610:5# config ip bgp local-as 40
```

- 4 Enable BGP on R3:

```
Passport-8610:5# config ip bgp enable
```

### Configure BGP peer interface for R3

- 1 Assign neighbor (RR1) as an R3 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.70.2 create
Passport-8610:5# config ip bgp neighbor 200.1.70.2
remote-as 40
Passport-8610:5# config ip bgp neighbor 200.1.70.2
admin-state enable
```

**2** Assign neighbor (R6) as an R3 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.30.2 create
Passport-8610:5# config ip bgp neighbor 200.1.30.2
remote-as 1000
Passport-8610:5# config ip bgp neighbor 200.1.30.2
admin-state enable
```

## Multiple route reflectors

Normally, in a route reflector cluster there is only one route reflector (RR) that is identified by the router ID. To increase resilience, a second route reflector can be installed (see [Figure 114 on page 275](#)).

When you install more than one route reflector in a cluster, the cluster must be configured with a 4-byte cluster ID. The cluster ID allows the route reflectors to recognize updates from other route reflectors in the same cluster. The cluster ID is also appended to all routes that are sent outside its cluster. If a route reflector receives an update that contains a cluster ID that is the same as the local customer ID, the update is dropped, preventing route loops.

The following sections provides a step-by-step procedure that shows how to configure RR1, RR2, R3, and R4 for the example shown in [Figure 114 on page 275](#).

### RR1 Configuration

To configure route reflector RR1, complete the following steps:

#### *Disable synchronization on RR1*

**1** Disable BGP on RR1:

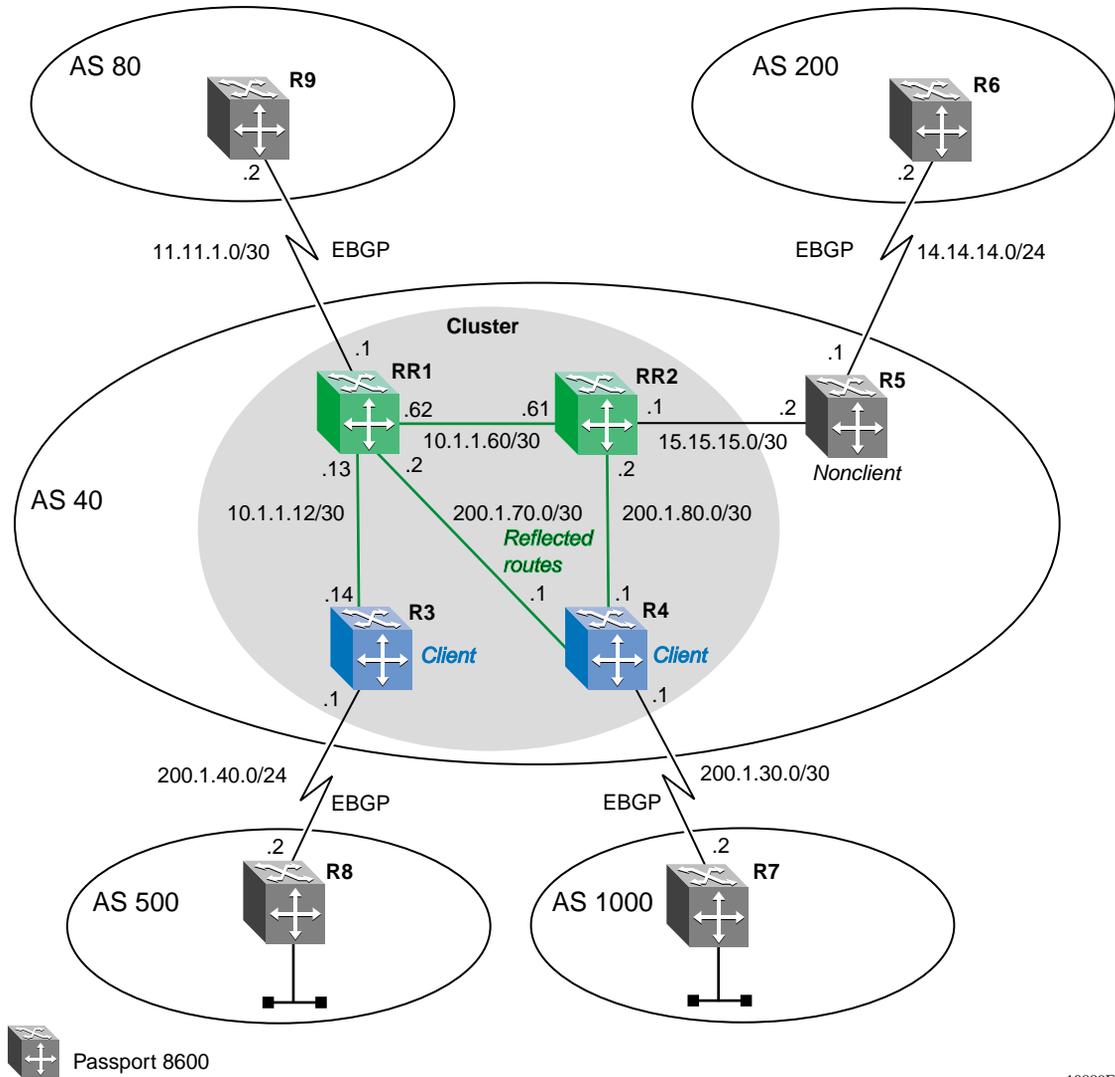
You must first disable the BGP administration state before you change the synchronization parameter value.

```
Passport-8610:5# config ip bgp disable
```

**2** Disable synchronization on RR1:

```
Passport-8610:5# config ip bgp synchronization disable
```

Figure 114 Multiple route reflector example



### 3 Assign RR1 to AS 40:

```
Passport-8610:5# config ip bgp local-as 40
```

**4** Enable route reflection on RR1:

```
Passport-8610:5# config ip bgp route-reflection enable
Passport-8610:5# config ip bgp cl-to-cl-reflection enable
```

**Note:** The `cl-to-cl-reflection` parameter lets you enable or disable the route reflector clients from distributing their respective BGP routes.

**5** Assign a 4-byte cluster ID to RR1:

The 4-byte cluster ID lets route reflectors recognize updates from *other* route reflectors in the same cluster.

```
Passport-8610:5# config ip bgp cluster-id 0.0.0.20 add
```

**6** Enable BGP on RR1:

```
Passport-8610:5# config ip bgp enable
```

*Configure BGP peer interfaces for RR1***1** Assign neighboring router (R9) as an RR1 peer:

```
Passport-8610:5# config ip bgp neighbor 11.11.1.2 create
Passport-8610:5# config ip bgp neighbor 11.11.1.2
remote-as 80
Passport-8610:5# config ip bgp neighbor 10.1.1.14
admin-state enable
```

**2** Assign neighboring router (RR2) as an RR1 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.61 create
Passport-8610:5# config ip bgp neighbor 10.1.1.61
remote-as 40
Passport-8610:5# config ip bgp neighbor 10.1.1.61
admin-state enable
```

**3** Assign neighboring router (R3) as an RR1 peer:

The following commands assign the R3 peer as an RR1 client.

```
Passport-8610:5# config ip bgp neighbor 10.1.1.14 create
Passport-8610:5# config ip bgp neighbor 10.1.1.14
remote-as 40
Passport-8610:5# config ip bgp neighbor 10.1.1.14
route-reflector-client enable
Passport-8610:5# config ip bgp neighbor 10.1.1.14
admin-state enable
```

**4** Assign neighboring router (R4) as an RR1 peer:

The following commands assign the R4 peer as an RR1 client.

```
Passport-8610:5# config ip bgp neighbor 200.1.70.1 create
Passport-8610:5# config ip bgp neighbor 200.1.70.1
remote-as 40
Passport-8610:5# config ip bgp neighbor 200.1.70.1
route-reflector-client enable
Passport-8610:5# config ip bgp neighbor 200.1.70.1
admin-state enable
```

**RR2 Configuration**

To configure route reflector RR2, complete the following steps:

*Disable synchronization on RR2***1** Disable BGP on RR2:

You must first disable the BGP administration state before you change the synchronization parameter value.

```
Passport-8610:5# config ip bgp disable
```

**2** Disable synchronization on RR2:

```
Passport-8610:5# config ip bgp synchronization disable
```

**3** Assign RR2 to AS 40:

```
Passport-8610:5# config ip bgp local-as 40
```

**4** Enable route reflection on RR2:

```
Passport-8610:5# config ip bgp route-reflection enable
Passport-8610:5# config ip bgp cl-to-cl-reflection enable
```

**Note:** The `cl-to-cl-reflection` parameter lets you enable or disable the route reflector clients from distributing their respective BGP routes.

**5** Assign a 4-byte cluster ID to RR2:

The 4-byte cluster ID lets route reflectors recognize updates from *other* route reflectors in the same cluster.

```
Passport-8610:5# config ip bgp cluster-id 0.0.0.20 add
```

**6** Enable BGP on RR2:

```
Passport-8610:5# config ip bgp enable
```

*Configure BGP peer interface for RR2***1** Assign neighboring router (RR1) as an RR2 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.62 create
Passport-8610:5# config ip bgp neighbor 10.1.1.62
remote-as 40
Passport-8610:5# config ip bgp neighbor 10.1.1.62
admin-state enable
```

**2** Assign neighboring router (R4) as an RR2 peer:

The following commands assign the R4 peer as an RR2 client.

```
Passport-8610:5# config ip bgp neighbor 200.1.80.1 create
Passport-8610:5# config ip bgp neighbor 200.1.80.1
remote-as 40
Passport-8610:5# config ip bgp neighbor 200.1.80.1
route-reflector-client enable
Passport-8610:5# config ip bgp neighbor 200.1.80.1
admin-state enable
```

**3** Assign neighboring router (R5) as an RR2 peer:

```
Passport-8610:5# config ip bgp neighbor 15.15.15.2 create
Passport-8610:5# config ip bgp neighbor 15.15.15.2
remote-as 40
Passport-8610:5# config ip bgp neighbor 15.15.15.2
admin-state enable
```

## R3 Configuration

To configure route reflector R3, complete the following steps:

*Disable synchronization on R3***1** Disable BGP on R3:

You must first disable the BGP administration state before you change the synchronization parameter value.

```
Passport-8610:5# config ip bgp disable
```

**2** Disable synchronization on R3:

```
Passport-8610:5# config ip bgp synchronization disable
```

**3** Assign R3 to AS 40:

```
Passport-8610:5# config ip bgp local-as 40
```

**4** Enable BGP on R3:

```
Passport-8610:5# config ip bgp enable
```

*Configure BGP peer interfaces for R3***1** Assign neighboring router (RR1) as an R3 peer:

```
Passport-8610:5# config ip bgp neighbor 10.1.1.13 create
Passport-8610:5# config ip bgp neighbor 10.1.1.13
remote-as 40
Passport-8610:5# config ip bgp neighbor 10.1.1.13
admin-state enable
```

**2** Assign neighboring router (R8) as an R3 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.40.2 create
Passport-8610:5# config ip bgp neighbor 200.1.40.2
remote-as 500
Passport-8610:5# config ip bgp neighbor 200.1.40.2
admin-state enable
```

**R4 Configuration**

To configure route reflector R4, complete the following steps:

*Disable synchronization on R4***1** Disable BGP on R4:

You must first disable the BGP administration state before you change the synchronization parameter value.

```
Passport-8610:5# config ip bgp disable
```

**2** Disable synchronization on R4:

```
Passport-8610:5# config ip bgp synchronization disable
```

**3** Assign R4 to AS 40:

```
Passport-8610:5# config ip bgp local-as 40
```

**4** Enable BGP on R4:

```
Passport-8610:5# config ip bgp enable
```

*Configure BGP peer interfaces for R4***1** Assign neighboring router (R7) as an R4 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.30.2 create
Passport-8610:5# config ip bgp neighbor 200.1.30.2
remote-as 1000
Passport-8610:5# config ip bgp neighbor 200.1.30.2
admin-state enable
```

**2** Assign neighboring router (RR1) as an R4 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.70.2 create
Passport-8610:5# config ip bgp neighbor 200.1.70.2
remote-as 40
Passport-8610:5# config ip bgp neighbor 200.1.70.2
admin-state enable
```

**3** Assign neighboring router (RR2) as an R4 peer:

```
Passport-8610:5# config ip bgp neighbor 200.1.80.2 create
Passport-8610:5# config ip bgp neighbor 200.1.80.2
remote-as 40
Passport-8610:5# config ip bgp neighbor 200.1.80.2
admin-state enable
```

<b>For more information about:</b>	<b>See:</b>
Route reflector concepts	<a href="#">"IBGP route reflection" on page 34.</a>
Configuring route reflection using Device Manager	<a href="#">"Configuring general parameters" on page 56.</a>
Configuring route reflection using the CLI	<a href="#">"Configuring general BGP parameters" on page 120.</a>

## Route flap dampening

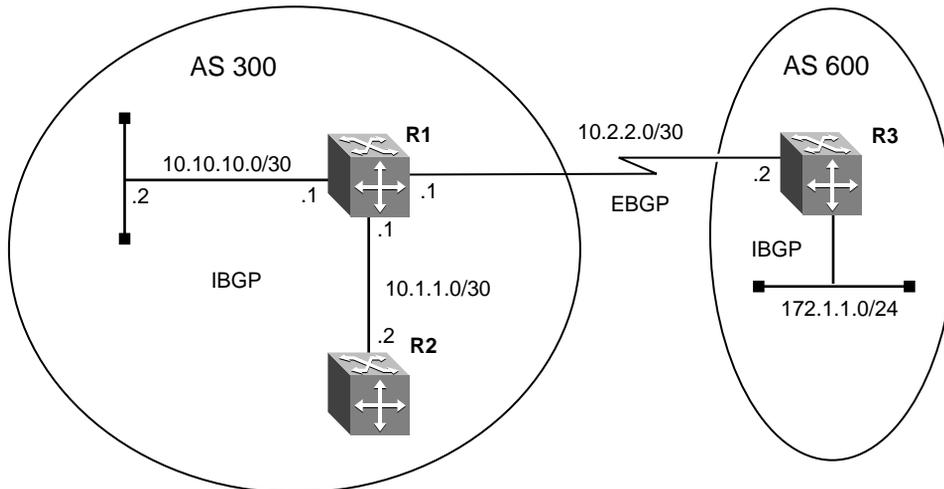
The frequent change of network reachability information that can be caused by an unstable route is commonly referred to as route flap. Route flap dampening suppresses the advertisement of the unstable route until the route becomes stable.



**Note:** Route flap dampening is applied only to routes that are learned via EBGP. Route flap dampening prevents routing loops and protects IBGP peers from having higher penalties for routes that are external to the AS.

The following steps demonstrate a simulation of a route flap occurrence between R1 and R3, shown in [Figure 115](#).

**Figure 115** Route flap dampening example



Passport 8600

10890FA

In this demonstration, you can see the affect of the simulation in the screen displays, before and after the simulation.

To simulate a route flap occurrence between R1 and R3, complete the following steps:

- 1 Enable BGP route flap damping on R1:

```
Passport-8610:1# config ip bgp flap-dampening enable
```

- 2 Enter the following *show* command on R1 to display the current flap damping configuration:

```
Passport-8610:1# show ip bgp flap-damp-config
```

Figure 116 displays the current output of the `show ip bgp flap-damp-config` command on R1.

**Figure 116** show ip bgp flap-damp-config command output

```
Passport-8610:1# show ip bgp flap-damp-config
*****
      Global Flap Dampening Configuration
*****
                                Status - enable
                                PolicyName - N/A
                                CutoffThreshold - 1536
                                ReuseThreshold - 512
                                Decay - 2
                                MaxHoldDown - 180
```

- 3 Enter the following command:

```
Passport-8610:1# show ip bgp dampened-paths 10.2.2.2
```

Figure 117 displays the output of the `show ip bgp dampened-paths 10.2.2.2` command.

Note that because the EBGp link between R1 and R3 is currently stable, the output shown in Figure 117 displays no dampened paths, at this time.

**Figure 117** show ip bgp dampened-paths 10.2.2.2 command output (Example 1)

```
Passport-8610:1# show ip bgp dampened-paths 10.2.2.2
Network/Mask      Peer Rem Addr   NextHop Address Org Loc Pref
-----
```

**4** Disable (and then enable) the admin-state on R3:

The following commands simulate a route flap situation on R3.

```
Passport-8610:1# config ip bgp neighbor 10.2.2.1
admin-state disable
Passport-8610:1# config ip bgp neighbor 10.2.2.1
admin-state enable
```

**5** Enter the following show command on R1:

```
Passport-8610:1# show ip bgp dampened-paths 10.2.2.2
```

[Figure 118](#) displays the output of the `show ip bgp dampened-paths 10.2.2.2`, *after* the route flap simulated by step 4.

**Figure 118** show ip bgp dampened-paths 10.2.2.2 command output (Example 2)

```
Passport-8610:1# show ip bgp dampened-paths 10.2.2.2
```

Network/Mask	Peer Rem Addr	NextHop Address	Org	Loc	Pref
172.1.1.0/24	10.2.2.2	N/A	IGP	0	

```
AS_PATH: no-AS_PATH-attr
MED:0
DAMPEN INFO:Penalty:1024 Count:1 Status:announced hist-del
time:set:180, remain:173
```

The following list describes the output of the screen display shown in [Figure 118](#):

- **Count:1** — indicates the number of times this route has flapped. If you disable/enable the admin-state on R3 a second time, the count will go up to 2.
- **Remain:173** — indicates the amount of hold down time left. Notice that when you display the flap damped configuration, the maximum hold down time displayed a value of 180 seconds—this counter is initially set upon a new route flap. This counter will continue to count down to zero unless there is another flap, in which case the counter will go back up to 180 and then count down again.
- **Penalty: 1024** — indicates that, if the penalty count is greater than the Cut Off Threshold, the route will be suppressed even if the route is up.

The following is a description of the algorithm that is used to control route flaps:

**1** When the route flaps the first time:

- a** A route history entry is created.

The penalty will be 1024.

- b** A timer is started (180 sec).

This timer is used to delete the history entry after the set time of 180 seconds, if the route does not flap again.

**2** When the route flaps a second time:

- a** The penalty is recalculated based on the decay function.

If the penalty is greater than the cut-off value (1536), the route is suppressed and a reuse time is calculated based on the reuse time function.

- b** The reuse timer starts.

When the reuse time expires, the suppressed route is announced again (the reuse time is recalculated if the route flaps again).

The penalty decays slower for withdrawn routes than for update routes. The route history entry is kept longer if the route is withdrawn.

For update history, the delete time is 90 seconds and the withdrawn history delete time is 180 seconds.

<b>For more information about:</b>	<b>See:</b>
Configuring flap dampening using Device Manager	<a href="#">“Configuring general parameters” on page 56.</a>
Configuring flap dampening using the CLI	<a href="#">“Configuring general BGP parameters” on page 120.</a>

## Appendix A

# Translating Cisco Systems-to-Nortel Networks command equivalents

This appendix shows how to translate Cisco Systems commands and functions into corresponding Nortel Networks equivalents.

This appendix includes the following topics:

- “[Configuration Commands](#),” next
- “[Operational Commands](#)” on page 298
- “[Route preferences](#)” on page 300

## Configuration Commands

[Table 33](#) lists the equivalent Nortel Networks command line interface (CLI) commands and Device Manager logical steps for the Cisco router configuration commands. In this table, **bold** text indicates user-supplied variables. Following the table is an itemized list that describes the corresponding items in the table.

**Table 33** Cisco Systems-to-Nortel Networks command equivalents

item	Cisco configuration	CLI commands	Device Manager logical steps
1	<pre>router bgp 333  neighbor 1.1.1.2 remote-as 444</pre>	<pre>config ip bgp &lt;enter&gt; local-as 333 enable <b>Note:</b> If you are changing the local-as, first disable BGP: config ip bgp disable config ip bgp neighbor 1.1.1.2 &lt;enter&gt; create remote-as 444 admin-state enable</pre>	<pre>IP_Routing&gt;BGP LocalAS: 333 AdminStatus: Enable  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 1.1.1.2 RemoteAs: 444 Insert</pre>

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
2	network 1.1.1.0 mask 255.255.255.0	config ip bgp network 1.1.1.0/24 add	IP_Routing>BGP>Nework>Insert NetworkAddr: 1.1.1.0 NetworkMask: 255.255.255.0
3	neighbor 1.1.1.1 distribute-list 5 out ...access list 5 deny 128.1.0.0 0.0.255.255 ...access list 5 permit 0.0.0.0 255.255.255.255	config ip prefix-list 128.1.0.0 add-prefix 128.1.0.0/16  config ip route-policy distribute <enter> seq 1 create enable action deny match-network 128.1.0.0  config ip route-policy distribute <enter> seq 2 create enable action permit  config ip bgp neighbor 1.1.1.1 route-policy out distribute add	IP_Routing>Policy>Prefix List>Insert Id: 1 Prefix: 128.1.0.0 PrefixMaskLen: 24  IP_Routing>Policy>Route Policy>Insert Id: 1 SequenceNumber: 1 Name: distribute Enable Mode: deny MatchNetwork: 128.1.0.0  IP_Routing>Policy>Route Policy>Insert Id: 1 SequenceNumber: 2 Name: distribute Enable Mode: permit  IP_Routing>BGP>Peers>Insert IpAddress: 1.1.1.1 RemoteAs: 444 RoutePolicyOut: distribute Insert

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
4	<pre>neighbor 1.1.1.1 route-map IncomingMap in ...route-map <b>IncomingMap</b> permit 10 match as-path 5 set local-preference 125 ...ip as-path access-list permit 333_444</pre>	<pre>config ip as-list 1 &lt;enter&gt; create add-as-path 1 permit "333 444"  config ip route-policy <b>IncomingMap</b> &lt;enter&gt; seq 1 create enable action permit match-as-path 1 set-local-pref 125  config ip route-policy Preference &lt;enter&gt; seq 2 create enable action permit  config ip bgp neighbor 1.1.1.1 route-policy in <b>IncomingMap</b> add</pre>	<pre>IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 1 MemberId: 1 Mode: <b>permit</b> AsRegularExpression: <b>333 444</b>  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert SequenceNumber: 1 Name: <b>IncomingMap</b> Enable Mode: <b>permit</b> MatchAsPath: 1 SetLocalPref: 125  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 2 Name: <b>IncomingMap</b> Enable Mode: <b>permit</b>  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 1.1.1.1 RemoteAs: 444 RoutePolicyIn: <b>IncomingMap</b> Insert</pre>

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
5	<pre>neighbor 1.1.1.1 route-map setASPath out ...route-map setASPath permit 10 set as-path prepend 123 123</pre>	<pre>config ip prefix-list 200.1.40.0 &lt;enter&gt; add-prefix 200.1.40.0/ 24  config ip as-list 1 create add-as-path 1 permit "123 123"  config ip route-policy setASPath &lt;enter&gt; seq 1 create enable action permit match-network 200.1.40.0 set-as-path 1  config ip bgp neighbor 1.1.1.1 route-policy out setASPath add</pre>	<pre>IP_Routing&gt; Policy&gt;Prefix List&gt;Insert Id:1 Prefix: 200.1.40.0 PrefixMaskLen: 24 Name: 200.1.40.0  IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 1 MemberId: 1 Mode: permit AsRegularExpression: 123 123  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 1 Name: setASPath Enable Mode: permit MatchNetwork: 200.1.40.0 SetAsPath: 1 Insert  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 1.1.1.1 RemoteAs: 444 RoutePolicyOut: setASPath Insert</pre>

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
6	<pre>neighbor 1.1.1.1 route-map AdvertiseMap out ...route-map <b>AdvertiseMap</b> permit 10 match ip address 1 set metric 100 route-map <b>AdvertiseMap</b> permit 20 set metric 50 ...access-list 1 permit 192.10.20.0 0.0.0.255</pre>	<pre>config ip prefix-list <b>192.10.20.0</b> &lt;enter&gt; add-prefix <b>192.10.20.0/24</b>  config ip route-policy <b>AdvertiseMap</b> &lt;enter&gt; seq 1 create enable action permit match-network <b>192.10.20.0</b> set-metric 100  config ip route-policy <b>AdvertiseMap</b> &lt;enter&gt; seq 2 create enable action permit set-metric 50  config ip bgp neighbor <b>1.1.1.1</b> route-policy out <b>AdvertiseMap</b> add</pre>	<pre>IP_Routing&gt; Policy&gt;Prefix List&gt;Insert Id: 1 Prefix: <b>192.10.20.0</b> PrefixMaskLen: 24 Name: <b>192.10.20.0</b>  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 1 Name: <b>AdvertiseMap</b> Enable Mode: <b>permit</b> MatchNetwork: <b>192.10.20.0</b> SetMetric: 1 Insert  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>1.1.1.1</b> RemoteAs: <b>444</b> RoutePolicyOut: <b>AdvertiseMap</b> Insert</pre>
7	<pre>neighbor 1.1.1.1 filter-list 5 in ...ip as-path access-list 5 permit ^1000\$ ^5000\$</pre>	<pre>config ip as-list 1 &lt;enter&gt; create add-as-path 1 permit ^1000\$ <b>Note:</b> Although this example shows the use of AS Path expressions, the AS Path expressions are not supported in release 3.3.  config ip as-list 2 &lt;enter&gt; create add-as-path 1 permit ^5000\$</pre>	<pre>IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 1 MemberId: 1 Mode: <b>permit</b> AsRegularExpression: ^1000\$  IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 2 MemberId: 1 Mode: <b>permit</b> AsRegularExpression: ^5000\$</pre>

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
		<pre> config ip route-policy <b>AS_Filter</b> &lt;enter&gt; seq 1 create enable action permit match-as-path 1 </pre>	<pre> IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 1 Name: <b>AS_Filter</b> Enable Mode: <b>permit</b> MatchAsPath: 1 </pre>
		<pre> config ip route-policy <b>AS_Filter</b> &lt;enter&gt; seq 2 create enable action permit match-as-path 2 </pre>	<pre> IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 2 Name: <b>AS_Filter</b> Enable Mode: <b>permit</b> MatchAsPath: 2 </pre>
		<pre> config ip route-policy <b>AS_Filter</b> &lt;enter&gt; seq 3 create action deny </pre>	<pre> IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 3 Name: <b>AS_Filter</b> Enable Mode: <b>deny</b> </pre>
		<pre> config ip bgp neighbor <b>1.1.1.1</b> route-policy in <b>AS_Filter</b> add </pre>	<pre> IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>1.1.1.1</b> RemoteAs: <b>444</b> RoutePolicyIn: <b>AS_Filter</b> Insert </pre>

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
8	<pre>neighbor 1.1.1.1 filter-list 10 out ...ip as-path access-list 10 deny 350_400_500 ip as-path access-list 10 permit 350_400</pre>	<pre>config ip as-list 1 &lt;enter&gt; create add-as-path 1 deny "350 400 500"  config ip as-list 1 &lt;enter&gt; add-as-path 2 permit "350 400"  config ip route-policy Deny_AS &lt;enter&gt; seq 1 create enable action permit match-as-path 1</pre>	<pre>IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 1 MemberId: 1 Mode: deny AsRegularExpression: 350 400 500  IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 1 MemberId: 2 Mode: permit AsRegularExpression: 350 400  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 1 Name: Deny_AS Enable Mode: permit MatchAsPath: 1  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 1.1.1.1 RemoteAs: 444 RoutePolicyIn: Deny_AS Insert</pre>

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
9	<pre>neighbor <b>MyPeers</b> peer-group neighbor <b>MyPeers</b> remote-as <b>333</b> neighbor <b>MyPeers</b> route-map <b>AdvertiseMap</b> out neighbor <b>MyPeers</b> route-map <b>FilterMap</b> in neighbor 1.1.1.1 peer-group <b>MyPeers</b> neighbor 2.2.2.2 peer-group <b>MyPeers</b></pre>	<pre>config ip bgp neighbor <b>MyPeers</b> &lt;enter&gt; remote-as 333 ebgp-multihop disable route-policy in <b>FilterMap</b> add route-policy out <b>AdvertiseMap</b> add  config ip bgp neighbor <b>1.1.1.1</b> peer-group <b>MyPeers</b> add  config ip bgp neighbor <b>2.2.2.2</b> peer-group <b>MyPeers</b> add</pre>	<pre>IP_Routing&gt;BGP&gt;Peer Groups&gt;Insert Index: 1 GroupName: <b>MyPeers</b> RemoteAS: 333 EbgpMultiHop: disable RoutePolicyIn: <b>FilterMap</b> RoutePolicyOut: <b>AdvertiseMap</b>  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>1.1.1.1</b> GroupName: <b>MyPeers</b> Insert  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>2.2.2.2</b> GroupName: <b>MyPeers</b> Insert</pre>
10	<pre>aggregate-address 195.89.8.0 255.255.248.0</pre>	<pre>Config ip bgp aggregate-address <b>195.89.8.0/20</b> add</pre>	<pre>IP_Routing&gt;BGP&gt;Aggregates&gt;Ins ert Address: <b>195.89.8.0</b> Mask: <b>255.255.248.0</b></pre>
11	<pre>aggregate-address 172.1.1.0 255.255.255.0 summary-only</pre>	<pre>Config ip bgp aggregate-address <b>195.89.8.0/20</b> add summary-only enable</pre>	<pre>IP_Routing&gt;BGP&gt;Aggregates&gt;Ins ert Address: <b>195.89.8.0</b> Mask: <b>255.255.248.0</b> SummaryOnly: Enable</pre>

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
12	<pre>router ospf 101  redistribute bgp  2000</pre>	<pre>config ip ospf  redistribute bgp &lt;enter&gt;  create  enable</pre> <p><b>Note:</b> Before you enable OSPF redistribution, make sure the Passport 8600 is configured for OSPF ASBR.</p> <pre>config ip ospf &lt;enter&gt;  as-boundary-router  enable  admin-state enable  area 0.0.0.0 create</pre>	<p>IP_Routing&gt;OSPF&gt;Redistribute&gt;Insert</p> <p>RouteSource: bgp</p> <p>Enable</p> <p>IP_Routing&gt;OSPF&gt;General</p> <p>RouterId: &lt;ipaddr&gt;</p> <p>AdminStat: enabled</p> <p>ASBdrRtrStatus: checked</p>
13	<pre>router bgp 2000  redistribute ospf  101  redistribute static</pre>	<pre>config ip bgp  redistribute ospf  &lt;enter&gt;  create  enable</pre> <pre>config ip bgp  redistribute static  &lt;enter&gt;  create  enable  apply</pre>	<p>IP_Routing&gt;BGP&gt;Redistribute&gt;Insert</p> <p>RouteSource: ospf</p> <p>Enable: enabled</p> <p>IP_Routing&gt;BGP&gt;Redistribute&gt;Insert</p> <p>RouteSource: static</p> <p>Enable: enabled</p>
14	<pre>timers bgp 60 180</pre>	<pre>config ip bgp neighbor  1.1.1.1 &lt;enter&gt;  admin-state disable  keepalive-time 60  hold-time 180  admin-state: enable</pre>	<p>IP_Routing&gt;BGP&gt;Peers&gt;1.1.1.1</p> <p>Enable: disable</p> <p>HoldTimeConfigured: 180</p> <p>KeepAliveConfigured: 60</p> <p>Enable: enable</p>
15	<pre>interface loopback0  ip address 1.1.1.1  255.255.255.255</pre>	<pre>config ip  circuitless-ip-int 1  &lt;enter&gt;  create 1.1.1.1/32</pre> <p><b>Note:</b> To enable circuitless-ip for OSPF distribution, enter</p> <pre>area &lt;ipaddr&gt;  ospf enable</pre>	<p>IP_Routing&gt;IP&gt;Circuitless IP&gt;Insert</p> <p>Interface: 1</p> <p>Ip Address: 1.1.1.1</p> <p>Net Mask: 255.255.255.255</p> <p>OSPF: enable (click on tab)</p>

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
16	ip subnet zero	Passport has no parameter for zero subnet, already enabled.	Passport has no parameter for zero subnet, already enabled.
17	router bgp 4001 Bgp confederation identifier 5 bgp confederation peers 4002 4003 4004 neighbor 1.2.3.4 remote-as 4002 neighbor 3.4.5.6 remote-as 510	config ip bgp <enter> local-as 4001 confederation identifier 5 add confederation peers 4002 4003 4004  config ip bgp neighbor 1.2.3.4 <enter> create remote-as 4002 admin-state enable  config ip bgp neighbor 3.4.5.6 <enter> create remote-as 510 admin-state enable	IP_Routing>BGP>Generals AdminStatus: disable LocalAS: 4001 ConfederationIdentifier: 5 ConfederationPeers: 4002 4003 4004 AdminStatus: enable  IP_Routing>BGP>Peers>Insert IpAddress: 1.2.3.4 RemoteAs: 4002  IP_Routing>BGP>Peers>Insert IpAddress: 3.4.5.6 RemoteAs: 510
18	router bgp 1000 neighbor 132.245.10.2 password bla4u00=2nkq	config ip bgp neighbor 132.245.10.2 <enter> password bla4u00=2nkq add MD5-authentication enable Password	IP_Routing>BGP>Peers>Insert IpAddress: 132.245.10.2 Password: bla4u00=2nkq MD5Authentication: enable
19	neighbor 1.1.1.1 remote-as 100 neighbor 1.1.1.1 remote-as 100 route-reflector-client	config ip bgp <enter> local-as 100 route-reflection enable cl-to-cl-reflection enable enable  config ip bgp neighbor 1.1.1.1 <enter> create remote-as 100 route-reflector-client enable admin-state enable	IP_Routing>BGP>Generals AdminStatus: disable LocalAS: 100 RouteReflectionEnable: enable ReflectorClientToClientReflection: enable AdminStatus: enable  IP_Routing>BGP>Peers>Insert IpAddress: 1.1.1.1 RemoteAs: 100 RouteReflectoinClient: checked off

**Table 33** Cisco Systems-to-Nortel Networks command equivalents (continued)

item	Cisco configuration	CLI commands	Device Manager logical steps
20	<pre>neighbor 5.5.5.5 remote-as 100 neighbor 1.1.1.1 remote-as 100 route-reflector-client bgp cluster-id 10</pre>	<pre>config ip bgp &lt;enter&gt; local-as 100 route-reflection enable cl-to-cl-reflection enable cluster-id 0.0.0.10 add del enable  config ip bgp neighbor 5.5.5.5 &lt;enter&gt; create remote-as 100 admin-state enable  config ip bgp neighbor 1.1.1.1 &lt;enter&gt; create remote-as 100 route-reflector-client enable admin-state enable</pre>	<pre>IP_Routing&gt;BGP&gt;Generals AdminStatus: disable LocalAS: 100 RouteReflectionEnable: enable RouteReflectorClusterId: 0.0.0.10 ReflectorClientToClientReflection: enable AdminStatus: enable  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 5.5.5.5 RemoteAs: 100  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 1.1.1.1 RemoteAs: 100 RouteReflectoinClient: checked off</pre>

## Interpreting Cisco Systems-to-Nortel Networks command equivalents

The numbers in the following list correspond to the item numbers in [Table 33](#). Each numbered item in this list describes the function of the commands in the corresponding row of [Table 33](#).

- 1 Enable the Border Gateway Protocol (BGP) routing process and identify the local router autonomous system (AS), 333. Activate a BGP session with peer router, IP address, 1.1.1.2 that belongs to AS 444. If the local and remote AS numbers are the same, the BGP session is internal, otherwise it is an external BGP session.
- 2 Advertise network 1.1.1.0 mask 255.255.255.0 and originate it from my AS. Note that network 1.1.1.0 must be present in the IP routing table for Cisco's BGP network command to advertise the route.

- 3 Deny incoming advertisement of network 128.1.0.0, mask 255.255.0.0 from peer IP address, 1.1.1.1, as specified by Cisco access list 5 or Nortel Networks policy name distribute.
- 4 Accept incoming advertisements, from peer 1.1.1.1, match on AS-Path that contain either AS “333 444” or 345 and set Local Preference to 125, as specified by Cisco route-map and Nortel Networks policy name IncomingMap.
- 5 Announce advertisements to peer 1.1.1.1 and append AS-Path <123 123> to all outgoing updates, as specified by Cisco route-map and Nortel Networks policy name setASPath.
- 6 Announce advertisement of network 192.10.20.0 mask 255.255.255.0 to peer IP address 1.1.1.1, setting multi-exit discriminator (MED) to 100 as specified by Cisco route-map and Nortel Networks policy name AdvertiseMap. In addition, advertise any other networks with MED set to 50.
- 7 Accept incoming advertisements from peer 1.1.1.1, of AS-Path that contain either exactly AS 1000 or 5000 as specified by Cisco as-path access-list 5 and Nortel Networks policy name AS\_Filter.
- 8 Announce advertisements to peer 1.1.1.1 if the update includes an AS-Path that matches <350 400> and deny updates of AS-Path that contain <350 400 500> as specified by Cisco as-path access-list 5 and Nortel Networks policy names Deny\_AS.
- 9 Create a peer group named MyPeers with the following elements: peer router AS is 333, advertise networks as specified by route-map AdvertiseMap and accept incoming networks as specified by FilterMap. Assign peer routers 1.1.1.1 and 2.2.2.2 to peer group MyPeers
- 10 Advertise the aggregate address 195.89.8.0 mask 255.255.248.0 (195.89.8.0/21) as well as the more specific addresses i.e. 195.89.8.0 - 195.89.15.0.
- 11 Advertise the aggregate address 195.89.8.0 mask 255.255.248.0 (195.89.8.0/21) only.
- 12 To redistribute BGP routes into OSPF.
- 13 To redistribute OSPF and static routes into BGP.

- 14** Keep-alive timer is used between BGP peers as a periodic check of the TCP connection between them. Hold-down timer is the amount of elapsed time before the BGP peering session is declared dead. RFC 1771 suggests values of 30 and 90 seconds respectively. Hold-down timer is suggested to be three times the amount of the keep-alive timer.
- 15** Cisco's loopback interface and Nortel Networks circuitless IP interface is useful in BGP environments to use as peer interfaces. It is highly recommended using loopback interfaces for BGP as it eliminates the dependency that would otherwise occur when you use the IP address of a physical interface.
- 16** Enable the use of subnet zero for interface addresses and routing updates.
- 17** Enable Confederations for IBGP full mesh reduction. In this example, the outside world sees this as a single AS, number 5, but within the AS it is divided into autonomous systems 4001, 4002, 4003 and 4004. This router's confederation ID is 4001. It has a peer 1.2.3.4 within its routing confederation domain and another peer 3.4.5.6 outside.
- 18** Enables MD5 authentication on the TCP connection between the two BGP peers (132.245.10.1 and 132.245.10.2). In this example, the MD5 key is **b1a4u00=2nkq**.
- 19** Enable Route Reflectors for IBGP full mesh reduction. The Passport 8600 is also configured to allow router reflector client to client route distribution.
- 20** Enable Route Reflectors with two route reflectors for redundancy. A cluster id must be configured when there are two or more router reflectors in a cluster.

## Operational Commands

Table 34 compares the corresponding Cisco Systems and Nortel Networks operational commands. Following the table is an itemized list that describes the function of the commands in the corresponding row of this table.

**Table 34** Cisco Systems-to-Nortel Networks operational commands

Item	Cisco	Nortel Networks
1	no synchronization	Synchronization disabled
2	Route reflector	Route reflector
3	Bgp damping	Bgp damping
4	Confederation	Confederation
<b>BGP monitoring commands</b>		
5	show ip route bgp	show ip bgp route
6	show ip bgp neighbors	show ip bgp sum
7	show ip bgp neighbors 1.1.1.2	show ip bgp neighbor info 1.1.1.2
8	show ip bgp neighbors 1.1.1.2	show ip bgp neighbor stats 1.1.1.2
9	show ip bgp neighbors 1.1.1.2	show ip bgp neighbor route 1.1.1.2
10	clear ip bgp neighbor-ip-address	config ip bgp enable disable config ip bgp neighbor 1.1.1.1 admin-state enable admin-state disable
11	show ip route	show ip route info
12	trace 1.1.1.1	Traceroute 1.1.1.1
13	debug ip bgp	You can use the local console port on the Passport 8600 to configure debug commands, which can display BGP state, events, and more. For more information about configuring debug commands, see <a href="#">"Configuring BGP debug commands" on page 177</a> .
		<ul style="list-style-type: none"> <li>To display bgp global debug messages, enter the following command: <b>config ip bgp global-debug mask &lt;value&gt;</b> Mask values include: none, all, error, packet, event, trace, warning, state, init, filter, and update.</li> </ul>

**Table 34** Cisco Systems-to-Nortel Networks operational commands (continued)

Item	Cisco	Nortel Networks
		<ul style="list-style-type: none"> <li>• To display specific debug messages for your global BGP neighbors, enter the following command:  <code>config ip bgp neighbor-debug-all mask &lt;value&gt;</code>  Mask values include: none, all, error, packet, event, trace, warning, state, init, filter, and update.</li> </ul>

## Interpreting Cisco Systems-to-Nortel Networks Operational equivalents

The following list describes the function of the Cisco Systems and Nortel Networks operational commands in the corresponding row of [Table 34](#).

- 1 Do not synchronize between BGP and IGP; this enables a router to advertise a BGP network to an external peer without having that network exist in the IP routing table.
- 2 Route reflection is a method to alleviate the need for “full mesh” IBGP by allowing an internal BGP speaker to reflect (or re-advertise) routes learned through an IBGP connection to another IBGP peer.
- 3 Minimize the instability caused by route flapping.
- 4 Confederations are used to reduce the number of peers in an AS by breaking the network into multiple (smaller) ASs.
- 5 Show BGP routing table.
- 6 Show status of BGP peers.
- 7 Show the router’s BGP timers. Within Cisco’s show ip bgp neighbor command the keep-alive, hold-down and external advertisement timers are displayed.
- 8 Display the router’s statistics.
- 9 Cisco’s show ip bgp neighbor command displays the router’s incoming and outgoing route filters. The Nortel Networks show ip bgp neighbor route command display incoming routes from peer 1.1.1.2.

- 10 Reset a neighbor's BGP connection.
- 11 Display the IP routing table.
- 12 Discover the routes the router's packets take when traveling to destination 1.1.1.1.
- 13 Display BGP updates/changes/events as they occur.

## Route preferences

[Table 35](#) compares Cisco Systems-to-Nortel Networks route preference values.

**Table 35** Cisco Systems-to-Nortel Networks route preference comparison

Route type	Cisco — Preference value	Nortel Networks — Preference value
Directly connected	0	0
Static	1	5
EBGP	20	45
OSPF Intra	110	20
OSPF Inter		25
BGP	20	30
RIP	120	100
OSPF External 1		120
OSPF External 2		125
IBGP	200	175

## Appendix B

# Translating Juniper Networks-to-Nortel Networks equivalents

This appendix shows how to translate Juniper Networks\* commands and functions into Nortel Networks equivalents.

This appendix includes the following topics:

- “Configuration Commands,” next
- “Operational Commands” on page 316
- “Route preferences” on page 318

## Configuration Commands

Table 36 lists the equivalent Nortel Networks command line interface (CLI) commands and Device Manager logical steps for the Juniper router configuration commands. In this table, **bold** text indicates user-supplied variables. Following the table is an itemized list that describes the corresponding items in the table.

**Table 36** Juniper Networks-to-Nortel Networks command equivalents

item	Juniper configuration	CLI commands	Device Manager logical steps
1	<pre>set routing-options autonomous-system 333</pre>	<pre>config ip bgp &lt;enter&gt; local-as 333 enable **</pre>	<pre>IP_Routing&gt;BGP LocalAS: 333 AdminStatus: Enable</pre>
	<pre>edit protocols bgp group ebgp &lt;enter&gt; set type external set peer-as 444 set neighbor 1.1.1.2</pre>	<pre>config ip bgp neighbor 1.1.1.2 &lt;enter&gt; create remote-as 444 admin-state enable</pre>	<pre>IP_Routing&gt;BGP&gt;Nework&gt;Insert NetworkAddr: &lt;ip address of direct interface&gt; NetworkMask: &lt;ip mask of direct interface&gt;</pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
2	<pre> protocols {   bgp {     export <b>direct</b>;   }   policy-options {     policy-statement <b>direct</b> {       term <b>dir_export</b> {         from protocol direct;         then accept;       }     }   } </pre>	<pre> config ip bgp <b>redistribute direct</b> &lt;enter&gt; create enable <b>metric 100</b> apply </pre>	<p>IP Routing&gt;BGP&gt;Redistribute&gt;Insert</p> <p>RouteSource: <b>direct</b></p> <p>Enable: <b>enable</b></p> <p>Metric: <b>100</b></p> <p>Insert</p>
3	<pre> protocols {   bgp {     group <b>ebgp</b> {       type external;       export <b>drop</b>;       peer-as <b>300</b>;       neighbor <b>1.1.1.1</b>;     }   }   policy-options {     policy-statement <b>drop</b> {       term <b>list</b> {         from {           protocol bgp;           route-filter <b>128.1.0.0/16</b> exact reject;         }         then accept;       }     }   } </pre>	<pre> config ip prefix-list <b>128.1.0.0</b> add-prefix <b>128.1.0.0/16</b>  config ip route-policy <b>distribute</b> &lt;enter&gt; seq 1 create enable action deny match-network <b>128.1.0.0</b>  config ip route-policy <b>distribute</b> &lt;enter&gt; seq 2 create enable action permit  config ip bgp neighbor <b>1.1.1.1</b> &lt;enter&gt; remote-as <b>300</b> route-policy out <b>distribute</b> add </pre>	<p>IP_Routing&gt;Policy&gt;Prefix List&gt;Insert</p> <p>Id: <b>1</b></p> <p>Prefix: <b>128.1.0.0</b></p> <p>PrefixMaskLen: <b>24</b></p> <p>IP_Routing&gt;Policy&gt;Route Policy&gt;Insert</p> <p>Id: <b>1</b></p> <p>SequenceNumber: <b>1</b></p> <p>Name: <b>distribute</b></p> <p>Enable</p> <p>Mode: <b>permit</b></p> <p>MatchNetwork: <b>128.1.0.0</b></p> <p>IP_Routing&gt;Policy&gt;Route Policy&gt;Insert</p> <p>Id: <b>1</b></p> <p>SequenceNumber: <b>2</b></p> <p>Name: <b>distribute</b></p> <p>Enable</p> <p>Mode: <b>deny</b></p> <p>IP_Routing&gt;BGP&gt;Peers&gt;Insert</p> <p>IpAddress: <b>1.1.1.1</b></p> <p>RemoteAs: <b>444</b></p> <p>RoutePolicyOut: <b>distribute</b></p> <p>Insert</p>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
4	<pre> policy-options {   policy-statement <b>IncomingMap</b> {     term <b>as</b> {       from {         neighbor <b>1.1.1.1</b>;         as-path <b>aslist</b>;       }       then {         local-preference <b>125</b>;       }     }   }   as-path <b>aslist 333-444</b>; } protocols {   bgp {     import <b>IncomingMap</b>;   } } </pre>	<pre> config ip as-list <b>1</b> &lt;enter&gt; create add-as-path <b>1</b> permit <b>333_444</b>  config ip route-policy <b>IncomingMap</b> &lt;enter&gt; seq <b>1</b> create enable action permit match-as-path <b>1</b> set-local-pref <b>125</b>  config ip route-policy Preference &lt;enter&gt; seq <b>2</b> create enable action permit  config ip bgp neighbor <b>1.1.1.1</b> route-policy in <b>IncomingMap</b> add </pre> <p><b>Note:</b> Although this example shows the use of AS Path expressions, the AS Path expressions are not supported in release 3.3.</p>	<pre> IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: <b>1</b> MemberId: <b>1</b> Mode: <b>permit</b> AsRegularExpression: <b>333_444</b>  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert SequenceNumber: <b>1</b> Name: <b>IncomingMap</b> Enable Mode: <b>permit</b> MatchAsPath: <b>1</b> SetLocalPref: <b>125</b>  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: <b>1</b> SequenceNumber: <b>2</b> Name: <b>IncomingMap</b> Enable Mode: <b>permit</b>  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>1.1.1.1</b> RemoteAs: <b>444</b> RoutePolicyIn: <b>IncomingMap</b> Insert </pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
5	<pre> policy-options {   policy-statement   setASPath {     term ASList {       from {         route-filter         200.1.40.0/24 exact;       }       then as-path-prepend       "123 123";     }   } }  protocols {   bgp {     group ebgp {       type external;       export setASPath;       peer-as 300;       neighbor 1.1.1.1;     }   } } </pre>	<pre> config ip prefix-list 200.1.40.0 &lt;enter&gt;   add-prefix   200.1.40.0/24  config ip as-list 1 &lt;enter&gt; create add-as-path 1 permit "123 123"  config ip route-policy setASPath &lt;enter&gt; seq 1 create enable action permit match-network 200.1.40.0 set-as-path 1  config ip bgp neighbor 1.1.1.1 &lt;enter&gt; remote-as 300 route-policy out setASPath add </pre>	<pre> IP_Routing&gt; Policy&gt;Prefix List&gt;Insert Id:1 Prefix: 200.1.40.0 PrefixMaskLen: 24 Name: 200.1.40.0  IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 1 MemberId: 1 Mode: permit AsRegularExpression: 123 123  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 1 Name: setASPath Enable Mode: permit MatchNetwork: 200.1.40.0 SetAsPath: 1 Insert  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 1.1.1.1 RemoteAs: 444 RoutePolicyOut: setASPath Insert </pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
6	<pre> policy-options {   policy-statement   <b>AdvertiseMap</b> {     term <b>seq1</b> {       from {         route-filter         <b>192.10.20.0/24</b> exact;       }       then {         metric <b>100</b>;         accept;       }     }     term <b>seq2</b> {       from {         route-filter <b>0.0.0.0/0</b>         orlonger;       }       then {         metric <b>50</b>;         accept;       }     }   } }  protocols {   bgp {     group <b>ebgp</b> {       type external;       export       <b>AdvertiseMap</b>;       peer-as <b>300</b>;       neighbor <b>1.1.1.1</b>;     }   } } </pre>	<pre> config ip prefix-list <b>192.10.20.0</b> &lt;enter&gt; add-prefix <b>192.10.20.0/24</b>  config ip route-policy <b>AdvertiseMap</b> &lt;enter&gt; seq 1 create enable action permit match-network <b>192.10.20.0</b> set-metric <b>100</b>  config ip route-policy <b>AdvertiseMap</b> &lt;enter&gt; seq 2 create enable action permit set-metric <b>50</b>  config ip bgp neighbor <b>1.1.1.1</b> &lt;enter&gt; remote-as <b>300</b> route-policy out <b>AdvertiseMap</b> add </pre>	<pre> IP_Routing&gt; Policy&gt;Prefix List&gt;Insert Id:1 Prefix: <b>192.10.20.0</b> PrefixMaskLen: <b>24</b> Name: <b>192.10.20.0</b>  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 1 Name: <b>AdvertiseMap</b> Enable Mode: <b>permit</b> MatchNetwork: <b>192.10.20.0</b> SetMetric: <b>100</b> Insert  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert t Id: 1 SequenceNumber: 2 Name: <b>AdvertiseMap</b> Enable Mode: <b>permit</b> SetMetric: <b>50</b> Insert  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>1.1.1.1</b> RemoteAs: <b>444</b> RoutePolicyOut: <b>AdvertiseMap</b> Insert </pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
7	<pre> policy-options { policy-statement <b>AS_Filter</b> {   term <b>aslist</b> {     from {       neighbor       10.10.10.2;       as-path <b>AS_3</b>;     }     then accept;   }   term <b>nothing_else</b> {     then reject;   } } as-path <b>AS_3</b> "(1000)   (5000)";  protocols {   bgp {     group <b>ebgp</b> {       type external;       import <b>AS_Filter</b>;       peer-as <b>300</b>;       neighbor <b>1.1.1.1</b>;     }   } } </pre>	<pre> config ip as-list 1 &lt;enter&gt; create add-as-path 1000 permit ^1000\$  config ip as-list 2 &lt;enter&gt; create add-as-path 5000 permit ^5000\$  config ip route-policy <b>AS_Filter</b> &lt;enter&gt; seq 1 create enable action permit match-as-path 1  config ip route-policy <b>AS_Filter</b> &lt;enter&gt; seq 2 create enable action permit match-as-path 2  config ip route-policy <b>AS_Filter</b> &lt;enter&gt; seq 3 create action deny  config ip bgp neighbor <b>1.1.1.1</b> &lt;enter&gt; remote-as <b>300</b> route-policy in <b>AS_Filter</b> add </pre> <p><b>Note:</b> Although this example shows the use of AS Path expressions, the AS Path expressions are not supported in release 3.3.</p>	<pre> IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 1 MemberId: 1000 Mode: permit AsRegularExpression: ^1000\$  IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 2 MemberId: 5000 Mode: permit AsRegularExpression: ^5000\$  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 1 Name: <b>AS_Filter</b> Enable Mode: permit MatchAsPath: 1  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 2 Name: <b>AS_Filter</b> Enable Mode: permit MatchAsPath: 2  IP_Routing&gt;Policy&gt;Route Policy&gt;Insert Id: 1 SequenceNumber: 3 Name: <b>AS_Filter</b> Enable Mode: deny  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>1.1.1.1</b> RemoteAs: <b>300</b> RoutePolicyIn: <b>AS_Filter</b> Insert </pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
8	<pre> protocols {   bgp {     group ibgp {       import <b>Deny_AS</b>;       neighbor 1.1.1.1;     }   }   policy-options {     policy-statement <b>Deny_AS</b> {       term <b>aslist1</b> {         from as-path 1;         then accept;       }       term <b>aslist2</b> {         from as-path 2;         then reject;       }     }   }   as-path 1 "<b>350 400</b>";   as-path 2 "<b>350 400 500</b>"; </pre>	<pre> config ip as-list 1 &lt;enter&gt; create add-as-path 1 deny "350 400 500"  config ip as-list 1 &lt;enter&gt; add-as-path 2 permit "350 400"  config ip route-policy <b>Deny_AS</b> &lt;enter&gt; seq 1 create enable action permit match-as-path 1  config ip bgp neighbor 1.1.1.1 route-policy in <b>Deny_AS</b> add </pre>	<pre> IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 1 MemberId: 1 Mode: <b>deny</b> AsRegularExpression: <b>350 400 500</b> IP_Routing&gt; Policy&gt;As Path List&gt;Insert Id: 1 MemberId: 2 Mode: <b>permit</b> AsRegularExpression: <b>350 400</b> IP_Routing&gt;Policy&gt;Route Poli cy&gt;Insert Id: 1 SequenceNumber: 1 Name: <b>Deny_AS</b> Enable Mode: <b>permit</b> MatchAsPath: 1 IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 1.1.1.1 RemoteAs: <b>444</b> RoutePolicyIn: <b>Deny_AS</b> Insert </pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
9	<pre> protocols {   bgp {     group <b>ibgp</b> {       type internal;       export <b>NHS</b>       peer-as <b>333</b>;       neighbor <b>1.1.1.1</b>;       neighbor <b>2.2.2.2</b>;      policy-options {       policy-statement <b>NHS</b> {         term <b>next-hop</b> {           from {             protocol bgp;           }           then {             next-hop self;           }         }       }     }   } } </pre>	<pre> config ip bgp neighbor <b>NHS</b> &lt;enter&gt; create remote-as <b>333</b> nexthop-self <b>enable</b> ebgp-multihop disable  config ip bgp neighbor <b>1.1.1.1</b> peer-group <b>NHS</b> add  config ip bgp neighbor <b>2.2.2.2</b> peer-group <b>NHS</b> add </pre>	<pre> IP_Routing&gt;BGP&gt;Peer Groups&gt;Insert Index: <b>1</b> GroupName: <b>MyPeers</b> EbgpMultiHop: disable  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>1.1.1.1</b> GroupName: <b>MyPeers</b> Insert  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>2.2.2.2</b> GroupName: <b>MyPeers</b> Insert </pre>
10	<pre> policy-options {   policy-statement <b>agg-add</b>   {     term <b>agg</b> {       from {         route-filter <b>195.89.8.0/20</b> orlonger;       }       then accept;     }   } } </pre>	<pre> config ip bgp aggregate-address <b>195.89.8.0/20</b> add </pre>	<pre> IP_Routing&gt;BGP&gt;Aggregates&gt;Ins ert Address: <b>195.89.8.0</b> Mask: <b>255.255.248.0</b> </pre>
11	<pre> routing-options {   aggregate {     route 195.89.8.0/20 passive;   } }  policy-options {   policy-statement <b>agg</b> {     from protocol aggregate;     then accept;   } } </pre>	<pre> config ip bgp aggregate-address <b>195.89.8.0/20</b> add summary-only enable </pre>	<pre> IP_Routing&gt;BGP&gt;Aggregates&gt;Ins ert Address: <b>195.89.8.0</b> Mask: <b>255.255.248.0</b> SummaryOnly: Enable </pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
12	<pre> protocols {   ospf {     export <b>bgp_routes</b>;     area <b>0.0.0.0</b> {       interface <b>20.1.1.1</b> {         metric <b>200</b>;       }     }   }   policy-options {     policy-statement <b>bgp_routes</b> {       from protocol bgp;       then accept;     }   } </pre>	<pre> config ip ospf redistribute bgp &lt;enter&gt; metric <b>200</b> create enable </pre> <p><b>Note:</b> Before you enable OSPF redistribution, be sure the Passport 8600 is configured for OSPF ASBR.</p> <pre> config ip ospf &lt;enter&gt; as-boundary-router enable admin-state enable area 0.0.0.0 create </pre>	<pre> IP_Routing&gt;OSPF&gt;Redistribute&gt;I nser RouteSource: bgp Enable </pre> <pre> IP_Routing&gt;OSPF&gt;General RouterId: &lt;ipaddr&gt; AdminStat: enabled ASBdrRtrStatus: checked </pre>
13	<pre> protocols {   bgp {     export <b>ospf_into_bgp</b>;   }   policy-options {     policy-statement <b>ospf_into_bgp</b> {       term <b>ospf-only</b> {         from protocol ospf;         then accept;       }     }   } </pre>	<pre> config ip bgp redistribute ospf &lt;enter&gt; create enable </pre>	<pre> IP_Routing&gt;BGP&gt;Redistribute&gt;In sert RouteSource: ospf Enable: enabled </pre> <pre> IP_Routing&gt;BGP&gt;Redistribute&gt;In sert RouteSource: static Enable: enabled </pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
14	<pre>protocols {   bgp {     group ebgp {       type external;       hold-time 180;       peer-as 300;       local-as 100;       neighbor 1.1.1.1;     }   } }</pre> <p><b>Note:</b> The Juniper software default for the Keepalive timer is one-third the HoldTime. For this example, the HoldTime is set for 180, so the KeepAlive default is 60.</p>	<pre>config ip bgp neighbor 1.1.1.1 &lt;enter&gt; admin-state disable keepalive-time 60 hold-time 180 admin-state: enable</pre>	<pre>IP_Routing&gt;BGP&gt;Peers&gt;1.1.1.1 Enable: disable HoldTimeConfigured: 180 KeepAliveConfigured: 60 Enable: enable</pre>
15	<pre>interfaces {   lo0 {     unit 0 {       family inet {         address 1.1.1.1/32;       }     }   } }</pre>	<pre>config ip circuitless-ip-int 1 &lt;enter&gt; create 1.1.1.1/32</pre> <p><b>** To enable circuitless-ip for ospf distribution, enter area &lt;ipaddr&gt; ospf enable</b></p>	<pre>IP_Routing&gt;IP&gt;Circuitless IP&gt;Insert Interface: 1 Ip Address: 1.1.1.1 Net Mask: 255.255.255.255 OSPF: enable (click on tab)</pre>
16	<p>Synchronization Disabled.</p> <p><b>Note:</b> With Juniper software, synchronization is disabled by default. There is no option to enable or disable synchronization.</p>	<pre>config ip bgp synchronization disable</pre>	<pre>IP_Routing&gt;BGP&gt;Generals&gt; Synchronization: disable</pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
17	<pre>routing-options {   autonomous-system <b>4001</b>;   confederation <b>5</b> members [ <b>4002 4003 4004</b> ]    protocols {     bgp {       group <b>1234</b> {         type external;         peer-as <b>4002</b>;         neighbor <b>1.2.3.4</b>;       }       group <b>3456</b> {         type external;         peer-as <b>510</b>;         neighbor <b>3.4.5.6</b>;       }     }   } }</pre>	<pre>config ip bgp &lt;enter&gt; local-as <b>4001</b> confederation identifier <b>5</b> add confederation peers "<b>4002 4003 4004</b>"  config ip bgp neighbor <b>1.2.3.4</b> &lt;enter&gt; create remote-as <b>4002</b> admin-state enable  config ip bgp neighbor <b>3.4.5.6</b> &lt;enter&gt; create remote-as <b>510</b> admin-state enable</pre>	<pre>IP_Routing&gt;BGP&gt;Generals AdminStatus: disable LocalAS: <b>4001</b> ConfederationIdentifier: <b>5</b> ConfederationPeers: <b>4002 4003 4004</b> AdminStatus: enable  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>1.2.3.4</b> RemoteAs: <b>4002</b>  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>3.4.5.6</b> RemoteAs: <b>510</b></pre>
18	<pre>set protocols bgp group <b>ebgp</b> authentication-key <b>bla4u00=2nkq</b></pre>	<pre>config ip bgp neighbor <b>132.245.10.2</b> &lt;enter&gt; password <b>bla4u00=2nkq</b> add MD5-authentication enable Password</pre>	<pre>IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>132.245.10.2</b> Password: <b>bla4u00=2nkq</b> MD5Authentication: <b>enable</b></pre>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
19	<pre> interfaces {   fe-1/1/0 {     unit 0 {       family inet {         address 10.10.10.1/ 30;       }     }   }   fe-1/1/1 {     unit 0 {       family inet {         address 10.10.10.13/30;       }     }   }   lo0 {     unit 0 {       family inet {         address 1.1.1.1/32;       }     }   }   routing-options {     static {       route 1.1.1.2/32 next-hop [ 10.10.10.2 10.10.10.14];     }   }   protocols {     bgp {       group ebgp {         type external;         multihop ttl 2;         local-address 1.1.1.1;         peer-as 300;         neighbor 1.1.1.2; </pre>	<pre> config ethernet 1/15 ip create 10.10.10.2/30 2078 &lt;enter&gt;  config ethernet 1/17 ip create 10.10.10.14/30 2079 &lt;enter&gt;  config ip circuitless-ip 1 &lt;enter&gt; create 1.1.1.2/32 add  config ip bgp neighbor 1.1.1.1 &lt;enter&gt; create remote-as 100 ebgp-multihop enable update-source-inter 1.1.1.2 add admin-state enable  config ip static-route &lt;enter&gt; create 1.1.1.1/32 next-hop 10.10.10.1 create 1.1.1.1/32 next-hop 10.10.10.13 </pre>	<p>Right-click port 1/15&gt;Edit IP Address&gt;Insert Ip Address: <b>10.10.10.2</b> Net Mask: <b>255.255.255.252</b> &gt;Insert</p> <p>Right-click port 1/17&gt;Edit IP Address&gt;Insert Ip Address: <b>10.10.10.14</b> Net Mask: <b>255.255.255.252</b> &gt;Insert</p> <p>IP Routing&gt;IP&gt;Circuitless IP&gt;Insert Interface: <b>1</b> Ip Address: <b>1.1.1.2</b> Net Mask: <b>255.255.255.255</b> &gt;Insert</p> <p>Ip Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: <b>1.1.1.1</b> RemoteAs: <b>100</b> EbgpMultiHop: enable UpdateSourceInterface: <b>1.1.1.2</b> &gt;Insert</p> <p>IP Routing&gt;IP&gt;Static Routes&gt;Insert Dest: <b>1.1.1.1</b> Mask: <b>255.255.255.255</b> NextHop: <b>10.10.10.1</b> &gt;Insert</p> <p>IP Routing&gt;IP&gt;Static Routes&gt;Insert Dest: <b>1.1.1.1</b> Mask: <b>255.255.255.255</b> NextHop: <b>10.10.10.13</b> &gt;Insert</p>

**Table 36** Juniper Networks-to-Nortel Networks command equivalents (continued)

item	Juniper configuration	CLI commands	Device Manager logical steps
20	<pre>routing-options {   autonomous-system 100;    protocols {     bgp {       group rr-cluster1 {         peer-as 100         local-address 5.5.5.4;         cluster 0.0.0.10         neighbor 5.5.5.5;       }       group rr-cluster2 {         peer-as 100         local-address 1.1.1.2;         cluster 0.0.0.10         neighbor 1.1.1.1;       }     }   } }</pre>	<pre>config ip bgp &lt;enter&gt; local-as 100 route-reflection enable cl-to-cl-reflection enable cluster-id 0.0.0.10 add del enable  config ip bgp neighbor 5.5.5.5 &lt;enter&gt; create remote-as 100 admin-state enable  config ip bgp neighbor 1.1.1.1 &lt;enter&gt; create remote-as 100 route-reflector-client enable admin-state enable</pre>	<pre>IP_Routing&gt;BGP&gt;Generals AdminStatus: disable LocalAS: 100 RouteReflectionEnable: enable RouteReflectorClusterId: 0.0.0.10 ReflectorClientToClientReflection: enable AdminStatus: enable  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 5.5.5.5 RemoteAs: 100  IP_Routing&gt;BGP&gt;Peers&gt;Insert IpAddress: 1.1.1.1 RemoteAs: 100 RouteReflectoinClient: checked off</pre>

## Interpreting Juniper Networks-to-Nortel Networks command equivalents

The numbers in the following list correspond to the item numbers in [Table 36](#). Each numbered item in this list describes the function of the commands in the corresponding row of [Table 36](#).

- 1 Enable the Border Gateway Protocol (BGP) routing process and identify the local router autonomous system (AS), 333. Activate a BGP session with peer router, IP address, 1.1.1.2 that belongs to AS 444. If the local and remote AS numbers are the same, the BGP session is internal, otherwise it is an external BGP session.
- 2 Advertise network 1.1.1.0 and 1.1.1.4 mask 255.255.255.252 that are direct interfaces on the Passport 8600 and originate it from my AS. Note that by default Juniper will advertise all learned routes and the BGP Network command is not used. A policy statement can be added, as shown in this

configuration example, in order for the Juniper router to advertise it's direct interfaces.

- 3 Deny incoming advertisement of network 128.1.0.0, mask 255.255.0.0 from peer IP address, 1.1.1.1, as specified by Juniper policy-statement drop or Nortel Networks policy name distribute.
- 4 Accept incoming advertisements, from peer 1.1.1.1, match on AS-Path that contain either AS "333 444" or 345 and set Local Preference to 125, as specified by Juniper policy-statement IncomingMap and Nortel Networks policy name IncomingMap.
- 5 Announce advertisements to peer 1.1.1.1 and append AS-Path <123 123> to all outgoing updates, as specified by Juniper policy-statement setASPath route-map and Nortel Networks policy name setASPath.
- 6 Announce advertisement of network 192.10.20.0 mask 255.255.255.0 to peer IP address 1.1.1.1, setting multi-exit discriminator (MED) to 100 as specified by Juniper policy-statement AdvertiseMap and Nortel Networks policy name AdvertiseMap. In addition, advertise any other networks with MED set to 50.
- 7 Accept incoming advertisements from peer 1.1.1.1, of AS-Path that contain either exactly AS 1000 or 5000 as specified by Juniper policy-statement AS\_Filter and Nortel Networks policy name AS\_Filter.
- 8 Announce advertisements to peer 1.1.1.1 if the update includes an AS-Path that matches <350 400> and deny updates of AS-Path that contain <350 400 500> as specified by Juniper policy-statement Deny\_AS and Nortel Networks policy name Deny\_AS.
- 9 Create a peer group named NHS with the following elements: nexthop-self enabled. Assign peer routers 1.1.1.1 and 2.2.2.2 to peer group MyPeers. Similar functionality is performed on Juniper by using the policy-statement NHS.
- 10 Advertise the aggregate address 195.89.8.0 mask 255.255.248.0 (195.89.8.0/21) as well as the more specific addresses i.e. 195.89.8.0 - 195.89.15.0.
- 11 Advertise the aggregate address 195.89.8.0 mask 255.255.248.0 (195.89.8.0/21) only.
- 12 To redistribute BGP routes into OSPF.
- 13 To redistribute OSPF into BGP.

- 14** Keep-alive timer is used between BGP peers as a periodic check of the TCP connection between them. Hold-down timer is the amount of elapsed time before the BGP peering session is declared dead. RFC 1771 suggests values of 30 and 90 seconds respectively. Hold-down timer is suggested to be three times the amount of the keep-alive timer.
- 15** The Juniper Network loopback interface and Nortel Networks circuitless IP interface is useful in BGP environments to use as peer interfaces. It is highly recommended using loopback interfaces for BGP as it eliminates the dependency that would otherwise occur when you use the IP address of a physical interface.
- 16** Disable synchronization on the Passport 8600. By default, synchronization is disabled on Juniper and there is no option to enable or disable this functionality.
- 17** Enable Confederations for IBGP full mesh reduction. In this example, the outside world sees this as a single AS, number 5, but within the AS it is divided into autonomous systems 4001, 4002, 4003 and 4004. This router's confederation ID is 4001. It has a peer 1.2.3.4 within its routing confederation domain and another peer 3.4.5.6 outside.
- 18** Enables MD5 authentication on the TCP connection between the two BGP peers (132.245.10.1 and 132.245.10.2). In this example, the MD5 key is bla4u00=2nkq.
- 19** Enable EBGp multihop load balancing. The EBGp peering is between the loopback interface on Juniper and the circuitless ip on Nortel. On each router, static routes to the remote peer's loopback address must be configured for each data link connection.
- 20** Enable Route Reflectors for IBGP full mesh reduction. A cluster id is always used by Juniper and must be configured on Passport 8600 when there are two or more router reflectors in a cluster.

## Operational Commands

Table 37 compares the corresponding Juniper Networks and Nortel Networks operational commands. Following the table is an itemized list that describes the function of the commands in the corresponding row of this table.

**Table 37** Juniper Networks-to-Nortel Networks operational commands

Item	Juniper	Nortel Networks
1	no synchronization	synchronization disabled
2	Route reflector	route reflector
3	bgp damping	bgp damping
4	confederation	confederation
<b>BGP monitoring commands</b>		
5	show route protocol bgp	show ip bgp route
6	show bgp summary	show ip bgp sum
7	show bgp neighbor 1.1.1.2	show ip bgp neighbor info 1.1.1.2
8	show bgp neighbor 1.1.1.2	show ip bgp neighbor stats 1.1.1.2
9	show route advertising-protocol bgp 1.1.1.2	show ip bgp neighbor route 1.1.1.2
10	clear bgp neighbor <ip address>	config ip bgp enable disable config ip bgp neighbor 1.1.1.1 admin-state enable admin-state disable
11	show route	show ip route info
12	traceroute 1.1.1.1	Traceroute 1.1.1.1

**Table 37** Juniper Networks-to-Nortel Networks operational commands (continued)

Item	Juniper	Nortel Networks
13	<p>a) show log messages</p> <p>b) configure the following:</p> <pre>[edit protocols bgp] set traceoptions file bgp-log size 1m files 10</pre> <p>then use the following command:</p> <pre>show log bgp-log</pre>	<p>You can use the local console port on the Passport 8600 to configure debug commands, which can display BGP state, events, and more. For more information about configuring debug commands, see <a href="#">“Configuring BGP debug commands” on page 177</a>.</p> <ul style="list-style-type: none"> <li>To display bgp global debug messages, enter the following command: <pre>config ip bgp global-debug mask &lt;value&gt;</pre> <p>Mask values include: none, all, error, packet, event, trace, warning, state, init, filter, and update.</p> </li> <li>To display specific debug messages for your global BGP neighbors, enter the following command: <pre>config ip bgp neighbor-debug-all mask &lt;value&gt;</pre> <p>Mask values include: none, all, error, packet, event, trace, warning, state, init, filter, and update.</p> </li> </ul>

## Interpreting Juniper Networks-to-Nortel Networks BGP Operational equivalents

The following list describes the functions of the Juniper Networks and the Nortel Networks operational commands in the corresponding row of [Table 37](#).

- 1 Do not synchronize between BGP and IGP; this enables a router to advertise a BGP network to an external peer without having that network exist in the IP routing table.
- 2 Route reflection is a method to alleviate the need for “full mesh” IBGP by allowing an internal BGP speaker to reflect (or re-advertise) routes learned through an IBGP connection to another IBGP peer.
- 3 Minimize the instability caused by route flapping.
- 4 Confederations are used to reduce the number of peers in an AS by breaking the network into multiple (smaller) ASs.
- 5 Show BGP routing table.

- 6 Show status of BGP peers.
- 7 Show the router's BGP neighbor information.
- 8 Display the router's statistics.
- 9 Juniper's show route advertising-protocol BGP command displays the router's incoming and outgoing routes. The Nortel Networks show ip bgp neighbor route command display incoming routes from peer 1.1.1.2.
- 10 Reset a neighbor's BGP connection.
- 11 Display the IP routing table.
- 12 Discover the routes the router's packets take when traveling to destination 1.1.1.1.
- 13 Display BGP updates/changes/events as they occur.

## Route preferences

Table 38 compares Juniper Networks-to-Nortel Networks route preference values.

**Table 38** Juniper Networks-to-Nortel Networks route preference comparison

Route type	Juniper — Preference value	Nortel Networks — Preference value
Directly connected	0	0
Static	5	5
EBGP	170	12
OSPF Intra	10	15
OSPF Inter	150	17
BGP	170	30
RIP	100	100
OSPF External 1	150	120
OSPF External 2	150	125
IBGP	170	200

---

## Glossary

---

This appendix provides a glossary of common terms used with the Border Gateway Protocol (BGP).

aggregate	A prefix length that is formed by combining several specific prefixes. The resulting prefix is used to combine blocks of address space into a single routing announcement.
AS (autonomous system)	A set of routers under a single technical administration, using a single IGP and common metrics to route packets within the AS, and using an EGP to route packets to other ASs.
AS confederation	A single logical AS that comprises multiple sub-ASs to ensure scalability.
ASN (autonomous system number)	A two-byte number that is used to identify a specific AS.
attribute	A unit of data that is used by BGP to describe any of the following prefixes: AS-PATH, LOCAL-PREF, NEXT-HOP, and so on.
BGP (Border Gateway Protocol)	An inter-domain routing protocol that provides loop-free inter-domain routing between autonomous systems (ASs) or within an AS.
BGP peer	A relationship that is formed between any two routers that open a TCP connection to each other for the purpose of exchanging routing informations.
BGP neighbor	BGP routers that have interfaces to a common network.
BGP session	An active connection between two routers running BGP.
BGP speaker	An entity within a BGP router that is used to communicate with other BGP speakers by establishing a peer-to-peer session.

CIDR (Classless Inter-Domain Routing)	A method for creating additional addresses on the Internet, which are given to Internet service providers (ISPs) that in turn delegate them to their customers. CIDR reduces the burden on Internet routers by aggregating routes so that one IP address represents thousands of addresses that are serviced by a major backbone provider.
Circuitless IP	A virtual interface that does not map to any physical interface. This interface is often called a <i>loopback</i> .
cluster	One or more route reflectors and their associated clients that form a relationship where the designated route reflectors provide route reflection for their clients, as well as nonclient peers.
community	A BGP attribute that contains a list of 32-bit values used to identify a route as belonging to a category of routes. All of the routes in the category are treated equally by routing policies.
dampen	Indicates that routes which exhibit instability are not advertised until the routes become stable for a minimum time period.
External BGP (EBGP)	A BGP session between two BGP speakers in different ASs.
Internal BGP (IBGP)	A BGP session between two BGP speakers in the same AS. IBGP is used to distribute routes within the AS that were learned from other sources, such as: EBGP, static routes, and so on.
mask	A bit string that is used along with an IP address to indicate the number of leading bits in the address that correspond with the network part.
multihomed AS	An AS that has multiple connections to one or more ASs and does not carry transit traffic.
next hop	The next hop to which a packet should be sent in order to advance the packet to the destination.

OSPF (Open Shortest Path First protocol)	A link-state protocol in the IP suite that enables routers in the same AS to exchange routing information by means of periodic updates. Each router periodically tests the status of the physical connection to each of its neighbors, and sends this information to its other neighbors. With this information, each router builds a <i>shortest-path tree</i> with itself as the tree root to identify the shortest path from itself to each destination, and to build its routing table.
Routing policy	Any form of routing that is influenced by factors other than the default algorithmically best route, such as the shortest or quickest path.
prefix	A group of contiguous bits, from 0 to 32 bits in length, that defines a set of addresses.
RIP (Routing Information Protocol)	A distance vector protocol in the IP suite, used by IP and IPX network-layer protocols, that enables routers in the same AS to exchange routing information by means of periodic updates. It is most often used as a very simple IGP within small networks.
route flapping	An instability that is associated with a prefix, where the associated prefix routes may exhibit frequent changes in availability over a period of time.
route reflector	A BGP speaker that advertises routes learned from its route reflector clients to other IBGP neighbors.
route reflector client	A BGP speaker that advertises its learned routes to a route reflector for readvertisement of its routes to the rest of the AS.
TCP (Transport Control Protocol)	The main reliable transport protocol used in the Internet (and with BGP).
transit AS	An AS that has multiple connections to one or more ASs and is used (with certain policy restrictions) to carry both transit and local traffic.
well-known attribute	A BGP attribute that is required to be known by all BGP implementations.



---

# Index

---

## A

- aggregate route, definition 31
- Aggregate routes
  - example 31
- AS Path Attribute 241
- AS Path Filtering 243
- AS, transit 25
- Attribute Type field 41
- authentication, MD5
  - described 44
- Autonomous system (AS)
  - confederations 25
  - definition 22
  - multihomed 25
  - route reflection 25
  - stub 25
  - transit 25

## B

- BGP
  - implementation notes 50
  - interior BGP (IBGP) 22
  - route reflector 32
- BGP debug commands
  - tips for using 177
- BGP path attributes 236
  - optional non-transitive 236
  - optional transitive 236
  - well-known discretionary 236
  - well-known mandatory 236
- BGP Peer Groups 233
- BGP peer/peer group debug commands 185

- BGP speaker
  - 23
  - peer-to-peer session 23
- BGP updates
  - update messages 38
- BGP/OSPF interaction 53

## C

- CIDR
  - configuration example 209
  - conversion 27
  - prefix 27
- Circuitless IP 47
  - configuring with CLI commands 173
- CLI
  - commands 114
  - roadmap of CLI commands 114
- CLI commands
  - config ip as-list 142
  - config ip bgp confederations 128
  - config ip bgp global-debug mask 178
  - config ip bgp global-debug mask all 179
  - config ip bgp global-debug mask error 181
  - config ip bgp global-debug mask event 182
  - config ip bgp global-debug mask filter 183
  - config ip bgp global-debug mask init 183
  - config ip bgp global-debug mask packet 181
  - config ip bgp global-debug mask update 184
  - config ip bgp global-debug mask warning 182
  - config ip bgp neighbor 129
  - config ip bgp neighbor neighbor-debug mask 185
  - config ip bgp neighbor neighbor-debug mask all 186

- config ip bgp neighbor neighbor-debug mask error 188
- config ip bgp neighbor neighbor-debug mask event 190
- config ip bgp neighbor neighbor-debug mask filter 192
- config ip bgp neighbor neighbor-debug mask packet 188
- config ip bgp neighbor neighbor-debug mask state 191
- config ip bgp neighbor neighbor-debug mask update 194
- config ip bgp neighbor neighbor-debug mask warning 190
- config ip bgp neighbor-debug-all mask 179
- config ip circuitless-ip-int 174
- config ip community-list 144
- configuring route redistribution parameters 136
- show commands 146
- show ip circuitless-ip-int info 176

Community Attribute 258

Confederations 264

- example 32

Configuration examples

- 195

- Adding preferences to specific routes 250

- AS path attribute 241

- AS path filtering 243

- basic example 196

- CIDR and aggregate address 209

- community attribute 258

- confederations 264

- EBGP load balancing 212

- EBGP Multihop 212

- load balancing 255

- Local preference 246

- MED attribute 253

- multiple route reflectors 274

- next-hop self 215

- path attributes 236

- peer groups 233

- route flap dampening 281

- route reflectors 269

- synchronization 215

- TCP MD5 authentication 230

conventions, text 18

customer support 20

## D

debug commands

- config ip bgp global-debug mask all 179

- config ip bgp global-debug mask error 181

- config ip bgp global-debug mask event 182

- config ip bgp global-debug mask filter 183

- config ip bgp global-debug mask init 183

- config ip bgp global-debug mask packet 181

- config ip bgp global-debug mask update 184

- config ip bgp global-debug mask warning 182

- config ip bgp neighbor neighbor-debug mask 185

- config ip bgp neighbor neighbor-debug mask all 186

- config ip bgp neighbor neighbor-debug mask error 188

- config ip bgp neighbor neighbor-debug mask event 190

- config ip bgp neighbor neighbor-debug mask filter 192

- config ip bgp neighbor neighbor-debug mask packet 188

- config ip bgp neighbor neighbor-debug mask state 191

- config ip bgp neighbor neighbor-debug mask update 194

- config ip bgp neighbor neighbor-debug mask warning 190

- config ip bgp neighbor-debug-all mask command 179

- global debug command output examples 179

- global debug commands 178

- tips for using 177

## E

Equal-cost multipath (ECMP) 43

**I**

IBGP (interior BGP) 34  
described 22  
implementation notes  
BGP 50  
IP Globals tab  
fields 105, 106, 110

**L**

Load balancing traffic 255  
Local Preference Attribute 246

**M**

Metric field 103  
Multi-exit discriminator (MED) attribute 253  
Multiple route reflectors 274

**N**

Network Layer Reachability Information field 43

**O**

Origin Attribute 237  
configuration example 237  
OSPF/BGP interaction  
described 53

**P**

Path Attributes field 40  
Peers 25  
product support 20  
publications  
hard copy 19

**R**

route balancing  
described 43

Route Flap Dampening 281  
route reflection  
concepts 34  
Route reflectors 269  
RoutePolicy field 103  
RouteSource field 102

**S**

supernet address 28  
Supernet advertisement  
aggregate address 26  
CIDR 26, 27  
support, Nortel Networks\* 20

**T**

technical publications 19  
technical support 20  
text conventions 18  
Total Path Attributes Length field 40  
traffic balancing  
described 43  
transit AS 25  
Translation  
Cisco Systems-to-Nortel Networks equivalents  
285  
Juniper Networks-to-Nortel Networks  
equivalents 301

**W**

Withdrawn Routes field 39  
Withdrawn Routes Length field 39