

Bay Networks

The Merged Company of SynOptics and Wellfleet

Customizing IP Services

Part No. 110079 A

Customizing IP Services

Router Software Version 8.10
Site Manager Software Version 2.10

Part No. 110079 Rev. A
February 1995



Bay Networks

The Merged Company of SynOptics and Wellfleet

Copyright © 1995 Bay Networks, Inc.

All rights reserved. Printed in USA. February 1995.

The information in this document is subject to change without notice. This information is proprietary to Bay Networks, Inc.

The software described in this document is furnished under a license agreement or nondisclosure agreement and may only be used in accordance with the terms of that license. The terms of the Software License are provided with the documentation.

Restricted Rights Legend

Use, duplication, or disclosure by the 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.

Notice for All Other Executive Agencies

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

Trademarks of Bay Networks, Inc.

ACE, BLN, BN, and Wellfleet are registered trademarks and AFN, AN, ASN, BCN, BCNX, BLNX, BNX, CN, FN, FRE, LN, PPX, Bay Networks, and the Bay Networks logo are trademarks of Bay Networks, Inc.

Third-Party Trademarks

3Com is a registered trademark of 3Com Corporation.

AIX, NetView, and IBM are registered trademarks of International Business Machines Corporation.

AppleTalk and EtherTalk are registered trademarks of Apple Computer, Inc.

AT&T and ST are registered trademarks of American Telephone and Telegraph Company.

DEC, DECnet, VAX, and VT100 are trademarks of Digital Equipment Corporation.

Distinct is a registered trademark and Distinct TCP/IP is a trademark of Distinct Corporation.

Fastmac and MADGE are trademarks of Madge Networks, Ltd.

Hayes is a registered trademark of Hayes Microcomputer Products, Inc.

HP is a registered trademark of Hewlett-Packard Company.

Intel is a registered trademark of Intel Corporation.

IPX, NetWare, and Novell are registered trademarks of Novell, Inc.

MCI is a registered trademark of MCI Communications Corporation.

Microsoft, MS, and MS-DOS are registered trademarks and Windows is a trademark of Microsoft Corporation.

Motif and OSF/Motif are registered trademarks of Open Software Foundation, Inc.

Motorola is a registered trademark of Motorola, Inc.

NetBIOS is a trademark of Micro Computer Systems, Inc.

Open Look and UNIX are registered trademarks of UNIX System Laboratories, Inc.

Sun and Solaris are registered trademarks and SPARCstation is a trademark of Sun Microsystems, Inc.

VINES is a registered trademark of Banyan Systems Incorporated.

X Window System is a trademark of the Massachusetts Institute of Technology.

Xerox is a registered trademark and XNS is a trademark of Xerox Corporation.

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

Bay Networks Software License

This Software License shall govern the licensing of all software provided to licensee by Bay Networks ("Software"). Bay Networks will provide licensee with Software in machine-readable form and related documentation ("Documentation"). The Software provided under this license is proprietary to Bay Networks and to third parties from whom Bay Networks has acquired license rights. Bay Networks will not grant any Software license whatsoever, either explicitly or implicitly, except by acceptance of an order for either Software or for a Bay Networks product ("Equipment") that is packaged with Software. Each such license is subject to the following restrictions:

1. Upon delivery of the Software, Bay Networks grants to licensee a personal, nontransferable, nonexclusive license to use the Software with the Equipment with which or for which it was originally acquired, including use at any of licensee's facilities to which the Equipment may be transferred, for the useful life of the Equipment unless earlier terminated by default or cancellation. Use of the Software shall be limited to such Equipment and to such facility. Software which is licensed for use on hardware not offered by Bay Networks is not subject to restricted use on any Equipment, however, unless otherwise specified on the Documentation, each licensed copy of such Software may only be installed on one hardware item at any time.
2. Licensee may use the Software with backup Equipment only if the Equipment with which or for which it was acquired is inoperative.
3. Licensee may make a single copy of the Software (but not firmware) for safekeeping (archives) or backup purposes.
4. Licensee may modify Software (but not firmware), or combine it with other software, subject to the provision that those portions of the resulting software which incorporate Software are subject to the restrictions of this license. Licensee shall not make the resulting software available for use by any third party.
5. Neither title nor ownership to Software passes to licensee.
6. Licensee shall not provide, or otherwise make available, any Software, in whole or in part, in any form, to any third party. Third parties do not include consultants, subcontractors, or agents of licensee who have licensee's permission to use the Software at licensee's facility, and who have agreed in writing to use the Software only in accordance with the restrictions of this license.

-
7. Third-party owners from whom Bay Networks has acquired license rights to software that is incorporated into Bay Networks products shall have the right to enforce the provisions of this license against licensee.
 8. Licensee shall not remove or obscure any copyright, patent, trademark, trade secret, or similar intellectual property or restricted rights notice within or affixed to any Software and shall reproduce and affix such notice on any backup copy of Software or copies of software resulting from modification or combination performed by licensee as permitted by this license.
 9. Licensee shall not reverse assemble, reverse compile, or in any way reverse engineer the Software. [Note: For licensees in the European Community, the Software Directive dated 14 May 1991 (as may be amended from time to time) shall apply for interoperability purposes. Licensee must notify Bay Networks in writing of any such intended examination of the Software and Bay Networks may provide review and assistance.]
 10. Notwithstanding any foregoing terms to the contrary, if licensee licenses the Bay Networks product "Site Manager," licensee may duplicate and install the Site Manager product as specified in the Documentation. This right is granted solely as necessary for use of Site Manager on hardware installed with licensee's network.
 11. This license will automatically terminate upon improper handling of Software, such as by disclosure, or Bay Networks may terminate this license by written notice to licensee if licensee fails to comply with any of the material provisions of this license and fails to cure such failure within thirty (30) days after the receipt of written notice from Bay Networks. Upon termination of this license, licensee shall discontinue all use of the Software and return the Software and Documentation, including all copies, to Bay Networks.
 12. Licensee's obligations under this license shall survive expiration or termination of this license.

Contents

Chapter 1

IP Concepts and Terminology

| | |
|--|------|
| IP Router Functions | 1-2 |
| IP Datagrams | 1-2 |
| IP Addresses | 1-3 |
| Subnet Addressing | 1-6 |
| Supernet Addressing and Classless Interdomain Routing (CIDR) | 1-9 |
| Autonomous Systems and Routing Protocols | 1-10 |
| Routing Information Protocol (RIP) | 1-11 |
| Open Shortest Path First (OSPF) Protocol | 1-12 |
| Border Gateway Protocol (BGP) | 1-12 |
| Exterior Gateway Protocol (EGP) | 1-13 |
| Static Routes | 1-13 |
| Route Preferences | 1-14 |
| Route Weights | 1-15 |
| IP Routing Policies and Filters | 1-16 |
| RFC Compliance | 1-17 |

Chapter 2

Customizing IP Routers and Interfaces

| | |
|--|------|
| Configuring the Router for IP Services | 2-2 |
| Configuring IP Interfaces | 2-2 |
| Multinet Interfaces | 2-4 |
| Specifying a Broadcast Address | 2-4 |
| Subnet Broadcast Addresses | 2-5 |
| Defining a Path to an Adjacent Host | 2-6 |
| Selecting an Address Resolution Protocol | 2-6 |
| Proxy ARP | 2-8 |
| Inverse ARP | 2-9 |
| HP Probe | 2-10 |
| X.25 DDN and X.25 PDN Address Resolution | 2-10 |
| Enabling Source Routing over Token Ring Networks | 2-11 |
| Configuring the Trivial File Transfer Protocol | 2-13 |
| Defining a Circuitless IP Interface | 2-14 |
| Configuring the Revised IP Security Option | 2-14 |
| Security Label Format | 2-15 |
| How RIPS0 Works on the Router | 2-17 |
| Inbound IP Datagrams | 2-17 |
| Forwarded IP Datagrams | 2-18 |
| Originated IP Datagrams | 2-18 |
| Unlabeled IP Datagrams | 2-19 |
| RIPS0 Example | 2-19 |

| | |
|--|------|
| Defining a Static Route | 2-21 |
| Defining a Black Hole for a Supernet | 2-21 |
| Configuring Router Discovery | 2-22 |
| Connecting the Router to a Blacker Front End | 2-22 |
| BFE Addressing | 2-24 |
| Editing IP Parameters | 2-25 |
| Editing IP Global Parameters | 2-26 |
| IP Global Parameter Descriptions | 2-28 |
| Editing IP Interface Parameters | 2-37 |
| IP Interface Parameter Descriptions | 2-39 |
| Deleting IP from an Interface | 2-56 |
| Configuring a Circuitless IP Interface | 2-57 |
| Configuring Static Routes | 2-58 |
| Adding a Static Route | 2-59 |
| Editing a Static Route | 2-61 |
| Deleting a Static Route | 2-64 |
| Configuring a Path to an Adjacent Host | 2-64 |
| Adding an Adjacent Host | 2-65 |
| Editing Adjacent Host Parameters | 2-67 |
| Deleting an Adjacent Host | 2-70 |
| Editing TFTP Parameters | 2-71 |
| TFTP Interface Parameter Descriptions | 2-72 |
| Configuring RIPS0 Support | 2-74 |
| RIPS0 Interface Parameter Descriptions | 2-75 |

| | |
|--|------|
| Configuring Router Discovery | 2-86 |
| Router Discovery Window Parameter Descriptions | 2-87 |
| Configuring Blacker Front-End Support | 2-90 |

Chapter 3

Customizing RIP Services

| | |
|---|-----|
| Routing Information Protocol (RIP) Overview | 3-1 |
| Editing RIP Parameters | 3-2 |
| Editing Routing Information Protocol (RIP) Interface Parameters | 3-2 |
| RIP Parameter Descriptions | 3-4 |

Chapter 4

Customizing OSPF Services

| | |
|---|------|
| Link States and Shortest Path Trees | 4-2 |
| Variable Length Subnet and Supernet Addresses | 4-2 |
| Configuring Network Support on an Interface | 4-3 |
| Defining a Routing Area | 4-5 |
| OSPF Router Classifications | 4-8 |
| Intra-area, Inter-area, and External Routing | 4-9 |
| Configuring a Boundary Router | 4-10 |
| Configuring a Virtual Link through a Transit Area | 4-10 |
| Configuring Cost Metrics | 4-11 |
| Defining a Summary Route | 4-12 |
| Enabling Authentication and Specifying a Password | 4-13 |

| | |
|---|------|
| Constructing an External Route Advertisement | 4-13 |
| Using the Route Weight as the Type 2 Metric | 4-14 |
| Generating and Matching an External Route Tag | 4-15 |
| Generating an External Route Tag for OSPF/BGP Interaction | 4-15 |
| Discovering and Configuring Neighbors | 4-16 |
| Electing a Designated and Backup Designated Router | 4-17 |
| Configuring the OSPF Primary and Backup Soloist | 4-17 |
| Configuring OSPF Message Logging | 4-18 |
| Putting the Pieces Together | 4-18 |
| For More Information about OSPF | 4-20 |
| OSPF Implementation Notes | 4-20 |
| Editing OSPF Parameters | 4-22 |
| Editing OSPF Global Parameters | 4-23 |
| OSPF Global Parameter Descriptions | 4-24 |
| Editing OSPF Area Parameters | 4-30 |
| Adding an Area | 4-31 |
| Editing an Area | 4-32 |
| Deleting an Area | 4-35 |
| Adding a Range to an Area | 4-35 |
| Editing an Area's Range | 4-38 |
| Deleting a Range from an Area | 4-40 |
| Editing OSPF Interface Parameters | 4-40 |
| Editing an OSPF Interface | 4-42 |
| Deleting OSPF from an Interface | 4-51 |
| Adding a Neighbor to an NBMA Interface | 4-51 |

| | |
|---|------|
| Editing a Neighbor | 4-53 |
| Deleting a Neighbor | 4-55 |
| Configuring OSPF Virtual Interfaces | 4-55 |
| Adding a Virtual Interface | 4-56 |
| Editing a Virtual Interface | 4-58 |
| Deleting a Virtual Interface | 4-63 |

Chapter 5

Customizing BGP Services

| | |
|--|------|
| BGP Features | 5-1 |
| Establishing a Peer-to-Peer Connection | 5-4 |
| BGP Messages | 5-6 |
| Open Message | 5-7 |
| Keepalive Message | 5-8 |
| Update Message | 5-8 |
| BGP-3 Update Message Format | 5-8 |
| BGP-4 Update Message Format | 5-10 |
| Notification Message | 5-12 |
| How BGP Selects the Best Path | 5-14 |
| AS Weight and Class Values | 5-14 |
| Routing Policies | 5-15 |
| Calculating the BGP-4 Local Preference Attribute | 5-15 |
| Best Route Calculation for Equal Routes | 5-17 |
| OSPF/BGP Interaction | 5-17 |
| Using IBGP in a Transit AS | 5-18 |

| | |
|---|------|
| Using IBGP in Intra-AS Routing | 5-19 |
| Configuring BGP Message Logging | 5-20 |
| For More Information about BGP | 5-20 |
| BGP Implementation Notes | 5-21 |
| Editing BGP Parameters | 5-22 |
| Editing BGP Global Parameters | 5-23 |
| BGP Global Parameter Descriptions | 5-24 |
| Editing BGP-3 Global Parameters | 5-29 |
| BGP-3 Global Parameter Descriptions | 5-30 |
| Editing BGP-4 Global Parameters | 5-31 |
| BGP-4 Global Parameter Descriptions | 5-32 |
| Configuring a BGP Peer Relationship | 5-32 |
| Adding a BGP Peer | 5-35 |
| Editing a BGP Peer Relationship | 5-37 |
| Deleting a BGP Peer | 5-44 |
| Configuring BGP AS Weights and Weight Classes | 5-45 |
| Specifying a Class and Adding a Weight Value to an AS | 5-46 |
| Editing the Weight Value Parameters of an AS | 5-51 |
| Deleting a Weight Value from an AS | 5-52 |
| Generating BGP Event Messages | 5-53 |
| BGP Debug Parameters Descriptions | 5-55 |
| Deleting BGP from the Router | 5-57 |
| Deleting BGP-3 from the Router | 5-57 |
| Deleting BGP-4 from the Router | 5-58 |

Chapter 6

Customizing EGP Services

| | |
|---|------|
| EGP Overview | 6-1 |
| Neighbor Acquisition Phase | 6-3 |
| Modes | 6-4 |
| Neighbor Reachability Phase | 6-7 |
| Network Reachability Phase | 6-10 |
| Modes | 6-12 |
| For More Information about EGP | 6-13 |
| EGP Implementation Notes | 6-13 |
| Editing EGP Parameters | 6-14 |
| Editing EGP Global Parameters | 6-15 |
| EGP Global Parameter Descriptions | 6-16 |
| Configuring EGP Neighbors | 6-17 |
| Adding an EGP Neighbor | 6-19 |
| Editing an EGP Neighbor | 6-21 |
| Deleting an EGP Neighbor | 6-24 |
| Deleting EGP from the Router | 6-24 |

Chapter 7

IP Multicasting

| | |
|--|-----|
| Host Groups | 7-1 |
| Multicast Networks and Multicast Source Networks | 7-2 |
| Internet Group Management Protocol | 7-2 |
| How IGMP Works | 7-3 |

| | |
|---|------|
| Distance Vector Multicast Routing Protocol | 7-4 |
| How DVMRP Works | 7-4 |
| Calculating a Route Metric | 7-6 |
| Comparing Routes | 7-7 |
| Creating a Shortest Path Tree | 7-8 |
| Identifying a Leaf Network | 7-8 |
| Aging a Route | 7-8 |
| Specifying a Threshold | 7-9 |
| Types of Multicast Support | 7-11 |
| Editing Multicasting Parameters | 7-12 |
| Editing DVMRP Global Parameters | 7-12 |
| DVMRP Global Configuration Parameter Descriptions | 7-13 |
| Editing DVMRP Circuit Parameters | 7-19 |
| DVMRP Circuit Parameter Descriptions | 7-20 |
| Editing DVMRP Tunnel Parameters | 7-23 |
| DVMRP Tunnel Parameter Descriptions | 7-24 |
| Adding a DVMRP Tunnel | 7-26 |
| Add Tunnel Parameters Descriptions | 7-27 |
| Editing IGMP Global Configuration Parameters | 7-28 |
| IGMP Global Configuration Parameters Description | 7-29 |
| Editing IGMP Entry Interface Parameters | 7-30 |
| IGMP Entry Interface Parameters Description | 7-31 |

Chapter 8

NetBIOS over IP

| | |
|--|------|
| Overview of NetBIOS Services | 8-2 |
| Customizing IP Support for NetBIOS | 8-3 |
| Configuring a Static NetBIOS Name | 8-5 |
| Configuring and Customizing a NetBIOS Cache | 8-6 |
| Aging a Cache Entry | 8-7 |
| Customizing a Cache Search | 8-7 |
| Adding a Traffic Filter to a NetBIOS Interface | 8-8 |
| Editing NetBIOS Parameters | 8-9 |
| Editing NetBIOS/IP Global Parameters | 8-9 |
| NetBIOS Global Parameters | 8-11 |
| Editing NetBIOS/IP Interface Table Parameters | 8-17 |
| NetBIOS Interface Parameter Descriptions | 8-18 |
| Editing NetBIOS/IP Static Entry Table Parameters | 8-20 |
| NetBIOS/IP Static Entry Table Parameter Descriptions | 8-22 |
| Adding a Statically Configured NetBIOS Name | 8-23 |
| NBIP Addresses Parameter Descriptions | 8-24 |

Chapter 9

IP Policies

| | |
|--|-----|
| IP Routing Table | 9-1 |
| Configuring Accept Policies | 9-5 |
| IP Accept Policy Parameters Descriptions | 9-7 |
| Common IP Accept Policy Parameters | 9-8 |

| | |
|---|------|
| RIP-Specific Accept Policy Parameters | 9-11 |
| OSPF-Specific Accept Policy Parameters | 9-12 |
| EGP-Specific Accept Policy Parameters | 9-13 |
| BGP-3-Specific Accept Policy Parameters | 9-15 |
| BGP-4-Specific Accept Policy Parameters | 9-19 |
| Configuring Announce Policies | 9-25 |
| IP Announce Policy Parameters | 9-26 |
| Common IP Announce Policy Parameters | 9-27 |
| RIP-Specific Announce Policy Parameters | 9-40 |
| OSPF-Specific Announce Policy Parameters | 9-42 |
| EGP-Specific Announce Policy Parameters | 9-45 |
| BGP-3-Specific Announce Policy Parameters | 9-48 |
| BGP-4-Specific Announce Policy Parameters | 9-53 |

Chapter 10

Import and Export Route Filters

| | |
|--|-------|
| RIP Route Filters | 10-2 |
| Configuring RIP Import Route Filters | 10-2 |
| Adding a RIP Import Route Filter | 10-3 |
| RIP Import Route Filter Parameter Descriptions | 10-5 |
| Editing a RIP Import Route Filter | 10-10 |
| Deleting a RIP Import Route Filter | 10-10 |
| Configuring RIP Export Route Filters | 10-11 |
| Adding a RIP Export Route Filter | 10-11 |
| RIP Export Route Filter Parameter Descriptions | 10-12 |

| | |
|--|-------|
| Editing a RIP Export Route Filter | 10-17 |
| Deleting a RIP Export Route Filter | 10-17 |
| OSPF Route Filters | 10-18 |
| Configuring OSPF Import Route Filters | 10-18 |
| Adding an OSPF Import Route Filter | 10-19 |
| OSPF Import Route Filter Parameter Descriptions | 10-21 |
| Editing an OSPF Import Route Filter | 10-27 |
| Deleting an OSPF Import Route Filter | 10-27 |
| Configuring OSPF Export Route Filters | 10-28 |
| Adding an OSPF Export Route Filter | 10-29 |
| OSPF Export Route Filter Parameter Descriptions | 10-31 |
| Editing an OSPF Export Route Filter | 10-35 |
| Deleting an OSPF Export Route Filter | 10-36 |
| BGP-3 Route Filters | 10-37 |
| Configuring BGP-3 Import Route Filters | 10-37 |
| Adding a BGP-3 Import Route Filter | 10-38 |
| BGP-3 Import Route Filter Parameter Descriptions | 10-40 |
| Editing a BGP-3 Import Route Filter | 10-46 |
| Deleting a BGP-3 Import Route Filter | 10-47 |
| Configuring BGP-3 Export Route Filters | 10-47 |
| Adding a BGP-3 Export Route Filter | 10-48 |
| BGP-3 Export Route Filter Parameter Descriptions | 10-51 |
| Editing a BGP-3 Export Route Filter | 10-57 |
| Deleting a BGP-3 Export Route Filter | 10-58 |

| | |
|--|-------|
| EGP Route Filters | 10-59 |
| Configuring EGP Import Route Filters | 10-59 |
| Adding an EGP Import Route Filter | 10-60 |
| EGP Import Route Filter Parameter Descriptions | 10-62 |
| Editing an EGP Import Route Filter | 10-65 |
| Deleting an EGP Import Route Filter | 10-66 |
| Configuring EGP Export Route Filters | 10-66 |
| Adding an EGP Export Route Filter | 10-67 |
| EGP Export Route Filter Parameter Descriptions | 10-69 |
| Editing an EGP Export Route Filter | 10-74 |
| Deleting an EGP Export Route Filter | 10-75 |

Index

Figures

| | |
|---|------|
| Figure 1-1. Network and Host Portions of IP Addresses | 1-5 |
| Figure 1-2. Internet Segmented into Three Autonomous Systems | 1-11 |
| Figure 2-1. IP Interface | 2-3 |
| Figure 2-2. Multinet Configuration | 2-4 |
| Figure 2-3. ARP Example | 2-7 |
| Figure 2-4. Proxy ARP Example | 2-9 |
| Figure 2-5. IP Routers Source Routing across a Token Ring Network | 2-12 |
| Figure 2-6. RIPS Security Label | 2-15 |
| Figure 2-7. RIPS Network | 2-20 |
| Figure 2-8. Blacker Front-End Network Configuration | 2-23 |
| Figure 2-9. Configuration Manager Window | 2-25 |
| Figure 2-10. Edit IP Global Parameters Window | 2-27 |
| Figure 2-11. IP Interfaces Window | 2-38 |
| Figure 2-12. IP Configuration Window | 2-57 |
| Figure 2-13. IP Static Routes Window | 2-58 |
| Figure 2-14. Add IP Static Route Window | 2-59 |
| Figure 2-15. IP Adjacent Hosts Window | 2-65 |
| Figure 2-16. IP Adjacent Host Configuration Window | 2-66 |
| Figure 2-17. Edit TFTP Parameters Window | 2-71 |
| Figure 2-18. IP Router Discovery Window | 2-86 |
| Figure 3-1. IP RIP Interfaces Window | 3-3 |
| Figure 4-1. Point-to-Multipoint Topology | 4-4 |
| Figure 4-2. OSPF Areas | 4-7 |
| Figure 4-3. Configurable Cost Metrics Usage Example | 4-12 |
| Figure 4-4. OSPF ASE Routes | 4-15 |

| | | |
|--------------|--|------|
| Figure 4-5. | Edit OSPF Global Parameters Window | 4-23 |
| Figure 4-6. | Primary Log Mask Window | 4-28 |
| Figure 4-7. | Backup Log Mask Window | 4-29 |
| Figure 4-8. | OSPF Area List Window | 4-31 |
| Figure 4-9. | OSPF Range List Window | 4-36 |
| Figure 4-10. | OSPF Range Area Window | 4-37 |
| Figure 4-11. | OSPF Interface List Window | 4-41 |
| Figure 4-12. | OSPF Neighbor List Window | 4-52 |
| Figure 4-13. | OSPF Neighbor Configuration Window | 4-52 |
| Figure 4-14. | OSPF Virtual Interface List Window | 4-56 |
| Figure 4-15. | OSPF Virtual Interface Configuration Window | 4-57 |
| Figure 5-1. | BGP Connection between Two Autonomous Systems Running OSPF | 5-2 |
| Figure 5-2. | Establishing and Confirming a Connection between BGP Peers | 5-5 |
| Figure 5-3. | Transit Autonomous System | 5-18 |
| Figure 5-4. | Edit BGP Global Parameters Window | 5-23 |
| Figure 5-5. | Edit BGP-3 Global Parameters Window | 5-29 |
| Figure 5-6. | BGP-4 Global Parameters | 5-31 |
| Figure 5-7. | IP Interface List for BGP Window | 5-33 |
| Figure 5-8. | BGP Peer List Window | 5-34 |
| Figure 5-9. | BGP Peer Parameters Window | 5-35 |
| Figure 5-10. | BGP AS Weight Parameters Window | 5-45 |
| Figure 5-11. | BGP AS Weights Window | 5-46 |
| Figure 5-12. | BGP Debug Parameters Window | 5-53 |
| Figure 5-13. | New BGP Debug Parameters Window | 5-54 |
| Figure 6-1. | EGP Connection between Two Autonomous Systems Running RIP | 6-2 |
| Figure 6-2. | Neighbor Acquisition Sequence | 6-6 |

| | | |
|-------------|--|------|
| Figure 6-3. | Neighbor Reachability Exchange Begins between Two EGP Neighbors | 6-9 |
| Figure 6-4. | Neighbor Reachability Is Established with Both Routers in the UP State | 6-10 |
| Figure 6-5. | Network Reachability Sequence between Two EGP Neighbors | 6-12 |
| Figure 6-6. | Edit EGP Global Parameters Window | 6-15 |
| Figure 6-7. | IP Interface List for EGP Window | 6-17 |
| Figure 6-8. | EGP Neighbors List Window | 6-18 |
| Figure 6-9. | EGP Neighbor Parameters Window | 6-19 |
| Figure 7-1. | Multicast Routers | 7-5 |
| Figure 7-2. | DVMRP Global Configuration Window | 7-12 |
| Figure 7-3. | DVMRP Circuit Parameters Window | 7-19 |
| Figure 7-4. | DVMRP Tunnel Parameters Window | 7-23 |
| Figure 7-5. | DVMRP Tunnel Address Window | 7-26 |
| Figure 7-6. | IGMP Global Configuration Parameters Window | 7-28 |
| Figure 7-7. | IGMP Entry Interface Parameters Window | 7-30 |
| Figure 8-1. | NetBIOS over IP | 8-1 |
| Figure 8-2. | Broadcasting a Name Query Request | 8-4 |
| Figure 8-3. | Returning a Unicast Name Query Response | 8-5 |
| Figure 8-4. | Edit NetBIOS/IP Global Parameters Window | 8-10 |
| Figure 8-5. | NetBIOS/IP Interface Table | 8-17 |
| Figure 8-6. | NetBIOS/IP Static Entry Table Window | 8-21 |
| Figure 8-7. | NBIP Addresses Window | 8-23 |
| Figure 9-1. | IP Routing Table | 9-2 |
| Figure 9-2. | Accept and Announce Policies | 9-3 |
| Figure 9-3. | BGP-3 Accept Policy Filters Window | 9-5 |
| Figure 9-4. | BGP-3 Accept IP Policy Filter Configuration Window | 9-6 |
| Figure 9-5. | BGP-3 Announce Policy Filters Window | 9-25 |

| | |
|--|-------|
| Figure 9-6. BGP-4 Announce IP Policy Filter Configuration Window | 9-26 |
| Figure 10-1. RIP Import Route Filters List Window | 10-2 |
| Figure 10-2. RIP Import Route Filter Configuration Window | 10-4 |
| Figure 10-3. RIP Import Route Filter Window | 10-4 |
| Figure 10-4. RIP Export Route Filters List Window | 10-11 |
| Figure 10-5. RIP Export Route Filter Configuration Window | 10-13 |
| Figure 10-6. RIP Export Route Filters Windows | 10-13 |
| Figure 10-7. OSPF Import Route Filters List Window | 10-19 |
| Figure 10-8. OSPF Import Route Filter Configuration Window | 10-20 |
| Figure 10-9. OSPF Import Route Filters Window | 10-22 |
| Figure 10-10. OSPF Export Route Filters List Window | 10-28 |
| Figure 10-11. OSPF Export Route Filter Configuration Window | 10-29 |
| Figure 10-12. OSPF Export Route Filters Window | 10-30 |
| Figure 10-13. BGP-3 Import Route Filters List Window | 10-38 |
| Figure 10-14. BGP-3 Import Route Filter Configuration Window | 10-39 |
| Figure 10-15. BGP-3 Import Route Filter Window | 10-40 |
| Figure 10-16. BGP-3 Export Route Filters List Window | 10-48 |
| Figure 10-17. BGP-3 Export Route Filter Configuration Window | 10-49 |
| Figure 10-18. BGP-3 Export Route Filter Window | 10-50 |
| Figure 10-19. EGP Import Route Filters List Window | 10-60 |
| Figure 10-20. EGP Import Route Filter Configuration Window | 10-61 |
| Figure 10-21. EGP Import Route Filter Window | 10-62 |
| Figure 10-22. EGP Export Route Filters List Window | 10-67 |
| Figure 10-23. EGP Export Route Filter Configuration Window | 10-68 |
| Figure 10-24. EGP Export Route Filter Window | 10-69 |

Tables

| | | |
|------------|---|------|
| Table 1-1. | Possible Subnet Masks for Class B and Class C Addresses | 1-8 |
| Table 1-2. | IP Router RFC Support | 1-17 |
| Table 2-1. | BFE Required X.25 Packet-Level Parameter Settings | 2-91 |
| Table 2-2. | BFE Required X.25 Network Service Record Parameter Settings | 2-94 |
| Table 4-1. | OSPF Router Classifications | 4-8 |
| Table 5-1. | BGP-3 Path Attributes | 5-9 |
| Table 5-2. | BGP-4 Optional Path Attributes | 5-11 |
| Table 5-3. | Notification Message Error Codes and Subcodes | 5-13 |
| Table 6-1. | Router Mode Determinator | 6-5 |
| Table 6-2. | UP and DOWN State Thresholds | 6-8 |
| Table 7-1. | Recommended Hop Metrics | 7-6 |
| Table 7-2. | Recommended TTL and Threshold Values | 7-10 |

About This Guide

If you are responsible for configuring and managing Wellfleet® routers, you need to read this guide.

This guide describes how to customize your router software for Internet Protocol (IP) services and the following IP protocols:

- ❑ Routing Information Protocol (RIP)
- ❑ Open Shortest-Path First (OSPF) Protocol
- ❑ Border Gateway Protocol, Version 3 (BGP-3)
- ❑ Border Gateway Protocol, Version 4 (BGP-4)
- ❑ Exterior Gateway Protocol (EGP)
- ❑ IP Multicasting Protocols
- ❑ NetBIOS over IP

Refer to this guide for

- ❑ An overview of the IP routing protocol, and instructions on editing IP global and interface parameters and configuring basic IP services.
- ❑ An overview of RIP, a description of how Wellfleet RIP routing services work, and instructions on editing RIP parameters and configuring RIP route filters.
- ❑ An overview of OSPF, a description of how Wellfleet OSPF routing services work, and instructions on editing OSPF parameters and configuring OSPF route filters.

-
- ❑ An overview of BGP, BGP-3 and BGP-4, a description of how Wellfleet BGP routing services work, and instructions on editing BGP parameters.
 - ❑ An overview of EGP, a description of how Wellfleet EGP routing services work, and instructions on editing EGP parameters.
 - ❑ An overview of IP multicasting services and instructions on editing multicasting parameters.
 - ❑ An overview of NetBIOS services, a description of how NetBIOS works over IP, and instructions for setting NetBIOS over IP parameters.
 - ❑ An overview of IP accept and announce policies and a description of IP policy parameters.
 - ❑ An overview of IP import and export filters and a description of IP import and export parameters.

For information and instructions about the following topics, see *Configuring Wellfleet Routers*.

- ❑ Initially configuring and saving an IP interface on which RIP, OSPF, BGP, and/or EGP are enabled
- ❑ Retrieving a configuration file
- ❑ Rebooting the router with a configuration file

Before You Begin

Before using this guide, you must complete the following procedures:

- Create and save a configuration file that contains at least one IP interface.
- Retrieve the configuration file in local, remote, or dynamic mode.

Refer to *Configuring Wellfleet Routers* for instructions.

How to Get Help

For additional information or advice, contact the Bay Networks Help Desk in your area:

| | |
|-------------------|-----------------|
| United States | 1-800-2LAN-WAN |
| Valbonne, France | (33) 92-966-968 |
| Sydney, Australia | (61) 2-903-5800 |
| Tokyo, Japan | (81) 3-328-0052 |

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 command syntax is ping <ip_address> , you enter ping 192.32.10.12 |
| arrow character (→) | Separates menu and option names in instructions. Example: Protocols→AppleTalk identifies the AppleTalk option in the Protocols menu. |
| brackets ([]) | Indicate optional elements. You can choose none, one, or all of the options. |
| user entry text | Denotes text that you need to enter. Example: Start up the Windows environment by entering the following after the prompt: win |
| command text | Denotes command names in text. Example: Use the xmodem command. |

| | |
|-------------------------|--|
| <i>italic text</i> | Indicates variable values in command syntax descriptions, new terms, file and directory names, and book titles. |
| screen text | Indicates data that appears on the screen. Example: Set Trap Monitor Filters |
| ellipsis points | Horizontal (. . .) and vertical (:) ellipsis points indicate omitted information. |
| quotation marks (" ") | Indicate the title of a chapter or section within a book. |
| vertical line () | Indicates that you enter only one of the parts of the command. The vertical line separates choices. Do not type the vertical line when entering the command. Example: If the command syntax is show at routes nets , you enter either show at routes or show at nets , but not both. |

Acronyms

| | |
|-------|--|
| ANSI | American National Standards Institute |
| ARP | Address Resolution Protocol |
| BGP | Border Gateway Protocol |
| CIDR | Classless Interdomain Routing |
| ATM | Asynchronous Transfer Mode |
| CMIP | Common Management Information Protocol |
| DVMRP | Distance Vector Multicast Routing Protocol |
| EGP | Exterior Gateway Protocol |
| FDDI | Fiber Distributed Data Interface |
| IEEE | Institute of Electrical and Electronic Engineers |
| IGMP | Internet Group Management Protocol |
| IGP | interior gateway protocol |
| ILI | intelligent link interface |
| IS-IS | Intermediate System to Intermediate System |
| MAC | media access control |
| MOP | Maintenance Operations Protocol |
| OSI | Open Systems Interconnection |

| | |
|---------------|--|
| OSPF | Open Shortest Path First |
| PVCs | permanent virtual circuits |
| QENET | Quad Ethernet Link Module |
| RIP | Routing Information Protocol |
| SMDS | Switched Multimegabit Data Services |
| SNAP | Subnetwork Access Protocol |
| SNMP | Simple Network Management Protocol |
| SRM | system resource modules |
| SVCs | switched virtual circuits |
| TCP/IP | Transmission Control Protocol/Internet Protocol |
| TFTP | Trivial File Transfer Protocol |

Chapter 1

IP Concepts and Terminology

The following sections introduce concepts and terminology used in this manual:

- “IP Router Functions” on page 1-2
- “IP Datagrams” on page 1-2
- “IP Addresses” on page 1-3
- “Autonomous Systems and Routing Protocols” on page 1-10
- “Route Preferences” on page 1-14
- “Route Weights” on page 1-15
- “IP Routing Policies and Filters” on page 1-16
- “RFC Compliance” on page 1-17

IP Router Functions

An IP router performs three basic functions:

- ❑ Acquires knowledge of other routers and hosts on the network.

IP routers use routing protocols—for example OSPF and BGP—to learn transmission paths (or routes) to other networks and to hosts residing on networks directly connected to the router.

- ❑ Stores network topology information about transmission paths in routing tables.
- ❑ Selects the best path, based on the information in its routing tables, for a particular *datagram* (a self-contained unit of data) to follow to reach its destination.

IP routers process each datagram individually. The datagram header provides the router with the destination IP address, as well as other routing information. Routers select a transmission path based on the IP address of the destination network, not of the destination host.

IP Datagrams

An IP datagram is the unit of data exchanged between IP modules. In addition to data, a datagram includes a header with fields that provide the following information used by IP routers:

- ❑ Type of Service

This field indicates the quality of service the datagram requires. The IP router inspects the Type of Service field to obtain information about the datagram's precedence and expected delay characteristics.

□ Time to Live

This field determines the datagram's lifetime in the Internet system. Each time an IP router processes the datagram header, it decrements the value in the Time to Live field by at least one. When the value reaches zero, the IP router discards the datagram, unless it is destined for the router itself, thus preventing undeliverable datagrams from looping endlessly through the network, consuming Internet resources.

□ Options

This field may or may not be present in a datagram; therefore, IP datagrams vary in length. There are three classes of Options:

- Security, which specifies security level and distribution restrictions.
- Timestamps, which is a 32-bit value measured in milliseconds since midnight universal time, or any other value if the high-order bit is set to 1.
- Special Routing, which specifies host-discovered paths to other hosts, or a specific path for the datagram to take.

□ Header Checksum

This field contains a value that the IP router calculates each time it processes a datagram's IP header. The algorithm used to calculate the checksum value is a 16-bit ones complement addition of the 16-bit words contained only within the IP header. The IP router discards datagrams received with an incorrect IP header checksum.

IP Addresses

An IP address consists of 32 bits having the form *network.host*. The network portion is a network number ranging from 8 to 24 bits. The host portion is the remaining 8 to 24 bits identifying a specific host on the network. The Internet Network Information Center (NIC) assigns the network portion of the IP address. Your network administrator assigns the host portion.


The NIC recognizes three primary classes of networks: A, B, and C. In addition, the NIC has recently identified two other classes: Class D for networks that support multicasting, which allows an IP datagram to be transmitted to a single multicast group consisting of hosts spread across separate physical networks, and Class E for experimental networks. The IP router does not fully support Class D or Class E networks.

Based on the size of the network, the NIC classifies a network as Class A, B, or C (the most common). The network class determines the number of bits assigned to the network and host portions of the IP address, as follows:

| Network Size | Class | Network Portion | Host Portion |
|---------------------|-------|-----------------|--------------|
| Over 65,534 hosts | A | 8 bits | 24 bits |
| 254 to 65,533 hosts | B | 16 bits | 16 bits |
| Under 254 hosts | C | 24 bits | 8 bits |

The position of the first bit set to 0 (whether it is the first, second, third, or fourth bit) in the first octet of an IP address indicates the network class (A, B, C, or D). If no bit is set to 0, it is a Class E network.

Figure 1-1 shows the placement of the first bit set to 0 for Class A, B, and C networks. The figure also shows how a network's class affects the network and host portion of the IP address.



| | First Octet | Range | Example | Network | Host |
|---------|-------------|---------|-------------|---------|------|
| Class A | | 1-127 | 25.0.0.1 | 25 | 1 |
| Class B | | 128-191 | 140.250.0.1 | 140.250 | 1 |
| Class C | | 192-223 | 192.2.3.1 | 192.2.3 | 1 |



Figure 1-1. Network and Host Portions of IP Addresses

You specify IP addresses in dotted decimal notation. To specify an IP address in dotted decimal notation, you convert each 8-bit octet of the IP address to a decimal number, and separate the numbers by decimal points.

For example, you specify the 32-bit IP address 10000000 00100000 00001010 10100111 in dotted decimal notation as 128.32.10.167. The most significant 2 bits (10) in the first octet indicate that the network is Class B; therefore, the first 16 bits compose the NIC-assigned network portion field. The third octet (00001010) and fourth octet (10100111) compose the host field.

Subnet Addressing

The concept of subnetworks (or subnets) extends the IP addressing scheme. Subnets are two or more physical networks that share a common network-identification field (the NIC-assigned network portion of the 32-bit IP address). Subnets allow an IP router to hide the complexity of multiple LANs from the rest of the internet.

With subnets, you partition the host portion of an IP address into a subnet number and a “real” host number on that subnet. The IP address is then defined by *network.subnet.host*. Routers outside the network do not interpret separately the subnet and host portions of the IP address.

Routers inside a network containing subnets use a 32-bit subnet mask that identifies the extension bits. In *network.subnet.host*, the *subnet.host* portion (or the local portion) contains an arbitrary number of bits. The network administrator allocates bits within the local portion to subnet and host, and then assigns values to subnet and host.

For example, the following is the IP address of a network that contains subnets: 10000000 00100000 00001010 10100111. You specify this address in dotted decimal notation as 128.32.10.167.

The second bit of the first octet is set to 0, indicating that the network is a Class B network. Therefore, the NIC-assigned network portion contains 16 bits, and the locally assigned local portion contains 16 bits.

The network administrator allocates the 16 bits in the local portion field as follows:

- ❑ Allocates the upper 8 bits (00001010) with a value of 10 to the subnet portion
- ❑ Allocates the lower 8 bits (10100111) with a value of 167 to the host portion

In other words, the 16-bit local portion field, together with the 16-bit network field, specify host 167 on Subnet 10 of network 128.32.

You now need a subnet mask to identify those bits in the 32-bit IP address that specify the network field and those bits that specify the subnet field. Like the IP address, you specify the subnet mask in dotted decimal notation.

You construct a subnet mask as follows:

- ❑ Assign a value of 1 to each of the 8, 16, or 24 bits in the network field.
- ❑ Assign a value of 1 to each bit in the subnet field.
- ❑ Assign a value of 0 to each bit in the host field.
- ❑ Convert the resulting 32-bit string to dotted decimal notation.

For example, to construct a subnet mask for the IP address described earlier (10000000 00100000 00001010 10100111), do the following:

1. Assign a value of 1 to each bit in the network field.

The position of the first bit set to 0 in the first octet of the IP address indicates that the network is Class B; therefore, the network field contains 16 bits: 11111111 11111111.

2. Assign a value of 1 to each bit in the subnet field.

The network administrator allocated the upper eight bits of the local portion to the subnet portion, as follows: 11111111.

3. Assign a value of 0 to each bit in the host field.

The network administrator allocated the lower eight bits of the local portion field to the host identification, as follows: 00000000.

4. Convert the resulting 32-bit string (11111111 11111111 11111111 00000000) to dotted decimal notation, as follows: 255.255.255.000.

Table 1-1 shows the range of possible subnet masks for Class B and Class C addresses, along with the number of bits that the mask allocates for a subnet address, the number of recommended subnets associated with the mask, and the number of hosts per subnet.

Table 1-1. Possible Subnet Masks for Class B and Class C Addresses

| Number of Bits | Subnet Mask | Number of Subnets (Recommended) | Number of Hosts per Subnet |
|-----------------------|--------------------|--|-----------------------------------|
| Class B | | | |
| 2 | 255.255.192.0 | 2 | 16,382 |
| 3 | 255.255.224.0 | 6 | 8,190 |
| 4 | 255.255.240.0 | 14 | 4,094 |
| 5 | 255.255.248.0 | 30 | 2,046 |
| 6 | 255.255.252.0 | 62 | 1,022 |
| 7 | 255.255.254.0 | 126 | 510 |
| 8 | 255.255.255.0 | 254 | 254 |
| 9 | 255.255.255.128 | 510 | 126 |
| 10 | 255.255.255.192 | 1,022 | 62 |
| 11 | 255.255.255.224 | 2,046 | 30 |
| 12 | 255.255.255.240 | 4,094 | 14 |
| 13 | 255.255.255.248 | 8,190 | 6 |
| 14 | 255.255.255.252 | 16,382 | 2 |
| Class C | | | |
| 2 | 255.255.255.192 | 2 | 62 |
| 3 | 255.255.255.224 | 6 | 30 |
| 4 | 255.255.255.240 | 14 | 14 |
| 5 | 255.255.255.248 | 30 | 6 |
| 6 | 255.255.255.252 | 62 | 2 |

Supernet Addressing and Classless Interdomain Routing (CIDR)

A *supernet* is a group of networks identified by contiguous network addresses. IP service providers can assign customers blocks of contiguous addresses to define supernets as needed.

Each supernet has a unique supernet address that consists of the upper bits shared by all of the addresses in the contiguous block. For example, consider the following block of contiguous 32-bit addresses (192.32.0.0 through 192.32.7.0 in decimal notation).

```

11000000 00100000 00000000 00000000
11000000 00100000 00000001 00000000
11000000 00100000 00000010 00000000
11000000 00100000 00000011 00000000
11000000 00100000 00000100 00000000
11000000 00100000 00000101 00000000
11000000 00100000 00000111 00000000

```

The supernet address for this block is 11000000 00100000 00000, the 21 upper bits shared by the 32-bit addresses.

A complete supernet address consists of an *address/mask* pair:

- *address* is the first 32-bit IP address in the contiguous block. In this example, the address is 11000000 00100000 00000000 00000000 (192.32.0.0 in decimal notation).
- *mask* is a 32-bit string containing a set bit for each bit position in the supernet part of the address. The mask for the supernet address in this example is 11111111 11111111 11111000 00000000 (255.255.248.0 in dotted decimal notation).

The complete supernet address in this example is 192.32.0.0/255.255.248.0.

Classless interdomain routing (CIDR) is an addressing scheme that employs supernet addresses to represent multiple IP destinations. Rather than advertise a separate route for each destination in a supernet, a router can use a supernet address to advertise a single

route — called an *aggregate* route — that represents all of the destinations. This reduces the size of the routing tables used to store advertised IP routes.

BGP-4 supports classless interdomain routing. OSPF supports classless routing within a domain.

Autonomous Systems and Routing Protocols

LANs and WANs interconnected by IP routers form a group of networks called an internet. For administrative purposes, an internet is divided into autonomous systems. An autonomous system is simply a collection of routers and hosts. Figure 1-2 depicts a sample internet segmented into three autonomous systems.

Routers inside an autonomous system use an interior gateway protocol (IGP) to communicate network topology changes to each other. Routers in separate autonomous systems use an exterior gateway protocol (EGP) to communicate. The IP router implements two dynamic IGPs: the Routing Information Protocol (RIP) and the Open Shortest Path First (OSPF) Protocol. The IP router implements two EGPs: the Border Gateway Protocol (BGP) and the Exterior Gateway Protocol (EGP).

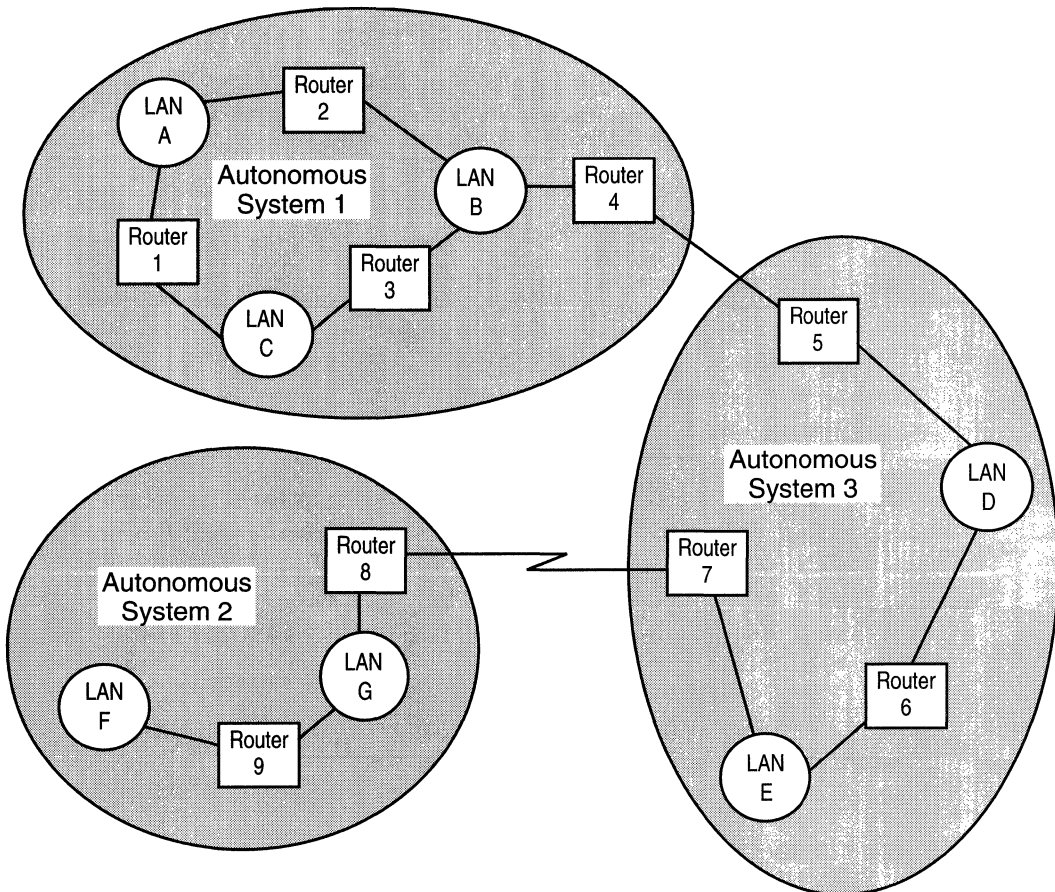


Figure 1-2. Internet Segmented into Three Autonomous Systems

Routing Information Protocol (RIP)

The Routing Information Protocol (RIP) is a distance-vector protocol that enables routers in the same autonomous system to exchange routing information by means of periodic RIP updates. Routers transmit their own RIP updates to neighboring networks and listen for RIP updates from the routers on those neighboring networks. Routers use the information in the RIP updates to keep their internal routing

tables current. For RIP, the “best” path to a destination is the shortest path (the path with the fewest hops). RIP computes distance as a metric, usually the number of hops (or routers) from the origin network to the target network.

For RIP configuration information, see Chapter 3.

Open Shortest Path First (OSPF) Protocol

The Open Shortest Path First (OSPF) protocol is an IGP intended for use in large networks. Using a link state algorithm, OSPF exchanges routing information between routers in an autonomous system. Routers synchronize their topological databases. Once the routers are synchronized and the routing tables are built, the routers will flood topology information only in response to some topological change. For OSPF, the “best” path to a destination is the path that offers the least cost metric delay. In OSPF, cost metrics are configurable, allowing you to specify preferred paths.

OSPF supports CIDR and can carry supernet advertisements within a routing domain.

For a more detailed overview and OSPF configuration information, see Chapter 4.

Border Gateway Protocol (BGP)

The Border Gateway Protocol (BGP) is an exterior gateway protocol used to exchange network reachability information with other BGP systems. BGP routers form relationships with other BGP routers. Using an entity called a BGP speaker, BGP routers transmit and receive current routing information over a reliable transport layer connection. Because a reliable transport mechanism is used, periodic updates are not necessary.

BGP updates contain “path attributes” that describe the route to some set of destination networks. When multiple paths are available, BGP compares these path attributes to choose the preferred path.

BGP-3 and BGP-4 are supported. BGP-4 is the border gateway protocol that supports CIDR.

For a more detailed overview and BGP configuration information, see Chapter 5.

Exterior Gateway Protocol (EGP)

The Exterior Gateway Protocol (EGP-2) is an exterior gateway protocol used to exchange network reachability information between routers in different autonomous systems. An IGP, such as RIP or OSPF, is used within an AS to facilitate the communication of routing information with the autonomous system. The routers that serve as the end points of a connection between two autonomous systems run an exterior gateway protocol, such as EGP-2.

Routers establish EGP neighbor relationships in order to periodically exchange reliable network reachability information. The router uses this information to maintain a list of gateways, the networks the gateways can reach, and the corresponding distances.

For a more detailed overview and EGP configuration information, see Chapter 6.

Static Routes

You can manually configure a route to another network and enter the route in the IP routing table. Such a route is called a static route.

For information about static routes and instructions on including a static route in the routing table, see “Defining a Static Route” on page 2-21.

Route Preferences

The IP router maintains an internal routing table. When making a decision on how to forward a datagram, the IP router consults the table to determine the specific route a datagram should take. A routing table can contain direct routes for the IP router's network interfaces, static routes, and the routes learned from RIP, OSPF, BGP, and/or EGP, if enabled (adjacent hosts are maintained in a separate table).

It is possible for a routing table to contain multiple routes to the same destination. In such a situation, IP uses (among other information) a preference value to determine which route to select. Preference values range from 1 to 16 (the higher the number, the greater the preference).

By default, RIP, BGP, EGP, and OSPF external routes have a preference value of 1. Static routes, direct routes, and OSPF intra-area and inter-area routes have a default preference of 16.

You can configure a preference value in the range of 1 to 16 for RIP, BGP, EGP, OSPF external, and static routes. The preference of direct routes and OSPF intra-area and inter-area routes cannot be user-configured.

To assign a greater or lesser preference to a static route, you supply a value when you define the route. For instructions, see the static route Preference parameter on page 2-64.

To assign a preference to a route learned by RIP, OSPF, BGP, and EGP, you configure an accept policy for the route. If an incoming route matches the policy, IP assigns the preference value you specify to the route and considers the route for possible inclusion in the routing table.

For instructions, see the Route Preference parameter on page 9-10.

Route Weights

Route weight calculation is an internal tool that IP uses to facilitate selection of the best route among alternative routes to the same destination. Route selection criteria are encoded into the route weight in a way that allows IP to compare routes simply by comparing their weight values, regardless of route sources.

Route weight calculation increases the efficiency of the route selection process and at the same time reduces the size of the routing database, since all route selection parameters for each route are encoded in a single integer — the weight value — rather than stored in separate variables.

Using selection criteria encoded in the route weight, IP chooses routes in the following order:

1. The route with the highest preference value (see “Route Preferences” on page 1-14)
2. A direct or OSPF intra-area route with the lowest metric

Note: Beginning with Wellfleet Router Software Version 8.00, a direct route (interface) that is part of an OSPF area is not automatically chosen over an OSPF intra-area route. As a result, it is possible to configure a slow direct link (for example, a backup dialup line) with a high metric value (`wfIpInterfaceCost`) and route packets to a fast link on another router in the same OSPF area. Direct routes that are not included in an OSPF area are assumed to have a metric of 0 and are always chosen over other routes.

3. A direct route with the lowest metric
4. An OSPF intra-area route with the lowest metric
5. An OSPF inter-area route with the lowest metric
6. An OSPF Type 1 external route with the lowest metric
7. A BGP route with the highest LOCAL_PREF value

8. A RIP route with the lowest metric
9. An EGP route with the lowest metric
10. A static route with the lowest metric
11. An OSPF Type 2 external route with a pre-8.00-style metric

Note: If OSPF is configured to propagate external routes using the route weight as the Type 2 metric, routes that are received as OSPF ASE Type 2 routes are evaluated according to their respective origins (for example, RIP or BGP).

IP Routing Policies and Filters

The IP router allows you to control the flow of routing data to and from the routing tables. This control is provided by two mechanisms:

- IP accept and announce policies
- Import and export filters

Note: Accept and announce policies provide a superset of the parameters provided by import and export filters. We currently support both IP policies and IP route filters. However, network administrators using import and export filters for routing table management should migrate as quickly as possible to IP policies. In a future release, support for import and export filters will be dropped.

IP accept policies (and the subset of parameters provided by import filters) govern the addition of new RIP-, OSPF-, BGP-, or EGP-derived routes to the routing tables. When RIP, OSPF, BGP, or EGP receives a new routing update, it consults its accept policies to validate the information before entering the update into the routing tables. Accept policies contain search information (to match fields in incoming routing updates) and action information (to specify the action to take with matching routes).

IP announce policies (and the subset of parameters provided by export filters) govern the propagation of RIP, OSPF, BGP, or EGP routing information. When preparing a routing advertisement, RIP, OSPF, BGP, or EGP consults its announce policies to determine whether the routes to specific networks are to be advertised and how they are to be propagated. Announce policies contain network numbers (to associate a policy with a specific network) and action information (to specify a route propagation procedure).

IP accept and announce policies and policy parameters are described in Chapter 9.

IP import and export filters and filter parameters are described in Chapter 10.

RFC Compliance

Table 1-2 lists the Internet Requests for Comments (RFCs) with which the IP router complies. This chapter assumes you are familiar with these RFCs.

Table 1-2. IP Router RFC Support

| RFC | Specifies |
|------------|--|
| 768 | User Datagram Protocol (UDP) |
| 783 | Trivial File Transfer Protocol (TFTP) |
| 791 | Internet Protocol (IP) |
| 792 | Internet Control Message Protocol (ICMP) |
| 826 | Address Resolution Protocol (ARP) |
| 950 | Internet subnetting procedures |
| 1009 | Internet gateways |
| 1058 | Routing Information Protocol (RIP) |
| 1063 | Maximum Transmission Unit (MTU) discovery option |

Table 1-2. IP Router RFC Support (*continued*)

| RFC | Specifies |
|------------|--|
| 1247 | Open Shortest Path First (OSPF) Protocol Version 2 |
| 1157 | Simple Network Management Protocol (SNMP) |
| 1188 | IP over FDDI networks |
| 1042 | IP over IEEE 802.x networks |
| 1027 | Proxy ARP |
| 1112 | Host Extensions for IP Multicasting |
| 1256 | ICMP Router Discovery Messages |
| 1267 | BGP-3 |
| 1403 | BGP OSPF Interaction |
| 1654 | BGP-4 |

Chapter 2

Customizing IP Routers and Interfaces

This chapter consists of the following sections:

- ❑ “Configuring the Router for IP Services” on page 2-2
- ❑ “Configuring IP Interfaces” on page 2-2
- ❑ “Specifying a Broadcast Address” on page 2-4
- ❑ “Defining a Path to an Adjacent Host” on page 2-6
- ❑ “Selecting an Address Resolution Protocol” on page 2-6
- ❑ “Enabling Source Routing over Token Ring Networks” on page 2-11
- ❑ “Configuring the Trivial File Transfer Protocol” on page 2-13
- ❑ “Defining a Circuitless IP Interface” on page 2-14
- ❑ “Configuring the Revised IP Security Option” on page 2-14
- ❑ “Defining a Static Route” on page 2-21
- ❑ “Defining a Black Hole for a Supernet” on page 2-21
- ❑ “Configuring Router Discovery” on page 2-22
- ❑ “Connecting the Router to a Blacker Front End” on page 2-22
- ❑ “Editing IP Parameters” on page 2-25

Configuring the Router for IP Services

You configure and customize the router for IP services by setting parameters on the Edit IP Global Parameters window. These parameters allow you to enable and disable IP on the router, specify whether the router forwards IP traffic to other routers, and supply aging, time-to-live, and other values. IP global parameters also allow you to help IP software preallocate system resources by providing the router with an estimate of the number of networks and hosts the router will be required to support.

For instructions on customizing IP services on the router, see “Editing IP Global Parameters” on page 2-26.

Configuring IP Interfaces

An IP network interface consists of a physical circuit configured with the appropriate data link and IP protocols. Each interface connects the router to one or more IP networks.

For example, the router in Figure 2-1 is configured with three IP interfaces. One of these interfaces is a point-to-point interface that connects the router to a single long-haul medium terminated by a host or another router. The other two interfaces are LAN interfaces that connect the router to an Ethernet or FDDI local area medium.

An IP interface can provide access to multiple networks. For example, in Figure 2-1, LAN interface 1 provides a connection to both LAN B and LAN C.

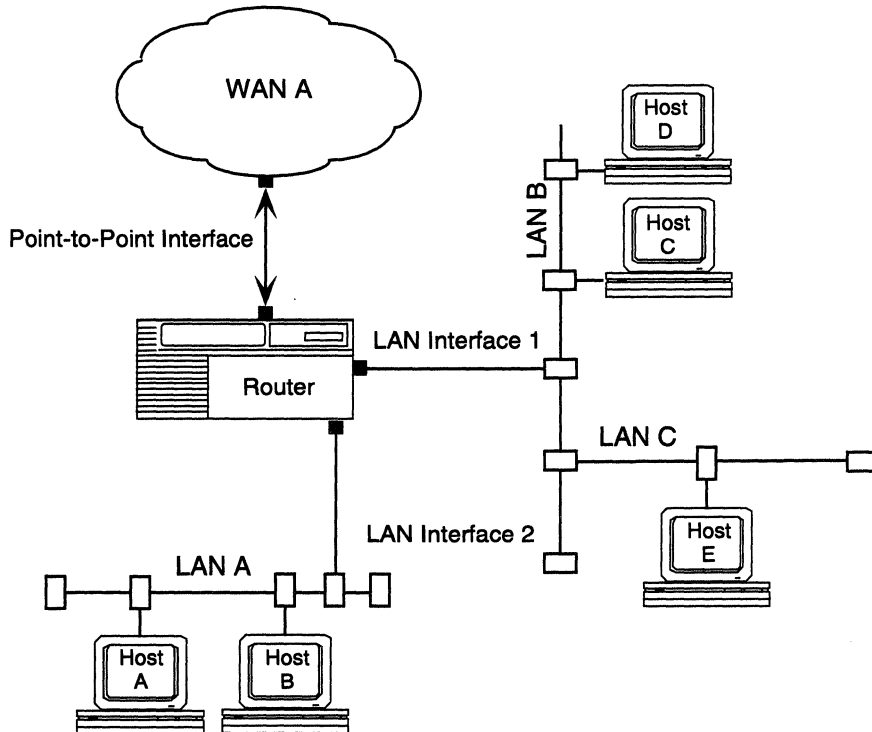


Figure 2-1. IP Interface

As part of the router configuration process, the network administrator associates a network with an interface by assigning the network's unique IP address to the circuit on which the interface is configured.

For instructions on customizing an IP interface, see “Editing IP Interface Parameters” on page 2-37.

Multinet Interfaces

The multinet capability allows you to assign multiple IP network/subnet addresses to a single circuit; each IP address represents a separate network interface on the circuit.

Multinet is commonly used in IP networks containing hosts that do not understand subnetting. For example, in Figure 2-2, hosts A, B, and C are connected by a router. Because the hosts do not understand subnetting, A, B, and C operate as if they are all on the same network. While A and C are on the same network, B is not. To facilitate connectivity between the three hosts, the router is configured with interfaces that connect three distinct subnets, as defined by the mask 255.255.255.0. In Figure 2-2, A and C are on a multinet interface.




Figure 2-2. Multinet Configuration

Specifying a Broadcast Address

Broadcasting occurs when the IP router transmits a single packet to every host on an attached network. To do so, it uses a broadcast address that refers to all hosts on the network. A broadcast address is simply an IP address that contains all 1s or all 0s in the host portion.

For example, if you have an IP network with IP address 10.3.45.12, you can configure a broadcast address for that network, as follows:

- Because the address is for a Class A network (the network portion is 1 byte), the host portion contains 3 bytes.
- Because the host portion of a broadcast address consists of all 1s or all 0s, the broadcast address for that network can be one of the following: 10.255.255.255, 10.0.0.0, 255.255.255.255, or 0.0.0.0.

Some networks do not support broadcasts; thus, configuring an IP broadcast address does not guarantee efficient broadcast delivery.

To configure a broadcast address on an IP interface, see the Broadcast Address parameter on page 2-40.

Subnet Broadcast Addresses

The way you configure a broadcast address for a subnet is different from the way you configure a broadcast address for a network. Because you extend the network portion of the IP address when you create subnets, you automatically take away from the host portion of the IP address. To configure a subnet broadcast, you take the subnet mask for that subnet and invert it. For example, if the subnet's IP address is 10.4.2.3, and the mask is 255.255.0.0, then the subnet broadcast address is either 10.4.255.255 or 10.4.0.0.

Note: IP permits an all-zero subnet address but discourages its use for the following reason. If an all-zero subnet address and an all-zero broadcast address are both valid, the router cannot distinguish an all subnets broadcast from a directed broadcast for the zero subnet. For details, see the Zero Subnet Enable parameter on page 2-34.

Defining a Path to an Adjacent Host

An adjacent host is a network device. This device may or may not be a router but must reside on a locally attached network. You configure a transmission path to an adjacent host if your topology includes a network or hosts that do not implement Address Resolution Protocol (ARP). In this situation, you need to configure an adjacent host to each non-ARP device that resides on a network directly connected to the router.

Also, if a local network does implement ARP, you may want to configure adjacent hosts to pre-empt the ARP process and, thereby, resolve the media address control (MAC) address.

To configure a connection to an adjacent host, see “Configuring a Path to an Adjacent Host” on page 2-64.

Selecting an Address Resolution Protocol

The IP router needs both a physical address and an IP address to transmit a datagram. In situations where the router knows only the network host’s IP address, the Address Resolution Protocol (ARP) enables the router to determine a network host’s physical address by binding a 32-bit IP address to a 48-bit MAC address. A router can use ARP across a single network only, and the network hardware must support physical broadcasts.




Figure 2-3. ARP Example

For example, in Figure 2-3, the router and Host C are on the same physical network. Both devices have an assigned IP address (the router's is 140.250.200.1 and Host C's is 140.250.200.4) and both devices have an assigned physical address (the router's is 00 00 A2 00 00 01 and Host C's is 00 00 A2 00 10 40).

In Figure 2-3, the router wants to send a packet to Host C, but only knows Host C's IP address. The router uses ARP to determine Host C's physical address, as follows:

1. The router broadcasts a special packet, called an ARP request, that asks IP address 140.250.200.4 to respond with its physical address.
2. All network hosts receive the broadcast request.
3. Only Host C responds with its hardware address.

The router maps Host C's IP address (140.250.200.4) to its physical address (00 00 A2 00 10 40) and saves the results in an address-resolution cache for future use.

Note: It is possible for the router to send out ARP requests even if ARP, which is a dynamically loaded module, is not currently loaded on the router. It is the responsibility of the network administrator to ensure that ARP is loaded correctly on a slot. To do this through Site Manager, select Events Manager→Options→Filters; then select LOADER and Debug, and do a File→Get Current Log File. Verify that ARP is loaded on a slot by locating the following message in the log:

```
#xx:01/01/95 10:10:55.00 DEBUG SLOT x LOADER CODE:33  
Loader service completed for ARP.EXE 0xxxxxxxxx
```

In addition to ARP, IP routers support the following address resolution schemes:

- Proxy ARP
- Inverse ARP
- HP Probe
- DDN and PDN

To select an address resolution protocol for an IP interface, see the Address Resolution parameter on page 2-44.

The following sections briefly describe the address resolution schemes that can be configured on an IP interface.

Proxy ARP

Proxy ARP allows a router to answer a local ARP request for a remote destination. For example, in Figure 2-4, Hosts B and C are located on the same network but on separate subnetworks. Hosts B and C do not understand subnetting. The router connecting the two physical networks knows which host resides on which network. The address mask is 255.255.255.000. In this example, one subnet is a remote network with respect to the other subnet.

Host B wants to talk to Host C, so Host B broadcasts an ARP request, which asks IP address 140.250.250.2 to respond with its physical address. The router captures Host B's ARP request and responds with its hardware address 00 00 A2 00 00 01 and Host C's IP address 140.250.250.2. Host B maps Host C's IP address 140.250.250.2 to the router's hardware address 00 00 A2 00 00 01.

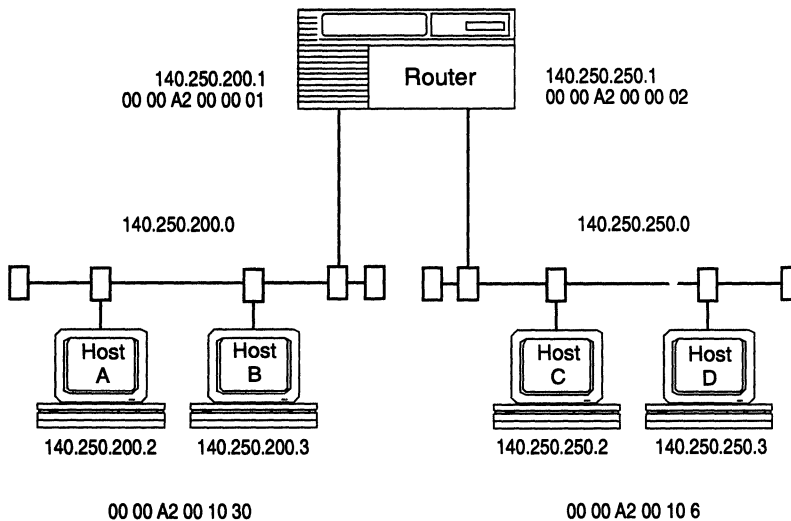


Figure 2-4. Proxy ARP Example

Inverse ARP

Inverse ARP enables the address resolution for Frame Relay interfaces. It is used to discover the IP address of the station at the remote end of the virtual circuit.

HP Probe

HP[®] Probe, a proprietary Hewlett-Packard protocol, is an address resolution mechanism that functions much like ARP to determine a network host's physical address when all it knows is the network host's IP address, by binding a 32-bit IP address to a 48-bit MAC address. We support the following HP Probe messages:

- Unsolicited Reply (incoming and outgoing)
- Name Request (incoming)
- Name Reply (outgoing)
- Virtual Address Reply (incoming and outgoing)
- Gateway Request (incoming)
- Gateway Reply (outgoing)

We support the concurrent operation of HP Probe and ARP.

X.25 DDN and X.25 PDN Address Resolution

For network interfaces that support the X.25 DDN service, we provide a DDN X.25 address resolution algorithm.

For network interfaces that support the X.25 PDN service, we provide an RFC 877-compliant address resolution mechanism.

Enabling Source Routing over Token Ring Networks

The IP router can route over token ring (TR) networks that contain one or more source routing bridges.

In a source routing network, every end station that sends out a frame supplies the frame with the necessary route descriptors so that it can be source routed across the network. Thus, in order for IP routers to route packets across a source-routing network, *they must act like end stations*; supplying route descriptors for each packet before they send it out onto the network.

With end-node support enabled, whenever an IP router receives a packet and determines that the packet's next hop is located across a source-routing network, the router does the following:

- Adds the necessary Routing Information Field (RIF) information to the packet's MAC header
- Sends the packet out onto the network where it is source routed toward the next hop

Upon receiving the packet from the token ring network, the peer router strips off the RIF field and continues to route the packet toward the destination network address (Figure 2-5).

You configure source-route end-node support on a per-circuit basis. See the TR Endstation parameter on page 2-48 for instructions and details.

Enabling Source Routing over Token Ring Networks






Figure 2-5. IP Routers Source Routing across a Token Ring Network

Configuring the Trivial File Transfer Protocol

The Trivial File Transfer Protocol (TFTP) is a TCP/IP standard protocol for transferring files with minimum capability and minimal overhead. TFTP is implemented on top of the unreliable, connectionless datagram delivery service and is used to move files between network devices.

TFTP was designed to be small and easy to implement. Because it is small, it is more restrictive, lacking most of the features of the File Transfer Protocol (FTP). TFTP provides inexpensive, unsophisticated file-transfer service only. It cannot list directories and provides no authentication.

TFTP runs on top of the User Datagram Protocol (UDP) and uses time-out and retransmission to ensure that data arrives. Each file transfer begins with a request to read or write to a file; this request also serves to ask for a connection. If the server grants the request, the connection is opened and the file is sent in fixed-length blocks (data packets) of 512 bytes. Each data packet contains one block of data and must be acknowledged by an acknowledgment packet before the next packet is sent. A data packet of less than 512 bytes terminates the transfer.

If a packet gets lost in the network, the intended recipient will time out and may retransmit its last packet (which can be data or an acknowledgment), causing the sender of the lost packet to retransmit the packet. Because the lock-step acknowledgment guarantees that all older packets have been received, the sender keeps one packet only on hand for transmission.

Both devices involved in a TFTP transfer are senders and receivers. One device sends data and receives acknowledgments; the other device sends acknowledgments and receives data.

The IP router includes a client and server implementation of the Trivial File Transfer Protocol, enabling the router to transmit and receive files across an Internet.

To specify the operating characteristics of TFTP on a router, see “Editing TFTP Parameters” on page 2-71.

Defining a Circuitless IP Interface

The IP router allows you to configure one circuitless IP interface. A circuitless IP interface specifies an IP address for the router without mapping the address to a specific circuit. If one or more of the router's IP interfaces become disabled, this feature ensures that the router is always reachable using the circuitless IP interface address, as long as a viable path to the router exists.

IP traffic is delivered to and transmitted from the circuitless interface in the same way as any other IP interface. In addition, the circuitless IP interface can receive packets from any application, including SNMP, BGP-3, TFTP, and Telnet.

When you configure a circuitless IP interface, note the following:

- You can configure one circuitless IP interface per router. Additional circuitless IP interfaces will not initialize.
- You must assign a unique IP address and subnetwork number to the circuitless IP interface.
- You *cannot* configure a circuitless IP interface in nonforwarding mode.

See "Configuring a Circuitless IP Interface" on page 2-57 for instructions.

Configuring the Revised IP Security Option

IP routers support the Department of Defense (DoD) Revised IP Security Option (RIPSO), as defined in RFC 1108) on a per-interface basis. While RIPSO RFC 1108 specifies both "basic" and "extended" security options, our implementation supports only the basic option.

RIPSO is a feature that allows end systems and intermediate systems (routers) to add labels to or process security labels in IP datagrams that they transmit or receive on an IP network. The labels specify security classifications (for example, Top Secret, Secret, Confidential,

and Unclassified, in descending order), which can be used to limit the devices that can access these labeled IP datagrams.

As a labeled IP datagram traverses an IP network, only those systems that have the proper clearance (that is, whose security classification range covers the classification specified by the datagram) should accept and forward the datagram.

Any system whose security classification range does not cover the classification specified by the security label should drop the datagram.

Note: RIPS0 does not include any method of preventing a system that does not support RIPS0 from simply accepting and forwarding labeled datagrams. Thus, in order for RIPS0 to be effective, *all* systems in a network must support RIPS0 and process IP datagrams as described.

For instructions on enabling RIPS0 support on an IP interface, see the Enable Security parameter on page 2-56. For complete information on RIPS0 parameters, see “Configuring RIPS0 Support” on page 2-74.

Security Label Format

A RIPS0 security label is 3 or more bytes long and specifies the security classification level and protection authority values for the datagram (Figure 2-6).

| Type | Length | Security Classification | Protection Authority | IP Datagram... |
|---------|---------|-------------------------|----------------------|----------------|
| 1 octet | 1 octet | 1 octet | 1 octet or more | |

Figure 2-6. RIPS0 Security Label

The format of the security label is as follows:

- ❑ Octet 1 contains a type value of $82_{(16)}$ identifying the basic security option format.
- ❑ Octet 2 specifies the length of the option (3 or more octets, depending on the presence or absence of authority flags).
- ❑ Octet 3 specifies the security classification levels for the datagrams. Valid security classification levels include
 - $3D_{(16)}$ Top Secret
 - $5A_{(16)}$ Secret
 - $96_{(16)}$ Confidential
 - $AB_{(16)}$ Unclassified
- ❑ Octet 4 and beyond identify the protection authorities under whose rules the datagram is classified at the specified level. (If no authorities have been identified, then this field is not used.)

The first 7 bits (0 through 6) are flags. Each flag represents a protection authority. The flags defined for Octet 4 are as follows:

| | | |
|-------|-----------------------|---|
| Bit 0 | GENSER | General Services (as per DoD 5200.28) |
| Bit 1 | SIOP-ESI | DoD (Organization of the Joint Chiefs of Staff) |
| Bit 2 | SCI | Central Intelligence Agency |
| Bit 3 | NSA | National Security Agency |
| Bit 4 | DOE | Department of Energy |
| Bit 5 | Reserved | |
| Bit 6 | Reserved | |
| Bit 7 | Termination indicator | |

Note: Bit 7 acts as a “more” bit, indicating that another octet (containing additional authority flags) follows.

How RIPS0 Works on the Router

When you configure RIPS0 on an IP interface, you specify the following conditions:

- ❑ A range of acceptable security levels for IP datagrams the interface receives and transmits
- ❑ A set of required and allowed authority values for IP datagrams the interface receives and transmits
- ❑ Whether inbound datagrams received on this interface require security labels
- ❑ Whether outbound datagrams transmitted on this interface (either forwarded or originated by the router) require security labels
- ❑ Whether datagrams received or transmitted on this interface should have their labels stripped

You also specify whether the router creates the following types of labels:

- ❑ An implicit label, which the router uses to label unlabeled inbound datagrams, when required
- ❑ A default label, which the router uses to label unlabeled outbound datagrams, when required
- ❑ An error label, which the router uses to label ICMP error messages associated with processing security options

The following sections describe how the router uses this information to handle labeled IP traffic.

Inbound IP Datagrams

When the router receives an IP datagram on a RIPS0 interface, it compares the security classification and authority values specified in the security label with those configured on the inbound interface.

If the interface does *not* require a security label for inbound IP datagrams, then the router accepts both unlabeled IP datagrams and

datagrams that meet the classification and authority rules described in the next paragraph.

If the interface *does* require a security label, then for the router to accept the datagram, the following RISPO conditions must be met:

- ❑ The datagram must be labeled.
- ❑ The security classification value in the datagram's label must be within the security level range configured for the interface.
- ❑ The authority flags in the datagram's label must include all of the flags required for the interface, and cannot contain any flags not allowed for the interface.

The router drops any datagrams not meeting these requirements and generates an ICMP error message.

On a *non-RISPO* interface, the router only accepts unlabeled IP datagrams and IP datagrams that are labeled as Unclassified with no authority flags set.

Forwarded IP Datagrams

When the router receives an IP datagram that needs forwarding on a RISPO interface, the router compares the security classifications and authority values specified in the security label with those configured on the outbound interface. So, before forwarding the datagram, the router

- ❑ Checks that all RISPO conditions are met (see above)
- ❑ Applies any out-bound specific configuration parameters

The router drops any datagrams not meeting these requirements and generates an ICMP error message.

Originated IP Datagrams

When the router originates a datagram, and the following conditions are true:

- ❑ The datagram needs forwarding out a RISPO interface

- The RIPS0 interface requires outbound labels for originated datagrams

the router labels the datagram with the default security label before transmitting it.

Unlabeled IP Datagrams

If the router receives an unlabeled IP datagram from an interface on which RIPS0 is *not* enabled (or on which labels are not required for inbound datagrams), and the IP datagram needs forwarding to an interface on which RIPS0 *is* enabled and labels are required for outbound datagrams, then the router labels the datagram using either an implicit label or default label as follows:

- If the inbound interface has an implicit label configured, then the router uses it to label the datagram.
- If the inbound interface does not have an implicit label configured, then the router labels the datagram with the default label configured for the outbound interface.

If the interface does not have an implicit or default label configured, then the datagram is simply dropped.

RIPS0 Example

The router in Figure 2-7 has RIPS0 configured on all three IP interfaces. The security ranges specified for each interface vary, as shown. (For simplicity, this example assumes that none of the interfaces require any authority flags on inbound and outbound traffic, but any flags that are present are acceptable.)

When host 1.1.0.1 broadcasts an all subnets broadcast IP datagram with the security level classification set to Secret, the router compares the datagram's classification with the range configured on inbound interface 1.1.0.2. Because Secret is within the range configured on the interface, the router accepts the datagram. In order to forward the datagram, the router does the following:

- ❑ Compares the datagram's security level, Secret, to the security level ranges configured on interface 1.2.0.2 and 1.3.0.2
- ❑ Forwards the datagram on interface 1.2.0.2, because Secret is within the security range configured on the interface
- ❑ Does *not* forward the datagram on interface 1.3.0.2, because Secret is outside of the security range configured on the interface

| Interface | Min. Security Classification | Max. Security Classification |
|-----------|------------------------------|------------------------------|
| 1.1.01 | Unclassified | Top secret |
| 1.2.02 | Secret | Top secret |
| 1.3.0.2 | Top secret | Top secret |

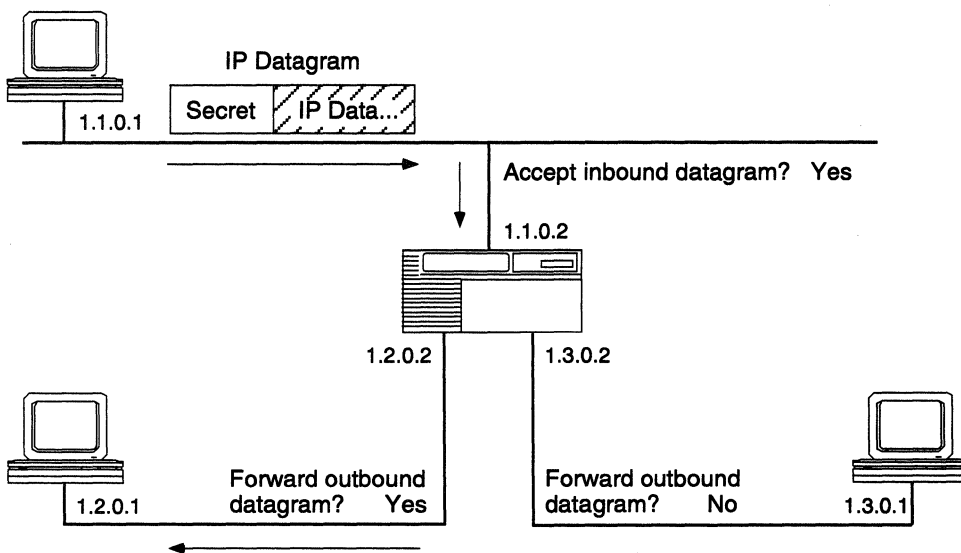


Figure 2-7. RIPS0 Network

Defining a Static Route

A static route is a manually configured route that specifies the transmission path a datagram must follow, based on the datagram's destination address. A static route specifies a transmission path to another network. You configure a static route if you want to restrict the paths that datagrams follow to paths you specifically configure.

Static routes remain in IP routing tables until you remove them. Note, however, that if the interface that was used to reach the next hop in the static route becomes disabled, the static route disappears from the IP routing table.

For instructions, see “Configuring Static Routes” on page 2-58.

Defining a Black Hole for a Supernet

A router that advertises an aggregate route by using a supernet address to represent multiple explicit routes must be able to discard packets that match the supernet address but that do not match any of the explicit routes.

For example, consider a router that advertises an aggregate route using the supernet address 192.32.0.0/255.255.248. The supernet address represents eight specific networks: 192.32.0.0 through 192.32.7.0. Once the aggregate route has been propagated, the router receives network traffic for each of these specific destinations.

At some point, the router loses connectivity to network 192.32.3.0, one of the networks in the supernet. The router continues to forward traffic that matches destinations 0.0 through 2.0 and 4.0 through 7.0. However, the router can no longer find a complete match in the routing table for the disconnected network, 3.0. The router must drop all traffic destined for 192.32.3.0.

To force the router to drop the packet for an unmatched destination, you configure a special type of static route for a supernet called a black hole. Specifically, you enter the supernet address/mask pair as the

Destination IP Address and Address Mask parameter values on the IP Static Routes window. To create the black hole, you enter the black hole encoding (255.255.255.255) as the Next Hop Addr and Next Hop Mask parameter values.

For details, see “Configuring Static Routes” on page 2-58.

Configuring Router Discovery

Before a host can send IP datagrams beyond its directly attached subnet, the host must discover the address of at least one operational router on that subnet. Router Discovery is an extension of the Internet Control Message Protocol (ICMP) that enables hosts attached to multicast or broadcast networks to discover the IP addresses of their neighboring routers.

Routers configured with Router Discovery periodically multicast a router advertisement from each of their multicast interfaces, announcing the IP address or addresses of that interface. Hosts discover the addresses of their neighboring routers by listening for these advertisements.

For information about Router Discovery parameters and settings, see “Configuring Router Discovery” on page 2-86.

Connecting the Router to a Blacker Front End

The Blacker Front End (BFE) is a classified encryption device used by hosts that want to communicate across unsecured, wide area networks. BFE devices are typically found in government networks (for example, DSNET), which handle sensitive data requiring a greater degree of security.

Blacker Front End support allows the router to connect to BFE devices. The BFE device, in turn, provides the router with encryption services while acting as the Data Communications Equipment (DCE) end of the connection between the router and the X.25 network (Figure 2-8).

Hosts using attached BFE devices can communicate with each other over an unsecured packet-switched network using data paths secured by the encryption services of the BFEs. These hosts are part of a *Red Virtual Network*. The packet-switched network that carries both the data secured by BFEs and any other unsecured data is known as the *Black network*.




Figure 2-8. Blacker Front-End Network Configuration

BFE devices receive authorization and address translation services from an Access Control Center residing on the Black network. The ACC makes access control decisions that determine which hosts are allowed to communicate with each other. A Key Distribution Center (KDC) residing on the Black network provides encryption keys and key management services. A BFE device uses these encryption keys for encrypting traffic between itself and other BFE devices.

The router-to-BFE interface is a modified version of the interface presented in the 1983 DDN X.25 Host Interface Specification. It supports data rates between 1200 b/s and 64 KB/s. In order to support BFE services, the interface must be configured to support IP with the Revised IP Basic Security Option (RIPSO) enabled. All IP datagrams transmitted on the interface must contain a RIPSO security label. The first option in each IP datagram header must be the Basic Security option.

For information about BFE parameters and settings, see “Configuring Blacker Front-End Support” on page 2-90.

BFE Addressing

You can enable BFE support on individual IP interfaces. When you enable BFE support, the router uses the BFE address resolution algorithm to map IP addresses to their corresponding X.121 addresses.

BFE IP-to-X.121 address translation differs from standard DDN address translation. Each physical router-to-BFE connection is identified by a BFE X.121 network address and a BFE IP address. The format of a BFE X.121 address is

zzzzzpddbbb

where:

| | |
|--------------|--------------------------------------|
| <i>zzzzz</i> | is zero |
| <i>p</i> | is the BCD encoding of the port ID |
| <i>ddd</i> | is the BCD encoding of the domain ID |
| <i>bbb</i> | is the BCD encoding of the BFE ID |

All BFE hosts are members of Class A IP networks. The format of a BFE IP address is as follows:

nnnnnnnn.Zpppddd.dddddbb.bbbbbbb

where:

| | |
|------------------|-----------------------------------|
| <i>nnnnnnnn</i> | identifies the network ID in bits |
| <i>Z</i> | is zero |
| <i>ppp</i> | is the port ID in bits |
| <i>ddddddddd</i> | is the domain ID in bits |
| <i>bbbbbbbbb</i> | is the BFE ID in bits |

BFE supports only physical addressing. It does not support either logical addresses or subaddresses.

Editing IP Parameters

This section describes how to edit, or customize, IP parameters.

Note: The instructions in this section assume that you have already configured at least one IP interface. If you have *not* yet configured an IP interface, or want to add additional IP interfaces, see *Configuring Wellfleet Routers* for instructions.

You access all IP parameters from the Configuration Manager window shown in Figure 2-9 (refer to *Configuring Wellfleet Routers* for instructions on accessing this window). For each IP parameter, this section describes the default setting, all valid setting options, the parameter function, instructions for setting the parameter, and the Management Information Base (MIB) object ID

The Technician Interface allows you to modify parameters by issuing **set** and **commit** commands with the MIB object ID. This process is equivalent to modifying parameters using Site Manager. For more information, refer to *Using Technician Interface Software*.

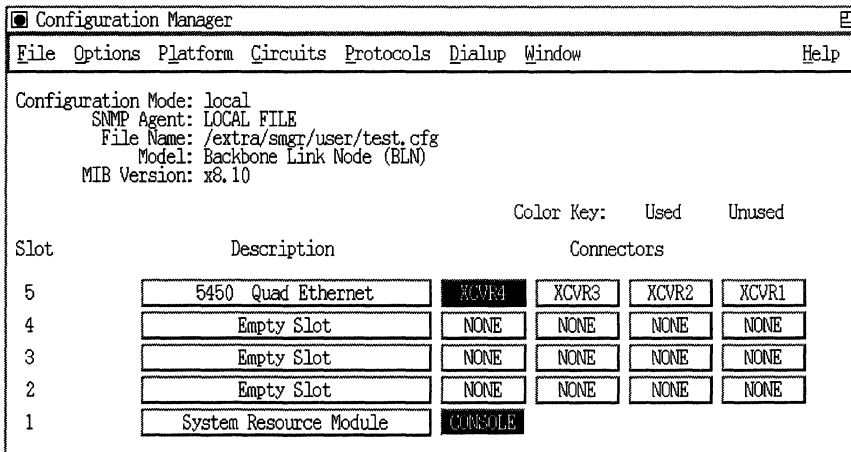


Figure 2-9. Wellfleet Configuration Manager Window

Editing IP Global Parameters

To edit IP global parameters, begin at the Configuration Manager window shown in Figure 2-9 and proceed as follows:

1. Select Protocols→IP→Global.

The Edit IP Global Parameters window appears (Figure 2-10).

2. Edit those parameters you wish to change.

Note: When you edit parameters in dynamic mode, the IP router restarts, causing Site Manager to lose its router connection temporarily, and to display a warning message. To verify that the change took effect, redisplay the IP Global Parameters window and inspect the setting.

All IP global parameters are described in the following section.

3. Click on OK to exit the window and save your changes when you are finished.

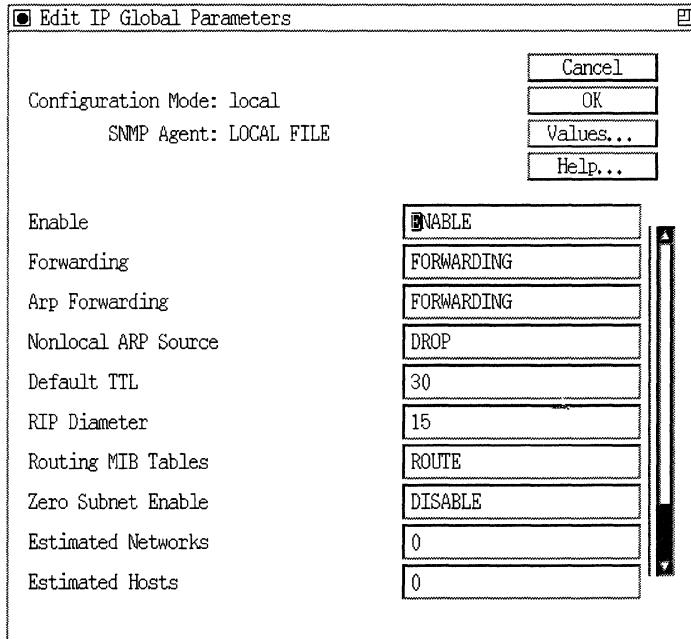


Figure 2-10. Edit IP Global Parameters Window

IP Global Parameter Descriptions

Use the following descriptions to set parameters on the IP Global Parameters window.

Parameter: **Enable**

Default: This parameter defaults to Enable when you add IP support to a circuit.

Options: Enable | Disable

Function: Specifies the state of the IP router software.

Instructions: Select Enable if you have previously disabled the IP router software and now wish to re-enable it. Select Disable to disable the IP router software.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.2



Warning In dynamic mode, when you set the global Enable parameter to Disable, you immediately prohibit all Site Manager communication with the router.

| | |
|-------------------|--|
| Parameter: | Forwarding |
| Default: | Forwarding |
| Options: | Forwarding Not Forwarding |
| Function: | Specifies whether the IP router forwards IP traffic that is not explicitly addressed to it. |
| Instructions: | <p>Select Forwarding if you want the IP router to route (forward) IP traffic. Forwarding configures the IP router to process all broadcast packets and all IP packets explicitly addressed to it, and to <i>route</i> all other IP packets. Select Not Forwarding if you want to provide IP management access (by means of TFTP and SNMP) to all active IP interfaces, but want to prohibit the IP router from forwarding IP traffic. You must specify an identical IP address and mask combination for each active IP interface that will provide management access. Not Forwarding configures the IP router to act as an IP host; it does not forward IP traffic, but still processes packets explicitly addressed to it. In Not Forwarding mode, only static routes and adjacent-host routes are allowed. No routing protocols are initiated.</p> <p>Because the IP router does not forward IP traffic in Not Forwarding mode, you must configure the router to <i>bridge</i> IP traffic not explicitly addressed to it. You must configure the Bridge for each circuit that conveys IP datagrams. The Bridge will then forward all IP datagrams that are not explicitly addressed to the router.</p> |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.1.4 |

Parameter: ARP Forwarding

Default: Forwarding

Options: Forwarding | Not Forwarding

Function: Specifies how ARP should act in relation to IP's forwarding state. Note that Forwarding means IP is in forwarding mode. If this parameter is set to Forwarding, then ARP packets are either consumed (if destined for the router) or dropped. If this parameter is set to Not Forwarding, ARP packets are consumed, if destined for the router, or bridged onto remaining ARP interfaces.

Instructions: Always set this parameter the same way you set IP Forwarding.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.1.1.3

Parameter: Nonlocal ARP Source

Default: Drop

Options: Drop | Drop and Log

Function: Determines what happens when IP encounters an invalid ARP source address. If the parameter is set to Drop and Log, IP logs an invalid ARP source address when processing an ARP request. If this parameter is set to Drop, IP does not log the invalid ARP source address. In either case, IP drops the invalid ARP request.

Instructions: If you want to log the invalid ARP source address, set the parameter to Drop and Log. Otherwise, set the parameter to Drop.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.1.1.4

Parameter: **Nonlocal ARP Destination**
Default: Drop
Options: Drop | Accept
Function: Determines whether IP drops ARP requests in which the source and destination addresses are located in different networks or subnetworks.
Instructions: To process ARP requests with source and destination addresses from different networks, set the parameter to Accept.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.1.1.5

Parameter: **Default TTL**
Default: 30
Range: 1 to 255 hops
Function: Specifies the starting value of the Time to Live (TTL) counter for each packet the router originates and transmits (called a source packet). When the router transmits a source packet, the TTL counter starts to decrement. Each router, or hop, that the packet traverses decrements the TTL counter by one. When the counter reaches zero, the router discards the packet unless it is destined for a locally attached network. The TTL counter prevents packets from looping endlessly through the network.
Instructions: Enter the maximum number of hops a source packet can traverse.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.5

Parameter: RIP Diameter

Default: 15

Range: 1 to 127

Function: Specifies the value, or hop count, the Routing Information Protocol (RIP) uses to denote infinity. In order for RIP to operate properly, every router within the network must be configured with an identical RIP Diameter value. If RIP is not enabled, the RIP Diameter parameter specifies the maximum number of hops within the autonomous system; if RIP is not enabled, the IP router still must understand network width.

Instructions: You must set the RIP Diameter parameter so that none of the interface cost, static cost, or route filter cost parameters exceed the RIP Diameter parameter. We recommend that you accept the default RIP Diameter value of 15.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.6

Parameter: Route Cache Interval

Default: 60 (seconds)

Options: Any value

Function: Specifies the interval at which aging routing entries are flushed from the forwarding tables.

Instructions: Select an interval to determine how often aging routing entries are to be flushed.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.7

Parameter: Routing MIB Table(s)

Note: This parameter is not valid in software versions later than Version 7.70.

Default: Route

Options: None | Route | Forward | Both

Function: Specifies which MIB routing tables IP maintains. IP uses these MIB routing tables only to store statistics; do not confuse them with the routing tables maintained to route packets. Maintaining both the Routing and Forwarding tables uses more memory than maintaining either. In the absence of variable-length subnet masks, these tables are identical. The routing table does not support variable-length subnet masks. This table is MIB-II compliant. The forwarding table does support variable-length subnet masks. It is not MIB-II compliant.

Instructions: Depending on your network requirements, select

- None to disable maintenance of both tables.
- Route if you are *not* using variable-length subnet masks.
- Forward if you are using variable-length subnet masks, and want to maintain statistics on them.
- Both if you are using other network management applications to manage the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.8.

Parameter: Zero Subnet Enable

Default: Disable

Options: Enable | Disable

Function: Specifies whether an interface address whose subnet portion is all zeros should be declared legal or not. If you set this parameter to Enable, then you can configure IP interfaces with a subnet ID of zero. Setting this parameter to Disable prevents you from doing so.

Instructions: Accept the default, Disable, if you do not have any interfaces that have a zero subnet ID. Otherwise, reset this parameter to Enable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.10

Note: The use of all-zero subnet addresses is discouraged for the following reason. If an all-zero subnet address and an all-zero broadcast address are both valid, the router cannot distinguish an all subnets broadcast from a directed broadcast for the zero subnet.

Parameter: Estimated Networks**Default:** 0**Range:** 0 to 2147483647**Function:** Allows the IP software to preallocate system resources based on the anticipated size of the routing table. Preallocation of memory increases the speed with which IP software can learn routes because it removes the overhead caused by dynamic memory allocation. Preallocation also makes better use of memory and reduces the amount of memory required.**Instructions:** Set to the number of networks (including unique subnets) that you expect. Avoid using a number that is excessively large. This will cause a wasteful overallocation of memory.

If you use the default value (0), IP software preallocates memory for 500 routing table entries.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.11

Parameter: Estimated Hosts

Default: 0

Range: 0 to 2147483647

Function: Allows the IP software to pre-allocate system resources based on the anticipated size of the routing table. Pre-allocation of memory increases the speed with which IP software can learn routes because it removes the overhead caused by dynamic memory allocation.

Instructions: Set to the number of hosts that you expect. Avoid using a number that is excessively large. This will cause a wasteful over-allocation of memory.

If you use the default value (0), IP software preallocates memory for 500 routing table entries.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.13

Parameter: Enable Default Route for Subnets

Default: Disable

Options: Enable | Disable

Function: Specifies whether the IP router uses a default route for unknown networks and subnets. When this parameter is set to Enable, the IP router uses a default route. When this parameter is set to Disable, the IP router does not use a default route.

Instructions: Accept the default, Disable, if you do not want the IP router to use a default route for unknown networks and subnets. Otherwise, reset this parameter to Enable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.1.14

| | |
|-------------------|---|
| Parameter: | Maximum Policy Rules |
| Default: | 32 |
| Options: | Any integer |
| Function: | Specifies the maximum number of policy rules that can be configured per policy type (Accept or Announce) per protocol. |
| Instructions: | To configure more than 32 Accept or Announce policy rules for a protocol, you must set this parameter to a larger value. IP will round the value up to the next multiple of 32. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.1.15 |

Editing IP Interface Parameters

To edit an IP interface, begin at the Configuration Manager window shown in Figure 2-9 and proceed as follows:

1. Select Protocols→IP→Interfaces.

The IP Interfaces window appears (Figure 2-11). It lists all IP interfaces configured on the router.

2. Click on the interface you want to edit.
3. Edit those parameters you wish to change.

If you are reconfiguring an interface in dynamic mode, see the warning below.

All IP interface parameters are described in the following section.

4. Click on Apply to implement your changes.
5. Click on Done to exit the window.

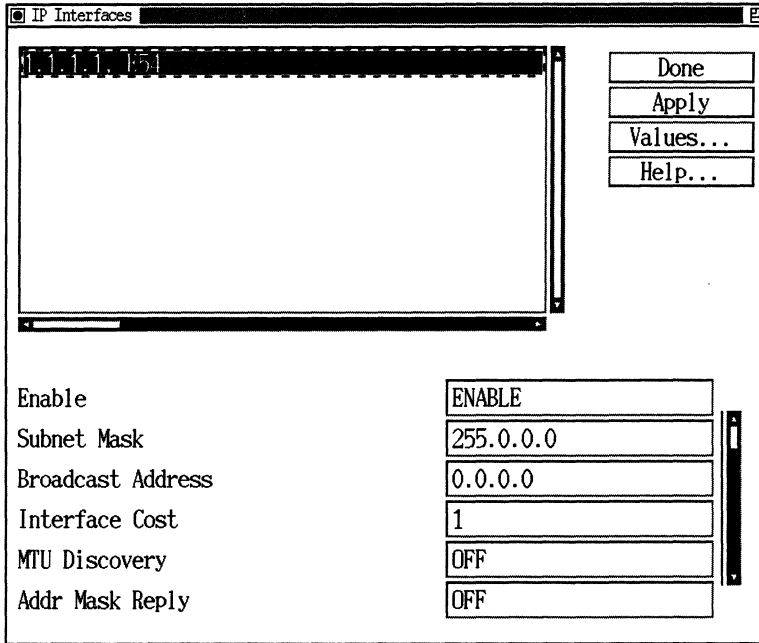


Figure 2-11. IP Interfaces Window



Warning When you reconfigure an interface in dynamic mode, IP restarts on that interface. Thus, if the interface you reconfigure is the interface that supports Site Manager's SNMP connection to the router, restarting IP on that interface will cause Site Manager to temporarily lose its router connection and to display a warning message. To verify that the change took effect, redisplay the IP Global Parameters window and inspect the setting.

Note: If you are configuring IP over an SMDS circuit, be sure to enter the correct addresses in the MAC Address, SMDS Group Address, and SMDS Arp Req Address parameter boxes displayed in this screen. These addresses are the same as those you entered in the Individual Address, Group Address, and ARP Address parameters of the SMDS Configuration window when you configured SMDS.

IP Interface Parameter Descriptions

Use the following descriptions to set parameters on the IP Interfaces window.

Parameter: **Enable**

Default: Enable

Options: Enable | Disable

Function: Enables or disables IP routing on this interface.

Instructions: Set to Disable to disable IP routing over this circuit.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.2

Parameter: **Subnet Mask**

Default: You specified the subnet mask when you added IP to the circuit.

Options: Depends on the class of the network to which the interface connects

Function: Specifies the network and subnetwork portion of the 32-bit IP address.

Instructions: Enter the subnet mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.6

Parameter: Broadcast Address

Default: You specified the Broadcast Address parameter when you added IP to the circuit.

Options: 0.0.0.0 or any IP address

Function: Specifies the broadcast address that the IP router uses to broadcast packets. Accepting 0.0.0.0 for the broadcast address specifies that the IP router will use a broadcast address with a host portion of all 1s. Accepting 0.0.0.0 does not configure the router to use the address 0.0.0.0 to broadcast packets. For example, if you have set the IP address to 123.1.1.1 and the subnet mask to 255.255.255.0, accepting the default value 0.0.0.0 configures the IP router to use the address 123.1.1.255 to broadcast packets. For the explicit broadcast address of all 1s, enter 255.255.255.255 for this parameter.

Instructions: Accept the default, 0.0.0.0, unless the calculated broadcast address (host portion) of all 1s is not adequate. If this is the case, then enter the appropriate IP broadcast address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.9

| | |
|-------------------|--|
| Parameter: | Interface Cost |
| Default: | 1 |
| Options: | 1 to the value of RIP Diameter (maximum 127) |
| Function: | Sets the cost of this interface. The interface cost is added to routes learned on this interface through RIP and is specified in subsequent RIP packets transmitted out other interfaces. |
| Instructions: | Enter the interface cost value (standard RIP implementation assigns a cost of 1); however, keep in mind that increasing this value causes the upper bound set by RIP Network Diameter to be attained more rapidly. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.4.1.8 |

Parameter: MTU Discovery

Default: Off

Options: On | Off

Function: Specifies whether the Reply MTU option (option 11 in RFC 1063) is enabled on this interface. When the option is enabled, this interface responds to Probe MTUs (option 12 in RFC 1063). A Probe MTU requests the minimum MTU (Maximum Transmission Unit) of all networks an IP datagram must traverse from source to destination. By enabling this interface to respond to Probe MTUs, you eliminate transit fragmentation and destination reassembly for datagrams destined for this interface, and therefore decrease network load.

Instructions: Select On to enable the Reply MTU option on this interface; select Off to disable the option on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.10

Parameter: Addr Mask Reply

Default: Off

Options: On | Off

Function: Specifies whether this interface generates ICMP (Internet Control Message Protocol) address-mask-reply messages in response to valid address-mask-request messages. The interface generates ICMP address-mask-reply messages in compliance with the relevant sections of RFCs 950 and 1009.

Instructions: Select On to enable ICMP address-mask-reply message generation on this interface. Select Off to disable ICMP address-mask-reply message generation on this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.11

Parameter: All Subnet Bcast

Default: Off

Options: On | Off

Function: Specifies whether or not the IP router floods ASB datagrams it receives out this interface. An ASB datagram has a destination address equal to the broadcast address for an entire network (all subnets). For example, if a network interface serves the subnet 128.10.2.1, with a subnet mask of 255.255.255.0, the IP router considers any datagram with a destination address of 128.10.255.255 or 128.10.0.0 to be an ASB datagram.

Instructions: Specify On if you want the IP router to flood ASBs out this interface; specify Off to restrict the router from flooding ASBs out this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.12

Parameter: Address Resolution

Default: ARP

Options: ARP | X.25_DDN | X.25_PDN | INARP |
ARPINARP | NONE | X.25_BFEDDN | PROBE |
ARPPROBE

Function: Indicates the address resolution for this interface. The default option ARP enables ARP on this interface. The option INARP (Inverse ARP) enables the address resolution for Frame Relay interfaces. It is used to discover the IP address of the station at the remote end of the virtual circuit.

Instructions: Depending on your network requirements, select
—INARP only when all Frame Relay stations support Inverse ARP.

—ARPINARP for your Frame Relay interfaces. ARPINARP enables both ARP and Inverse ARP.

—X.25_DDN for your X.25 DDN interfaces.

—X.25_PDN for your X.25 PDN interfaces.

—PROBE to enable HP Probe on the interface.

—ARPPROBE to enable both ARP and HP Probe

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.13

| | |
|-------------------|---|
| Parameter: | Proxy |
| Default: | Off |
| Options: | On Off |
| Function: | Specifies whether this interface uses Proxy ARP to respond to ARPs for a remote network. |
| Instructions: | Select On to enable Proxy ARP on this interface. In order to enable Proxy ARP, you must have set the ARP parameter to Enable for this interface. When you enable Proxy ARP, the IP router assumes responsibility for IP datagrams destined for the remote network. Select Off to disable Proxy ARP on this interface. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.4.1.14 |

Parameter: Host Cache

Default: Off

Options: Off | 120 | 180 | 240 | 300 | 600 | 900 | 1200 (seconds)

Function: Specifies whether the IP router times out entries in the address-resolution cache for this interface, and specifies the timeout interval in seconds if the interface does time out entries. The address-resolution cache contains host physical addresses learned by means of ARP or Proxy ARP. A host entry is timed out (deleted) if the IP router sends no traffic destined for that host within the specified aging period.

Instructions: Select Off to disable timeout on this interface; the IP router does not time out address-resolution cache entries. Select one of the other options to enable timeout with a timeout interval equal to the value you select (for example, 120 seconds); the IP router removes address-resolution cache entries that have not been accessed within the specified number of seconds. Once an entry is removed, the IP router must use ARP to re-acquire the physical-level address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.15

| | |
|-------------------|--|
| Parameter: | Checksum |
| Default: | On |
| Options: | On Off |
| Function: | Specifies whether UDP checksum processing is enabled on this interface. |
| Instructions: | Select On to enable UDP checksum processing for the interface; all outgoing and incoming UDP datagrams are subject to checksumming. You should select On in virtually all instances. Select Off to disable UDP checksum processing and provide backward compatibility with UNIX BSD 4.1. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.4.1.16 |

Parameter: MAC Address

Default: None

Options: 0 | a user-specified MAC address | if the interface is on an SMDS circuit, the entire E.164 address—for example, E1 617 275 5000 FFFF

Function: Specifies a MAC (media access control) address for this IP interface. The IP router will use its IP address and this MAC address when transmitting and receiving packets on this interface.

Instructions: Enter 0 to configure the IP router to use its IP address and the circuit's MAC address when transmitting packets on this interface. Enter your own MAC address to configure the IP router to use its IP address and the specified MAC address when transmitting packets on this interface.

To configure this parameter for a multinet or multigroup configuration, refer to *Customizing SMDS Services*.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.17

Parameter: TR Endstation

Default: Off

Options: On | Off

Function: Specifies source routing over token ring selection.

Instructions: Use the On option to enable the parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.64

Parameter: Redirects**Default:** Enable**Options:** Enable | Disable**Function:** Indicates whether or not this interface sends out ICMP redirects.

ICMP redirects are messages sent by the router to alert a host that it should be using a different path to route data.

Instructions: Reset to Disable if you do not want this interface to send out redirects. For example, in a Frame Relay network, two stations on the same network may not be directly connected if the network is not fully meshed. Thus, in this case, you would set Redirects to Disable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.70

Parameter: **Enet Arp Encaps**

Default: ARP Ethernet

Options: ARP Ethernet | ARP SNAP | ARP Both |
 Probe LSAP | ARP Ethernet/Probe LSAP |
 ARP SNAP/Probe LSAP | ARP Both/Probe LSAP

Function: Defines the data-link encapsulation to use for
 ARP and HP Probe packets generated at this
 interface if the underlying medium is Ethernet.
 This parameter is ignored if the underlying
 medium is anything other than Ethernet.

Instructions: Depending on the selection you have made for the
 ARP Resolution parameter (ARP, Probe, or ARP/
 Probe), select the appropriate encapsulation
 option. If your address-resolution scheme is ARP
 only, select Ethernet encapsulation, SNAP
 encapsulation, or Ethernet/SNAP encapsulation.
 If your resolution scheme is HP Probe only, select
 LSAP encapsulation. If your resolution scheme is
 ARP/Probe, select Ethernet/LSAP encapsulation,
 SNAP/LSAP encapsulation, or Ethernet/SNAP/
 LSAP encapsulation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.71

Parameter: SMDS Group Address**Default:** None**Options:** A complete SMDS E.164 address specified by the SMDS subscription agreement that you have with your SMDS provider**Function:** Provides a MAC-layer multicast address for this IP interface in an SMDS network. This parameter is displayed only if this is an SMDS circuit.**Instructions:** Enter an entire E.164 address — for example, E1 617 555 1212 FFFF.

To configure this parameter for a multinet or multigroup configuration, refer to *Customizing SMDS Services*.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.65**Parameter: SMDS Arp Req Address****Default:** None**Options:** A complete SMDS E.164 address specified by the SMDS subscription agreement that you have with your SMDS provider**Function:** Provides an address resolution multicast address for this IP interface in an SMDS network. This parameter is only displayed if this is an SMDS circuit.**Instructions:** Enter an entire E.164 address — for example, E1 617 555 1212 FFFF.

To configure this parameter for a multinet or multigroup configuration, refer to *Customizing SMDS Services*.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.66

| | |
|-------------------|--|
| Parameter: | WAN Broadcast |
| Default: | 0 |
| Options: | Any decimal number |
| Function: | Provides a broadcast address for this IP interface in a Frame Relay network. If you enter a value for this parameter, the Frame Relay switch, rather than the router, will broadcast the message. This parameter is displayed only if this is a Frame Relay circuit. |
| Instructions: | Enter the broadcast address provided by the Frame Relay subscription agreement. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.4.1.67 |
| | |
| Parameter: | WAN Multicast #1 |
| Default: | 0 |
| Options: | Any decimal number |
| Function: | Provides a multicast address for this IP interface that will send messages to all OSPF routers in a Frame Relay network. If you enter a value for this parameter, the Frame Relay switch, rather than the router, will send the message to all OSPF routers. This parameter has meaning only if OSPF has been added to this interface. |
| Instructions: | Enter the multicast address for all OSPF routers as provided by the Frame Relay subscription agreement. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.4.1.68 |

Parameter: WAN Multicast #2**Default:** 0**Options:** Any decimal number**Function:** Provides a multicast address for this IP interface that will send messages to all OSPF designated routers in a Frame Relay network. If you enter a value for this parameter, the Frame Relay switch, rather than the router, will send the message to all OSPF designated routers. This parameter has meaning only if OSPF has been added to this interface.**Instructions:** Enter the multicast address for all OSPF designated routers as provided by the Frame Relay subscription agreement.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.4.1.69

Parameter: Slot Mask

Default: Slot-mask bit set to 1 (enabling circuitless IP interface support) for every router slot running IP

Options: For each slot in the router, Site Manager allows you to set the slot-mask bit to 1 (circuitless IP interface support enabled) or 0 (circuitless IP interface support disabled)

Function: Specifies whether circuitless IP interface support is enabled or disabled on each slot in the router.

Instructions: If you have configured a circuitless IP interface and do not wish it to run on certain slots, set the slot-mask bit to 0 on those slots. Be certain to keep the slot-mask-bit set to 1 on at least one slot running IP; otherwise, the circuitless IP interface will not initialize. Setting the slot-mask bit parameter to 1 on an empty slot, a slot containing a system resource module, or a slot with no IP support, does not affect the circuitless IP interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.75

Parameter: Max Forwarding Table Size**Default:** 128 entries**Range:** 64 to 1048 entries**Function:** Specifies the maximum number of entries allowed in the forwarding table at one time.**Instructions:** Specify a forwarding table size for each interface.

This parameter controls the number of destinations that are cached in the forwarding table on this receiving interface. When this interface receives an IP packet, the router looks up the destination in the forwarding table. Therefore, an interface that receives packets that are destined for a large number of different destinations may benefit from a larger forwarding table. The larger the number of entries, the more likely it is that the destination will already be in the forwarding table and the faster the route lookups will be for those destinations. Configuring a forwarding table size that is larger than necessary reduces the total amount of memory usable by other applications. Configuring a routing table too small can affect overall router performance. A check of the number of cache hits and misses will help determine the optimal size of the forwarding table. For debugging purposes, if you see the `wfIpInterfaceCacheMisses` statistic going up at an alarming rate, you should consider increasing the table size. However, an occasional cache miss does not warrant an increase in table size.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.104

Parameter: **Enable Security**

Default: Disable

Options: Enable | Disable

Function: Specifies whether Revised IP security option (RIPSO) is enabled for the interface.

Instructions: If you do not support RIPSO on your network, simply accept the default setting, Disable. If you are configuring RIPSO support, set this parameter to Enable. Then see “Configuring RIPSO Support” for instructions on setting the rest of the RIPSO parameters that you must configure.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.1.76

Note: Once you set the Enable Security parameter to Enable, you can access the rest of the RIPSO parameters. If you do not enable this parameter, Site Manager does not activate the RIPSO parameters.

Deleting IP from an Interface

To delete IP from an interface on which it is currently configured, begin at the Configuration Manager window and proceed as follows:

1. Click on the connector from which you want to delete IP services.
2. Click on Edit Circuit.
3. Select Protocols→Add or Delete.

The Select Protocols window appears. The IP button is highlighted to show that IP is enabled on the circuit.

4. Click on IP to deselect it.
5. Click on OK to exit the window.
6. Select File→Exit to exit the Circuit Definition window and return to the Configuration Manager window.

Configuring a Circuitless IP Interface

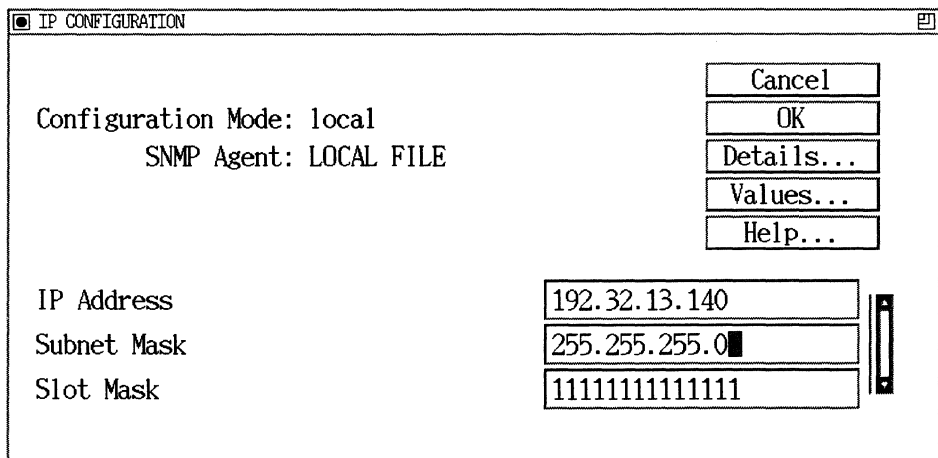
Note: The IP router supports one circuitless IP interface.

To configure a circuitless IP interface, begin at the Configuration Manager window shown in Figure 2-9 and proceed as follows:

1. Select Protocols→IP→Circuitless IP→Create to display the IP Configuration window (see Figure 2-12).
2. Specify the IP address and subnet mask for the interface (these parameters are required); optionally, specify the Slot Mask parameter.

The section “IP Interface Parameter Descriptions” on page 2-39 describes these parameters.

3. Click on OK to save the circuitless IP interface and return to the Configuration Manager window.



The screenshot shows a window titled "IP CONFIGURATION". Inside the window, the following text is displayed:

Configuration Mode: local
SNMP Agent: LOCAL FILE

On the right side of the window, there are five buttons stacked vertically: Cancel, OK, Details..., Values..., and Help....

Below the text, there are three input fields with labels to their left:

| | |
|-------------|----------------|
| IP Address | 192.32.13.140 |
| Subnet Mask | 255.255.255.0 |
| Slot Mask | 11111111111111 |

Figure 2-12. IP Configuration Window

Configuring Static Routes

To add, edit, or delete static routes, begin at the Configuration Manager window shown in Figure 2-9 and proceed as follows:

1. Select Protocols→IP→Static Routes.

The IP Static Routes window appears (Figure 2-13). It lists all static routes configured on the router. You add, edit, and delete static routes from this window.

2. Add, edit, or delete static routes as described in the following sections.

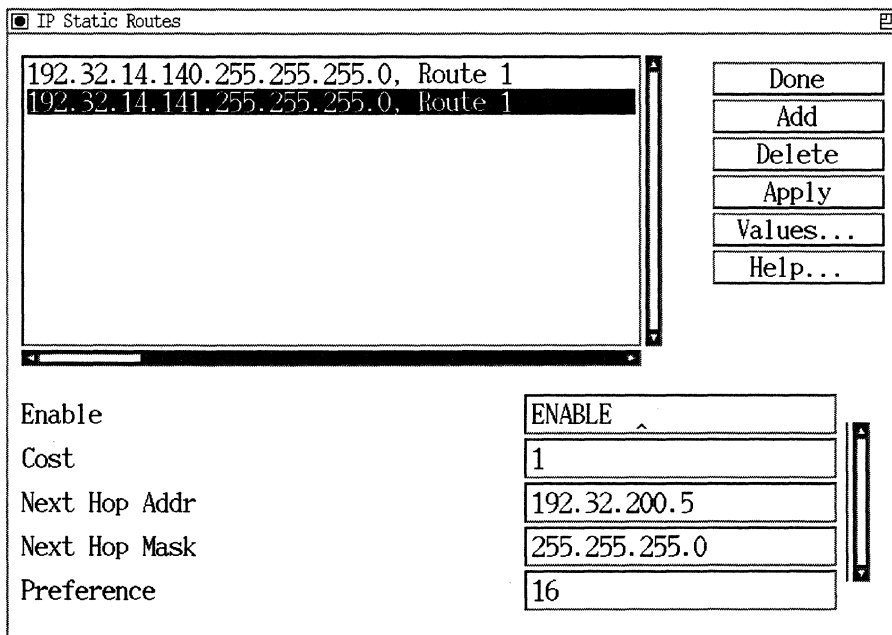


Figure 2-13. IP Static Routes Window

Adding a Static Route

To add a static route, begin at the IP Static Routes window and proceed as follows:

1. Click on Add.

The Add IP Static Route window (see Figure 2-14) appears.

2. Define the Destination IP Address and Address Mask parameters; then click on OK.

These two static route parameters are described following these instructions.

The IP Static Route window now lists the static route you configured. Note that the Enable, Cost, Next Hop Addr, Next Hop Mask, and Preference parameters display the default values for the static route. To edit these parameters, see “Editing a Static Route.”

Repeat Steps 1 and 2 to add additional static routes.

3. Click on OK to exit the IP Static Routes window.




Figure 2-14. Add IP Static Route Window

Parameter: Destination IP Address

Default: None

Options: Any valid IP network address

Function: Specifies the IP address of the network to which you want to configure the static route. Specifies a supernet for which you want to configure a black hole static route.

Instructions: Enter the destination IP address in dotted decimal notation. To configure a default route, enter 0.0.0.0. To configure a black hole static route, enter a supernet address. You can configure up to 12 static routes to the same destination.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.3

Parameter: Address Mask

Default: None

Options: Based on the network class of the IP address you specified at the Destination IP Address parameter.

Function: Specifies the subnet mask of the destination network. Specifies the supernet mask of the supernet for which you want to configure a black hole static route.

Instructions: Enter the subnet or supernet mask in dotted decimal notation. To configure a default route, enter 0.0.0.0. To configure a black hole static route, enter a supernet mask.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.1.4

Editing a Static Route

You can edit the Enable, Cost, Next Hop Addr, Next Hop Mask, and Preference parameters for a static route. To edit these parameters, begin at the IP Static Routes window shown in Figure 2-13 and proceed as follows:

1. Click on the static route you want to edit.
2. Edit those parameters you want to change.

The following section describes the static route parameters.

3. Click on Apply to implement your changes.
4. Click on Done to exit the window.

Note: You cannot reconfigure the Destination IP Address or Address Mask parameters for a static route. To change these parameters, you must delete the static route and add a new route with the proper information. See “Deleting a Static Route” for instructions.

Static Route Parameter Descriptions

Use the following descriptions to set parameters on the IP Static Routes window.

Parameter: **Enable**

Default: This parameter defaults to Enable when you configure the static route.

Options: Enable | Disable

Function: Specifies the state (active or inactive) of the static route record in the IP routing tables.

Instructions: Select Disable to make the static route record inactive in the IP routing table; the IP router will not consider this static route. Select Enable to make the static route record active again in the IP routing table.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.2

Parameter: **Cost**

Default: 1

Range: 1 to the value of the RIP Diameter parameter (maximum 126)

Function: Specifies the number of router hops a datagram can traverse before reaching the destination IP address. The IP router uses Cost when determining the best route for a datagram to follow.

Instructions: Enter the number of router hops.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.5.5

Parameter: Next Hop Addr**Default:** 0.0.0.0**Options:** Any valid IP address**Function:** Specifies the IP address of the next-hop router.
Defines a “black hole” route for a supernet.**Instructions:** Enter the IP address in dotted decimal notation.
To configure a black hole static route, enter
255.255.255.255.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.5.6**Parameter: Next Hop Mask****Default:** 0.0.0.0**Options:** Any valid subnet mask address**Function:** Specifies the subnet mask of the next hop router.
The parameter also defines a “black hole” route for
a supernet.**Instructions:** Enter the subnet mask in dotted decimal notation.
To configure a black hole static route, enter
255.255.255.255.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.5.7

| | |
|-------------------|--|
| Parameter: | Preference |
| Default: | 16 |
| Range: | 1 to 16 |
| Function: | Specifies a weighted value (from 1 to 16, with 16 being the most preferred) that the IP router uses to select a route when its routing tables contain multiple routes to the same destination. |
| Instructions: | Enter a value from 1 to 16 for this static route. To configure a black hole static route, enter the maximum preference value. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.5.8 |

Deleting a Static Route

To delete a static route, begin at the IP Static Routes window and proceed as follows:

1. Click on the static route you want to delete.
2. Click on Delete.

The IP Static Routes window no longer displays the static route.

3. Click on Done to exit the IP Static Routes window.

Configuring a Path to an Adjacent Host

To add, edit, or delete a transmission path to an adjacent host, begin at the Configuration Manager window shown in Figure 2-9 and proceed as follows:

1. Select Protocols→IP→Adjacent Hosts.

The IP Adjacent Hosts window appears (Figure 2-15). It lists all adjacent hosts configured on the router. You add, edit, and delete adjacent hosts from this window.

2. Add, edit, or delete adjacent hosts as described in the following sections.

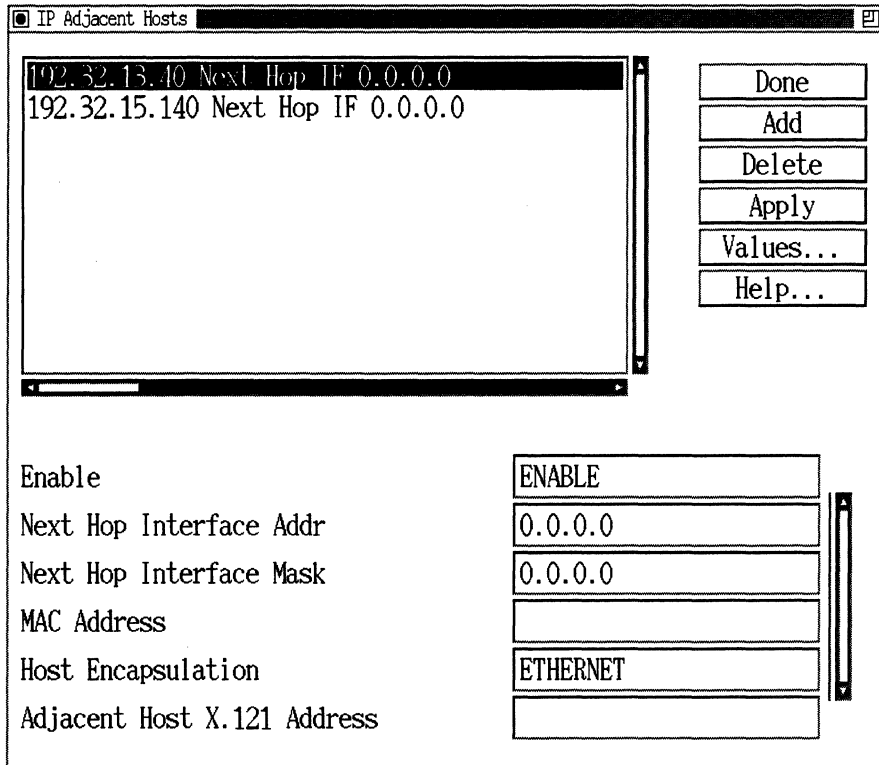


Figure 2-15. IP Adjacent Hosts Window

Adding an Adjacent Host

To add an adjacent host, begin at the IP Adjacent Hosts window and proceed as follows:

1. Click on Add.

The IP Adjacent Host Configuration window (see Figure 2-16) appears.

2. Specify the IP address for the adjacent host; then click on the OK button.

The IP address parameter is described following these instructions.

The IP Adjacent Hosts window now lists the adjacent host you configured. Note that the Enable, Next Hop Interface Addr, Next Hop Interface Mask, MAC Address, Adjacent Host X.121 Address, and Host Encapsulation parameters display the default values for the adjacent host. To edit these parameters, see “Editing Adjacent Host Parameters” on page 2-67.

Repeat Steps 1 and 2 to add additional adjacent hosts.

3. Click on Done to exit the window.




Figure 2-16. IP Adjacent Host Configuration Window

| | |
|-------------------|--|
| Parameter: | IP Address |
| Default: | None |
| Options: | Any valid IP address |
| Function: | Specifies the IP address of the device for which you want to configure an adjacent host. |
| Instructions: | Enter the IP address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.6.1.3 |

Editing Adjacent Host Parameters

You can edit the **Enable**, **Next Hop Interface Addr**, **Next Hop Interface Mask**, **MAC Address**, **Adjacent Host X.121 Address**, and **Host Encapsulation** parameters for an adjacent host.

To edit these parameters, begin at the **IP Adjacent Hosts** window shown in **Figure 2-15** and proceed as follows:

1. Click on the adjacent host that you want to edit.
2. Edit those parameters you want to change.

All adjacent host parameters that you can edit are described in the following section.

3. Click on **Apply** to implement your changes.
4. Click on **Done** to exit the **IP Adjacent Hosts** window.

Note: You cannot change the adjacent host's IP address. If you wish to change this parameter, you must delete the adjacent host and configure a new adjacent host with the proper IP address. See "Deleting an Adjacent Host" for instructions.

Adjacent Host Parameter Descriptions

Use the following descriptions to set parameters on the Adjacent Hosts window.

Parameter: **Enable**

Default: Enable

Options: Enable | Disable

Function: Specifies the state (active or inactive) of the adjacent host in the IP routing tables.

Instructions: Select Disable to make the adjacent host record inactive in the IP routing table; the IP router will not consider this adjacent host.

Select Enable to make the adjacent host record active again in the IP routing table.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.2

Parameter: **Next Hop Interface Addr**

Default: 0.0.0.0

Options: Valid IP address

Function: Specifies the IP address of the router's network interface to the adjacent host.

Instructions: Enter the IP address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.4

Parameter: Next Hop Interface Mask

- Default:** 0.0.0.0
- Options:** Based on the network class of the IP address specified at the Next Hop Interface Addr parameter
- Function:** Specifies the subnet mask of the IP address specified for the Next Hop Addr parameter.
- Instructions:** Enter the subnet mask in dotted decimal notation.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.6.1.5

Parameter: MAC Address

- Default:** None
- Options:** Any valid MAC address
- Function:** Specifies the MAC address of the adjacent host. This value can be a 48-bit Ethernet (or 64-bit SMDS) address or an ATM VPI/VCI address.
- Instructions:** Enter the MAC address as a 12-digit hexadecimal number. Enter an ATM address in the form *Virtual Path Identifier/Virtual Channel Identifier* — for example, 0/32.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.6.1.6

Parameter: Host Encapsulation

Default: Ethernet

Options: Ethernet | SNAP | PDN | DDN

Function: Specifies the adjacent host's encapsulation method.

Instructions: Select Ethernet or SNAP (Service Network Access Point) if you are defining a point-to-point network interface, or if the adjacent host resides on an Ethernet. For an X.25 interface, select PDN or DDN.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.7

Parameter: Adjacent Host X.121 Address

Default: None

Options: Any valid X.121 address

Function: Specifies the X.121 address of the adjacent host. Set this parameter only if this is a PDN/X.25, DDN/X.25, or BFE/X.25 connection.

Instructions: Enter the appropriate X.121 address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.6.1.9

Deleting an Adjacent Host

To delete an adjacent host, begin at the IP Adjacent Hosts window shown in Figure 2-15 and proceed as follows:

1. Click on the adjacent host you want to delete.
2. Click on Delete.
3. Click on Done to exit the IP Adjacent Hosts window.

Editing TFTP Parameters

To edit TFTP parameters for IP, begin at the Configuration Manager window shown in Figure 2-9 and proceed as follows:

1. Select Protocols→IP→TFTP.

The Edit TFTP Parameters window appears (Figure 2-17).

2. Edit those parameters you wish to change.

All TFTP interface parameters are described in the following section.

3. Click on OK to save your changes and exit the window.

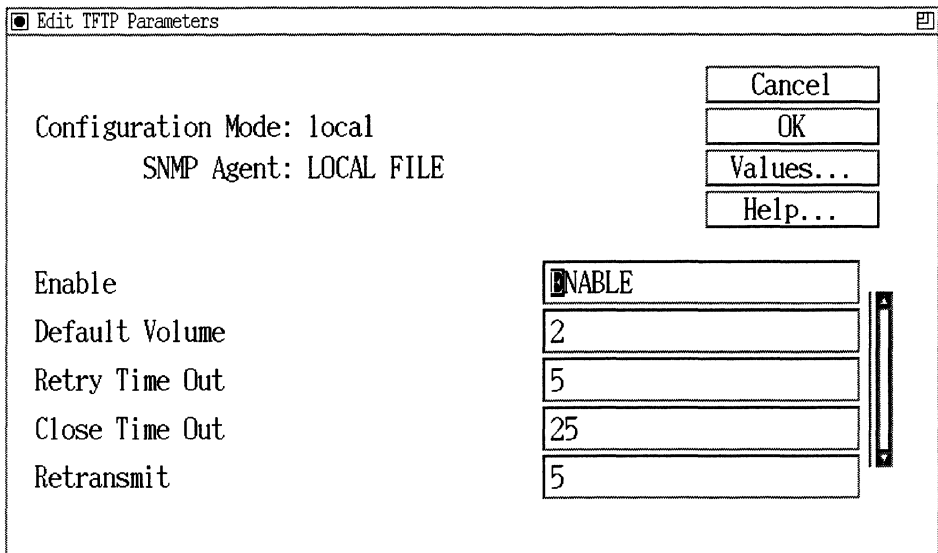


Figure 2-17. Edit TFTP Parameters Window

TFTP Interface Parameter Descriptions

Use the following descriptions to set TFTP interface parameters.

Parameter: **Enable**

Default: Enable

Options: Enable | Disable

Function: Specifies whether TFTP is enabled for the IP router.

Instructions: Select Enable to enable TFTP for the IP router. Because TFTP allows write access to the router's file system, We recommend that you do not enable TFTP in network environments in which you are concerned with security. Select Disable to disable TFTP for the IP router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.1

Parameter: **Default Volume**

Default: 2

Options: 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14

Function: Specifies which of the router's slots will be used, by default, for all TFTP GETs and PUTs.

Instructions: Specify the appropriate slot number. If you are configuring an AN, you must specify slot 1.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.2

Parameter: Retry Time Out

Default: 5 seconds

Options: Any number of seconds

Function: Specifies the number of seconds TFTP waits for an acknowledgment before retransmitting the last packet.

Instructions: Specify the appropriate number of seconds.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.4

Parameter: Close Time Out

Default: 25 seconds

Options: Any number of seconds

Function: Specifies the number of seconds TFTP waits, after it has successfully received a file, to make sure that the sender has received the last acknowledgment.

Instructions: Specify the appropriate number of seconds.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.5

Parameter: Retransmit

Default: 5 retransmissions

Options: Any number of retransmissions

Function: Specifies the number of times TFTP retransmits an unacknowledged message before abandoning the transfer attempt.

Instructions: Specify the number of retransmissions.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.6.6

Configuring RIPS0 Support

To configure RIPS0 support on an IP interface, begin at the Configuration Manager window shown in Figure 2-9 and proceed as follows:

1. Select **Protocols→IP→Interfaces**.
The IP Interfaces window (see Figure 2-1) appears.
2. Click on the IP interface on which you want to enable RIPS0.
3. Scroll through the IP interface parameters until you can access the **Enable Security** parameter.
4. Set the **Enable Security** parameter to **Enable**.
5. Set the remaining RIPS0 parameters.
6. Click on **Apply** to implement your changes.
7. Click on **Done** to exit the window.

RIPSO Interface Parameter Descriptions

Use the following descriptions to set RIPSO parameters.

| | |
|-------------------|---|
| Parameter: | Enable Security |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Enables or disables IP security options for this interface. |
| Instructions: | Set to Disable if you want to disable IP security options. If you set this parameter to Disable, then the router accepts only the following IP datagrams: —Labeled IP datagrams with the classification level set to Unclassified and no authority flags set. —Unlabeled IP datagrams |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.4.76 |

Parameter: Strip Security

Default: None

Options: None | Incoming | Outgoing | All

Function: Specifies the type of IP datagrams from which the router should remove the IP security options.

Instructions: Select the type of IP datagrams from which you want IP security options to be removed as follows:

—None: The router leaves IP security options on all inbound and outbound IP datagrams intact.

—Incoming: The router strips the IP security option from each incoming IP datagram, after checking the IP datagram against the interface's security configuration.

—Outgoing: The router strips the IP security option from each outgoing IP datagram, before checking each datagram against the interface's security configuration.

—All: The router strips the IP security options from both incoming and outgoing IP datagrams; incoming datagrams after checking each against this interface's security configuration and outgoing datagrams before checking each against the interface's security configuration.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.77

Note: If you set the Strip Security parameter to Outgoing or All, then you must set the Require Out Security parameter to None. (Similarly, if you set the Require Out Security parameter to Forwarded, Originated, or All, then you must set the Strip Security parameter to None or Incoming.)

Parameter: Require Out Security**Default:** All**Options:** None | Forwarded | Originated | All**Function:** Specifies which type of outbound datagrams require IP security labels.**Instructions:** Select a Require Out Security type as follows:

—None: The router forwards unlabeled IP datagrams unchanged on this interface. In addition, those IP datagrams that it originates and transmits do not require labels.

—Forwarded: The router requires all IP datagrams it forwards on this interface (not those it originates) to contain basic IP security options. If the datagram already contains an IP security label, the router forwards the datagram unchanged. If the datagram is unlabeled, the router adds the implicit or default label to the datagram before forwarding it.

—Originated: The router specifies basic IP security options for all IP datagrams it originates and transmits on this interface. The router adds the default label to IP datagrams it originates and transmits on this interface.

—All: The router requires all datagrams (both those that it forwards and those it originates) on this interface to contain basic IP security options. It supplies the implicit or default label for those datagrams that do not already contain one.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.78

Note: If you set the Require Out Security parameter to Originated or All, then you must enable the Default Label and Error Label parameters.

Parameter: Require In Security

Default: All

Options: None | All

Function: Specifies which type of incoming IP datagrams require security labels.

Instructions: Select an In Security type that matches your network requirements, as follows:

—None: The router does not require inbound IP datagrams to contain labels.

—All: The router requires all inbound IP datagrams received on this interface to contain basic IP security options.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.79

Parameter: Min Level

Default: Unclassified

Options: Unclassified | Confidential | Secret | Top Secret

Function: Specifies the minimum security level that the router allows for inbound or outbound IP datagrams. This parameter, together with the Max Level parameter, specifies the range of classification levels that the router will accept and process. The router drops IP datagrams it receives on this interface that are below the minimum level specified here.

Instructions: Select a minimum security level for this interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.80

Parameter: Max Level

Default: Top Secret

Options: Unclassified | Confidential | Secret | Top Secret

Function: Specifies the maximum security level that the router allows for inbound or outbound IP datagrams. This parameter, together with the Min Level parameter, specifies the range of classification levels that the router accepts. The router drops IP datagrams it receives or transmits on this interface that are above the maximum level specified here.

Instructions: Select a maximum security level for this interface. The maximum level must be greater than or equal to the minimum level.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.81

Parameter: Must Out Authority

Default: No authority flags selected

Options: No authority flags selected | GENSER | SIOPESE | SCI | NSA | DOE

Function: Specifies which authority flags *must* be set in the protection authority field of all outbound datagrams.

Instructions: Select all of those authority flags that the router must set in all outbound IP datagrams it transmits on this interface. If you do not select any authority flags (the default setting) then the router does not set any protection authority flags in outbound IP datagrams.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.82

Parameter: May Out Authority

Default: ANY

Options: GENSER | SIOPESE | SCI | NSA | DOE

Function: Specifies which authority flags *may* be set in the protection authority field of all outbound datagrams. The authorities you specify here must be a superset of the authorities you specify for the Must Out Authority parameter.

Instructions: The default setting specifies that any of the authority flags may be set. Either accept the default setting, or reset and select only those authority flags that are appropriate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.83

Parameter: Must In Authority

Default: No authority flags selected

Options: No authority flags selected | GENSER | SIOPESE | SCI | NSA | DOE

Function: Specifies which authority flags *must* be set in the protection authority field of inbound IP datagrams.

Instructions: Select all of those authority flags that must be set in inbound IP datagrams received on this interface. If you do not select any authority flags (the default setting) then the router does not require a datagram to have authority flags set, but still accepts the datagram if any flags are set.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.84

Parameter: May In Authority

Default: Any

Options: Any | GENSER | SIOPESE | SCI | NSA | DOE

Function: Specifies which authority flags *may* be set in the protection authority field of inbound IP datagrams. The authorities you specify here must be a superset of the authorities you specify for the Must In Authority parameter.

Instructions: The default setting specifies that any of the authority flags may be set. Either accept the default setting, or reset and select only those authority flags that are appropriate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.85

Parameter: Implicit Label

Default: Enable

Options: Enabled | Disabled

Function: If enabled, the router uses the Implicit Authority and Implicit Level fields to create an implicit label. The router supplies the implicit label to unlabeled inbound datagrams received by this interface. If disabled, the router does not supply implicit labels for this interface.

Instructions: Accept the default, Enable, to allow the router to supply implicit labels for unlabeled inbound datagrams.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.86

Parameter: Implicit Authority

Default: No authority flags selected

Options: No authority flags selected | GENSER | SIOPESE
SCI | NSA | DOE

Function: Specifies the authority flags that the router sets when it supplies implicit security labels for unlabeled inbound IP datagrams.

Instructions: Select all of those authority flags that the router should set when it supplies an implicit security label. The set of authority flags you specify here must include the set of authority flags you specified for the Must In Authority parameter, and cannot include any of the flags you did *not* specify for the May In Authority parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.87

Parameter: Implicit Level

Default: Unclassified

Options: Unclassified | Confidential | Secret | Top Secret

Function: Specifies the security level that the router sets when it supplies implicit security labels for unlabeled, inbound IP datagrams.

Instructions: Specify a level within the range specified by the Min Level and Max Level parameters.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.88

Parameter: Default Label**Default:** Enable**Options:** Enabled | Disabled**Function:** If enabled, the router uses the Default Authority and Default Level fields to create a default label. The router supplies the default label to unlabeled outbound datagrams originated or forwarded out this interface. If disabled, the router does not supply default labels for this interface.**Instructions:** To allow the router to supply default labels for unlabeled outbound datagrams, accept the default Enable.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.4.89**Parameter: Default Authority****Default:** No authority flags selected**Options:** No authority flags selected | GENSER | SIOPEI
SCI | NSA | DOE**Function:** Specifies the authority flags that the router uses when it supplies default security labels to unlabeled outbound IP datagrams.**Instructions:** Select those authority flags that the router should set when it supplies default security labels. The set of authority flags you specify must include the set of authority flags specified for the Must Out Authority parameter, and cannot include any of the flags you did *not* specify for the May Out Authority parameter.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.4.90

Parameter: Default Level

Default: Unclassified

Options: Unclassified | Confidential | Secret | Top Secret

Function: Specifies the security level that the router sets when it supplies default security labels to unlabeled outbound IP datagrams.

Instructions: Specify a default level within the range specified by the Min Level and Max Level parameters.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.91

Parameter: Error Label

Default: Enable

Options: Enabled | Disabled

Function: If enabled, the router uses the Error Authority and Min Level fields to create an error label. The router supplies the error label to outbound ICMP error datagrams. If disabled, the router does not supply error labels for this interface.

Instructions: To allow the router to supply error labels for outbound ICMP error datagrams, accept the default, Enable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.4.92

| | |
|-------------------|---|
| Parameter: | Error Authority |
| Default: | No authority flags selected |
| Options: | No authority flags selected GENSER SIOPESE SCI NSA DOE ALL |
| Function: | Specifies the authority flags that the router uses when it supplies error security labels to outbound ICMP error datagrams. |
| Instructions: | Select those authority flags that the router should set when it supplies error security labels to outbound ICMP error datagrams. The set of authority flags you specify here must include the set of authority flags you specified for the Must Out Authority parameter, and cannot include any of the flags you did <i>not</i> specify for the May Out Authority parameter. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.4.93 |

Configuring Router Discovery

To configure Router Discovery, begin at the Configuration Manager window (Figure 2-9) and proceed as follows:

1. Select Protocols→IP→Router Discovery. The IP Router Discovery window appears (Figure 2-18). Edit the options from this window

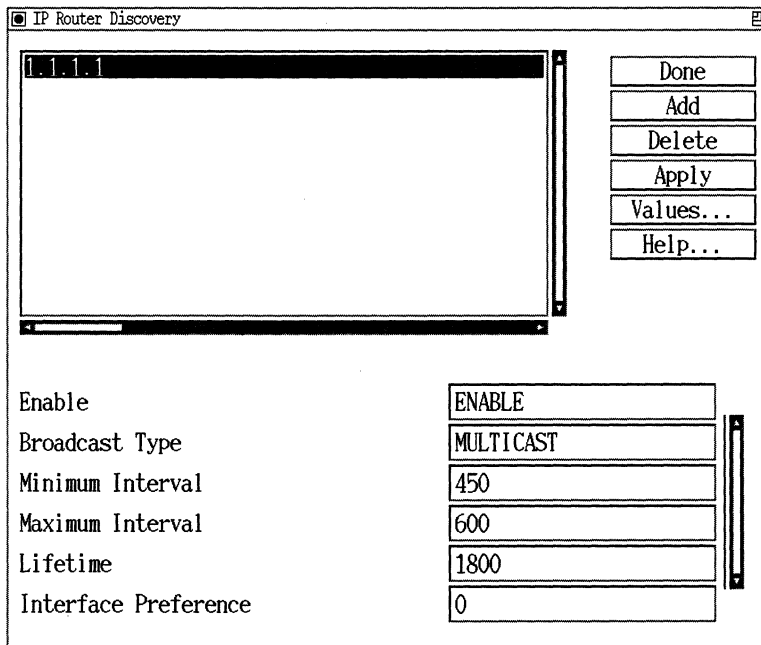


Figure 2-18. IP Router Discovery Window

Router Discovery Window Parameter Descriptions

Use the following descriptions to set Router Discovery parameters.

Parameter: **Enable**

Default: Enable

Options: Enable | Disable

Function: Disables and enables Router Discovery on this interface.

Instructions: If you configured this interface with Router Discovery, use this parameter to disable Router Discovery.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.2

Parameter: **Broadcast Type**

Default: Multicast

Options: Multicast | Local | Direct

Function: Specifies the type of broadcast to use in sending advertisements.

Instructions: Use Multicast wherever possible; that is, on any link where all listening hosts support IP multicast.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.5

Parameter: Minimum Interval

Default: 450

Options: A value specifying the number of seconds

Function: Specifies the minimum time interval between advertisements.

Instructions: Specify a value that is no less than 3 seconds and less than the value you set for the Maximum Interval parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.6

Parameter: Maximum Interval

Default: 600

Options: A value specifying the number of seconds

Function: Specifies the maximum time interval between advertisements.

Instructions: Specify a value that is no less than 4 seconds, greater than the value you specified for the Minimum Interval parameter, and no greater than 1800 seconds.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.7

Parameter: Lifetime**Default:** 1800**Options:** A value specifying the number of seconds**Function:** Specifies the maximum length of time that the advertised addresses are to be considered as valid router addresses by hosts, in the absence of further advertisements.**Instructions:** Specify a value that is no less than the value you set for the Maximum Interval parameter and no greater than 9000 seconds.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.17.1.8**Parameter: Interface Pref****Default:** Null**Options:** A numeric value**Function:** Specifies the preferability (a higher number indicates more preferred) of the address as a default router address, relative to other router addresses on the same subnet.**Instructions:** Enter a value indicating the relative preferability of the router address.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.17.1.9

Configuring Blacker Front-End Support

Configuring BFE support on an IP interface requires you to

- ❑ Configure an X.25 interface that conforms to the BFE requirements described in this section.
- ❑ Enable the IP routing protocol on the interface.
- ❑ Enable RIPS0 support on the interface.

Before you begin the procedures described in this section, we recommend that you have the following guides available for reference:

- ❑ *Configuring Wellfleet Routers*
- ❑ The appropriate protocol manual

To configure BFE support on an IP interface, begin at the Configuration Manager window and perform the following procedures:

1. Configure an X.25 interface.

When you initially configure packet level parameters for the X.25 interface, make certain to:

- Set the Network Address Type parameter to BFE_NETWORK.
- Set the DDN IP Address parameter to the IP address that is assigned to your BFE connection.

2. Add network service record(s) to the X.25 interface.
3. Enable the IP routing protocol on the X.25 interface.
4. Edit the packet layer parameters for the X.25 interface so that they match the settings specified in Table 2-1.
5. Edit the network service record parameters for the X.25 interface so that they match the settings specified in Table 2-2.

Remember to set the DDN BFE parameter to Enable.

6. Configure IP security options (RIPS0) on the interface. IP security must be enabled and labels are required on all outbound data.

For instructions on performing Steps 1 through 3, see the *Configuring Wellfleet Routers* guide. For instructions on performing Steps 4 and 5, see *Configuring X.25 Services*. For instructions on performing Step 6, see the section “Configuring the Revised IP Security Option” on page 2-14.

Note: Generally, the synchronous line parameter settings are the same for both a DDN X.25 link and a BFE X.25 link. However, if your operating environment has specific needs, you may want to edit synchronous line parameters. See the appropriate protocol manual for instructions.

Table 2-1. BFE Required X.25 Packet-Level Parameter Settings

| X.25 Parameter | BFE Required Setting |
|-------------------------------------|---|
| Enable | Enable |
| Network Address Type | BFE_NETWORK |
| PDN X.121 Address | Parameter is ignored. |
| DDN IP Address | Specify the IP address assigned to your BFE connection. |
| Sequence Size | MOD8 |
| Restart Procedure Type | DTE_RESTART |
| Default Tx/Rx Window Size | BFE Range is 1-7. This setting should match the default value configured in the BFE. |
| Default Tx/Rx Packet Length | BFE options include 128, 256, 512, and 1024. This setting should match the default value configured in the BFE. |
| Incoming Logical Channel Count | Zero (0). BFE does not support the one-way logical channel incoming facility. |
| Incoming LCN Start | Parameter is ignored. |
| Bidirectional Logical Channel Count | Any valid non-zero setting. |
| Bidirectional Logical Start | Any valid non-zero setting. |

| X.25 Parameter | BFE Required Setting |
|--|--|
| Outgoing Logical Channel Count | Zero (0). BFE does not support the one-way logical channel outgoing facility. |
| Outgoing LCN Start | Parameter is ignored. |
| T1 Timer, T2 Timer, T3 Timer, T4 Timer | BFE has no special requirements for any of these four parameters. |
| Flow Control Negotiation | Set to On if you do not want to use the default values configured in the BFE for this link. |
| Max Window Size | BFE Range is 1-7. If you specify any other setting than the default value configured in the BFE, set Flow Control Negotiation to On. |
| Max Packet Length | BFE options include 128, 256, 512, and 1024. If you specify any other value than the default value configured in the BFE, then set Flow Control Negotiation to On. (If IP interface is configured to support multiple IP security levels, then set to 1024.) |
| Trans/Recv Throughput Class | Parameter is ignored. |
| Max Throughput Class | Parameter is ignored. |
| Flow Control Negotiation | Set to On if you do not want to use the default values configured in the BFE for this link. |
| Throughput Class Negotiation | Off |
| Network User Identification | Off |
| Incoming Calls Accept | On |
| Outgoing Calls Accept | On |
| Fast Select Accept | Off |
| Reverse Charge Accept | Off |
| Fast Select | Off |
| Reverse Charging | Off |
| CUG Selection | Null |

| X.25 Parameter | BFE Required Setting |
|-------------------------|-----------------------------|
| CUG Outgoing Access | Null |
| CUG Bilateral Selection | Null |
| RPOA Selection | Off |
| Charging Information | Off |
| Transit Delay | Off |
| Full Addressing | On |
| Acceptance Format | Basic |
| Release Format | Basic |
| CCITT Conformance | DXE1980 |
| Network Standard | DOD |

Table 2-2. BFE Required X.25 Network Service Record Parameter Settings

| X.25 Parameter | BFE Required Setting |
|------------------------|---|
| Enable | Enable |
| Type | DDN |
| Connection ID | Parameter is ignored. |
| Remote IP Address | Specify the IP address of the remote system. |
| Remote X.121 Address | Parameter is ignored. |
| Broadcast | Parameter is ignored. |
| Max Connections | Any valid setting. |
| Precedence | Any valid setting. The BFE will accept, but not act on the DDN Precedence facility. |
| Max Idle | Any valid setting. |
| Call Retry | Any valid setting. |
| Flow Facility | Set to On if you want to use a value other than the default window size and packet size configured in the BFE. |
| Window Size | BFE Range is 1-7. If you if you want to use a value other than the default window size configured in the BFE, set Flow Facility to On. |
| Packet Size | BFE options include 128, 256, 512, and 1024. If you want to use a value other than the default packet size configured in the BFE, set Flow Facility to On. (If IP interface is configured to support multiple IP security levels, then set to 1024.) |
| Fast Select Request | Off |
| Fast Select Accept | Off |
| Reverse Charge Request | Off |
| Reverse Charge Accept | Off |
| User Facility | Null |
| DDN BFE | Enable |

Chapter 3

Customizing RIP Services

This chapter describes the Routing Information Protocol (RIP) and tells you how to edit RIP parameters. It contains

- A RIP overview
- Instructions for editing RIP parameters

Routing Information Protocol (RIP) Overview

The Routing Information Protocol (RIP) is a distance-vector protocol that lets routers in the same autonomous system exchange routing information by means of periodic RIP updates.

Routers transmit their own RIP updates to neighboring networks and listen for RIP updates from the routers on those neighboring networks. Routers use the information in the RIP updates to update their internal routing tables. For RIP, the “best” path to a destination is the shortest path (the path with the fewest hops). RIP computes distance as a metric, usually the number of hops (or routers) from the origin network to the target network.

Editing RIP Parameters

This section describes how to edit, or customize, RIP parameters for IP interfaces on which you enabled RIP.

Note: The instructions in this section assume that you have already configured at least one IP interface on which you enabled RIP (RIP interface). If you have *not* yet configured a RIP interface, or want to add additional IP interfaces, see the *Configuring Wellfleet Routers* guide for instructions.

You access all RIP parameters from the Configuration Manager window shown in Figure 2-9 (refer to the *Configuring Wellfleet Routers* guide for instructions on accessing this window).

For each RIP parameter, this chapter describes the default setting, all valid setting options, the parameter function, and instructions for setting the parameter.

Editing Routing Information Protocol (RIP) Interface Parameters

Once you have enabled RIP on an IP interface, you can edit the RIP parameters for the interface.

Note: If you want to change the RIP Diameter Value for the IP Router, refer to “Editing IP Global Parameters” on page 2-26.

To edit RIP parameters, begin at the Configuration Manager window and proceed as follows:

1. Select the Protocols→IP→RIP Interfaces option.

The RIP Interfaces window appears (see Figure 3-1). It lists all RIP interfaces configured on the router.

2. Click on the RIP interface you want to edit.
3. Edit those parameters you wish to change.

All RIP interface parameters are described in the following section.

4. Click on Apply to implement your changes.
5. Click on Done to save your changes and exit the window.

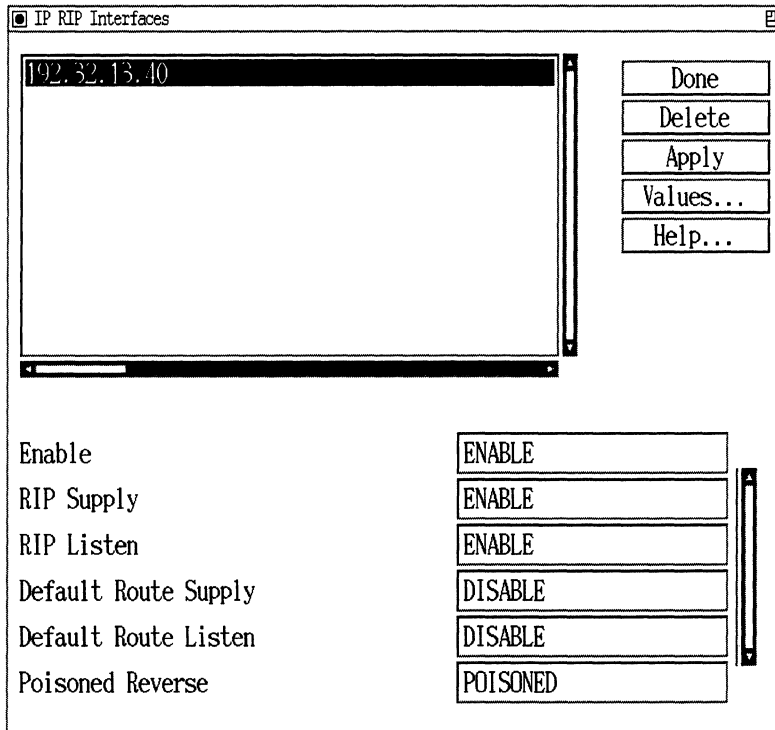


Figure 3-1. IP RIP Interfaces Window

RIP Parameter Descriptions

This section describes how to set all parameters shown on the IP RIP Interfaces window.

| | |
|-------------------|--|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Specifies whether the Routing Information Protocol (RIP) is enabled on this interface. |
| Instructions: | Select Enable to enable RIP on this interface. Select Disable to disable RIP on this interface. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.17.1.2 |
| | |
| Parameter: | RIP Supply |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Specifies whether the interface transmits periodic RIP updates to neighboring networks. |
| Instructions: | Select Enable to configure the interface to transmit RIP updates. You must select Enable if you want the interface to supply default route information. Select Disable to prohibit the interface from transmitting RIP updates. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.17.1.5 |

Parameter: **RIP Listen**

Default: Enable

Options: Enable | Disable

Function: Specifies whether this interface listens to RIP updates from neighboring networks.

Instructions: Select Enable to configure this interface to listen to RIP updates, and thus, add received routing information to its internal routing table.

If you set RIP Listen to Enable, a route filter can still prohibit the interface from updating its internal routing tables.

Select Disable to configure the interface to ignore RIP updates from neighboring routers. Thus, the interface does not add received routing information to its internal routing table.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.6

Parameter: Default Route Supply

Default: Disable

Options: Enable | Disable

Function: Specifies whether or not the interface advertises a default route in RIP updates sent to neighboring networks. In order to advertise a default route, you must have either statically configured a default route, or the router must have learned the default route (0.0.0.0). When a router does not know the direction of a particular address, it uses the default route as the destination.

Instructions: Select Enable to configure the interface to advertise the default route. If you set Default Route Supply to Enable, you must also set RIP Supply to Enable.

Select Disable to configure the interface not to advertise the default route.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.7

Parameter: Default Route Listen

Default: Disable

Options: Enable | Disable

Function: Specifies whether or not IP adds default route information to its internal routing table.

Instructions: Select Enable to configure the RIP interface to listen for and potentially add the default route (0.0.0.0) information to its internal routing table. Note that you must also enable RIP Listen on this interface.

Select Disable to prohibit the RIP interface from adding the default route (0.0.0.0) information to its internal routing table.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.17.1.8

Parameter: Poisoned Reverse**Default:** Poisoned**Options:** Poisoned | Actual | Split**Function:** Specifies how the RIP interface advertises routes it learns from an adjacent network in periodic updates subsequently sent to that network.**Instructions:** Select **Poisoned** to configure this RIP interface to implement **Poisoned Reverse**. When **poisoned reverse** is enabled, the RIP interface advertises routes learned from an adjacent network in RIP updates subsequently sent to that network with a hop count of **RIP Network Diameter** plus one; thus declaring the destination unreachable. **Poisoned Reverse** can speed up the convergence of the network routing tables.Select **Actual** to configure this RIP interface to advertise routes with the learned cost.Select **Split** to configure this RIP interface to implement a **split horizon**. When a **split horizon** is configured, the RIP interface omits routes learned from a neighbor in RIP updates subsequently sent to that neighbor.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.17.1.9



Chapter 4

Customizing OSPF Services

Open Shortest Path First (OSPF) is an internal gateway protocol for use in large networks. This chapter describes OSPF and tells you how to edit OSPF parameters.

- “Link States and Shortest Path Trees” on page 4-2
- “Variable Length Subnet and Supernet Addresses” on page 4-2
- “Configuring Network Support on an Interface” on page 4-3
- “Defining a Routing Area” on page 4-5
- “Configuring a Boundary Router” on page 4-10
- “Configuring a Virtual Link through a Transit Area” on page 4-10
- “Configuring Cost Metrics” on page 4-11
- “Defining a Summary Route” on page 4-12
- “Enabling Authentication and Specifying a Password” on page 4-13
- “Constructing an External Route Advertisement” on page 4-13
- “Discovering and Configuring Neighbors” on page 4-16
- “Configuring the OSPF Primary and Backup Soloist” on page 4-17
- “Configuring OSPF Message Logging” on page 4-18
- “Putting the Pieces Together” on page 4-18
- “Editing OSPF Parameters” on page 4-22

Link States and Shortest Path Trees

OSPF is a link state protocol. A router running a link state protocol periodically tests the status of the physical connection — the link — to each of its neighbor routers and sends this information to its other neighbors. A link state protocol does not require each router to send its entire routing table to its neighbors. Instead, each router floods only link state change information throughout the system (a system, in this case, may be the autonomous system, or a subset of the autonomous system called an area). This process is referred to as the synchronization of the routers' topological databases.

With the link information, each router builds a shortest path tree with itself as the root of the tree. It then can identify the shortest path from itself to each destination, and build its routing table. Once the routers are synchronized and the routing tables are built, the routers flood topology information only in response to some topological change (a disabled router, a downed line, and so on). How this reduces network routing traffic becomes clear when you consider that a distance vector routing protocol requires its routers to automatically flood their entire routing table every 30 seconds, regardless of whether there has been topological change.

Using a link state algorithm, OSPF exchanges routing information between routers in an autonomous system. OSPF responds quickly to topological changes, calculating new loop-free paths using only a small amount of routing protocol traffic.

Variable Length Subnet and Supernet Addresses

Subnetting and supernetting strategies allow a number of networks to be aggregated by address and appear to be one single network. Subnets and supernets are identified by variable length addresses. The address length is indicated by a mask.

OSPF understands IP subnet and supernet addresses and variable length masks. Supernetting allows OSPF to support Classless Interdomain Routing by performing classless routing within an IP domain.

Configuring Network Support on an Interface

OSPF provides interfaces to four types of networks:

- ❑ Point-to-point networks
- ❑ Broadcast networks
- ❑ Nonbroadcast multiaccess networks
- ❑ Point-to-multipoint networks

A point-to-point network joins a single pair of OSPF routers. An example of such a network would be a network of synchronous lines.

A broadcast network supports multiple routers and can address a single physical message to all attached routers. Examples of a such network are Ethernet, FDDI, and Token Ring.

A nonbroadcast multiaccess network supports multiple routers and cannot address a single physical message to all routers. Examples of such a network are Frame Relay and X.25 networks.

A point-to-multipoint network supports multiple routers in star Frame Relay topologies. The OSPF point-to-multipoint network interface is a Bay Networks proprietary solution for routers running OSPF in star Frame Relay topologies.

In Figure 4-1, for example, four AFN™ routers are connected by Frame Relay links to a BCN™ router. The AFN routers are the spokes of the topology, and the BCN router is the hub. All of the routers are running OSPF. The BCN router is connected to the Frame Relay network over a permanent virtual circuit (PVC) in group mode. The AFN routers are connected over PVCs in direct or group mode. (For details on Frame Relay, see *Customizing Frame Relay Services*.)

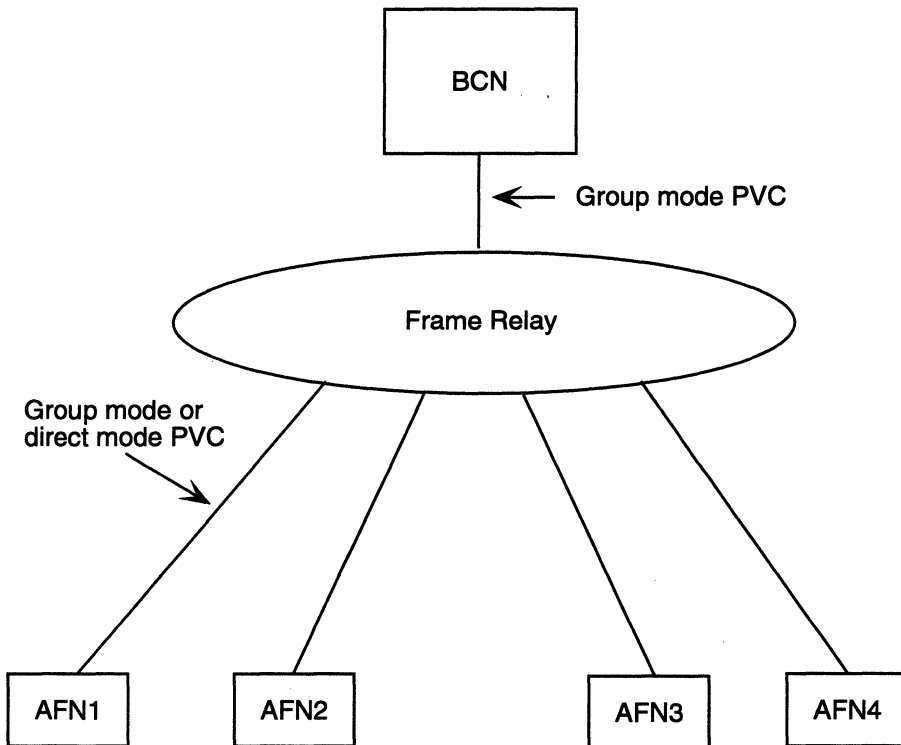


Figure 4-1. Point-to-Multipoint Topology

OSPF point-to-multipoint interfaces provide an efficient means to connect routers in a star topology. The routers are configured as follows:

1. The hub of the star topology — the BCN router in Figure 4-1 — is configured with a point-to-multipoint interface to the PVC and is set to be the OSPF designated router in the network. The Router Priority Parameter is set to a value greater than 0.
2. Each spoke of the star — the AFN routers in Figure 4-1 — is configured with a point-to-multipoint interface to the PVC and is made ineligible to become the designated router. The Router Priority parameter on each AFN is set to 0.

When the spokes of the topology (the AFN routers) are computing routes through the other spokes, the next hop is forced to be the hub (the BCN router). The hub can then forward the packet to the correct spoke.

Running OSPF with point-to-multipoint network interfaces addresses two problems: how to keep to a minimum the number of subnets required and the number of interfaces required to support communications within the star topology. With point-to-multipoint interfaces, each star topology requires only one subnet, rather than one subnet for each PVC. Secondly, the hub needs to support only one interface for each star, rather than one interface for each PVC. This reduces the demand for resources on the router.

To specify a network type on an OSPF interface, see the Type parameter on page 4-43.

Defining a Routing Area

OSPF allows the autonomous system to be subdivided into areas. An area consists of groups of contiguous networks and hosts, and any router having an interface to any of the networks in the group. Each area maintains a separate copy of the basic routing algorithm and shares an identical topological database with every other router in the area.

Each OSPF AS contains a central area — called the backbone — that is responsible for connecting all other areas and distributing routing information between them. It consists of the networks that are *not* included in any other area.

Like the other areas, the backbone

- ❑ Must be contiguous
- ❑ Uses link state information to create an SPF tree and, from that, build a routing table
- ❑ Has a topology that is invisible to all other areas and does not know the topology of other areas

The backbone's main function is to distribute routing information between all of the other areas in the autonomous system. It must always take the area ID of 0.0.0.0.

Another type of area that OSPF supports is the stub area. These areas are dependent on default routes to get out of the AS area; external routes are not flooded into or throughout stub areas. This not only reduces bandwidth overhead but also the internal router's topological database size and, in turn, its memory requirements.

You can configure a stub area when there is just one point of exit from an area. This is usually characterized by a leaf or branch node (a router that has one connection to a LAN and one connection to the rest of the world). You can also configure a stub area when the choice of exit does not need to be made on a per-external-destination basis (you can get to all AS external routes through all of the stub's area border routers). Default routing must be used in a stub area (the stub area's area border router must advertise a default route into the stub area). All routers in an area must agree on whether the area has been configured as a stub area.

Figure 4-2 shows an OSPF autonomous system divided into three areas and a backbone.

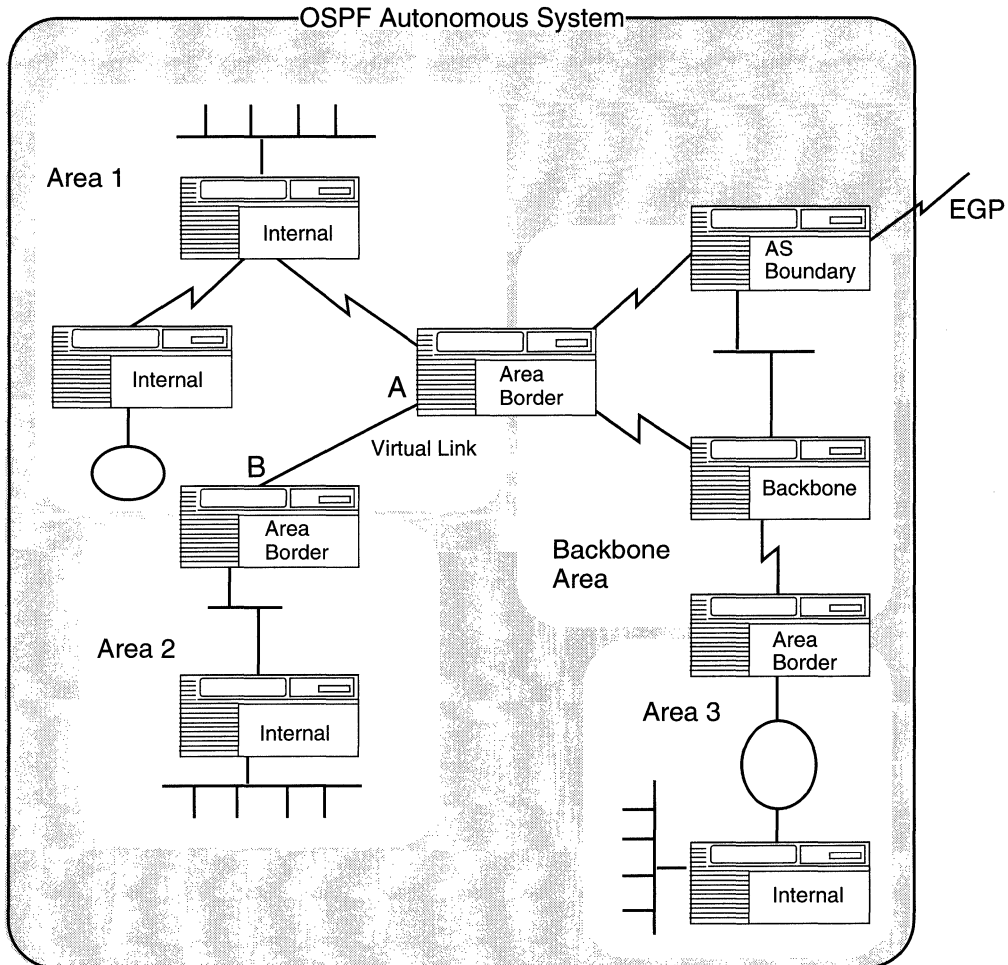


Figure 4-2. OSPF Areas

Dividing the autonomous system into areas reduces the level of protocol traffic by reducing the amount of flooded topology change information. That is, when there is a topological change, the router floods this information to other routers in its area only, rather than to every router in the autonomous system.

Dividing the autonomous system into areas also provides a level of security, in that the physical topology of an area is invisible to nonarea residents. Conversely, routers that reside within a single area know nothing of the physical topology external to that area.

For instructions on how to define an OSPF area, see “Editing OSPF Area Parameters” on page 4-30.

OSPF Router Classifications

Segmentation of the AS into areas results in four functional classifications of routers (see Table 4-1). Note that the four classifications overlap; a router can fall into more than one classification.

Table 4-1. OSPF Router Classifications

| Router Type | Description/Function |
|--------------------|---|
| Internal Router | The internal router resides within an area. All of its directly connected networks belong to the same area. Routers with only backbone interfaces also fall into this category. Each internal router runs a single copy of the basic routing algorithm. |
| Area Border Router | The area border router attaches to more than one area, and runs multiple copies of the basic routing algorithm — one copy for each area to which it is attached. An area border router distributes topological information about each of its attached areas to the backbone; then, the backbone distributes that same information to other areas. |
| Backbone Router | The backbone router is any router that has an interface to the backbone, including all routers that have an interface to more than one area (area border router). Backbone routers with all interfaces connected to the backbone are considered to be internal routers. |

Table 4-1. OSPF Router Classifications (*continued*)

| Router Type | Description/Function |
|--------------------|---|
| AS Boundary Router | The AS boundary router is the autonomous system's link to other routing domains. The AS boundary router exchanges router information with routers belonging to other routing domains. Such a router has AS external routes that are advertised throughout the autonomous system. The path to each AS boundary router is known to every other router in the autonomous system. |

Intra-area, Inter-area, and External Routing

There are two types of routing within the autonomous system: intra-area routing and inter-area routing. Intra-area routing takes place between a source and destination that reside in the same area. Inter-area routing takes place between a source and destination residing in different areas. Inter-area routing always involves area border routers, which are responsible for providing the source area with information about the topology of some other area or areas in the AS.

A third type of routing, external routing, takes place between a source and destination that are in different routing domains. External routing always involves an AS boundary router, which is responsible for providing the AS in which it resides information about other routing domains. It floods this information throughout its own AS, excepting all stub areas. Paths to AS boundary routers are summarized by nonstub area border routers.

Configuring a Boundary Router

External routes are learned and propagated by AS boundary routers. These routers run not only OSPF (on interfaces internal to the AS), but may also run some exterior gateway protocol (on the interface that connects to another AS), such as EGP. The following list defines routes that OSPF considers external routes:

- ❑ A route to a destination outside the AS
- ❑ A static route
- ❑ A default route
- ❑ A route derived by RIP
- ❑ A directly connected network not running OSPF

Boundary routers learn and propagate external routes. These external routes can be tagged if you configure IP policies or route filters to identify the system that is delivering the routes to the OSPF system. The AS boundary router floods an External Links Advertisement describing these routes throughout the entire AS.

To configure a boundary router, see the AS Boundary Router parameter on page 4-25.

To define the characteristics of the external route advertisements that a boundary router injects into an OSPF domain, see “Constructing an External Route Advertisement” on page 4-13.

Configuring a Virtual Link through a Transit Area

Every area border router must connect directly to the backbone. If an area border router does not connect physically to the backbone, a virtual link must be configured from it to the backbone. These virtual links are configured like a point-to-point connection between one area border router and another. The area through which the virtual link is configured is called the transit area. Once configured, these virtual links are considered part of the backbone.

In the network shown in Figure 4-2, a virtual link needs to be configured from area border router A, through Area 1 (the transit area), to area border router B. This is necessary to restore the contiguity of the backbone, because area border router B does not have a direct physical connection to the backbone. Every area border router must belong to the backbone area.

To configure a virtual link through a transit area and define a virtual interface for a border router, see “Configuring OSPF Virtual Interfaces” on page 4-55.

Configuring Cost Metrics

In contrast to RIP (a distance vector routing protocol), which considers only a hop count in calculating the best path, OSPF considers a cost metric that you assign to a path.

OSPF recognizes that a simple hop count takes no account of reliability, bandwidth, delay, or the actual dollar cost of using a path. Passing through an extra hop to get to a 1.54-MB T1 channel, for instance, may be more efficient than traversing a shorter, but slower route. For OSPF, the best path is the one that offers the least-cost metric delay. With the Wellfleet implementation of OSPF, every path automatically takes on a cost metric value of 1. You must configure cost metrics if you want to specify a preferred path. To specify a preferred path, you would allow the preferred path to retain the cost metric value of 1, and then assign higher cost metric values to the less preferred paths.

Figure 4-3 shows the benefit of using configurable cost metrics. Assigning the 56-KB line a cost metric value of 10 forces OSPF to choose the faster T1 line path as the best path, despite the extra hop, when transmitting a packet from host A to host B.




Figure 4-3. Configurable Cost Metrics Usage Example

For instructions on specifying a cost metric for an OSPF interface, see the Metric Cost parameter on page 4-49.

Defining a Summary Route

Area border routers generate summary advertisements. In OSPF, you can configure a corresponding subnet mask for an advertised route that indicates an address range being described by the route. The area border router sends a summary link advertisement for the single address/mask pair, rather than sending one summary link advertisement for every network defined by the address/mask pair, thus reducing the amount of routing traffic. For example, a summary advertisement for the destination 140.191.0.0 with a mask of 255.255.0.0 actually is describing a single route to the collection of destinations 140.191.0.0 to 140.191.255.255. When a packet is forwarded, it is always forwarded to the network that is the best (longest or most specific) match for the packet's destination.

For instructions, see “Adding a Range to an Area” on page 4-35.

Enabling Authentication and Specifying a Password

OSPF provides a measure of security through the use of passwords. If an area is configured to use authentication, all OSPF interfaces configured in that area must be configured with a password. The password must be identical on each interface connected to the same network. Different networks can have different passwords.

In such an area, a router that receives a packet verifies the password before doing anything else with the packet. Unauthorized routers are not allowed to communicate with the OSPF system.

For instructions on enabling authentication in an area, see the **Authentication Type** parameter on page 4-33.

For instructions on specifying a password for an OSPF interface, see the **Password** parameter on page 4-50.

Constructing an External Route Advertisement

OSPF boundary routers inject AS external (ASE) route advertisements into the OSPF domain. (For information on configuring a router as an OSPF boundary router, see “Configuring a Boundary Router” on page 4-10.)

The advertisement includes a Type 1 or Type 2 metric. The Type 1 metric is equivalent to the metric of the non-OSPF route. The Type 2 metric is either the metric of the non-OSPF route or the weight value calculated for that route (see “Using the Route Weight as the Type 2 Metric” on page 4-14).

By default, the boundary router generates a Type 2 metric for BGP, EGP, or RIP routes. For routes from all other sources, the boundary router generates a Type 1 metric. You can construct an OSPF announce policy to override the default metric type. For details, see the **Type** parameter on page 9-42.

Using the Route Weight as the Type 2 Metric

Beginning with Version 8.00, the network administrator has the option of generating OSPF AS external (ASE) routes that use the route weight as the Type 2 metric.

Figure 4-4, for example, shows three routers — A, B, and C — in an OSPF domain. Router A and Router B are both configured to generate ASE routes using the route weight as the Type 2 metric.

1. Router A learns a route to destination X via EGP.
2. Router A advertises the route to Router C as an OSPF ASE route. The Type 2 metric in the advertisement contains the route weight value calculated for the EGP route to destination X.
3. Router B learns a route to destination X via BGP.
4. Router B advertises the route to Router C as an OSPF ASE route. The Type 2 metric in the advertisement contains the route weight value calculated for a BGP route.
5. To determine the preferable route, Router C compares the Type 2 metrics — the EGP route weight and the BGP route weight.
6. Router C selects the BGP route — the route with the lower weight.

Note: The route weight will appear to be a greater value than the route's original metric. For this reason, all routers advertising a particular network must use the same metric type — Type 1 or Type 2. If not, the router that receives the advertisements may choose the wrong route.

For information about route weights, see “Route Weights” on page 1-15. For instructions on configuring a router to use the route weight as the OSPF metric, see the ASE Metric Support parameter on page 4-26.

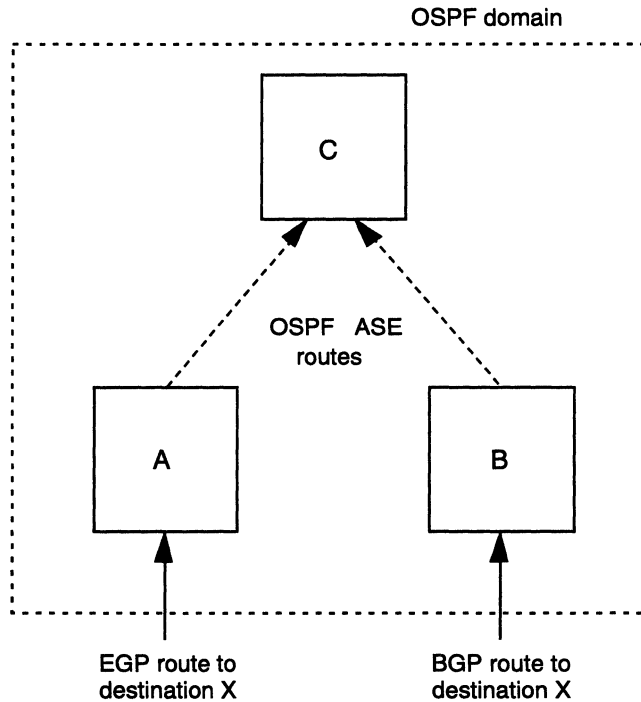


Figure 4-4. OSPF ASE Routes

Generating and Matching an External Route Tag

OSPF AS external route advertisements include a user-configurable external route tag. For instructions on specifying a tag value, see the `announce Tag` parameter on page 9-43. For instructions on matching a tag value, see the `accept policy Tag` parameter on page 9-13.

Generating an External Route Tag for OSPF/BGP Interaction

The network administrator can configure a router to automatically generate the external route tag according to RFC 1403, "OSPF/BGP Interaction." For instructions, see the `Tag Generation Method` parameter on page 4-29. You can also configure an OSPF announce

policy to enable automatic tag generation. For instructions, see the Automatic Tag parameter on page 9-43.

Routers that export OSPF routes into BGP use the OSPF external route tag to set the Origin and AS Path attributes in the BGP update message.

Discovering and Configuring Neighbors

OSPF neighbors are any two routers that have an interface to the same network. In each OSPF network, routers use the Hello protocol to discover their neighbors and maintain neighbor relationships. On a broadcast or point-to-point network, the Hello protocol dynamically discovers neighbors; however, on a nonbroadcast multiaccess network, you must manually configure neighbors.

The Hello protocol is responsible for ensuring that communication between neighbors is bidirectional. Periodically, OSPF routers send out hello packets over all interfaces. Included in these hello packets are

- ❑ The router's priority
- ❑ The router's hello timer and dead timer value
- ❑ A list of routers that have sent this router hello packets on this interface
- ❑ The router's choice for designated router and backup designated router

Bidirectional communication is determined when one router sees itself listed in the neighbor's hello packet.

For instructions on setting the characteristics of the Hello protocol on an OSPF interface, see the Hello Interval parameter on page 4-46, the Dead Interval parameter on page 4-47, and the Poll Interval parameter on page 4-48.

To configure a neighbor on a nonbroadcast multiaccess network, see "Adding a Neighbor to an NBMA Interface" on page 4-51.

Neighbors may form a relationship called an adjacency for the purpose of exchanging routing information. When two routers form an adjacency, the routers go through a process to synchronize their topological databases. When their databases are synchronized, the routers are said to be fully adjacent. From this point on, only routing change information is passed between the adjacencies, thus conserving bandwidth.

All routers connected by a point-to-point network or a virtual link will always form an adjacency. Also, every router on a multi-access network forms an adjacency relationship with the designated router and backup designated router.

Electing a Designated and Backup Designated Router

To further reduce the amount of routing traffic, the Hello protocol elects a designated router and a backup designated router on each multiaccess network. Instead of neighboring routers forming adjacencies and swapping link state information with each other (which on a large network can mean a lot of routing protocol traffic), all routers on the network form an adjacency with the designated router and the backup designated router only and send link state information to them. The designated router then redistributes the information from each router to every other router.

In case the designated router goes down, a backup designated router is always elected at the same time that the designated router is elected. Its responsibility is to take over all of the designated router's functions should the designated router fail.

Configuring the OSPF Primary and Backup Soloist

The OSPF protocol is implemented by a single process running on a single slot of a router. If the slot on which the OSPF soloist is running goes down, the router will attempt to run OSPF on another slot. By default, the router uses any available slot. To specify a slot or slots, see the OSPF Slot parameter on page 4-26.

The OSPF backup soloist provides a method of preserving information learned from the network in the event of an OSPF crash or slot removal.

Each time the OSPF soloist is restarted, all of the routing information is lost and must be relearned from the network.

The time-consuming and resource-intensive process of relearning routing information can be avoided by enabling the OSPF backup soloist. In the event of a crash or slot removal, transition between the OSPF primary and backup soloist occurs without relearning routing information from the network.

To disable and enable the OSPF backup soloist, see the Backup Enable parameter on page 4-27.

Configuring OSPF Message Logging

Two special Site Manager windows allow you to customize how much message logging you want from OSPF.

For instructions, see the Primary Log Mask parameter on page 4-27 and the Backup Log Mask parameter on page 4-28.

Putting the Pieces Together

An OSPF autonomous system consists of multiple areas and a backbone. Each area is a contiguous group of hosts and networks and routers that have interfaces to those networks. The backbone consists of networks not included in any area, routers attached to those networks, and routers attached to more than one network.

Within each area and within the backbone reside four classifications of routers: internal routers, backbone routers, area border routers, and AS boundary routers. These classifications are functional and can overlap.

All areas in the autonomous system must be physically contiguous with the backbone or, if not contiguous, have a virtual link to the backbone. An area containing border routers that are configured to create a virtual link between another area and the backbone is called a transit area.

OSPF supports interfaces to four types of network: point-to-point networks, broadcast networks, nonbroadcast multicast networks, and point-to-multipoint networks. OSPF also supports IP subnetting and supernetting, address ranges, and special areas called stubs that rely on default routing.

There are three categories of OSPF routing: intra-area routing, inter-area routing, and external routing. Inter-area routing occurs when source and destination reside in the same area. Inter-area routing occurs when source and destination reside in different areas within the same AS. External routing occurs when source and destination reside in different ASs or when source or destination reside on a RIP network within the AS.

All routers in an OSPF area must have databases that are synchronized for that area. First, the routers in the area use the Hello protocol to discover their neighbors— each router sends periodic hello packets out all interfaces and checks to see itself listed in the hello packets it receives from other routers. Next, it forms an adjacency relationship with certain neighbors or, on a multi-access network, with the designated router and backup designated router. This relationship is established to facilitate the distribution of routing information. All routing protocol packets, except for the hello packet, are sent over adjacencies.

By issuing link state advertisements, adjacent routers synchronize their area topology databases to facilitate routing between sources and destinations within the area. To route beyond the area, a router depends on area border routers. These border routers advertise topology information to the backbone; the backbone, in turn, advertises the information to all other areas, thus facilitating routing between different areas. Each AS boundary router exchanges information with routers from other autonomous systems or with routers from RIP

networks within the same autonomous system. Each AS boundary router receives routes from external networks — for example, RIP or EGP networks — which it advertises throughout the autonomous system. Each router in the area knows the path to every boundary router, thus facilitating routing to external networks.

For More Information about OSPF

If you would like more information about OSPF, refer to the following documents:

Moy, J. "OSPF Version 2." RFC 1247, Network Information Center (NIC), SRI International, Menlo Park, California, July 1991.

Comer, Douglas E. *Internetworking with TCP/IP, Volume I: Principle, Protocols, and Architecture*. 2d ed. Englewood Cliffs, New Jersey, Prentice Hall, Inc., 1991.

Perlman, Radia. *Interconnections: Bridges and Routers*. Reading, Massachusetts, Addison-Wesley Publishing Company, 1992.

OSPF Implementation Notes

This section provides some suggestions to help you when configuring your OSPF network. The Wellfleet OSPF implementation does not restrict you to these suggestions, but we are providing them as guidelines.

- ❑ Keep the same password throughout an area, or even throughout the entire OSPF AS, if possible.
- ❑ Use the default timers, unless you are running 9.6-KB sync lines. In this case, double the default timers on both ends of the link.
- ❑ Use address ranges if your network is a subnetted network.
- ❑ Keep all subnets within one area. If you cross areas, you cannot configure summaries.

- ❑ Make sure the AS Border Router parameter is enabled if the router has any non-OSPF interfaces, and you want that information propagated.
- ❑ You must configure virtual links for each area border router that does not reside within or directly interface to the backbone (see Figure 4-2). Every area border router must have a configured path to the backbone. See “Configuring OSPF Virtual Interfaces” on page 4-55.
- ❑ Rather than just a hop count, OSPF considers the cost of a path when choosing the best path. Each interface, however, is assigned the default cost 1 for the path to which it interfaces. If you have a preferred path, you must edit the Metric Cost parameter for your interfaces. You will need to assign a higher metric cost for those paths which are *not* preferred paths. See “Editing OSPF Interface Parameters” on page 4-40.
- ❑ If you have any devices in your network running OSPF, and are now adding a Bay Networks router, you must make sure that the router’s timer values coincide with the timers in your other devices. Determine the timer values of the other devices, and change the router’s timer values to match them. See “Editing OSPF Interface Parameters” on page 4-40.
- ❑ If there is a topology change (for example, if you add an area, combine two areas, move routers, and so on), you must reconfigure the appropriate OSPF elements (OSPF area ranges/interfaces/neighbors/virtual links, and so on).

Editing OSPF Parameters

This section describes how to edit, or customize, OSPF parameters for IP interfaces on which you enabled OSPF support.

Note: The instructions in this section assume that you have already configured at least one IP interface with OSPF support (OSPF interface) on the router. If you have *not* yet configured an OSPF interface, or want to add additional interfaces, see the *Configuring Wellfleet Routers* guide for instructions.

You access all OSPF parameters from the Configuration Manager window (refer to the *Configuring Wellfleet Routers* guide for instructions on accessing this window).

For each OSPF parameter, this chapter describes the default setting, all valid setting options, the parameter function, and instructions for setting the parameter.

OSPF parameters are described in the following sections:

- “Editing OSPF Global Parameters” on page 4-23
- “Editing OSPF Area Parameters” on page 4-30
- “Editing OSPF Interface Parameters” on page 4-40
- “Configuring OSPF Virtual Interfaces” on page 4-55

Editing OSPF Global Parameters

When you edit OSPF global parameters, you are editing parameters that affect OSPF on the entire router.

To edit OSPF global parameters, begin at the Configuration Manager window and complete the following steps:

1. Select Protocols→IP→OSPF→Global.

The Edit OSPF Global Parameters window appears (Figure 4-5).

2. Edit those parameters you wish to change.

All OSPF global parameters are described following these instructions.

3. Click on OK to save your changes and exit the window.

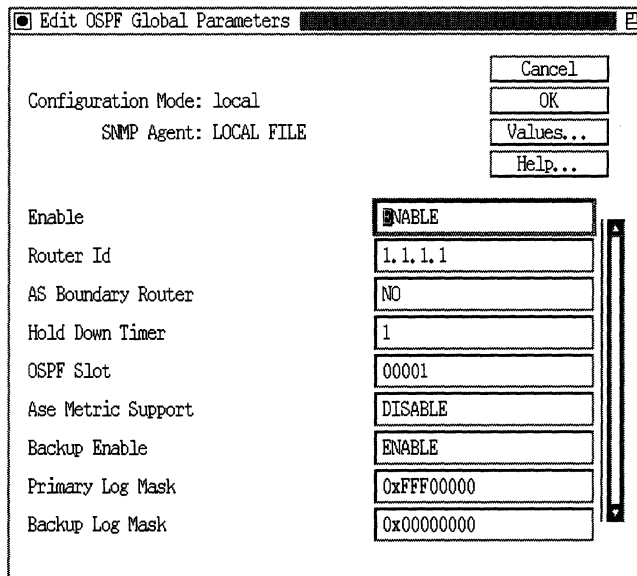


Figure 4-5. Edit OSPF Global Parameters Window

OSPF Global Parameter Descriptions

This section describes how to set all parameters shown on the Edit OSPF Global Parameters window.

| | |
|-------------------|--|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | This parameter allows you to globally enable or disable OSPF on all router interfaces. |
| Instructions: | Set to Disable if you want to disable OSPF for the entire router. Set to Enable if you previously disabled OSPF on the router and now wish to re-enable it. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.1.2 |
| | |
| Parameter: | Router ID |
| Default: | The IP address of the first OSPF circuit configured on this router. |
| Options: | Any IP address; preferably, one of the router's IP interface addresses |
| Function: | This IP address uniquely identifies this router in the OSPF domain. By convention, and to ensure uniqueness, one of the router's IP interface addresses should be used as the router ID. The router ID will determine the designated router on a broadcast link if the priority values of the routers being considered are equal. The higher the router ID, the greater its priority. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. See note below. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.1.4 |

Note: If both OSPF and BGP are running on the router, then the OSPF router ID must be identical to the BGP identifier.

Parameter: AS Boundary Router

Default: No

Options: Yes | No

Function: Indicates whether or not this router functions as an AS boundary router. Only AS boundary routers are allowed to convert non-OSPF routes into OSPF routes so that they can be passed along throughout the OSPF routing domain. The router can be an AS boundary router if one or more of its interfaces is connected to a non-OSPF network (for example, RIP, BGP, or EGP).

Instructions: Set this parameter to Yes if this router functions as an AS boundary router. Otherwise, accept the default value, No.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.7

Parameter: Hold Down Timer

Default: 1

Range: 0 to 10 seconds

Function: Prevents the algorithm from running more than once per hold down time. Its purpose is to free up the CPU. Note that a value of 0 means there is no hold down time.

Instructions: Either accept the default value of 1 second, or enter a new value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.9

Parameter: OSPF Slot

Default: All slots

Options: Any slot on the router

Function: Indicates which slot(s) the OSPF soloist is eligible to run on. If the slot on which the OSPF soloist is running goes down, the router will attempt to run OSPF on another slot specified by the OSPF Slot parameter.

Instructions: Select all of the appropriate slots.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.10

Note: Use caution when selecting the slot(s) on which OSPF may run. If you choose an empty slot, and it is the only slot you choose, OSPF will not run; if you choose a slot that becomes disabled, and it is the only slot you choose, OSPF will not restart.

Parameter: ASE Metric Support

Default: Disable

Options: Enable | Disable

Function: Causes the router to use the route weight as the OSPF metric in OSPF ASE Type 2 advertisements.

Instructions: Disable ASE metric support if the router is to interoperate with routers using a pre-8.00 OSPF version. The new metric is not compatible with the pre-8.00 metric.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.11

Parameter: Backup Enable

Default: Disable

Options: Enable | Disable

Function: Enables or disables the backup OSPF soloist's backup link state database. When the parameter is set to disabled, the OSPF backup soloist will not maintain a copy of the OSPF link state database.

Instructions: Select the default, Disable, if you do not want to back up the OSPF soloist.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.12

Parameter: Primary Log Mask

Default: Log all messages

Options: See Figure 4-6

Function: Specifies which OSPF log messages should be logged in the primary log.

Instructions: Highlight the line entry for Primary Log Mask in the Edit OSPF Global Parameters window and click on Values. The Primary Log Mask window appears (Figure 4-6). Choose the log messages that you wish to enter into the primary log by clicking on their buttons. Then click on OK.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.13




Figure 4-6. Primary Log Mask Window

Parameter: Backup Log Mask

Default: Log no messages

Options: See Figure 4-7

Function: Specifies which OSPF log messages should be logged in the backup log.

Instructions: Highlight the line entry for Backup Log Mask in the Edit OSPF Global Parameters window and click on Values. The Backup Log Mask window appears (see Figure 4-7). Choose the log messages that you wish to enter into the backup log by clicking on their buttons. Then click on OK.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.14




Figure 4-7. Backup Log Mask Window

Parameter: Tag Generation Method

Default: Zero

Options: Zero | Autotag | Proprietary

Function: Specifies the method of OSPF external tag field generation.

Instructions: Set the parameter to Autotag if you want OSPF to generate a tag value according to RFC 1403, "OSPF/BGP Interaction."

Use the default to insert 0 into the tag field.

The Proprietary option is reserved for debugging purposes.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.15

Editing OSPF Area Parameters

To edit OSPF Area Parameters, begin at the Configuration Manager window and proceed as follows:

1. Select the Protocols→IP→OSPF→Areas option.

The OSPF Area List window appears (Figure 4-8). It lists the areas currently configured on the router.

2. Perform any of the functions listed below, which are described in the following sections:
 - Add an area
 - Edit an area
 - Delete an area
 - Add a range to an area
 - Edit an area's range
 - Delete a range from an area

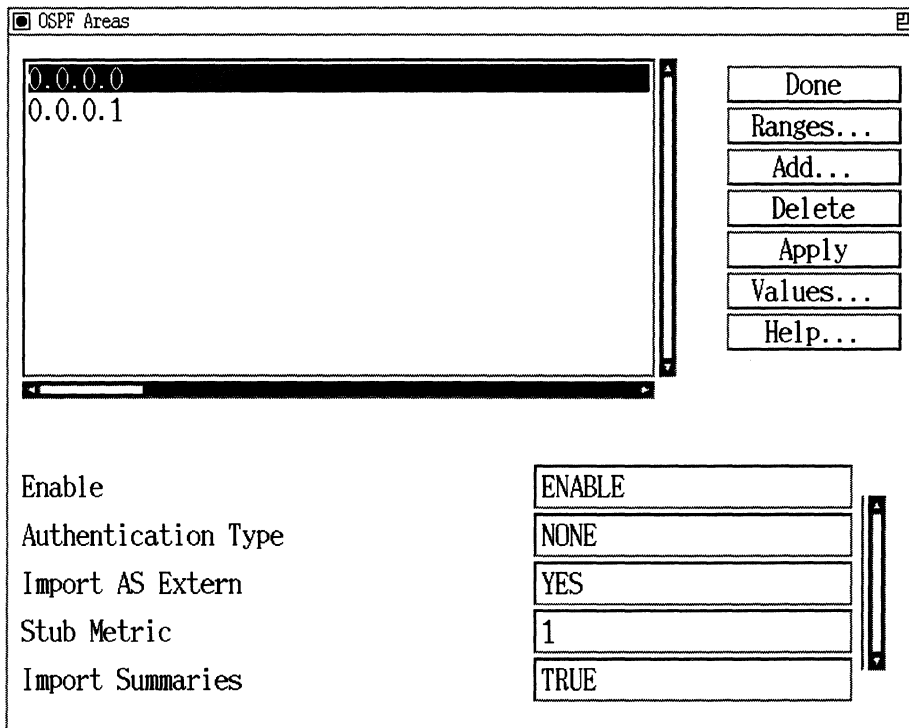


Figure 4-8. OSPF Area List Window

Adding an Area

To add an OSPF area, begin at the OSPF Area List window shown in Figure 4-8 and complete the following steps:

1. Click on Add.

The OSPF Area Configuration window appears.

2. Enter the area address of the new area at the OSPF Area parameter.

3. Click on Done to save your changes and exit.

The OSPF Area List window now lists the area you added.

OSPF Area Configuration Parameter Description

| | |
|-------------------|---|
| Parameter: | OSPF Area |
| Default: | 0.0.0.0 |
| Options: | Any 4-octet number in dotted decimal notation. |
| Function: | Identifies the area ID, which is the OSPF area to which this interface belongs. |
| Instructions: | Enter the appropriate area ID in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.2.1.4 |

Note: The backbone area ID is always 0.0.0.0.

Editing an Area

After you add an area, you may edit any of your area's default parameters. To edit an area, begin at the OSPF Area List window shown in Figure 4-8 and complete the following steps:

1. Click on the area you want to edit.
2. Edit the OSPF area parameters.
The OSPF area parameters that you can edit are described following these instructions.
3. Click on Apply to implement your changes.
4. Click on Done to exit the window.

OSPF Area Parameter Descriptions

This section describes how to set all OSPF area parameters.

| | |
|-------------------|---|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Allows you to enable and disable this area. This parameter is useful if you want to temporarily disable an area rather than delete it. |
| Instructions: | Set this parameter to Disable if you want to disable this area. Set this parameter to Enable if you previously disabled the area and now wish to re-enable it. This will cause OSPF to restart. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.1.2 |
| | |
| Parameter: | Authentication Type |
| Default: | None |
| Options: | None Simplepassword |
| Function: | Enables or disables password authentication for the area. If you select Simplepassword (enabling password authentication), only those routers sharing the correct password will be able to communicate with each other. If you accept the default, None, password authentication is disabled for this area. |
| Instructions: | Either accept the default value, None, to disable password authentication, or select Simplepassword to enable password authentication. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.1.5 |

Parameter: Import AS Extern

Default: Yes

Options: Yes | No

Function: Indicates whether or not this area imports AS external link state advertisements. If this area does *not* import AS external link state advertisements, it is a stub area. If it does import AS external link state advertisements, it is not a stub area.

Instructions: Set to No if this area functions as a stub area. Otherwise, accept the default value, Yes.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.6

Parameter: Stub Metric

Default: 1

Range: 1 to 255

Function: When an area border router is connected to a stub area, it generates a default link summary into the area specifying a default route. The stub metric is the cost of that route. By default, Stub Metric equals 1. This parameter has meaning only when the Import AS Extern parameter is set to No.

Instructions: Either accept the Stub Metric default value, 1, or supply the appropriate Stub Metric value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.1.7

| | |
|-------------------|--|
| Parameter: | Import Summaries |
| Default: | True |
| Options: | True False |
| Function: | Specifies whether network summaries are flooded into a stub area. This variable has meaning only if the Import AS Extern parameter is set to No. |
| Instructions: | Set to False if Import AS Extern is set to No <i>and</i> you do not want network summaries imported into the stub area. Otherwise, set to True. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.1.8 |

Deleting an Area

Sometimes, as the result of a topology change, you may want to delete an area. To delete an area, begin at the OSPF Area List window shown in Figure 4-8 and complete the following steps:

1. Click on the area you want to delete.
2. Click on Delete.
3. Click on OK to confirm the deletion.

The area no longer appears in the OSPF Area List window.

4. Click on Done to exit the window.

Adding a Range to an Area

Ranges are address/mask pairs that let you group subnetted networks that reside in the same area, and to have that group be advertised by *one* network summary advertisement. Otherwise, a summary advertisement would be generated for each subnet in the area.

To add a range to an area, begin at the OSPF Area List window shown in Figure 4-8 and complete the following steps:

1. Click on the area for which you want to define a range.
2. Click on Ranges.

The OSPF Range List window appears (Figure 4-9).

3. Click on Add.

The OSPF Range Area window appears (Figure 4-10).

4. Specify the Range Net and Range Mask parameters.

These parameters are described following these instructions.

5. Click on OK.
6. Click on Done to exit the window.




Figure 4-9. OSPF Range List Window

OSPF Range Parameter Descriptions

This section describes how to set OSPF range parameters.




Figure 4-10. OSPF Range Area Window

| | |
|-----------------------|---|
| Parameter: | Range Net |
| Default: | None |
| Options: | Any network number |
| Function: | Allows you to assign a single network address to a group of subnets. This network address, together with the subnet mask you provide, specifies the subnets to be grouped in this area range. Just one link summary advertisement will be generated for all subnets in this range, rather than one link summary advertisement for each of the subnets included in that network. |
| Instructions: | Enter the appropriate network number in dotted decimal notation. |
| MIB Object ID: | 1.3.6.14.1.18.3.5.3.2.3.4.1.5 |

Parameter: Range Mask

Default: None

Options: Any address mask

Function: This parameter, together with Range Net, indicates all of the networks that belong to this range. The range mask is not restricted to the natural address class mask for the address supplied at Range Net.

In this example, Range Net is 128.1.0.0 and Range Mask is 255.255.0.0. That means that the link summary advertisement generated will summarize networks 128.1.0.0 to 128.1.254.254.

Instructions: Enter the appropriate subnet mask in dotted decimal notation.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.6

Note: When setting up your OSPF network, keep all subnetted networks in the same area.

Editing an Area's Range

Once you add a range to an area, you can edit the Enable and Mask parameters for the range.

To edit a range, begin at the OSPF Area List window shown in Figure 4-8 and complete the following steps:

1. Click on the area for which you want to edit a range.
2. Click on Ranges.

The OSPF Range List window appears (Figure 4-9).

3. Click on the range that you want to edit.
4. Edit the Enable or Mask parameter, or both.

The Enable and Mask parameters are described following these instructions.

5. Click on Apply to implement your changes.
6. Click on Done to exit the window.

Parameter: Enable

Default: Enable

Options: Enable | Disable

Function: Enables or disables this range for the specified area. This parameter is useful if you want to disable the range, rather than delete it.

Instructions: Set this parameter to Disable if you want to disable this range. Set the parameter to Enable if you previously disabled this range and now wish to re-enable it.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.2

Parameter: Mask

Default: None

Options: Any address mask

Function: This parameter allows you to change the mask portion of this area range. Mask, together with Range Net, indicates all of the networks that belong to this range. Mask is not restricted to the natural address class mask for the address supplied at Range Net.

In this example, Range Net is 128.1.0.0 and Range Mask is 255.255.0.0. That means that the link summary advertisement generated will summarize networks 128.1.0.0 to 128.1.255.255.

Instructions: Enter the appropriate address mask in dotted decimal notation.

MIB Object ID: 1.3.6.14.1.18.3.5.3.2.3.4.1.6

Deleting a Range from an Area

If you no longer want a range to be associated with an area, you can delete it.

To delete a range, begin at the OSPF Area List window and complete the following steps:

1. Click on the area for which you want to delete a range.
2. Click on Ranges. The OSPF Range List window appears (Figure 4-9).
3. Click on the range you want to delete.
4. Click on Delete. This range no longer appears in the OSPF Range List window.
5. Click on Done to save your changes and exit the window.

Editing OSPF Interface Parameters

All OSPF interfaces assume certain default values when you first configure them. You can, however, change these defaults by editing the interface-specific parameters. The changes you make affect only the interface you select.

To edit OSPF interface parameters, begin at the Configuration Manager window and complete the following steps:

1. Select Protocols→IP→OSPF→Interfaces.
The OSPF Interface List window appears (Figure 4-11).
2. Perform any of the functions listed below, which are described in the following sections.
 - “Editing an OSPF Interface” on page 4-42
 - “Deleting OSPF from an Interface” on page 4-51
 - “Adding a Neighbor to an NBMA Interface” on page 4-51
 - “Editing a Neighbor” on page 4-53
 - “Deleting a Neighbor” on page 4-55

Note: When you reconfigure an interface in dynamic mode, OSPF restarts on all interfaces; the only exception to this is when you dynamically change the Transit Delay, Hello Interval, Retransmission Interval, or Dead Interval timers.

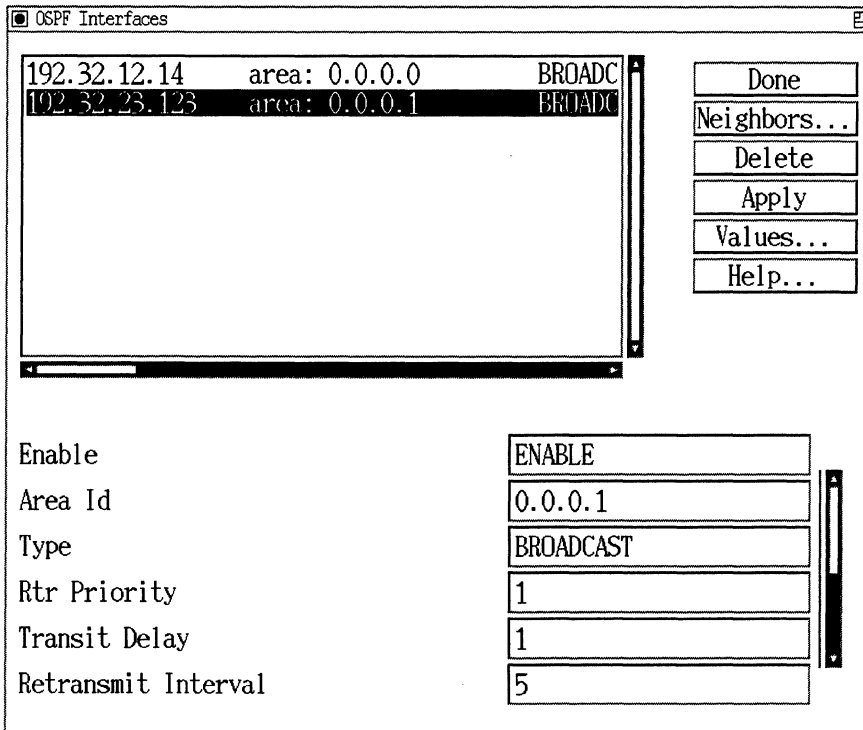


Figure 4-11. OSPF Interface List Window

Editing an OSPF Interface

To edit OSPF interface parameters, begin at the OSPF Interface List window (see Figure 4-11) and complete the following steps:

1. Click on the interface you want to edit.
2. Edit those parameters you wish to change.

All OSPF interface parameters are described following these instructions.

3. Click on Apply to implement your changes.
4. Click on Done to exit the window.

Parameter: Enable

Default: Enable

Options: Enable | Disable

Function: This parameter indicates whether or not OSPF is enabled on this interface. The default value Enable indicates that neighbor relationships may be formed on this interface, and that this interface will be advertised as an internal route to some area. The value Disable indicates that this is not an OSPF interface.

Instructions: Set this parameter to Disable if you do not want OSPF enabled on the interface. Or, set it to Enable if you previously disabled OSPF on this interface, and now wish to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.2

| | |
|-------------------|---|
| Parameter: | Area ID |
| Default: | 0.0.0.0 |
| Options: | Any 4-octet number in dotted decimal notation |
| Function: | This parameter identifies the area to which this interface belongs. |
| Instructions: | Enter the appropriate area ID in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.5.1.6 |

Note: Note that area ID 0.0.0.0 is used only for the OSPF backbone.

| | |
|-------------------|--|
| Parameter: | Type |
| Default: | Broadcast |
| Options: | Broadcast NBMA (nonbroadcast multiaccess) point-to-point point-to-multipoint |
| Function: | Indicates this interface's type (the type of network to which it is attached). Set this parameter to Broadcast if this network is a broadcast LAN, such as Ethernet. Set it to NBMA for an X.25 or similar type of interface. Set it to point-to-point for a synchronous, point-to-point interface. Set it to point-to-multipoint for a star Frame Relay topology. |
| Instructions: | Set this parameter to match this interface type. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.5.1.7 |

Note: If you set the Type parameter to NBMA, you need to configure neighbors manually.

Parameter: Rtr Priority

Default: 1

Range: 0 to 255

Function: Indicates the priority of this interface. The router priority value is used in multi-access networks (Broadcast, NBMA, or point-to-multipoint), for the election of the designated router. If this parameter is set to 0, this router is not eligible to become the designated router on this particular network.

In the case of equal Rtr Priority values, the router ID will determine which router will become the designated router. However, if there already is a designated router on the network when you start this router, it will remain the designated router no matter what your priority or router ID.

Instructions: Set the router priority to a value between 0 and 255, or accept the default value, 1.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.8

Parameter: Transit Delay

Default: 1 second

Range: 1 to 360 seconds

Function: Indicates the estimated number of seconds it takes to route a packet over this interface.

Instructions: Either accept the default value of 1 second, or enter some slightly higher number for slower speed serial lines, for example, 15 to 20 seconds for a 19.8KB line.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.9

Parameter: Retransmit Interval**Default:** 5 seconds**Range:** 1 to 360 seconds**Function:** Indicates the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. The Retransmit Interval value is also used when retransmitting OSPF packets. Although the default value is 5, we suggest the following values for Retransmit Interval.

| Network Type | Suggested Retransmit Interval |
|---------------------|--------------------------------------|
| Broadcast | 5 seconds |
| Point-to-point | 10 seconds |
| NBMA | 10 seconds |
| Point-to-multipoint | 10 seconds |

Instructions: Either accept the default value of 5 seconds, or set the retransmit interval to some slightly higher number for slower-speed serial lines.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.3.5.1.10

Parameter: Hello Interval

Default: 10 seconds

Range: 1 to 360 seconds

Function: Indicates the number of seconds between the hello packets that the router sends on the interface. Although the default value is 10 seconds, we suggest the following values for Hello Interval.

| Network Type | Suggested Hello Interval |
|---------------------|--------------------------|
| Broadcast | 10 seconds |
| Point-to-point | 15 seconds |
| NBMA | 20 seconds |
| Point-to-multipoint | 15 seconds |

Instructions: Either accept the default value of 10 seconds, or set the hello interval to some higher number for slower-speed serial lines.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.11

Note: The Hello Interval value must be the same for all routers attached to the same network.

Parameter: Dead Interval**Default:** 40 seconds**Range:** 1 to 2000 seconds**Function:** Indicates the number of seconds that a router's hello packets have not been seen before its neighbors declare the router down. The Dead Interval value should be some multiple of the Hello Interval. We suggest the following values for Dead Interval.

| Network Type | Suggested Dead Interval |
|---------------------|-------------------------|
| Broadcast | 40 seconds |
| Point-to-point | 60 seconds |
| NBMA | 80 seconds |
| Point-to-multipoint | 60 |

Instructions: Either accept the default value of 40 seconds, or set the dead interval some higher number for slower-speed serial lines.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.3.5.1.12**Note:** The Dead Interval value must be the same for all routers attached to the same network.

Parameter: **Poll Interval**
Default: 120 seconds
Range: 1 to 2000 seconds
Function: Indicates the largest number of seconds allowed between hello packets sent to an inactive non-broadcast multi-access neighbor.
Instructions: Either accept the default value of 120 seconds, or set Poll Interval to some slightly higher number for slower speed serial lines.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.13

Parameter: Metric Cost**Default:** 1**Options:** 1 to 65535**Function:** Indicates the cost of using this type of service on this interface. We suggest the following values for Metric Cost.

| Network Type/Bit Rate | Suggested Metric Cost |
|-----------------------|-----------------------|
| > = 100 Mb/s | 1 |
| Ethernet/802.3 | 10 |
| E1 | 48 |
| T1 | 65 |
| 64 Kb/s | 1562 |
| 56 Kb/s | 1785 |
| 19.2 Kb/s | 5208 |
| 9.6 Kb/s | 10416 |

Metric Cost allows you to configure preferred paths. If you do want to configure a preferred path, allow that path to retain the default value of 1, or assign it a relatively low metric cost. Then, assign the less preferred paths a higher Metric Cost value.

Instructions: Either accept the default value, 1, or enter a larger number for a slower path or a backup route.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.3.5.1.16

Parameter: Password

Default: None

Options: Any ASCII string up to eight characters long

Function: Specifies the password used for this area. You can specify a password up to eight ASCII characters in length that will appear in the authentication field of all OSPF packets across this interface. Password is valid only when Authentication Type is set to Simplepassword.

Instructions: Enter the appropriate password.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.17

Note: All routers in the same area must either have no Authentication, or have the same Password.

Parameter: MTU Size

Default: 1

Options: 1 | 2 | 3 to 1000

Function: Specifies the maximum transmission unit (MTU) size of OSPF updates on this interface.

Instructions: Accept the default value, 1, to use the IP MTU size for that physical interface. Enter 2 to send packets no larger than the IP MTU size for Ethernet. Enter a number from 3 to 1000 to specify an MTU size directly; the number you enter must be less than the IP MTU size for that physical interface.

Note: When running OSPF over a synchronous/PPP link, set the MTU size to a value less than the sync MTU size (1200). This allows all OSPF routes to be learned over the link.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.5.1.29

Deleting OSPF from an Interface

To delete OSPF from an interface on which it is currently configured, begin at the Configuration Manager window and proceed as follows:

1. Click on the connector from which you want to delete OSPF services.
2. Click on Edit Circuit.
3. Select Protocols→Add or Delete.

The Select Protocols window appears. The OSPF button is highlighted to show that OSPF is enabled on the circuit.

4. Click on OSPF to deselect it.
5. Click on OK to exit the window.
6. Select File→Exit to exit the Circuit Definition window and return to the Configuration Manager window.

Adding a Neighbor to an NBMA Interface

In an NBMA network, neighbors are not learned dynamically. For each neighbor on the network, you need to enter its IP address.

Note: You configure neighbors for NBMA interfaces only (those where the interface's Type parameter is set to NBMA).

To add a neighbor to an NBMA interface, begin at the OSPF Interface List window (refer to Figure 4-11) and complete the following steps.

1. Click on the interface to which you want to add a neighbor.
2. Click on Neighbors.

The OSPF Neighbor List window appears (Figure 4-12).

3. Click on Add.

The OSPF Neighbor Configuration window appears (Figure 4-13).

4. Enter the appropriate neighbor address.

5. Click on OK.
6. Click on Done to exit the window.

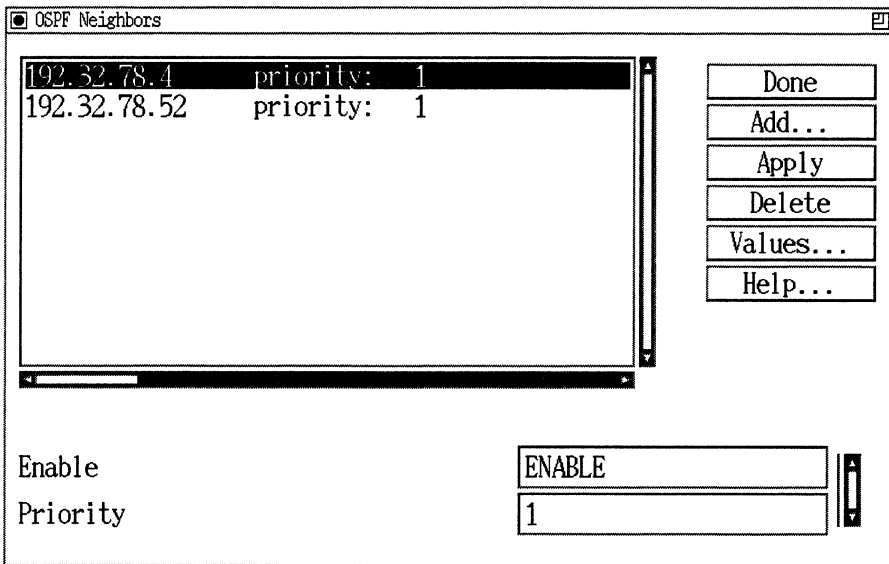


Figure 4-12. OSPF Neighbor List Window

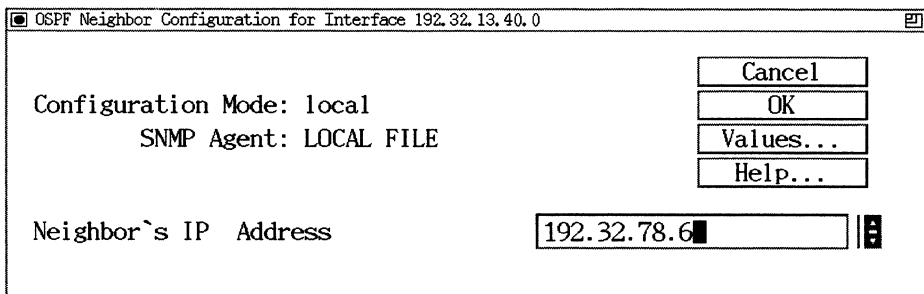


Figure 4-13. OSPF Neighbor Configuration Window

| | |
|-------------------|--|
| Parameter: | Neighbor Address |
| Default: | None |
| Options: | IP address of neighbor |
| Function: | Indicates by IP address a nonbroadcast multi-access neighbor for this interface. |
| Instructions: | Enter the appropriate IP address of the nonbroadcast multi-access neighbor in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.7.1.4 |

Editing a Neighbor

Once you have configured the neighbors for an NBMA interface, you can change them.

To edit a neighbor, begin at the OSPF Interface List window shown in Figure 4-11 and complete the following steps.

1. Click on the interface for which you want to edit a neighbor.
2. Click on Neighbors.

The OSPF Neighbor List window appears (Figure 4-12).

3. Click on the neighbor that you want to edit.
4. Edit the Enable and Priority parameters, which are described in the following section.
5. Click on Apply to implement your changes.
6. Click on Done to exit the window.

OSPF Neighbor Parameter Descriptions

This section describes how to set all OSPF neighbor parameters.

Parameter: Enable

Default: Enable

Options: Enable | Disable

Function: Allows you to enable and disable this neighbor configuration for this interface. This parameter is useful if you want to temporarily disable a neighbor configuration rather than delete it.

Instructions: Set to Disable if you want to disable this neighbor configuration. Or, set to Enable if you previously disabled this neighbor configuration and now wish to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.2

Parameter: Priority

Default: 1

Range: 0 to 255

Function: Indicates the priority of this neighbor, with 255 indicating the highest priority. The neighbor priority value is used in multi-access networks for the election of the designated router. If this parameter is set to 0, this router is not eligible to become the designated router on this particular network.

Instructions: Either accept the default neighbor Priority value of 1, or enter some other value between 0 and 255.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.9

Deleting a Neighbor

To delete a neighbor from an NBMA interface, begin at the OSPF Interface List window shown in Figure 4-11 and complete the following steps:

1. Select the interface from which you want to delete a neighbor.
2. Click on Neighbors.

The OSPF Neighbor List window appears (see Figure 4-12).

3. Click on the neighbor that you want to delete.
4. Click on Delete.

The OSPF neighbor interface is removed from the list.

5. Click on Done to save your changes and exit the window.

Configuring OSPF Virtual Interfaces

To add, edit, or delete OSPF virtual interfaces, begin at the Configuration Manager window and proceed as follows:

1. Select Protocols→IP→OSPF→Virtual Interfaces.

The OSPF Virtual Interface List window appears (Figure 4-14).

2. Perform any of the functions listed below, which are described in the following sections.
 - To add a virtual interface, see “Adding a Virtual Interface” on page 4-56.
 - To edit a virtual interface, see “Editing a Virtual Interface” on page 4-58.

- To delete a virtual interface, see “Deleting a Virtual Interface” on page 4-63.

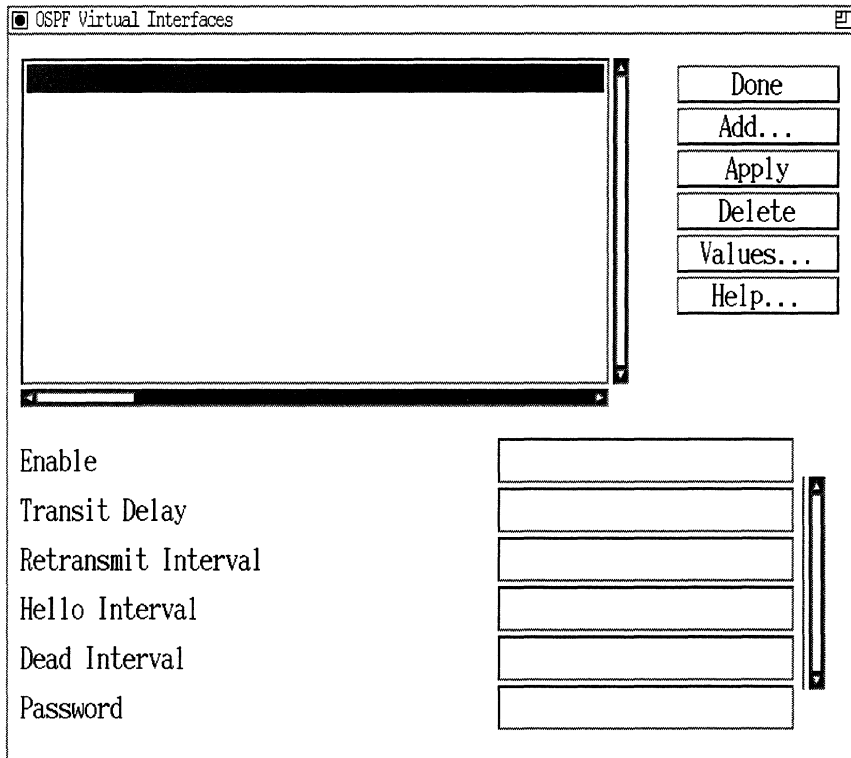


Figure 4-14. OSPF Virtual Interface List Window

Adding a Virtual Interface

To add a virtual interface, begin at the OSPF Virtual Interface List window shown in Figure 4-14 and complete the following steps:

1. Click on Add.

The OSPF Virtual Interface Configuration window appears (see Figure 4-15).

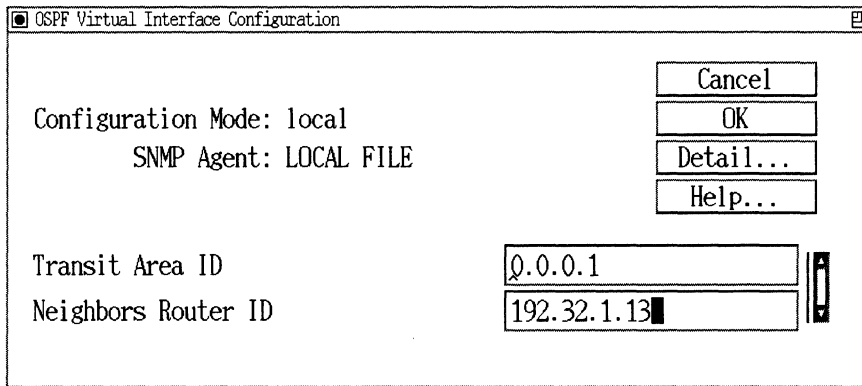


Figure 4-15. OSPF Virtual Interface Configuration Window

2. Specify the Transit Area ID and Neighbors Router ID parameters.

The Transit Area ID and Neighbors Router ID parameters are described following these instructions.

3. Click on OK.
4. Click on Done to save your changes and exit the window.

Parameter: Transit Area ID

Default: None

Options: Any area ID

Function: Identifies the transit area through which this virtual link is configured.

Instructions: Enter the appropriate area ID in dotted decimal notation. The transit area must contain the neighboring router identified in the Neighbors Router parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.6.1.4

| | |
|-------------------|---|
| Parameter: | Neighbors Router ID |
| Default: | None |
| Options: | Any IP address |
| Function: | Identifies the interface at the other end of this virtual link. |
| Instructions: | Enter the appropriate IP address. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.7.1.5 |

Editing a Virtual Interface

To edit the default parameters for a virtual interface, begin at the OSPF Virtual Interface List window shown in Figure 4-14 and complete these steps:

1. Click on the virtual interface that you want to edit.
2. Edit those parameters that you want to change.

All OSPF virtual interface parameters that you can edit are described following these instructions.

3. Click on Apply to implement your changes.
4. Click on Done to save your changes and exit the window.

Note: When you reconfigure a virtual interface in dynamic mode, OSPF restarts on that interface. The only exception to this rule is when you change the Hello Interval, Retransmit Interval, or Dead Interval timers.

OSPF Virtual Interface Parameter Descriptions

This section describes how to set all virtual interface parameters that you can edit.

| | |
|-------------------|--|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Enables or disables this virtual link. This parameter is useful when you want to temporarily disable a virtual link rather than delete it. |
| Instructions: | Set to Disable to turn off this virtual link. Or, set to Enable if you previously disabled this virtual link and now wish to re-enable it. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.7.1.2 |

| | |
|-------------------|--|
| Parameter: | Transit Delay |
| Default: | 1 second |
| Range: | 1 to 360 seconds |
| Function: | Indicates the estimated number of seconds it takes to transmit a link state update packet over this interface. |
| Instructions: | Either accept the default value of 1 second, or enter a new value between 1 and 360 seconds. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.7.1.6 |

Parameter: Retransmit Interval

Default: 5 seconds

Range: 1 to 360 seconds

Function: Indicates the number of seconds between link state advertisement retransmissions for adjacencies belonging to this interface. The Retransmit Interval value is also used when retransmitting database description and link state request packets. This value should be well over the expected round trip time. Although the default value is 10, we suggest the following values for Retransmit Interval.

| Network Type | Suggested Retransmit Interval |
|--------------------------|--------------------------------------|
| Broadcast | 10 seconds |
| Point-to-point | 15 seconds |
| NBMA | 15 seconds |
| Point-to-mul- tipoint | 15 seconds |

Instructions: Either accept the default value of 10 seconds, or set the retransmit interval to some other value between 1 and 360 seconds.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.7

Parameter: Hello Interval**Default:** 15 seconds**Range:** 1 to 360 seconds**Function:** Indicates the number of seconds between the hello packets that the router sends on the interface. Although the default value is 15 seconds, we suggest the following values for Hello Interval.

| Network Type | Suggested Hello Interval |
|---------------------|---------------------------------|
| Broadcast | 10 seconds |
| Point-to-point | 15 seconds |
| NBMA | 20 seconds |
| Point-to-multipoint | 15 seconds |

Instructions: Either accept the default value of 15 seconds, or set the hello interval to some other value between 1 and 360 seconds.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.3.7.1.8

Note: The Hello Interval value must be the same for the virtual neighbor and for all routers attached to the same network.

Parameter: Dead Interval

Range: 60 seconds

Options: 1 to 2000 seconds

Function: Indicates the number of seconds that a router's hello packets have not been seen before its neighbors declare the router down. The Dead Interval value should be some multiple of the Hello Interval. Although the default value is 60 seconds, we suggest the following values for Dead Interval.

| Network Type | Suggested Dead Interval |
|---------------------|-------------------------|
| Broadcast | 40 seconds |
| Point-to-point | 60 seconds |
| NBMA | 80 seconds |
| Point-to-multipoint | 60 seconds |

Instructions: Either accept the default value of 60 seconds, or enter some other value for Dead Interval.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.3.7.1.9

Note: The Dead Interval value must be the same for all routers attached to the same network.

| | |
|-------------------|--|
| Parameter: | Password |
| Default: | None |
| Options: | Any ASCII text string up to eight characters long. |
| Function: | Specifies the password used for this area. You can specify a password up to eight ASCII characters in length that will appear in the authentication field of all OSPF packets across this interface. Password is valid only when Authentication Type is set to Simplepassword. |
| Instructions: | Enter the appropriate password. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.3.7.1.10 |

Note: All routes in the same area must either have no authentication, or have the same password.

Deleting a Virtual Interface

To delete a virtual interface, begin at the OSPF Virtual Interface List window shown in Figure 4-14 and complete the following steps.

1. Select the virtual interface that you want to delete.
2. Click on Delete.

The virtual interface no longer appears on the OSPF Virtual Interfaces window.

3. Click on Done to save your changes and exit the window.



Chapter 5

Customizing BGP Services

This chapter describes the Bay Networks implementation of the Border Gateway Protocol (BGP) and shows you how to edit parameters for BGP Version 3 and Version 4 (BGP-3 and BGP-4). The chapter contains the following sections:

- “BGP Features” on page 5-1
- “Establishing a Peer-to-Peer Connection” on page 5-4
- “BGP Messages” on page 5-6
- “How BGP Selects the Best Path” on page 5-14
- “OSPF/BGP Interaction” on page 5-17
- “Using IBGP in a Transit AS” on page 5-18
- “Using IBGP in Intra-AS Routing” on page 5-19
- “Configuring BGP Message Logging” on page 5-20
- “Editing BGP Parameters” on page 5-22

BGP Features

BGP is an exterior gateway protocol primarily used to exchange network reachability information with other BGP systems in other autonomous systems. (Autonomous systems and interior and exterior gateway protocols are described in Chapter 1.)

Figure 5-1 shows two autonomous systems: AS1 and AS2. Networks within AS 1 and AS2 are connected by routers running an interior gateway protocol — in this case, OSPF. AS 1 and AS 2 are connected by routers that run an exterior gateway protocol — BGP — in addition to OSPF.

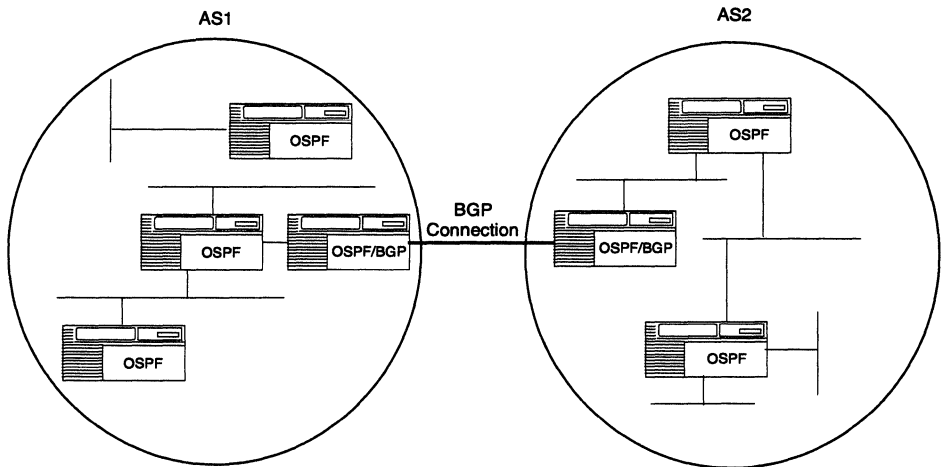


Figure 5-1. BGP Connection between Two Autonomous Systems Running OSPF

A BGP router employs a BGP speaker, which is an entity within the router that transmits and receives BGP messages and acts upon them. BGP routers form neighbor relationships with other BGP routers. BGP runs over the LAN and WAN media/protocols that IP runs over; this includes Ethernet, Token Ring, Sync, Wellfleet Proprietary Sync, Frame Relay, SMDS, X25 (DDN, PDN, Pt-to-Pt), ATM PVC, FDDI, T1, E1, HSSI, and PPP.

An autonomous system can include one or more BGP speakers that provide external route information for the networks within the AS. An AS containing a single BGP speaker with a single external BGP connection is a stub AS. The BGP speaker is providing external route information for the networks contained within its AS only.

BGP features include**□ TCP support**

The neighbors communicate over a reliable transport layer connection (TCP), so that BGP can assume that its communication with other BGP routers is reliable. This eliminates the need to implement the update, retransmission, acknowledgment and sequencing that are necessary with EGP.

□ Multiple path elimination

A BGP speaker can announce only routes that it actually uses. Therefore, a border router that learns multiple paths to an external destination must choose only one of those routes for further advertisement into the AS or to other BGP peers.

□ Authentication

BGP provides support for multiple authentication schemes. A scheme is identified in the Open message and each subsequent message on that TCP connection must contain a marker field that complies with the scheme. However, only the default authentication scheme (none) has been developed at this time.

□ AS Path attribute

Each BGP route contains a list of the autonomous systems that it has traversed. This allows a BGP speaker to eliminate looped routes. If a BGP speaker sees its own AS listed in a route, then there is a loop, and the route is not used.

□ Routing policy support

Each routing update contains information upon which hop-by-hop policies can be applied. For example, policies can be defined based on the information contained in a route's AS_PATH attribute. BGP can favor routes based on AS count, or the presence of a certain AS in the path. Conversely, it can also avoid routes that contain a certain AS in the path, or that originate in a certain AS.

Bay Networks supports two versions of BGP, BGP-3 and BGP-4.

BGP-3 assumes that each advertised network is a natural class network (A, B, or C) based on its high-order bits. BGP-3 cannot advertise subnets or supernets. In contrast, 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. This allows for both supernet and subnet advertisement. The supernet advertisement is what makes classless inter-domain routing (CIDR) deployment possible.

When BGP peers establish communications, they negotiate the version of BGP that they use to exchange routing information. If you add both BGP-3 and BGP-4 to an IP circuit, the router first attempts to use BGP-4. If the BGP peer is not a BGP-4 speaker, the router uses BGP-3.

The network administrator can control the way the router negotiates the BGP version with a BGP peer. For instructions, see the Min BGP Version parameter on page 5-38 and the Max BGP Version parameter on page 5-38.

Establishing a Peer-to-Peer Connection

A BGP speaker forms neighbor relationships with other BGP speakers. This happens when a BGP speaker establishes a TCP connection to a BGP peer (which is simply the BGP speaker at the other end of the connection), based on local configuration information.

A BGP speaker that wants to initiate peer-to-peer connections periodically issues an Open message. BGP speakers respond to connection requests by returning an Open message. In Figure 5-2, for example, BGP Speaker A sends an Open message to BGP Speaker B to request a connection; BGP Speaker B responds by sending an Open message to BGP Speaker A.

All BGP speakers respond to connection requests from other speakers. Site Manager allows you to specify whether BGP also issues connection requests and, if so, how frequently. (See Connect Retry Timer parameter on page 5-40.)

BGP speakers use the exchange of Open messages to negotiate the characteristics of the peer-to-peer connection — for example, to determine whether the speakers will use BGP-3 or BGP-4.

Once Open messages have been exchanged, each speaker then sends a Keepalive message to confirm the BGP connection. A neighbor relationship now exists between the two BGP peers. BGP peers periodically issue a Keepalive message to maintain the connection. Site Manager allows you to specify how often the BGP speaker issues a Keepalive message on a peer-to-peer connection.

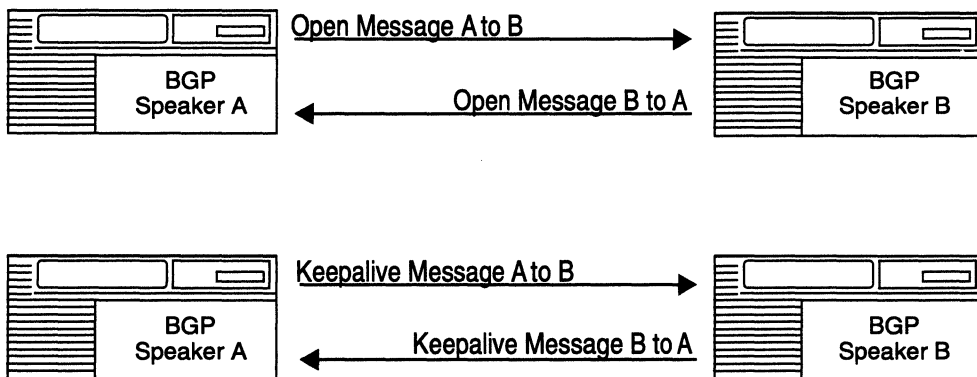


Figure 5-2. Establishing and Confirming a Connection between BGP Peers

Once a connection is established, the BGP speaker uses one or more Update messages to send the entire IP routing table (compliant to local BGP export policies). BGP, however, does *not* require the entire routing table to be sent again. Therefore, the BGP speaker must keep a current version of the routing information received from of all of its peers for as long as the connection to each peer is valid. This information will be updated via Update messages whenever changes occur.

Site Manager allows you to specify how long the BGP speaker waits for an Update message (or a Keepalive message) before terminating the connection. (See BGP Interval Timer parameter on page 5-26 and Holdtime parameter on page 5-41.)

If a condition occurs that causes a BGP speaker to terminate a peer-to-peer connection, the BGP speaker issues a Notification message, specifying the reason. The connection is immediately terminated.

For more information on BGP messages, see “BGP Messages” on page 5-6.

For complete information on how to configure a router to establish, maintain, and terminate peer-to-peer connections and send and receive Update messages, see “Configuring a BGP Peer Relationship” on page 5-32.

BGP Messages

BGP uses four different message types: Open, Keepalive, Update, and Notification.

All of the messages share a common BGP message header made up of the following three fields:

□ **Marker**

This field is used for authentication. Currently, this field can only be set to all 1s, specifying the null authentication scheme, the only authentication scheme yet defined.

□ **Length**

This field indicates the total length of the message in octets. The value of this field must be between 19 octets (header only message) and 4096 octets.

□ **Type**

This field indicates the type of message:

- 1 — Open message
- 2 — Update message
- 3 — Notification message
- 4 — Keepalive message

Open Message

The Open message is used to establish a BGP connection between two BGP speakers. In addition to the message header, the Open message includes the following fields:

- ❑ The BGP Version. Bay Networks currently supports BGP Version 3 and Version 4.
- ❑ The AS Number, which provides the autonomous system number of the transmitting BGP speaker.
- ❑ The Holdtime, which indicates the maximum number of seconds that can elapse between the receipt of a Keepalive and/or Notification and/or Update message. If this timer expires, the receiver assumes the connection is down.

Over switched virtual circuits, a Holdtime of zero can be used. This disables the transmission of periodic Keepalive messages on the connection so that the virtual circuit can go idle. While this can result in cost savings for line usage, it also makes it more difficult for either BGP speaker to determine, in a timely manner, if the BGP connection has gone down.

- ❑ The BGP Identifier, which is the IP address of the transmitting BGP speaker. A BGP speaker sets the BGP identifier to the IP address of one of its interfaces. It uses the same identifier in all the Open messages on every connection.

Note: If OSPF is running on the same router, and BGP-3 routes are advertised as OSPF external routes, the BGP identifier and the OSPF router ID must be identical.

- ❑ The Authentication Code, which indicates the authentication mechanism in use. Currently, only the null authentication mechanism is defined. Therefore, this field must be set to 0.
- ❑ Authentication Data is a field with variable length and contents depending on the value of the Authentication Code field. Currently, this field must be nil, because only the null authentication mechanism is defined.

Keepalive Message

The Keepalive message has two functions. First, it is used as a confirmation to the Open message when a connection between two BGP speakers is being established. Second, it is used to keep the hold timer from expiring and the connection from going down when there has been no other BGP message sent over the connection for a while. The periodic transmission of the Keepalive message is regulated by the Keepalive timer.

A Keepalive message consists only of the common BGP header.

Update Message

The Update message is used to transfer current routing information between BGP peers. It describes routes from the transmitting BGP speaker to some number of destination networks. Each destination is listed, and the path to the set of destinations is described using path attributes.

When a BGP speaker receives a route in an Update message, it applies any local routing policies to determine whether the router will use the route and whether it will propagate the route to other routers. Then, if the route can be used, it is compared against routes from other protocols and possibly included in the forwarding table.

BGP-3 Update Message Format

In addition to the message header, a BGP-3 Update message includes the following fields:

- **Total Path Attributes Length.** This field indicates the total length of the path attributes field.
- **Path Attributes.** This field is a variable-length sequence of path attributes. Each attribute entry consists of an attribute value and a field describing the attribute. Table 5-1 lists the mandatory and optional BGP-3 path attributes.

Table 5-1. BGP-3 Path Attributes

| Attribute | Description |
|--------------------|--|
| AS Path | Mandatory attribute containing a list of the ASs that must be traversed to reach the given destinations. |
| Origin | Mandatory attribute containing one of the following values: IGP (the path is valid all the way to the IGP of the originating AS), EGP (the path was advertised using EGP by the last AS in the AS path), or Incomplete (the path is valid only to the last AS in the AS Path). |
| Next Hop | Mandatory attribute that defines the IP address of the router to use as a next hop for the advertised destinations. |
| Inter-AS attribute | Optional attribute used to choose between paths to the destinations listed. |
| Unreachable | Discretionary attribute used to indicate destinations that have become unreachable. |

- The Networks. This field indicates the destinations being described by the path attributes.

You set values in BGP-3 accept and announce policy parameters to match and, in some cases, override the attribute values contained in inbound and outbound update messages.

For details about BGP-3 accept policy parameters, see

- Originating AS parameter on page 9-17
- Route Origin parameter on page 9-17

For details about BGP-3 announce policy parameters, see

- ❑ Inter-AS Metric Selector parameter on page 9-50
- ❑ Specific Inter-AS Metric parameter on page 9-50
- ❑ Origin parameter on page 9-51
- ❑ AS Path Override parameter on page 9-51
- ❑ Next Hop parameter on page 9-52

Note: For BGP-3, only natural class networks or the default route (0.0.0.0) can be advertised. BGP-3 assumes that each advertised network is a natural class network (A, B, or C) based on its high-order bits. It cannot advertise subnets or supernets.

BGP-4 Update Message Format

The BGP-4 update message has the same format and contains the same mandatory attributes as the BGP-3 update message with the following additions.

In place of the Unreachable attribute that BGP-3 includes as part of the path attribute description, the BGP-4 update includes an Unreachable field. This field specifies destinations that have become unreachable.

In place of the BGP-3 optional attributes, a BGP-4 update message can include the optional attributes described in Table 5-2 .

Table 5-2. BGP-4 Optional Path Attributes

| Attribute | Description |
|--------------------------|---|
| Multi-Exit Discriminator | Optional attribute used to choose between paths to the destinations listed. |
| Local Preference | Optional attribute allowing AS border routers to indicate the preference they have assigned to a chosen route when advertising it to IBGP peers. |
| Atomic Aggregate | Optional attribute used to ensure that certain network layer reachability information (NLRI) is not deaggregated. |
| Aggregator | Optional attribute identifying which AS performed the most recent route aggregation. The 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. |

You set values in BGP-4 accept and announce policy parameters to match and, in some cases, override the attribute values contained in inbound and outbound update messages.

For details about BGP-4 accept policy parameters, see

- ❑ Originating AS parameter on page 9-21
- ❑ Route Origin parameter on page 9-21
- ❑ Aggregator AS List parameter on page 9-22,
- ❑ Aggregator Router List parameter on page 9-22
- ❑ Local Preference parameter on page 9-23

For details about BGP-4 announce policy parameters, see

- ❑ Multi-Exit Discriminator parameter on page 9-55
- ❑ Multi-Exit Discriminator Value parameter on page 9-56,
- ❑ Origin parameter on page 9-56,
- ❑ AS Path parameter on page 9-57,
- ❑ Local Preference Override parameter on page 9-57
- ❑ Local Preference Value parameter on page 9-58
- ❑ Next Hop parameter on page 9-58,
- ❑ Atomic parameter on page 9-59

Notification Message

The notification message is sent whenever a condition is detected that causes a BGP speaker to terminate a connection. The BGP connection is closed after the notification is transmitted. In addition to the message header, the Notification message includes the following fields:

- ❑ The Error Code, which indicates the type of notification.
- ❑ The Error Subcode, which further specifies the reported error conditions.
- ❑ Error codes and their associated subcodes are described in Table 5-3.

Table 5-3. Notification Message Error Codes and Subcodes

| Error Code | Associated Error Subcode |
|---------------------------------------|---|
| Message Header Error (1) | (1) Connection not synchronized (2) Bad Message Length (3) Bad Message Type |
| Open Message Error (2) | (1) Unsupported version number (2) Bad Peer AS (3) Bad BGP Identifier (4) Unsupported Authentication Code (5) Authentication Failure (6) Unacceptable Hold Time |
| Update Message Error (3) | (1) Malformed attribute list (2) Unrecognized well-known attribute (3) Missing well-known attribute (4) Attribute flags error (5) Attribute length error (6) Invalid Origin attribute (7) AS routing loop (8) Invalid Next Hop attribute (9) Optional attribute error (10) Invalid Network field (11) Malformed AS_PATH |
| Hold Timer Expired (4) | No subcodes |
| Finite State Machine Error (5) | No subcodes |
| Cease (6) | No subcodes |

How BGP Selects the Best Path

A BGP speaker must, at times, evaluate and compare different paths to a destination network to determine the best path. Because all border routers must provide the same view of the AS to external ASs, having a selection strategy that is consistent in the router, and which can be consistent across all border routers, is very important in BGP. To select the best available path, BGP uses AS weights and classes and IP policies. To compare IBGP routes, BGP-4 can also calculate and use a Local Preference value. These mechanisms are described in the following sections.

AS Weight and Class Values

You can assign a weight value to any AS number. An assigned weight can range from 1 to 15 plus an infinity value. Weights provide a way either to prefer or to avoid routes that pass through certain ASs. The weights of each AS in a path are added, and the path with the smallest total weight is the preferred path. Any path containing an AS weight of infinity will be avoided.

When a BGP router receives a new route, it is evaluated against any existing accept policies. If after this evaluation, the path still is to be used, the total weight of the path is calculated.

AS weights should be configured the same on all BGP routers in an AS.

AS weight classes allow a network administrator to assign multiple weight values to the same AS. This feature allows the administrator to consider an AS path differently for different networks. For example, consider a situation in which two networks — 192.32.1.0 and 192.32.2.0 — are both reachable by two paths. The first path to each network shares a common AS — AS 5. The second path to each network also shares a common AS — AS 10. If the administrator for some reason wants to favor AS 5 in the path to 192.32.1.0 and AS 10 in the path to 192.32.2.0, he or she can assign one AS weight class to 192.32.1.0 and another weight class to 192.32.2.0.

For instructions on assigning weight and class values to an autonomous system, see “Configuring BGP AS Weights and Weight Classes” on page 5-45

Routing Policies

BGP accept and announce policies complement the AS weight feature to provide a mechanism to configure hop-by-hop routing policies. These policies govern which routes are used by a router, and which are propagated to other routers.

Note: By default, an external BGP-3 or BGP-4 speaker will neither advertise any routes to a peer, nor inject any routes into its IGP. Route policies must be configured to enable any route advertisement.

So that every BGP border router within an AS comes to the same decision in constructing path attributes for an external path, route policies must be coordinated between all of the BGP speakers within an AS. It is suggested that the accept and announce policies on all IBGP connections accept and propagate all routes. On external BGP connections, consistent routing policy decisions should be made.

Note: In addition to announce and accept policies, Bay Networks supports import and export filters for BGP-3. Import and export filters provide a subset of the parameters provided by the policies. In a future release, support for import and export filters will be dropped.

Calculating the BGP-4 Local Preference Attribute

BGP-4 update messages include a Local Preference attribute that allows an AS border router to assign a preference value to a route when advertising it to IBGP peers. The calculation of the Local Preference attribute is implementation-specific. A higher value indicates that the route is more preferred.

You can configure a BGP-4 accept or announce policy to override the value in the Local Preference attribute. For details and instructions, see the Accept Local Preference parameter on page 9-23 and the Announce Local Preference Override parameter on page 9-57.

The router uses the following equations to calculate a value for the Local Preference attribute:

$$\text{local preference} = 8191 - \text{origin value} - \text{AS path weight}$$

where *origin value* is 0 for routes with an Origin Path attribute of IGP and 4096 otherwise and *AS path weight* is a sum of weight values associated with AS numbers listed in the route's AS Path attribute. These weight values can be configured and default to 8.

A steep penalty is applied to routes that are advertised with an ORIGIN attribute other than IGP — that is, EGP or Incomplete.

For an OSPF internal route or a direct route, the Local Preference attribute is set to

$$\text{local preference} = (8191 + 256 - (\text{metric} \& 255))$$

where *metric* is the OSPF metric for an OSPF route or the configured cost for a direct route.

For a RIP route, an EGP route, an OSPF ASE route, or a static route, the local preference attribute is set to

$$\text{local preference} = (256 - \text{metric})$$

where *metric* is the RIP metric for a RIP route, the EGP metric for an EGP route, the OSPF metric for an OSPF ASE route, or the configured cost for a static route.

Note that Local Preference values for OSPF internal routes and direct routes are higher than the Local Preference values calculated for BGP routes.

Best Route Calculation for Equal Routes

The following eight rules (tiebreakers) are used to choose between two equal BGP routes:

1. Choose the route with the lower route weight.
2. Choose the route with the higher Local Preference attribute.
3. Choose the route with the lower Inter-AS Metric attribute (if both routes include this optional attribute).
4. Choose the route with the lower interior cost to the Next Hop.
5. Choose external BGP over IBGP.
6. Choose the route with the lower BGP identifier
7. Choose the route with the lower BGP connection remote address.
8. Choose the route with the lower BGP connection local address.

OSPF/BGP Interaction

RFC 1403 recommends certain interaction when OSPF is the IGP within an autonomous system. For routers running both protocols, the OSPF router ID and the BGP Identifier must be an IP address and must be identical. A route policy must be configured in order to allow BGP advertisement of OSPF routes.

For more information, see “Generating an External Route Tag for OSPF/BGP Interaction” on page 4-15.

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.

Using IBGP in a Transit AS

If an AS has more than one BGP speaker, it can provide transit service between multiple networks outside the AS. An AS that provides such a service for BGP speakers is known as a transit AS (see Figure 5-3).

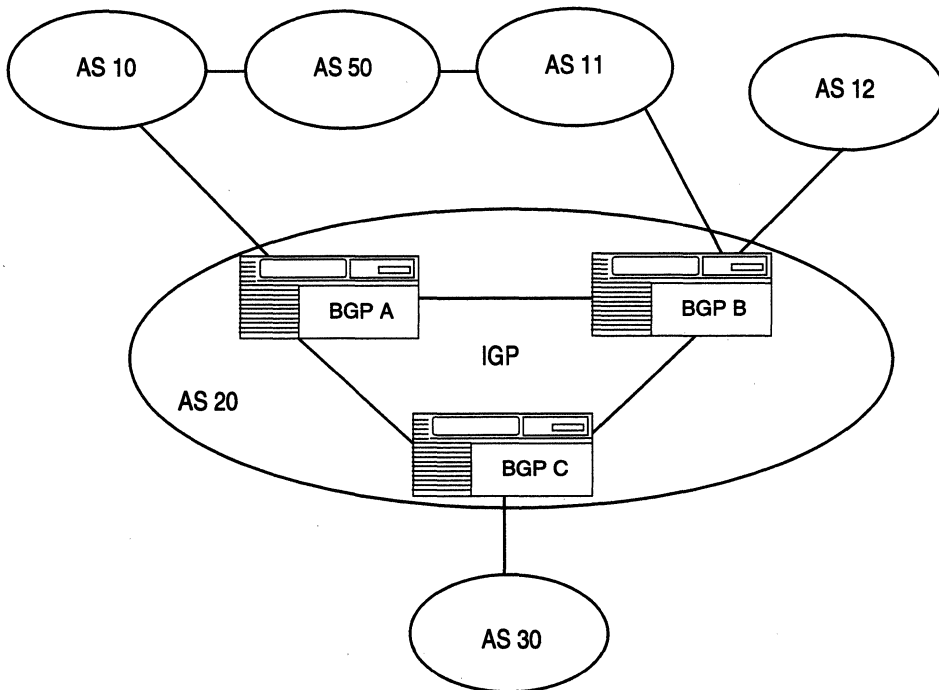


Figure 5-3. Transit Autonomous System

It is important that there is a consistent view of routing within the transit AS. This view is provided by whatever IGP the AS is running. It is also important that routes exterior to the AS are consistent. This can be accomplished by having all of the BGP speakers within the AS that connect to exterior ASes maintain direct connections with each other. This is known as internal BGP (IBGP). The speakers then agree upon which border routers will serve as exit/entry points for particular networks outside the AS. All internal routers must be updated with

this transit information before transit service is advertised to other ASs.

In Figure 5-3, Autonomous System 20 is the transit AS. It is providing information about its internal networks, as well as transit networks, to the remaining ASs. The IBGP connections between BGP routers A, B, and C are necessary to provide consistent information to the ASs.

When setting up IBGP connections, consider using a circuitless IP interface. Doing so separates the connection endpoint from any of the physical interfaces on the router. For example, if the BGP connection is configured over a physical interface and that interface becomes disabled, the IBGP connection will become disabled as well. However, if you configure the BGP connection on a circuitless interface, then as long as there is a valid path between the peer routers, the BGP connection will stay alive.

Using IBGP in Intra-AS Routing

Because situations will arise where OSPF is not the IGP within some autonomous systems, and because BGP does not interact well with IGP protocols other than OSPF, Bay Networks implements IBGP Intra-AS Routing.

With IBGP intra-AS routing, an AS need not propagate BGP routes into the AS. Instead, all routers in the AS must run IBGP to each border router. The IBGP information is used in conjunction with the IGP route to the authoring BGP border router to determine the next hop to use for external networks.

No BGP information is carried by the IGP. Each router uses IBGP exclusively to determine reachability to external networks. When an IBGP update for a network is received, it can be passed on to IP for inclusion in the forwarding tables only if a viable IGP route to the correct border gateway is available.

For instructions on configuring a router for IBGP intra-AS routing, see the IBGP Intra AS Routing parameter on page 5-25 and the From Protocols parameter on page 5-26.

Configuring BGP Message Logging

Site Manager allows you to control the event messages that BGP sends to the log file by specifying

- Local and remote address of a peer-to-peer session or sessions
- Message severity level: fault, warning, information, trace, or debug, or all levels
- BGP message type: Open, Update, Notification, or Keepalive

Use BGP message logging parameters to limit the volume of debug-level messages that BGP generates and logs. If you allow BGP to log all debug-level events, the messages that BGP generates will quickly overrun and overwrite the log file.

For instructions on configuring BGP message logging on the router, see “Generating BGP Event Messages” on page 5-53.

For More Information about BGP

For more information about BGP, refer to the following documentation:

Lougheed, K. and Rekhter, Y. “A Border Gateway Protocol 3.” RFC 1267, Network Information Center (NIC), SRI International, Menlo Park, California, October 1991.

Perlman, Radia. *Interconnections: Bridges and Routers*. Reading, Massachusetts: Addison-Wesley Publishing Company, 1992.

Rekhter, Y. “Application of the Border Gateway Protocol in the Internet.” RFC 1268, Network Information Center (NIC), SRI International, Menlo Park, California, October 1991.

Varadhan, K. “BGP OSPF Interaction.” RFC 1364, Network Information Center (NIC), SRI International, Menlo Park, California, September 1992.

Willis, S. and Burruss, J. "Definition of Managed Objects for the Border Gateway Protocol (Version 3)." RFC 1269, Network Information Center (NIC), SRI International, Menlo Park, California, October 1991.

BGP Implementation Notes

This section provides you with some guidelines that you should follow when you configure BGP. If you do not follow these guidelines, BGP will either 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 multi-homed AS (one that contains more than one exit point), we strongly encourage you to use OSPF for your IGP and BGP for your sole exterior gateway protocol, or use intra-AS IBGP routing.

If OSPF is the IGP, you should also use the default OSPF tag construction. Using EGP or modifying the OSPF tags makes network administration and proper construction of BGP path attributes more difficult.

- ❑ For any router supporting both BGP and OSPF, the OSPF router ID and the BGP identifier must be the same.

Editing BGP Parameters

Note: The instructions in this section assume that you have already configured at least one IP interface with BGP support (BGP interface). If you have *not* yet configured a BGP interface, or want to add additional BGP interfaces, see *Configuring Wellfleet Routers* for instructions.

You access all BGP parameters from the Configuration Manager window. Refer to *Configuring Wellfleet Routers* for instructions on accessing this window.

The following sections show you how to configure BGP parameters:

- “Editing BGP Global Parameters” on page 5-23
- “Editing BGP-3 Global Parameters” on page 5-29
- “Editing BGP-4 Global Parameters” on page 5-31
- “Configuring a BGP Peer Relationship” on page 5-32
- “Configuring BGP AS Weights and Weight Classes” on page 5-45
- “Generating BGP Event Messages” on page 5-53
- “Deleting BGP from the Router” on page 5-57

Each BGP parameter description includes the default setting, all valid setting options, the parameter function, and instructions for setting the parameter.

Editing BGP Global Parameters

When you edit the BGP global parameters, you are editing parameters that affect BGP on the entire router.

To edit BGP global parameters, begin at the Configuration Manager window and complete the following steps:

1. Select Protocols→IP→BGP→BGP Global.

The Edit BGP Global Parameters window appears (Figure 5-4).

2. Edit those parameters you want to change. BGP global parameters are described following these instructions.
3. Click on OK to save your changes and exit the window.

| Parameter | Value |
|---------------------------|---------|
| BGP Enable | ENABLE |
| BGP Identifier | 1.1.1.1 |
| BGP Local As | 1 |
| BGP Intra-AS | ENABLE |
| BGP From Protocols | BGP |
| BGP Interval Timer | 5 |
| BGP Collision Detect | ENABLE |
| Multi-hop Ebgp Connection | DISABLE |

Figure 5-4. Edit BGP Global Parameters Window

BGP Global Parameter Descriptions

This section describes how to set all BGP global parameters.

| | |
|-------------------|---|
| Parameter: | BGP Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Globally enables or disables BGP on all router interfaces. |
| Instructions: | Set to Disable if you want to disable BGP for the entire router. Set to Enable if you previously disabled BGP and now want to re-enable it. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.5.1.1.2 |
| | |
| Parameter: | BGP Identifier |
| Default: | None |
| Options: | An IP address of an IP interface on this router |
| Function: | Identifies the BGP router. There is no default for this parameter. You must use an IP address of one the router's IP interfaces. |
| Instructions: | Either accept the current BGP identifier or enter a new IP address. The BGP identifier must be one the router's IP interfaces. If both BGP and OSPF are running on the router, then the BGP identifier must be identical to the OSPF router ID. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.5.1.1.4 |

Parameter: BGP Local AS

Default: None

Range: 1 to 65535

Function: Identifies the autonomous system to which this BGP router belongs.

Instructions: Either accept the current BGP Local AS value or enter a new value for this parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.5

Parameter: IBGP Intra AS Routing

Default: Enable

Options: Enable | Disable

Function: Specifies whether BGP will perform intra-AS IBGP routing.

Instructions: Either accept the default value for Intra AS Routing or set to a different value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.8

Parameter: From Protocols

Default: BGP

Options: BGP | ALL

Function: Controls (if Intra-AS routing is enabled) the types of routes that BGP advertises in any IBGP sessions.

Instructions: Select BGP to propagate only advertised routes learned from external BGP peers. Select ALL to propagate routes learned from all route sources (excluding IBGP and OSPF inter-area and intra-area routes, which are never advertised with IBGP).

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.9

Parameter: BGP Interval Timer

Default: 5 seconds

Range: 1 to 2147483647

Function: Specifies the minimum time interval, in seconds, between injections of external BGP routes into the IP routing table.

Instructions: Accept the default or enter a nonzero value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.10

Parameter: Collision Detect**Default:** Enable**Options:** Enable | Disable**Function:** Specifies whether redundant BGP connections to the same router will be detected and disallowed.**Instructions:** If you want only one BGP connection to the same router to be maintained, use the default. If you want to allow redundant connections, enter Disable.

Collision detection is based on router ID. If two BGP peers have multiple physical connections and want to establish a BGP session across each physical connection, you must disable this parameter. The advantage of a configuration with multiple physical connections is redundancy. The disadvantage is that such a configuration results in multiple copies of each route.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.16

Parameter: Multi-hop EBGW Connection

Default: Disable

Options: Enable | Disable

Function: Specifies whether BGP allows multihop connections to an external BGP peer.

Instructions: By default, BGP enforces the rule that requires an external BGP peer to be located on a directly attached network. Use this parameter to override the restriction.



Warning Enabling multihop BGP connections is dangerous since it can cause BGP speakers to establish a BGP connection that traverses a third-party AS, which may violate policy considerations and may also introduce forwarding loops.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.1.6

Editing BGP-3 Global Parameters

When you edit the BGP-3 global parameters, you are editing parameters that affect BGP-3 on the entire router.

To edit BGP-3 global parameters, begin at the Configuration Manager window and complete the following steps:

1. Select Protocols→IP→BGP→BGP-3 Global.

The Edit BGP-3 Global Parameters window appears (Figure 5-5).

2. Edit those parameters you want to change. The BGP-3 global parameters are described following these instructions.
3. Click on OK to save your changes and exit the window.

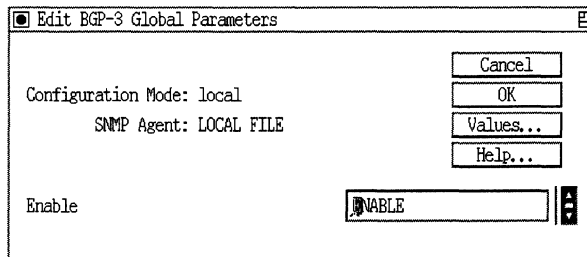


Figure 5-5. Edit BGP-3 Global Parameters Window

BGP-3 Global Parameter Descriptions

This section describes how to set all BGP-3 global parameters.

| | |
|-------------------|--|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Globally enables or disables BGP-3 on all router interfaces. |
| Instructions: | Set to Disable if you want to disable BGP-3 for the entire router. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.5.2.1.2 |

Editing BGP-4 Global Parameters

When you edit the BGP-4 global parameters, you are editing parameters that affect BGP-4 on the entire router.

To edit BGP-4 global parameters, begin at the Configuration Manager window and complete the following steps:

1. Select Protocols→IP→BGP→BGP-4 Global.

The BGP-4 Global Parameters window appears (Figure 5-6).

2. Edit those parameters you want to change. BGP-4 global parameters are described following these instructions.
3. Click on OK to save your changes and exit the window




Figure 5-6. BGP-4 Global Parameters

BGP-4 Global Parameter Descriptions

This section describes how to set all BGP-4 global parameters.

| | |
|-------------------|--|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Globally enables or disables BGP-4 on all router interfaces. |
| Instructions: | Set to Disable if you want to disable BGP-4 for the entire router. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.5.3.1.2 |

Configuring a BGP Peer Relationship

When you configure BGP peers, you are setting parameters that affect the formation of BGP peer relationships on a particular IP interface.

To configure BGP peers, begin at the Configuration Manager window and complete the following steps:

1. Select Protocols→IP→BGP→Peers.

The IP Interface List for BGP window appears (Figure 5-7). This window lists all IP interfaces on which you can enable BGP peers.

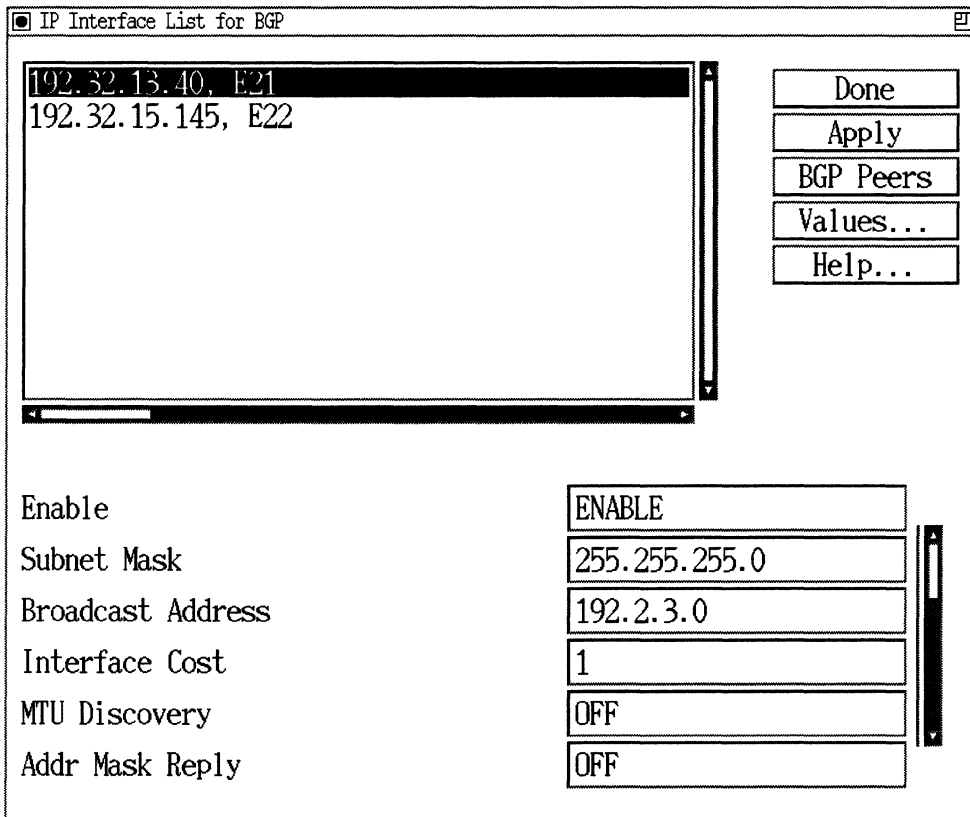


Figure 5-7. IP Interface List for BGP Window

2. Click on the IP interface for which you want to edit BGP peer parameters.
3. Click on BGP Peers.

The BGP Peer List window appears (Figure 5-8). It shows all of the neighbors configured for the IP interface that you selected in Step 2.

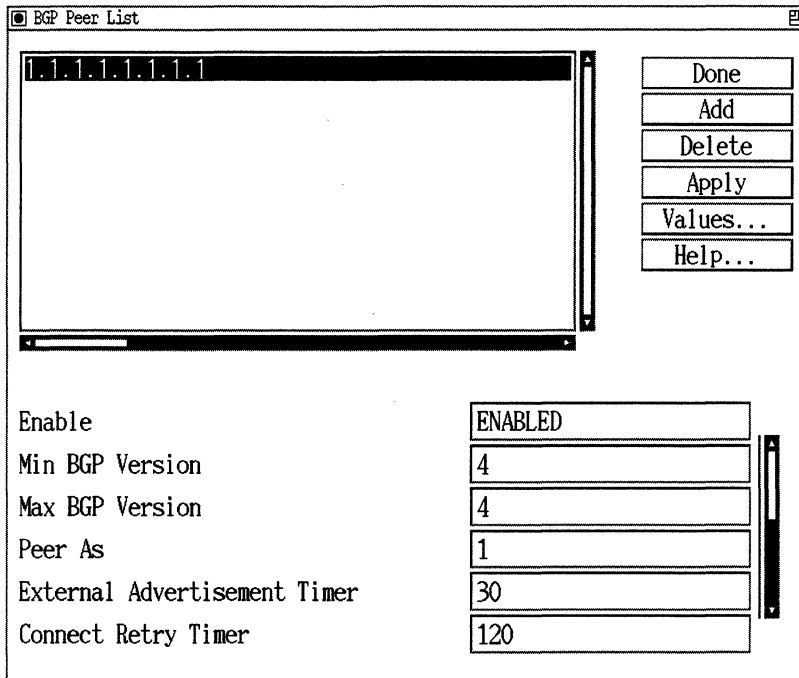


Figure 5-8. BGP Peer List Window

Add a BGP peer to the IP interface, edit parameters associated with a specific BGP neighbor, or delete a BGP peer from the IP interface as described in the following sections:

Adding a BGP Peer

To add a BGP peer to an IP interface, begin at the BGP Peer List window shown in Figure 5-8 and complete the following steps:

1. Click on Add.

The BGP Peer parameters window appears (Figure 5-9).

2. Set the BGP peer configuration parameters. The BGP peer configuration parameters are described following these instructions.
3. Click on OK.

The BGP Peer List window now lists the BGP peer you added. If you click on the peer, the default values for the rest of the peer parameters are shown at the bottom of the window (Enable, Min BGP Version, Max BGP Version, Remote AS, External Advertisement Timer, Connect Retry Timer, Holdtime, Keepalive Timer, and Path Attribute Table Switch). To edit these parameters, see “Editing a BGP Peer Relationship” on page 5-37.

BGP PEER

Configuration Mode: local
SNMP Agent: LOCAL FILE

Peer Address: 192.32.13

Peer AS: 1

Local Address: 192.32.13.40

Buttons: Cancel, OK, Values..., Help...

Figure 5-9. BGP Peer Parameters Window

BGP Peer Parameter Descriptions

This section describes how to set BGP peer configuration parameters shown on the BGP Peer Parameters window.

Parameter: Peer Address
Default: None
Options: Any IP address
Function: Specifies the IP address of the interface on the remote side of this BGP peer connection.
Instructions: Enter the IP address in dotted decimal notation. If the peer is in a remote AS, the address must be on the same subnet as the local interface.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.

Parameter: Peer AS
Default: None
Range: 1 to 65535
Function: Identifies the autonomous system to which the BGP router at the remote end of this BGP peer connection belongs.
Instructions: Enter the appropriate AS number.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.10

| | |
|-------------------|--|
| Parameter: | Local Address |
| Default: | None |
| Options: | Any IP Address |
| Function: | Specifies the IP address of the interface on the local side of this BGP peer connection. |
| Instructions: | Enter the appropriate address. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.5.1.2. |

Editing a BGP Peer Relationship

Note: You *cannot* reconfigure the Local Address or Peer Address for a BGP peer. To change these parameters, you must delete the peer and add a new peer with the proper information. See “Deleting a BGP Peer” on page 5-44 for instructions.

To edit a BGP peer, begin at the BGP Peer List window shown in Figure 5-8, and complete the following steps:

1. Click on the peer for which you want to edit parameters.

When you do this, all of the parameters shown at the bottom of the window will reflect the current values for the peer you selected.

2. Edit those parameters you want to change. The BGP peer parameters that you can edit are described following these instructions.

3. Click on Apply to implement your changes.

Repeat steps 1 through 3 to edit any other peers you want to change; remember to click on the Apply button each time.

4. Click on Done to exit the window.

Parameter: Enable

Default: Enable

Options: Enable | Disable

Function: Enables or disables a BGP peer relationship with the specified IP address.

Instructions: Set this parameter to Disable if you want to temporarily disable this peer relationship, rather than delete it. Or, set it to Enable if you previously disabled this peer relationship and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.2

Parameter: Min BGP Version

Default: 4

Options: 3 or 4

Function: Specifies the minimum acceptable BGP version to run on this peer connection.

Instructions: Specify BGP-3 or BGP-4.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.8

Parameter: Max BGP Version

Default: 4

Options: 3 or 4

Function: Specifies the maximum acceptable BGP version to run on this peer connection.

Instructions: Specify BGP-3 or BGP-4.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.9

| | |
|-------------------|--|
| Parameter: | Remote AS |
| Default: | None |
| Range: | 1 to 65535 |
| Function: | Identifies the autonomous system to which the BGP router at the remote end of this BGP peer connection belongs. |
| Instructions: | Either accept the current value or enter a new one. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.10 |
| | |
| Parameter: | External Advertisement Timer |
| Default: | 5 seconds |
| Range: | 1 to 2147483647 |
| Function: | Specifies the minimum number of seconds allowed between BGP updates for this peer connection. |
| Instructions: | Either accept the current External Advertisement Timer value, or enter a value greater than zero seconds. |
| | <p>The external advertisement interval controls how often the IP routing table is examined for changes. BGP update messages for routes that originate external to this AS will be issued no faster than the number of seconds you specify with this parameter.</p> |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.11 |

Parameter: Connect Retry Timer

Default: 120 seconds

Options: 0 to 2147483647

Function: Specifies the maximum number of seconds allowed between TCP connection attempts for this peer connection.

Instructions: Either accept the current Connect Retry Timer value or set this parameter to some value. A value of 0 indicates that no active attempt to establish a BGP connection to the peer is to be done. Incoming calls from the peers will be accepted.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.12

Parameter: Holdtime**Default:** 90 seconds**Options:** Zero, or any decimal number greater than 2**Function:** Specifies the holdtime that will be inserted into an Open message. Upon receipt of the peer's Open message, the lesser of the two hold times will be used (this must be at least 3 seconds). There are two exceptions:

- If one peer sends a zero Holdtime, then the non-zero Holdtime is used.
- If both peers send zero Holdtimes, then no Holdtime is used.

The calculated hold time is the amount of time the either peer will wait for a Keepalive or Update message before declaring the connection down.

Instructions: Either accept the current Holdtime Timer value or set the parameter to zero or some value greater than 2 seconds.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.13

Parameter: Keepalive Timer

Default: 30 seconds

Options: Any decimal number

Function: Specifies how often Keepalive messages will be sent across this peer connection.

If a hold time of zero is negotiated, no periodic Keepalive messages are sent. Otherwise, the Keepalive timer is set to the smaller of this configured value and one-third of the hold time.

Instructions: Either accept the current Keepalive value or set this parameter to some value greater than zero.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.15

Parameter: Min AS Origination Interval

Default: 15 seconds

Options: A value greater than 0

Function: Determines the minimum amount of time that must elapse between successive advertisements of Update messages that report changes within the advertising BGP speaker's own autonomous system.

Instructions: Enter a value greater than 0 seconds.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.30

Parameter: Local AS to Advertise to Peer**Default:** Null**Range:** 1 to 65535**Function:** Specifies the AS number that is sent in an open message to this peer.**Instructions:** Enter an AS number. To specify the AS number you set with the BGP Local AS parameter, use the default, null.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.31**Parameter: Max Update Size****Default:** 800 bytes**Range:** 64 to 4096 bytes**Function:** Specifies the maximum size (in bytes) of Update messages that are sent to this peer.**Instructions:** Use the default or specify a size. Note that, if the size of the Update message that is used to advertise a single route is greater than the configured message size, the actual message size can exceed the configured value.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.32

Parameter: Route Echo Switch

Default: Enable

Options: Enable | Disable

Function: Controls the way the router echoes a BGP route that is selected for forwarding. (Echoing in this case means advertising the route back to the peer from which it was received.) If enabled, the router advertises the route back as reachable and includes the local AS. If disabled, the router echoes the route as UNREACHABLE/withdrawn.

Instructions: If the peer router saves routes that contain its own AS number and is running short of memory, send an UNREACHABLE echo.

A BGP speaker that participates in inter-AS multicast routing must advertise a route it receives from one of its external peers. If the router stores the route in its routing table, it must also advertise it back to the peer from which the route was received. For a BGP speaker that does participate in inter-AS multicast routing, such echoing is optional.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.2.1.33

Deleting a BGP Peer

To delete a BGP peer from an IP interface, begin at the BGP Peers List window shown in Figure 5-8, and complete the following steps:

1. Click on the peer that you want to delete.
2. Click on Delete.

The peer you selected is deleted.

3. Click on Cancel to exit the window.

Configuring BGP AS Weights and Weight Classes

When you configure BGP AS weights and weight classes, you are affecting the way BGP selects routes.

To configure BGP AS weights, begin at the Configuration Manager window and proceed as follows:

1. Select Protocols→IP→BGP→Weights.

The BGP AS Weight Parameters window appears (see Figure 5-10). This window lists all ASs to which a weight value has been assigned.

2. Add a weight to an AS, delete a weight from an AS, or change the weight parameter values of an AS as described in the following sections:

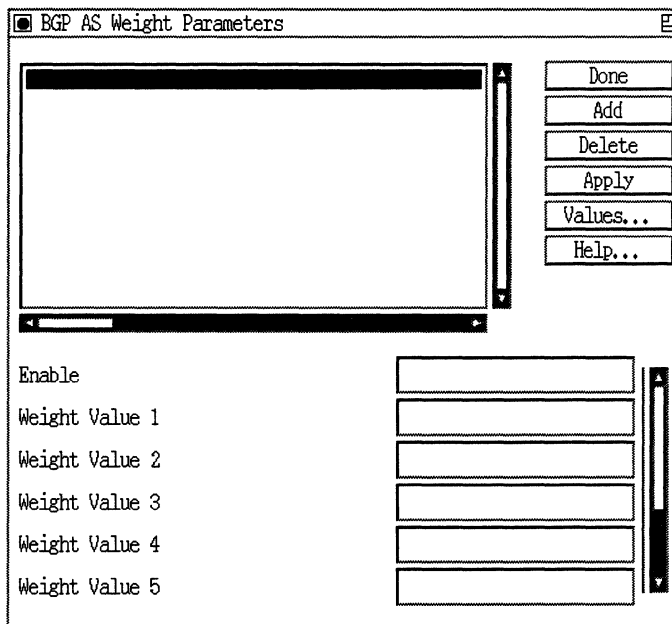


Figure 5-10. BGP AS Weight Parameters Window

Specifying a Class and Adding a Weight Value to an AS

To specify a class value and add a weight value to an AS, begin at the BGP AS Weight Parameters window shown in Figure 5-10 and complete the following steps:

1. Click on Add.
The BGP AS Weights window appears (Figure 5-11).
2. Specify the AS and Weight parameters.
3. Click on OK to save your changes and exit the window.

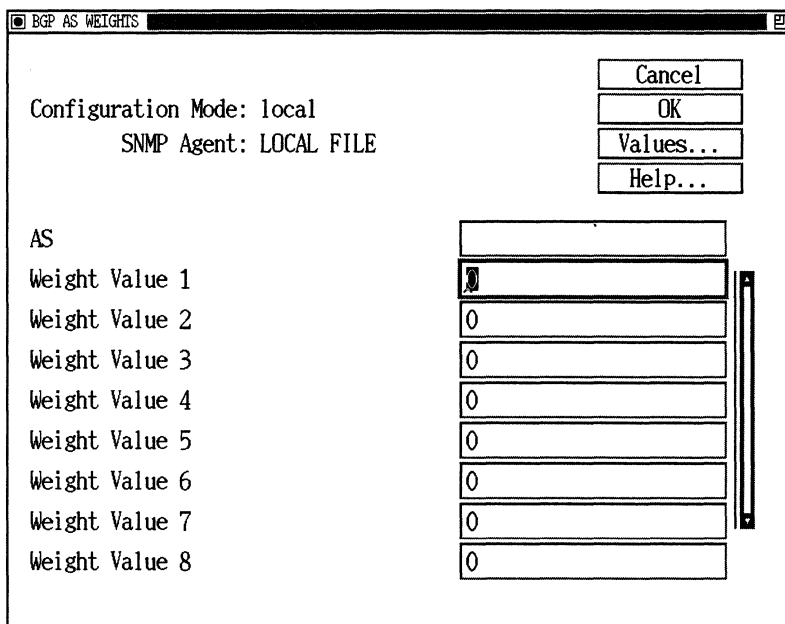


Figure 5-11. BGP AS Weights Window

BGP Weight Parameter Descriptions

This section describes how to set all parameters shown on the BGP Weight Parameters window.

Parameter: AS
Default: None
Range: 1 to 65535
Function: Identifies the autonomous system to which you want to assign a weight.
Instructions: Enter the appropriate AS number.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.4

Parameter: Weight Value 1
Default: 8
Range: 1 to 15, plus the infinity value of 16
Function: Specifies the Class 1 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.
Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.5

Parameter: Weight Value 2

Default: 8

Range: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 2 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.6

Parameter: Weight Value 3

Default: 8

Range: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 3 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.7

Parameter: Weight Value 4**Default:** 8**Range:** 1 to 15, plus the infinity value of 16**Function:** Specifies the Class 4 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.**Instructions:** Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.8**Parameter: Weight Value 5****Default:** 8**Range:** 1 to 15, plus the infinity value of 16**Function:** Specifies the Class 5 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.**Instructions:** Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.9

Parameter: Weight Value 6

Default: 8

Range: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 6 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.10

Parameter: Weight Value 7

Default: 8

Range: 1 to 15, plus the infinity value of 16

Function: Specifies the Class 7 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection.

Instructions: Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.11

| | |
|-------------------|--|
| Parameter: | Weight Value 8 |
| Default: | 8 |
| Range: | 1 to 15, plus the infinity value of 16 |
| Function: | Specifies the Class 8 weight value to add to this AS. This weight value is added to the other AS weight values in a route to determine the preference of the route and aid in route selection. |
| Instructions: | Either accept the current AS weight value or enter a new value. Any route that traverses an AS with an AS weight of 16 (infinity) will not be used. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.12 |

Editing the Weight Value Parameters of an AS

To edit the weight value of an AS, begin at the BGP AS Weight Parameters window shown in Figure 5-10 and complete the following steps:

1. Click on the AS for which you want to edit the weight value parameters.

When you do this, the parameters shown at the bottom of the BGP AS Weight Parameters window reflect the current values for the AS you selected.

2. Edit those parameters you want to change.

The Enable parameter is described following these instructions; see “BGP Weight Parameter Descriptions” on page 5-47 for instructions on setting the Weight parameter.

3. Click on Apply to implement your changes.

Repeat Steps 1 through 3 to edit any other AS you want to change; remembering to click on the Apply button each time.

4. Click on Done to exit the window.

| | |
|-------------------|--|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Enables or disables a weight assignment for a particular AS. |
| Instructions: | Set to Disable to disable the weight assignment for this AS; set to Enable if you previously disabled this weight assignment and now want to re-enable it. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.5.1.3.1.2 |

Deleting a Weight Value from an AS

To delete a weight value from an AS, begin at the BGP AS Weight Parameters window shown in Figure 5-10, and complete the following steps:

1. Click on the AS for which you want to delete the weight value.
2. Click on Delete.
3. Click on Done to exit the window.

Generating BGP Event Messages

To control the generation of BGP event messages:

1. Select Protocols→IP→BGP→Debug.

The BGP Debug Parameters window appears (see Figure 5-12).

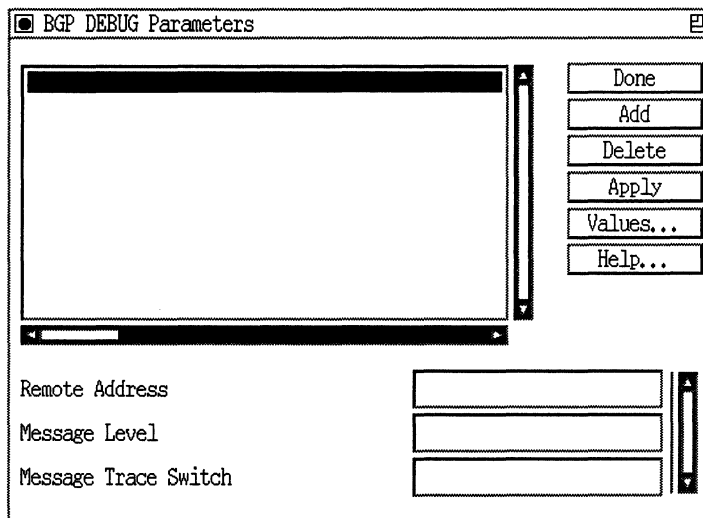


Figure 5-12. BGP Debug Parameters Window

Click on Add. The New BGP Debug Parameters window appears (see Figure 5-13).

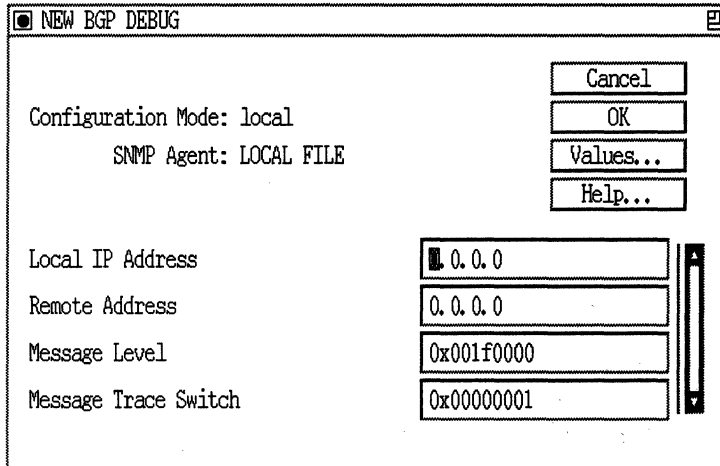


Figure 5-13. New BGP Debug Parameters Window

2. Edit the parameters to specify a connection and indicate the level of information you need.
3. Click on Done to exit the window.

BGP Debug Parameters Descriptions

This section describes how to set all parameters shown on the New BGP Debug Parameters window and the BGP Debug Parameters window.

Parameter: Local IP Address

Default: Null

Options: An IP address

Function: Specifies a BGP peer's local address

Instructions: Enter 0.0.0.0 to obtain event messages about all connections to a peer with the specified remote address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.5.1.2

Parameter: Remote Address

Default: Null

Options: An IP address

Function: Specifies a BGP peer's remote address

Instructions: Enter 0.0.0.0 to obtain event messages about all connections to peers using the specified local address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.5.1.3

Parameter: Message Level

Default: ALL

Options: ALL | DEBUG | INFO | WARNING
FAULT | TRACE

Function: Specifies the severity level of event messages required.

Instructions: Select the default to obtain event messages of all levels.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.5.1.4

Parameter: Message Trace Switch

Default: DISABLED

Options: DISABLED | OPEN | UPDATE
NOTIFICATION | KEEPALIVE

Function: Specifies whether or not BGP messages on the specified connection are logged and, if so, which messages are logged.

Instructions: Use the default or select a BGP message type.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.5.1.5.1.5

Deleting BGP from the Router

You can delete BGP from all router circuits on which it is currently enabled.

To delete BGP, begin at the Configuration Manager window, and complete the following steps.

1. Select the Protocols→IP→BGP→Delete BGP option.

A pop-up window appears prompting “Do you really want to delete BGP?”

2. Click on OK.

You are returned to the Configuration Manager window. BGP is removed from all circuits on the router.

Deleting BGP-3 from the Router

You can delete BGP-3 from all router circuits on which it is currently enabled. To delete BGP-3, begin at the Configuration Manager window and complete the following steps.

1. Select the Protocols→IP→BGP→Delete BGP-3 option.

A pop-up window appears prompting, “Do you really want to delete BGP-3?”

2. Click on OK.

You are returned to the Configuration Manager window. BGP-3 is removed from all circuits on the router.

Deleting BGP-4 from the Router

You can delete BGP-4 from all router circuits on which it is currently enabled. To delete BGP-4, begin at the Configuration Manager window and complete the following steps.

1. Select the Protocols→IP→BGP→Delete BGP-3 option.

A pop-up window appears prompting, “Do you really want to delete BGP-4?”

2. Click on OK.

You are returned to the Configuration Manager window. BGP-4 is removed from all circuits on the router.

Chapter 6

Customizing EGP Services

This chapter explains how to configure the Exterior Gateway Protocol (EGP). It contains:

- “EGP Overview” on page 6-1
- “EGP Implementation Notes” on page 6-13
- “Editing EGP Parameters” on page 6-14

EGP Overview

EGP-2 is an exterior gateway protocol used to exchange network reachability information between routers in different autonomous systems (AS). An AS is a group of networks and routers that share routing information using one or more Interior Gateway Protocols (IGP). An IGP, such as RIP or OSPF, is used within an AS to facilitate the communication of routing information with the AS. The routers that serve as end points of a connection between two autonomous systems run an exterior gateway protocol, such as EGP-2 (see Figure 6-1).

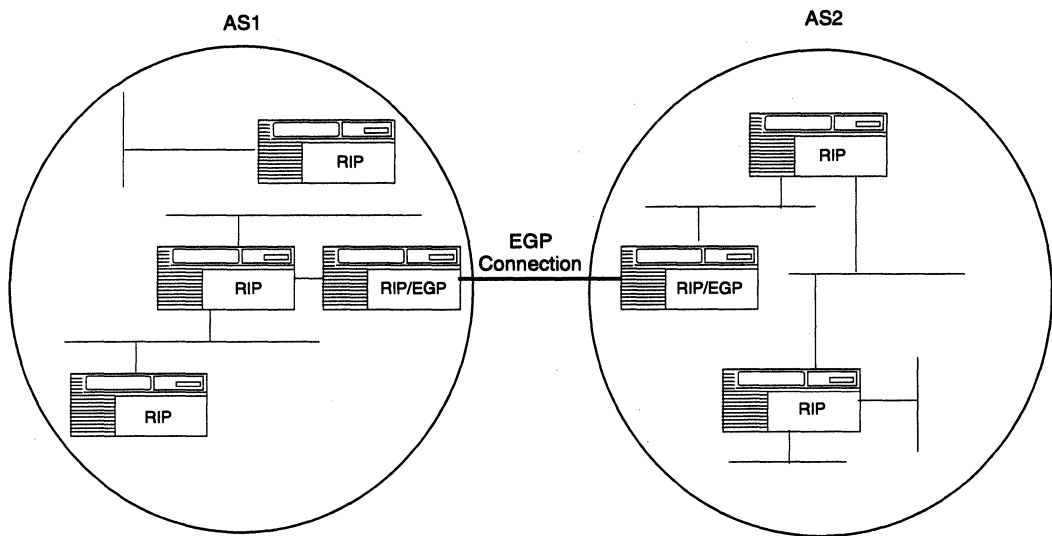


Figure 6-1. EGP Connection between Two Autonomous Systems Running RIP

Bay Network's implementation of EGP complies with RFCs 827 and 904. It runs over the same LAN and WAN media/protocols that IP runs over, including Ethernet, Token Ring, Synchronous, Wellfleet Proprietary Synchronous, Frame Relay, SMDS, X.25 (DDN, PDN, Pt-to-Pt), ATM PVC, FDDI, T1, E1, HSSI, and PPP.

Note: EGP assumes that each advertised network is a natural class network (A, B, or C) based on its high order bits. EGP cannot advertise or interpret subnets or supernets.

An EGP router has the following capabilities:

- ❑ It acquires EGP neighbors.
- ❑ It determines neighbor reachability.
- ❑ It exchanges network reachability information.

Each of these capabilities has an associated phase in EGP. The phases include the Neighbor Acquisition phase, the Neighbor Reachability phase, and the Network Reachability phase, respectively. The following three sections explain each phase.

Neighbor Acquisition Phase

This portion of EGP is responsible for forming neighbor relationships between routers that are peers. Routers that are peers each have an interface to a common network. One router attempts to acquire a peer router. If the peer agrees to be acquired, the two routers form a neighbor relationship. They then negotiate the mode of operation and the polling modes.

Certain messages that are used in the Neighbor Acquisition phase included the following:

- Neighbor Acquisition Request Command

This is the message that one router sends to another to request the formation of a neighbor relationship. The requesting router includes its:

- Autonomous system number
- Acquisition mode
- Hello interval it will accept from the peer
- Poll interval it will accept from the peer

- Neighbor Acquisition Confirm Response

This message is sent in response to a Neighbor Acquisition Request when the router agrees to being acquired; that is, it is willing to form the neighbor relationship. The responding router includes its:

- Autonomous system number
- Acquisition mode
- Hello interval it will accept from the peer
- Poll interval it will accept from the peer

□ Neighbor Acquisition Refuse Response

This message is sent in response to a Neighbor Acquisition Request when the router does *not* agree to being acquired; that is, it will not form the neighbor relationship. The status field of the Neighbor Acquisition Refuse message header supplies the reason for the refusal.

□ Neighbor Acquisition Cease Command

When two routers have an established neighbor relationship, either of the routers may send a Neighbor Acquisition Cease Command to the other to end the relationship. The status field of the Neighbor Acquisition Cease message header supplies the reason for ending the neighbor relationship.

□ Neighbor Acquisition Cease Ack Response

This message is sent in response to a Neighbor Acquisition Cease Command and indicates that the peer received and accepts the message.

Modes

Once two routers agree to form a neighbor relationship, they must then negotiate modes. Remember that in the Acquisition Request message, the requesting neighbor supplies its acquisition mode, and in the Acquisition Confirm Response message the responding router supplies its acquisition mode. The acquisition mode is configured for each router, and can be either active, passive or both. Ultimately, however, one of the routers must become the active router, and the other router must become the passive router.

The router that becomes the active router will later be responsible for Hello packets and Poll requests specified by the Hello Interval and the Poll Interval, respectively. The passive router just responds to the active router with I-H-U and Routing Update messages.

According to EGP, the routers' modes are determined as shown in Table 6-1.

Table 6-1. Router Mode Determinator

| Router A | Router B | Resulting Modes |
|----------|----------|--|
| Active | Passive | Router A is active; Router B is passive. |
| Passive | Passive | Not allowed |
| Active | Active | The router with the lower autonomous system number becomes active, the other becomes the passive router. |
| Both | Active | Router A is passive; Router B is active. |
| Both | Passive | Router A is active; Router B is passive. |
| Both | Both | The router with the lower autonomous system number becomes active, the other becomes the passive router. |

Table 6-1 shows all possible acquisition mode combinations, when you configure the EGP neighbors at each end of a connection. However, it is recommended that one router be configured in the Active acquisition mode, and the other in the Passive acquisition mode.

As an example of a neighbor acquisition, consider the Routers A and B in Figure 6-2. Router A attempts to acquire Router B by sending an Acquisition Request message to Router B. Router B agrees to form the neighbor relationship with Router A by responding with an Acquisition Confirm Message.

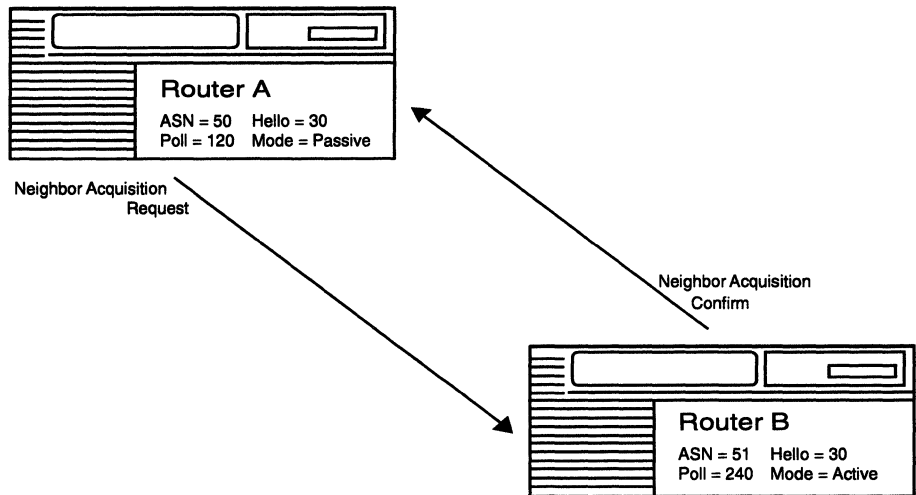


Figure 6-2. Neighbor Acquisition Sequence

Router B becomes the Active router because its configured acquisition mode was Active and Router A's configured mode was passive (refer to Table 6-1). This means that Router B, as the Active router, will later be responsible for sending Hello packets and Poll commands, and Router A will respond to Router B.

Had Router B sent an Acquisition Refuse Response, no relationship would have been formed. Also, at anytime after the neighbor relationship is formed, either Router A or Router B could send an Acquisition Cease Command. This would terminate the neighbor relationship between them.

Neighbor Reachability Phase

This portion of EGP is responsible for monitoring and maintaining an established EGP neighbor relationship between two routers. Its purpose is to ensure that the neighbors are operational and can provide reliable network reachability information.

Two neighbors will be able to exchange network reachability information only if they are both in the UP state and know that they are both in the UP state. This is the point at which neighbor reachability is positively determined.

Whether a router is in the UP or DOWN state is indicated in the status field of the Hello and I-H-U messages. Following is a description of these two messages. Poll and Update messages are also sent during the neighbor reachability phase, but will be discussed in the next section.

□ Neighbor Reachability Hello Command

This command is sent by the active neighbor to the passive neighbor to determine whether the passive neighbor is functioning. The frequency of the active router's Hello command transmissions is dictated by the passive router's configured Hello interval. The passive router specifies, in the Neighbor Acquisition Confirm response, an interval at which it is willing to respond to Hello commands. The active router can send Hello commands less frequently than the specified Hello interval, but not more frequently. The passive neighbor determines reachability by the status field in the active neighbor's Hello command.

□ Neighbor Reachability I-H-U Response

This response is sent by the passive neighbor in response to a Hello command. If the status field in the I-H-U is UP, then the active neighbor determines that the passive neighbor is reachable.

As stated previously, Hello commands and I-H-U messages are used to determine neighbor reachability. A neighbor is reachable when it moves to the UP state, which is indicated in the Status field of these two messages. A neighbor will move to the UP state only when it has

received a certain number of reachability indicators within a specified time interval. Similarly, a neighbor will move to the DOWN state when it has *not* received a certain number of reachability indicators within that same specified time. These UP and DOWN state thresholds differ for Active and Passive routers (see Table 6-2).

Table 6-2. UP and DOWN State Thresholds

| Mode | UP Threshold | DOWN Threshold | Specified Time Interval |
|---------|--------------|----------------|-------------------------|
| Active | 3 | 1 | T x 5 |
| Passive | 1 | 0 | T x 5 |

T is the agreed upon Hello interval for this neighbor relationship. If, after 5 Hello Intervals, the number of reachability indicators is 3 for an Active router or 1 for a passive router, the neighbor is considered UP. If, after 5 Hello Intervals, the number of reachability indicators is 1 for an Active router or 0 for a passive router, the neighbor is considered DOWN.

Figure 6-3 shows two routers that already have formed an EGP neighbor relationship in the Neighbor Acquisition phase, and are now attempting to determine neighbor reachability. Router B, the Active neighbor, will use the Hello and Poll intervals provided by Router A, the Passive neighbor. The Hello interval is 30 and the Poll interval is 120.

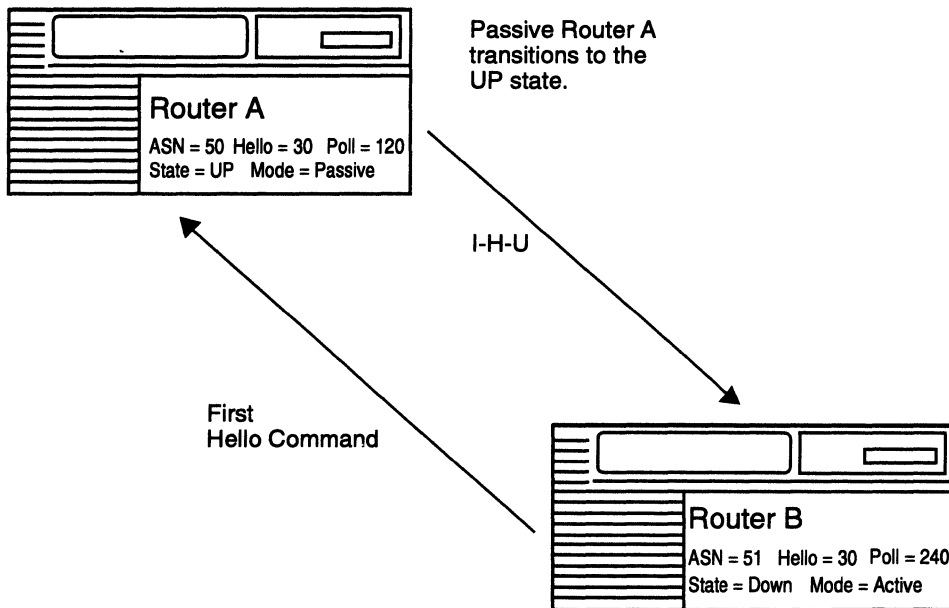


Figure 6-3. Neighbor Reachability Exchange Begins between Two EGP Neighbors

When Router B sends its first Hello command, Router A transitions to the UP State. Router A responds to the Hello command with an I-H-U; however, Router B does not yet transition to the UP state. As an Active router, it must receive 3 I-H-U's within a specified time (in this case 2.5 minutes, or 5 x 30 seconds) before transitioning to the UP state. Upon receipt of the third I-H-U within the specified time interval, Router B transitions to the UP state (see Figure 6-4). At this point, neighbor reachability is established.

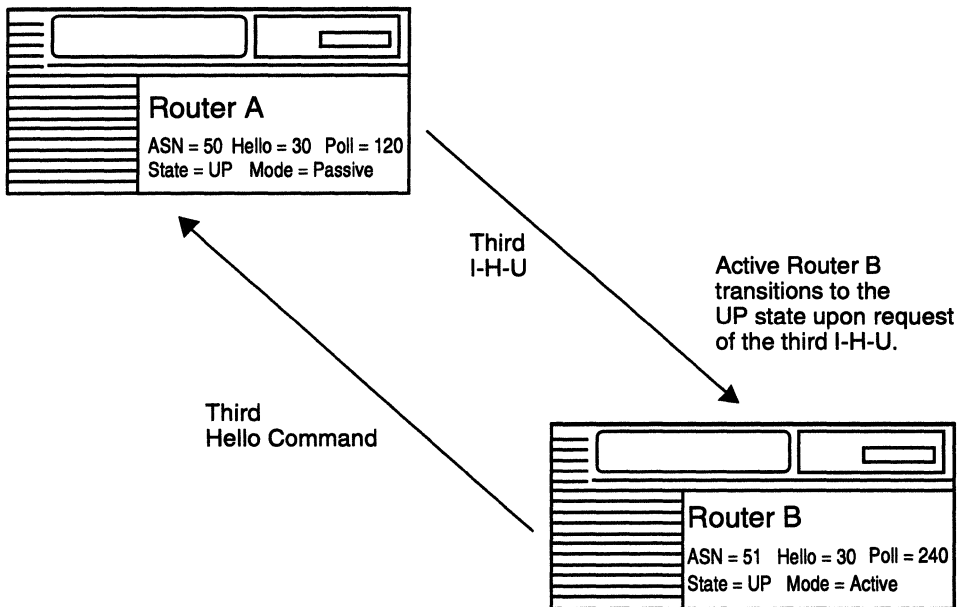


Figure 6-4. Neighbor Reachability Is Established with Both Routers in the UP State

Network Reachability Phase

This portion of EGP is responsible for determining what networks are reachable through two EGP neighbors; that is, it provides the network reachability information. This information provides a list of gateways, the networks those gateways can reach, and their associated distances.

Two neighbors determine network reachability by exchanging Poll commands and Routing Update responses as described below:

- Poll Command

The Active neighbor sends a Poll command to a passive neighbor that it already knows to be reachable. The Poll command requests routing information from the Passive neighbor.

- Routing Update Response

The Routing Update Response is the message that contains the routing information (the list of gateways on the common network, the networks they can reach, and associated distances). Both Active and Passive neighbors can send Routing Update messages. The Active neighbor usually sends a Routing Update response after it sends a Poll command. The Passive neighbor usually sends a Routing Update response in response to a Poll command.

Although the Routing Update Response is typically sent as a response, each router is allowed to send one Unsolicited Routing Update packet between Poll intervals. This Unsolicited Routing update is sent either upon a transition to the UP state, or when there is a neighbor reachability change.

The Poll command and Routing Update Response both use an IP Source Network Field. The IP Source Network Field contains the IP address of the network to which both EGP neighbors have an interface. From this network, all distances to reachable networks (contained in Routing Update Responses) are measured.

Figure 6-5 shows the typical Network Reachability sequence between two routers that have established an EGP neighbor relationship, and have determined neighbor reachability through the exchange of Hello and I-H-U messages.

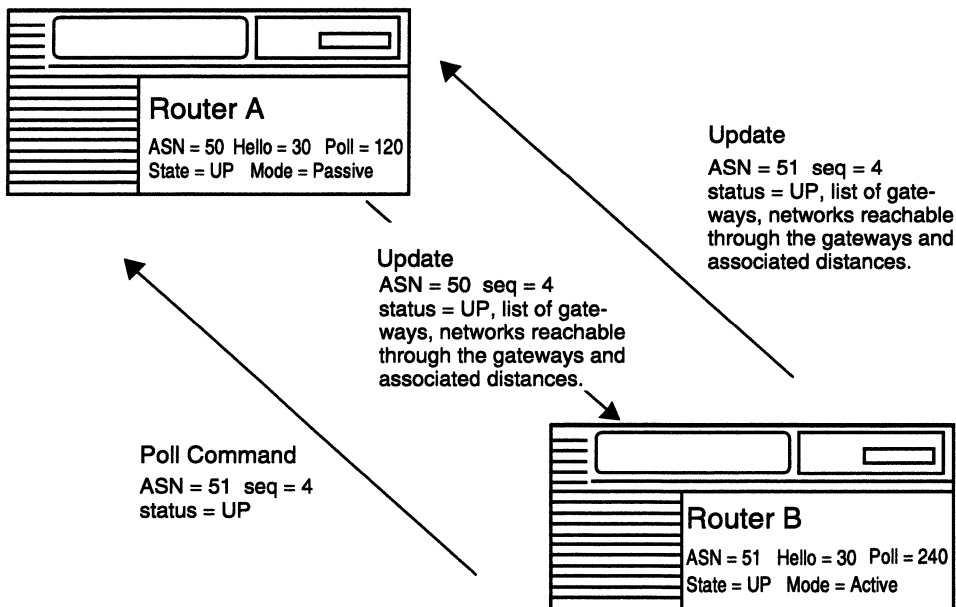


Figure 6-5. Network Reachability Sequence between Two EGP Neighbors

Modes

The EGP router can be configured to operate in one of two gateway modes for any given IP interface:

- ❑ **Non-Core**

When the router is configured as a noncore gateway, the AS to which it belongs acts as a stub AS. It advertises and forwards only traffic that originated or is destined for a network within its AS.

- ❑ **Core**

When the router is configured as a core gateway, the AS to which it belongs acts as a transit AS. In the core mode, it can advertise and forward traffic to networks reachable interior or exterior to its local AS.

The default gateway mode is core mode. If the EGP router is reconfigured to run in noncore mode, the Site Manager *automatically* configures EGP export route filters on that IP interface to suppress OSPF external routes to EGP and the advertisement of any networks learned by EGP.

For More Information about EGP

For more information about EGP, refer to the following documentation:
Comer, Douglas, E. *Networking With TCP/IP, Volume I*. 2d ed.
Englewood, Calif.: Prentice-Hall Inc., 1991.

Mills, D. L. "Exterior Gateway Protocol Formal Specification." RFC 904, Network Information Center (NIC), SRI International, Menlo Park, California, April 1984.

Perlman, Radia. *Interconnections: Bridges and Routers*. Reading, Massachusetts: Addison-Wesley Publishing Company, 1992.

Rosen, Eric, C. "Exterior Gateway Protocol (EGP)." RFC 827, Network Information Center (NIC), SRI International, Menlo Park, California, October 1982.

EGP Implementation Notes

This section provides you with some guidelines that must be followed when you configure EGP. If you do not follow these guidelines, EGP will become disabled on the interfaces involved.

- Autonomous system numbers must be between 1 and 65535.
- Two autonomous systems connected by an EGP link must have different autonomous system numbers.
- The Remote IP Address cannot be the same as any of the Local IP Interface Addresses.
- The Remote IP Address must be on the same subnet as one of the Local IP interfaces.

- ❑ EGP does not have any loop avoidance techniques — avoid loop topologies; otherwise, you will have to configure EGP route filters to counter the redundancies.

Editing EGP Parameters

This section describes how to edit, or customize, EGP parameters.

Note: The instructions in this section assume that you have already configured at least one IP interface with EGP support (EGP interface). If you have *not* yet configured an EGP interface, see the *Configuring Wellfleet Routers* guide for instructions.

You access all EGP parameters from the Configuration Manager window (refer to the *Configuring Wellfleet Routers* guide for instructions on accessing this window).

For each EGP parameter, this chapter describes the default setting, all valid setting options, the parameter function, and instructions for setting the parameter.

Editing EGP Global Parameters

When you edit the EGP global parameters, you are editing parameters that affect EGP on the entire router.

To edit EGP global parameters, begin at the Configuration Manager window (see Figure 1-11) and complete the following steps:

1. Select the Protocols→IP→EGP→Global option.

The Edit EGP Global Parameters window appears (Figure 6-6).

2. Edit those parameters you wish to change.

The EGP global parameters are described following these instructions.

3. Click on OK to exit the window and save your changes when you are finished.

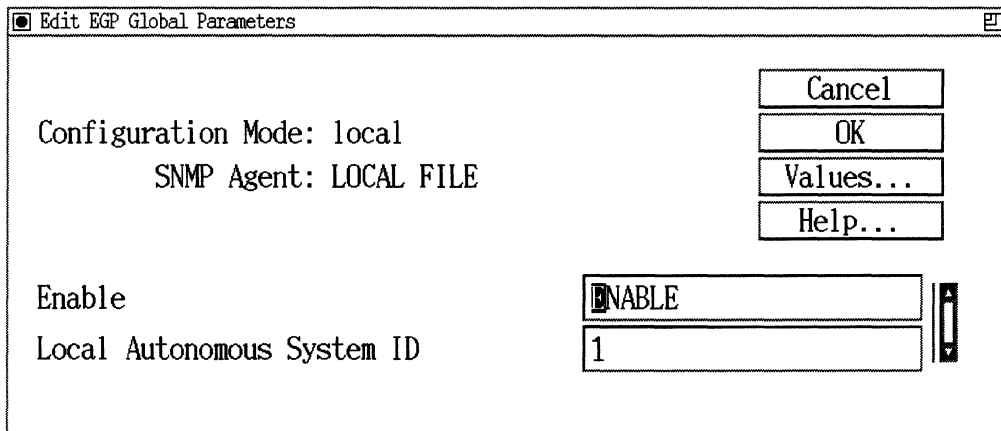


Figure 6-6. Edit EGP Global Parameters Window

EGP Global Parameter Descriptions

This section describes how to set all EGP global parameters.

Parameter: **Enable**

Default: Enable

Options: Enable | Disable

Function: This parameter allows you to globally enable or disable EGP on all router interfaces.

Instructions: Set to Disable if you want to disable EGP for the entire router. Set to Enable if you previously disabled EGP and now wish to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.4.1.2

Parameter: **Local Autonomous System ID**

Default: None

Range: 1 to 65535

Function: Identifies the local autonomous system (the AS to which this router belongs), by the NIC-assigned decimal number. There is no default for this parameter.

Instructions: Either accept the current value for Local Autonomous System ID, or enter a new value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.4.1.7

Configuring EGP Neighbors

When you configure EGP neighbors, you are setting parameters that affect the formation of EGP neighbor relationships on a particular IP interface.

To configure EGP Neighbor parameters, begin at the Configuration Manager window and complete the following steps:

1. Select the Protocols→IP→EGP→Neighbors option.

The IP Interface List for EGP window appears (Figure 6-7). It lists all IP interfaces on which EGP has been enabled.

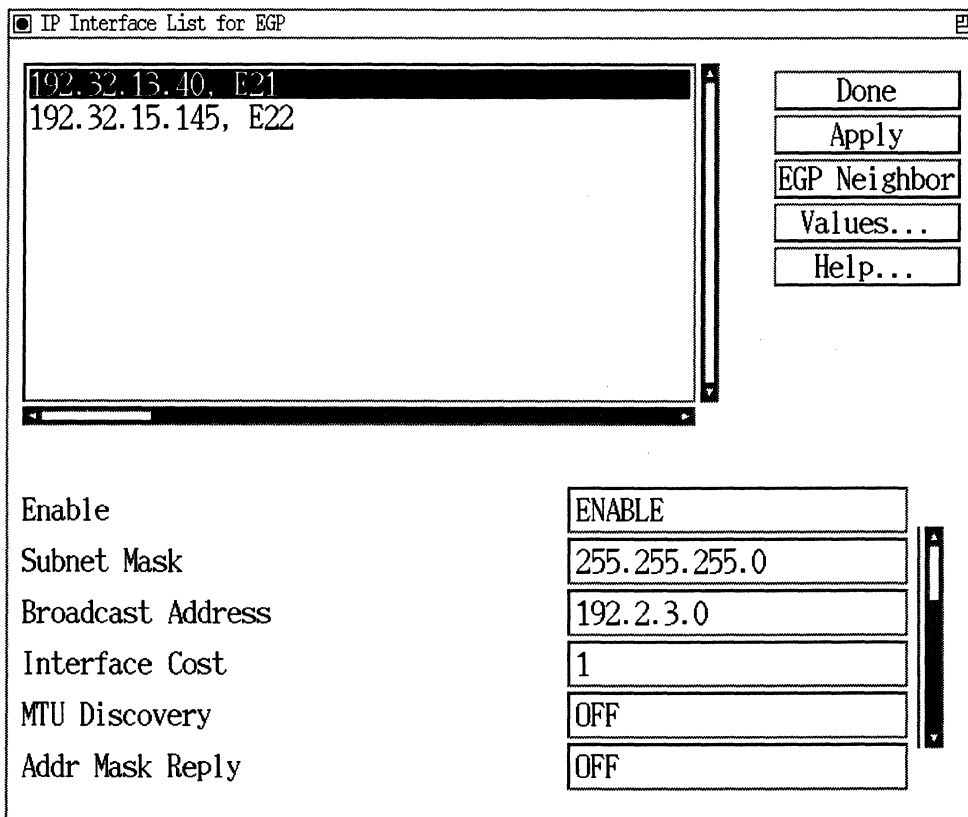


Figure 6-7. IP Interface List for EGP Window

2. Click on the IP interface for which you wish to edit EGP neighbor parameters.
3. Click on EGP Neighbor.

The EGP Neighbors List window appears (Figure 6-8). It shows all of the neighbors configured for the IP interface that you selected in Step 2. In this example, neighbors have not yet been configured for the chosen interface.

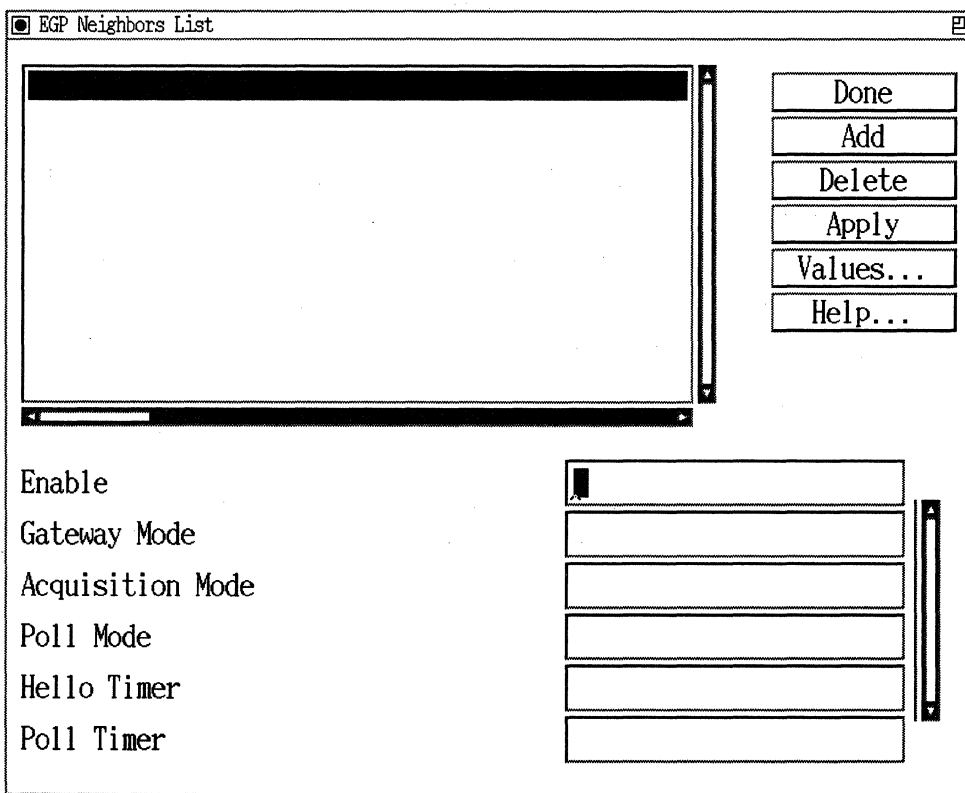


Figure 6-8. EGP Neighbors List Window

4. Add an EGP neighbor, edit parameters associated with a specific EGP neighbor, or delete an EGP neighbor from the IP interface as described in the following sections.

Adding an EGP Neighbor

To add an EGP neighbor to an IP interface, begin at the EGP Neighbors List window shown in Figure 6-8, and complete the following steps:

1. Click on Add.

The EGP Neighbor parameters window appears (see Figure 6-9).

2. Set the Remote Autonomous System IP Address and Gateway Mode parameters.

These EGP neighbor configuration parameters are described following these instructions.

3. Click on OK.

The neighbor you just added now appears in the scroll box in the EGP Neighbors List window.

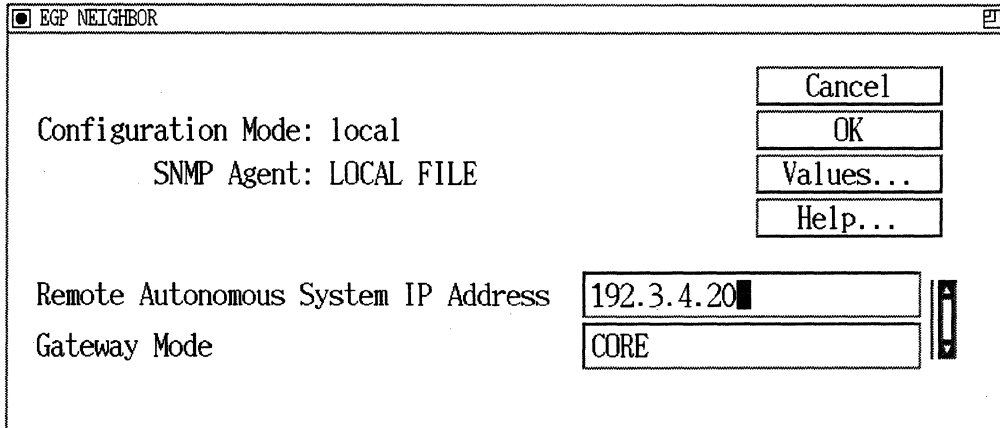


Figure 6-9. EGP Neighbor Parameters Window

EGP Neighbor Parameter Descriptions

This section describes how to set the EGP neighbor configuration parameters.

| | |
|-------------------|--|
| Parameter: | Remote Autonomous System IP Address |
| Default: | None |
| Options: | Any IP address |
| Function: | Specifies the IP address of the remote router that will form an EGP neighbor relationship with this router. |
| Instructions: | Enter the IP address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.4.3.1.4 |
| | |
| Parameter: | Gateway Mode |
| Default: | Core |
| Options: | Core Non Core |
| Function: | Specifies the gateway mode for this EGP neighbor. If you choose Core, the default, the local AS to which this EGP neighbor belongs will act as a transit AS. That is, it will advertise networks that reside within the AS as well as external networks. If you choose Non Core, the AS to which this EGP neighbor belongs will act as a stub AS. That is, it will only advertise networks that reside within the AS. |
| Instructions: | Set this parameter to either Core or Non Core, depending on how you want this EGP neighbor to function. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.4.3.1.5 |

Editing an EGP Neighbor

To edit an EGP neighbor, you must begin at the EGP Neighbors List window shown in Figure 5-9, and complete the following steps:

1. Click on the neighbor for which you want to edit parameters from the Neighbors List window.

When you do this, all of the parameters shown at the bottom of the window will reflect the *current* values for the neighbor you selected.

2. Edit those parameters you want to change.

All EGP neighbor parameters that you can edit are described following these instructions.

3. Click on Apply to implement your changes.

Repeat Steps 1 through 3 to edit any other neighbors you wish to change, remembering to click on Apply each time.

4. Click on Done to exit the window.

| | |
|-------------------|---|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Enables or disables an EGP neighbor relationship with the specified IP address. |
| Instructions: | Set this parameter to Disable if you want to temporarily disable this neighbor relationship, rather than delete it. Or, set it to Enable if you previously disabled this neighbor relationship, and now wish to re-enable it. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.4.3.1.2 |

Parameter: Acquisition Mode

Default: Passive

Options: Passive | Active

Function: Specifies which of the two neighbors initiates EGP connections. The router in the Active mode is the initiator.

Instructions: Set this parameter to Active if you want the local EGP neighbor to be the initiator of EGP connections. Otherwise, accept the default value, Passive.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.4.3.1.7

Parameter: Poll Mode

Default: Both

Options: Active | Passive | Both

Function: Specifies the type of neighbor reachability algorithm this local EGP neighbor executes. In the Active mode, a router sends Hello and Poll commands to request reachability status from its neighbor. In the Passive mode, a router responds to Hello and Poll commands with I-H-U and Update messages.

Instructions: Accept the default value, Both, or set to either Active or Passive (depending on the neighbor reachability algorithm you want this router to execute.)

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.4.3.1.8

Parameter: Hello Timer**Default:** 60 seconds**Range:** 30 to 120 seconds**Function:** Specifies the number of seconds between the local EGP neighbor's EGP Hello message retransmissions. This variable represents the RFC 904 t1 timer.**Instructions:** Accept the default value of 60 seconds for the Hello Timer; or, set this parameter to some value between 30 and 120 seconds.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.4.3.1.9**Parameter: Poll Timer****Default:** 180 hundredths of a second**Range:** 120 to 480 hundredths of a second.**Function:** Specifies the time period, in hundredths of a second, between the local EGP neighbor's EGP Poll message retransmissions. This variable represents the RFC 904 t2 timer.**Instructions:** Either accept the default value of 180 hundredths of a second for the Hello Timer; or, set this parameter to some value between 120 and 480 hundredths of a second.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.4.3.1.10

Deleting an EGP Neighbor

To delete an EGP neighbor from an IP interface, begin at the EGP Neighbors List window shown in Figure 6-8, and complete the following steps:

1. Click on the neighbor that you want to delete.
2. Click on Delete.

The neighbor you specified is deleted.

3. Click on Done to exit the window.

Deleting EGP from the Router

You can delete EGP from all router circuits on which it is currently enabled. To delete EGP, begin at the Configuration Manager window and complete the following steps.

1. Select the Protocols→IP→BGP→Delete EGP option.

A window appears prompting, “Do you really want to delete EGP?”

2. Click on OK.

You are returned to the Configuration Manager window. EGP is removed from all circuits on the router.

Chapter 7

IP Multicasting

An IP unicast router is a router that can forward unicast datagrams — datagrams that bear a unicast IP destination address. Each unicast datagram is delivered to a single destination. A multicast router is a router that forwards unicast datagrams and can also forward IP multicast datagrams — datagrams that bear a multicast IP address. Each multicast datagram is delivered to a host group, a set of zero or more hosts designated by the address.

This chapter contains the following sections describing Wellfleet support for IP multicast routing:

- “Host Groups” on page 7-1
- “Multicast Networks and Multicast Source Networks” on page 7-2
- “Internet Group Management Protocol” on page 7-2
- “Distance Vector Multicast Routing Protocol” on page 7-4
- “Types of Multicast Support” on page 7-11
- “Editing Multicasting Parameters” on page 7-12

Host Groups

Multicasting defines two categories of host groups: permanent and transient. A permanent host group has a well-known, administratively assigned IP multicast group address. It is the address, not the membership, that is permanent and defines the group. A permanent

host group can consist of zero members. A transient host group, on the other hand, exists only as long as it has members that need its services. IP addresses in the multicast range that are not reserved for permanent groups are available for dynamic assignment to transient host groups.

An IP host group places no restrictions on its membership. Host members can reside anywhere; they can join and leave the group at any time; and they can be members of more than one group at the same time. In order to receive a multicast message from a host group, a host must be a member of the group; however, a host need not be a member of a group to send a multicast message to its members.

In general, hosts that are members of the same group are located on different networks. However, a range of multicast addresses (224.0.0.x) is reserved for groups that are locally scoped. All message traffic for these hosts remains on the local network. Hosts that belong to a group in this address range and that reside in different networks will not receive each other's message traffic.

Multicast Networks and Multicast Source Networks

A multicast network is a network that can support the sending and receiving of multicast datagrams. The hosts on this network may or may not be members of various multicast host groups.

A multicast source network is a network containing hosts that can (but may or may not) send multicast packets. These hosts may or may not ever be members of a host group.

Internet Group Management Protocol

Any host system on any IP network can send a message to a multicast group using the group's IP multicast address. To receive a message addressed to a multicast group, however, the host must be a member of the group and reside on a network where that group is registered with a local multicast router.

The Internet Group Management Protocol (IGMP) allows a host to register its local network with the local router to receive any datagrams sent to this router and targeted to a specific IP multicast address.

How IGMP Works

A multicast router periodically sends IGMP host membership queries to its attached local networks. Hosts on the networks respond with host membership reports, one report for each supported multicast group. If multiple multicast routers exist on the network, the router with the lowest IP address takes on the job of sending out host membership queries. If at least one host on the local network specifies that group in a report, the router will forward to that network all datagrams bearing the group's multicast address.

Upon initialization, the host may immediately send out a report for each of its supported multicast groups. The router accepts and processes these asynchronous reports the same way it accepts requested reports.

Once in a steady state, hosts and routers communicate in a way that minimizes the exchange of the queries and reports.

A host that receives a query delays its reply by a random interval and listens for a reply from any other host in the same host group. Consider a network that includes two host members — host A and host B — of the same multicast group. The router sends out a host membership query on the local network. Host A and Host B both receive the query and listen on the network for a host membership report. Host B's delay time expires first, so it responds to the query with a membership report. Hearing the response, Host A does not send a report of its own for the same group.

For instructions on configuring IGMP on a circuit, see “Editing IGMP Global Configuration Parameters” on page 7-28 and “Editing IGMP Entry Interface Parameters” on page 7-30.

Distance Vector Multicast Routing Protocol

The Distance Vector Multicast Routing Protocol (DVMRP) provides a mechanism for routers to propagate multicast datagrams. Multicast routers running DVMRP exchange routing information, store the information in routing tables, and use the information to ensure that a multicast datagram is propagated to every network that needs this datagram, doing so in a manner that minimizes the number of excess copies sent to any particular network.

How DVMRP Works

DVMRP routers build and maintain their routing tables by exchanging routing information with their DVMRP neighbors. In a DVMRP environment, neighbors are multicasting routers that are connected in either of two ways: directly or by means of a path — or tunnel — between multicast routers through a unicast network.

In Figure 7-1, for example, multicasting Router A has two neighbors, Router B and Router C. Router A and Router B are connected to each other through network 6 by direct physical links to interfaces a1 and b1, respectively. Router A and Router C are connected to each other through a tunnel (and a unicast router) between interfaces a2 and c1.

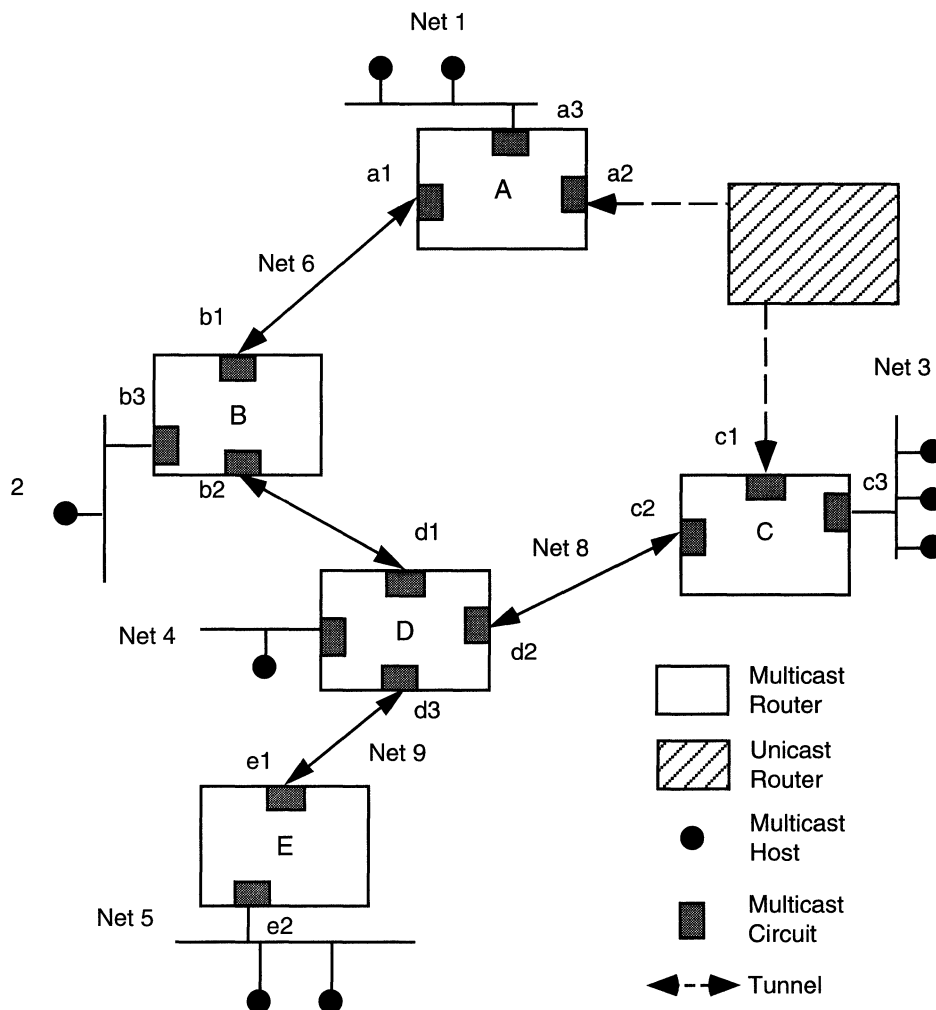


Figure 7-1. Multicast Routers

At startup, a DVMRP multicasting router:

1. Initializes its routing table with information on all of its local networks
2. Sends out a probe for all routes on each of its multicast interfaces (both physical circuits and tunnels)

3. Receives reports from its neighbors containing the routing information (including route costs)

In Figure 7-1, for example, Router D becomes active and issues routing probes on four multicasting interfaces. Router D receives reports from its multicasting neighbors, Router B, C, and E.

A router will not send out route reports on an interface until it knows (by means of received probes or reports) that it has a neighboring multicast router on that interface. It will continue to send probes periodically on an interface until it knows that a neighbor exists on that interface.

Calculating a Route Metric

Each interface — either a physical interface to a local network or a tunnelled interface to a remote multicasting router — is configured with a metric that indicates the cost of the hop. A route metric is the sum of all the interface (hop) metrics from a given route source to a given router. (Currently, *mrouted* restricts a route to a total metric value of 31 or less.) Table 7-1 lists the recommended hop metrics for topologies that will link up to the Internet Multicast Backbone (MBone).

Table 7-1. Recommended Hop Metrics

| Hop | Metric |
|---|---------------------------|
| LAN, or tunnel across a single LAN | 1 |
| Multihop tunnel | 2 or 3 |
| Serial link, or tunnel across a serial link | 1 |
| Backup tunnel | primary tunnel metric + 1 |

Comparing Routes

A router that receives multiple route reports for the same multicasting source network compares the cost specified in each (based on the metric field) and stores information from the report with the lowest cost in its routing table.

In Figure 7-1, for example, router D receives two reports for the network connected to multicasting router A, one from router B and one from router C. Using the metrics contained in the route reports, router D determines that the cost of the tunnelled route is greater than the cost of the route that uses direct physical connections. Router D discards the route received from router C and stores the route received from router B.

Router D then declares router B to be the next hop neighbor and interface a1 to be the next hop interface. Once a next hop neighbor has been declared for a route, the route updates received from that neighbor for that route take precedence until either the route times out or another router advertises a better metric for that route.

Periodically, each multicasting router issues full or partial routing information on each DVMRP circuit, using DVMRP report messages. This routing information represents the sending router's cost to reach the specified network (the cost is the sum of the hop metrics along the shortest path to the given source network). When DVMRP receives a DVMRP report from another router, it re-examines its routing table to determine if the shortest path information needs updating. Specifically, DVMRP looks in the routing table for an entry describing a route to the same source network. If one exists, DVMRP compares the cost of the two routes. DVMRP stores the route with the lower cost in its routing table. (Other received routing information is used in the construction of a shortest path tree as described in "Creating a Shortest Path Tree" on page 7-8.)

Creating a Shortest Path Tree

Route information used by DVMRP is independent of any other routing information used by the router — for example, routes provided by OSPF. The purpose of this routing information is to create a shortest path tree entry in the routing table for the propagation of multicast datagrams. The shortest path tree entry indicates the interface local to a network or tunnel on the local router that provides the shortest path from a particular source network to that router. A shortest path tree entry also indicates those interfaces local to networks or tunnels that are on the shortest path from neighboring local network routers to that source network.

In Figure 7-1, for example, the routing table on router D includes a shortest path route to the network connected to router A. The route indicates that circuit d1 provides the shortest path for router D to that network. Router E considers the network between itself and router D to be on shortest path to the network connected to router A. Router D has an interface — d3 — that is part of the shortest path from router E to the network connected to router A.

If neighboring routers have the same metric to a given source network, the router with the higher IP address will be responsible for propagating multicast traffic originating from that source network onto the network or tunnel that is common to these neighboring routers.

Identifying a Leaf Network

A network that is not on the shortest path from a multicast router to a source network is considered to be a leaf. In Figure 7-1, the network connected to router E is a leaf network.

Aging a Route

When a router adds or updates a route, it runs aging timers that control the useful life of the route. The route expiration timer is used to time out a route so that it is no longer used by this router in routing decisions. The garbage timer is used to time out a route so that it is no

longer propagated by this router in route updates; once the route expires, it is advertised as unreachable until it is garbage collected or until it receives a route report advertising reachability. If a router has not received any reports from a neighbor before the neighbor timer expires, that neighbor is considered to be down.

The router uses a leaf timeout timer to determine whether or not a network or tunnel local to a given interface is considered to be part of the shortest path to a given source network by any other local network routers.

If the local interface has not received during this time a route report for a given source network, this network or tunnel and its local interface are considered not to lie in the shortest path for any local network routers: in other words, the local network is not part of the shortest path to that specific source network.

Specifying a Threshold

Threshold values control the scope of datagram delivery. Threshold is the minimum IP TTL required for a multicast datagram to be forwarded out a given interface. The interface compares the TTL value of each multicast datagram to be forwarded with the threshold configured for that interface.

For *mrouted* compatibility, multicast datagrams originated by the router have a TTL of 1. These datagrams are not compared against the TTL. Unicast datagrams originated by the router (for example, for route reports issued via a tunnel) have a TTL of 255. A datagram that is to be forwarded through a tunnel is first compared against the threshold and, if accepted, is then encapsulated in an IP datagram with a TTL of 64.

Table 7-2 lists the originating TTL values that are recommended for certain types of multicast applications and the threshold values recommended for routers in order to permit the forwarding of packets from these applications. These values are recommended for topologies that will hook up to the MBone.

Table 7-2. Recommended TTL and Threshold Values

| Multicast Application | TTL | Threshold |
|-----------------------------------|------------|------------------|
| IETF channel 1 low-rate GSM audio | 255 | 224 |
| IETF channel 2 low-rate GSM audio | 223 | 192 |
| IETF channel 1 PCM audio | 191 | 160 |
| IETF channel 2 PCM audio | 159 | 128 |
| IETF channel 1 video | 127 | 96 |
| IETF channel 2 video | 95 | 64 |
| Local event audio | 63 | 32 |
| Local event video | 31 | 1 |

Types of Multicast Support

Using Site Manager, the network administrator can specify various types of multicast support for a circuit. (The first two and the last three are mutually exclusive.)

- ❑ IGMP host membership queries enabled. For a circuit that connects the router to a network with a host that may become a member of one or more multicasting groups. The network administrator configures IGMP on the circuit and enables host queries by setting the Interface Query Rate parameter to a nonzero value.
- ❑ IGMP host membership queries disabled. For a circuit that connects the router to a network that has no multicasting hosts. The network administrator configures IGMP on the circuit but disables queries by setting the Interface Query Parameter to zero.
- ❑ IGMP/DVMRP support with circuit routing support. For a circuit that connects the router to a network that requires the propagation of multicast datagrams. The network configures IGMP (as described above) and DVMRP on the circuit and enables the circuit for routing.
- ❑ IGMP/DVMRP support for tunnels with circuit routing support. For a circuit that links the router to one or more remote multicasting routers via tunnels and that also propagates multicast datagrams. The network administrator configures IGMP (as above) and DVMRP on the circuit and uses the DVMRP Tunnel Parameters window to configure one or more tunnels on the circuit.
- ❑ IGMP/DVMRP support for tunnels without circuit routing support. For a circuit that links the router to one or more remote multicasting routers via tunnels but that should not be allowed to propagate multicast datagrams. The network configures IGMP and DVMRP on the circuit, uses the DVMRP Tunnel Parameters window to configure one or more tunnels on the circuit, and sets the Route Enable parameter in the circuit entry to Disabled.

Editing Multicasting Parameters

The following sections describe and show you how to set DVMP and IGMP multicasting parameters.

Editing DVMRP Global Parameters

To edit DVMRP global parameters:

1. Select IP→Multicast→DVMRP→Global. The DVMRP Global Configuration window appears (see Figure 7-2).
2. Edit the parameters as described in the following section.
3. Click on Save.

| Parameter | Value |
|---------------------------|---------|
| Enable | DISABLE |
| Full Update Interval | 60 |
| Triggered Update Interval | 5 |
| Leaf Timeout | 200 |
| Neighbor Timeout | 140 |
| Route Expiration Timeout | 200 |
| Garbage Timeout | 340 |
| Estimated Routes | 25 |
| Neighbor Probe Interval | 190 |
| Route Switch Timeout | 140 |

Figure 7-2. DVMRP Global Configuration Window

DVMRP Global Configuration Parameter Descriptions

Use this section as a guide for setting DVMRP global parameters.

Parameter: **Enable**

Default: Enable

Options: Enable | Disable

Function: Enables and disables DVMRP support on the router.

Instructions: To disable DVMRP once you have configured it on the router, specify Disable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.2

Parameter: **Full Update Interval**

Default: 60 seconds

Range: 10 to 2000 seconds

Function: Specifies, in seconds, how often routing messages containing complete routing tables are sent.

Instructions: Determine the full update interval you require and specify a value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.4

Parameter: Triggered Update Interval

Default: 5 seconds

Range: 5 or more seconds

Function: Specifies, in seconds, the minimum amount of time between triggered updates.

Instructions: Triggered updates are sent in the period between full updates. Issuing a full update restarts the triggered update timer. Therefore, the triggered update interval you specify must be shorter than the full update interval you have specified with the Full Update Interval parameter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.5

Parameter: Leaf Timeout

Default: 200 seconds

Range: 25 to 4000 seconds

Function: Specifies, in seconds, a value for the virtual interface hold down timer.

Instructions: Determine the virtual hold down timer interval you require and specify a value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.6

Parameter: Neighbor Timeout

- Default:** 140 seconds
- Range:** 40 to 8000 seconds
- Function:** Specifies, in seconds, how long a connection with a router neighbor is considered active without receiving a subsequent probe or report from the neighbor.
- Instructions:** Determine a neighbor timeout period and specify a value.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.12.1.7

Parameter: Route Expiration Timeout

- Default:** 200 seconds
- Range:** 20 to 4000 seconds
- Function:** Specifies, in seconds, how long a route is considered valid without the receipt of a subsequent update indicating the route is reachable.
- This value represents the duration of time that this route will be used. Upon expiration of this timer, this route is advertised as unreachable until is refreshed or until is garbaged.
- Instructions:** Enter a value that represents the duration of time this route will be used without being refreshed.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.12.1.8

Parameter: Garbage Timeout

Default: 340 seconds

Range: 40 to 8000 seconds

Function: Specifies, in seconds, the duration of time that this route will be included in routing updates without the receipt of a subsequent update indicating that the route is reachable.

The difference between this value and the Route Expiration Timeout value represents the duration of time that the route will be advertised as unreachable without subsequent refreshment.

Instructions: Enter a Garbage Timeout value that is greater than the value you specified for Route Expiration Timeout to allow for sufficient time for the route to be advertised as unreachable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.9

Parameter: Estimated Routes**Default:** 25 routes**Range:** An integer of 10 or greater**Function:** Specifies the estimated number of routes.**Instructions:** Enter a value that the router can use for preallocating routing tables. For an Mbone deployment, a value of 1400 or higher is recommended.

Note that routes are kept on a per source network basis, independent of multicast groups. This number must include a route for every network that is local to a circuit configured for multicasting. This is to allow the router to utilize memory efficiently; exceeding this size during router operation will not cause an error but may cause the router to consume more memory than is required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.1.10**Parameter: Neighbor Probe Interval****Default:** 190 seconds**Range:** 10 to 8000 seconds**Function:** Specifies how often to send a probe on virtual interfaces from which no neighbors have been heard.**Instructions:** If your neighbor is running DVMRP *mrouted*, ensure that your probe interval value matches the value used by the neighbor. For interaction with newer versions of *mrouted*, a value of 10 is recommended.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.12.1.11

| | |
|-------------------|---|
| Parameter: | Route Switch Timeout |
| Default: | 140 seconds |
| Range: | 20 to 2000 seconds |
| Function: | Specifies how long to wait, without receiving a subsequent route update from the original neighbor, before switching to a different neighbor advertising equal cost for this route. |
| Instructions: | If your neighbor is running DVMRP <i>mrouterd</i> , the recommended value is 140 seconds. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.12.1.12 |

Editing DVMRP Circuit Parameters

DVMRP is configured on a per circuit basis. To edit DVMRP circuit parameters:

1. Select **IP→Multicast→DVMRP→Circuit**. The DVMRP Circuit Parameters window appears (see Figure 7-3).
2. Edit the parameters as described in the following section.
3. Click on **Done**.

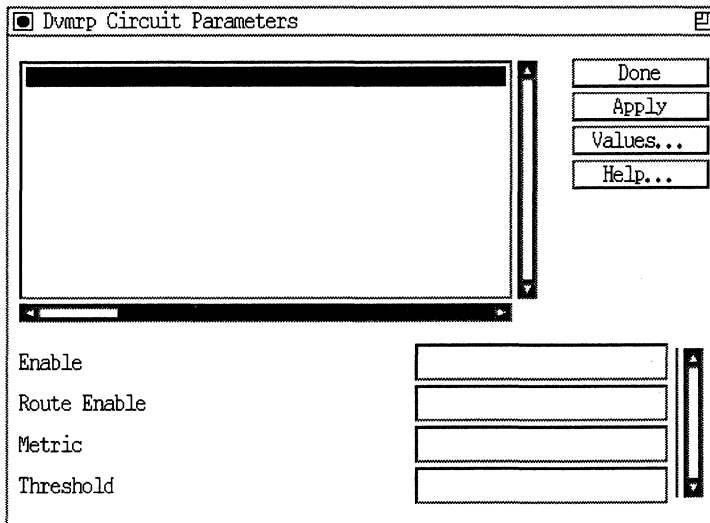


Figure 7-3. DVMRP Circuit Parameters Window

DVMRP Circuit Parameter Descriptions

Use this section as a guide for setting DVMRP circuit parameters.

| | |
|-------------------|---|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Enables or disables DVMRP on this circuit. |
| Instructions: | If you have configured DVMRP on this circuit, enter Disable to disable it. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.12.2.1.2 |

Parameter: Route Enable**Default:** Enable**Options:** Enable | Disable**Function:** Enables or disables this circuit for routing.**Instructions:** Specify Enabled if you want this circuit to be used to propagate routing information and if you want information about the source network associated with this circuit incorporated into routing updates.

Specify Enabled if you want multicast datagrams to be forwarded on this circuit in “native mode” — that is, as multicast datagrams. You can configure tunnels on this circuit.

Specify Disabled if you want this circuit to exist only to support unicast tunnels. If you specify Disabled, all other DVMRP circuit parameters are ignored. The source network associated with this circuit is not incorporated into the routing updates.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.5

Parameter: Metric

Default: 1

Range: 1 to 31

Function: Specifies the cost of this interface.

Instructions: Determine the cost that you want to assign to this hop and enter a value. We recommend the following values:

| Hop | Metric |
|---|---------------------------|
| LAN, or tunnel across a single LAN | 1 |
| multihop tunnel | 2 or 3 |
| serial link, or tunnel across a serial link | 1 |
| backup tunnel | primary tunnel metric + 1 |

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.6

Parameter: Threshold

Default: 1 hop

Range: 1 to 254 hops

Function: Specifies a time to live (TTL) value for the interface. This value is the minimum IP TTL required for a multicast datagram to be forwarded out this interface.

Instructions: Use this parameter to control the scope of the datagrams. If the IP TTL is less than the Threshold value you specify, the datagram is dropped by the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.2.1.7

Editing DVMRP Tunnel Parameters

To edit DVMRP tunnel parameters:

1. Select IP→Multicast→DVMRP→Tunnel. The DVMRP Tunnel Parameters window appears (see Figure 7-4).
2. Edit tunnel parameters as described in the following section.
3. Click on Done.

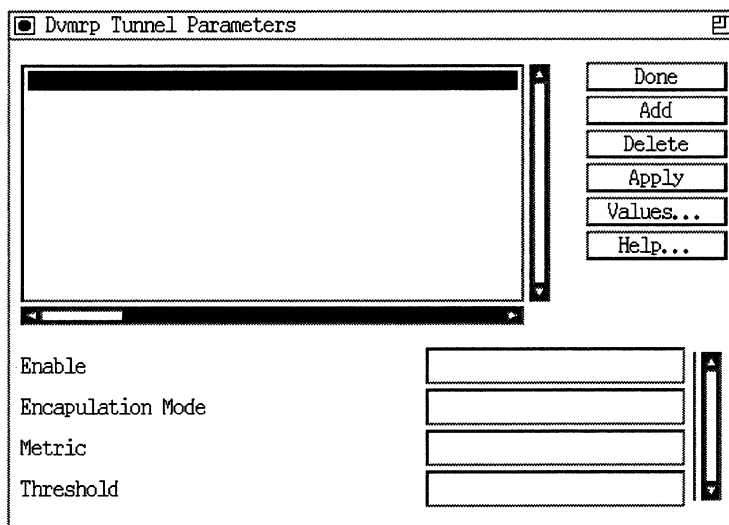


Figure 7-4. DVMRP Tunnel Parameters Window

DVMRP Tunnel Parameter Descriptions

Use this section as a guide for setting DVMRP tunnel parameters.

Parameter: **Enable**
Default: Enable
Options: Enable | Disable
Function: Enables or disables this tunnel interface.
Instructions: If you have configured this tunnel, specify Disable to disable the tunnel.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.2

Parameter: **Encapsulation Mode**
Default: IPINIP
Options: IPINIP | LSSR
Function: Specifies whether tunneled datagrams are encapsulated within an IP datagram or loosely encapsulated using the LSSR option.
Instructions: See RFC 1075 for information about the LSSR option, which is provided for backward compatibility.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.6

Parameter: Metric

Default: 1

Range: 1 to 31

Function: Specifies the cost of this tunnel.

Instructions: Determine the cost you want to assign to this hop and enter a value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.7

Parameter: Threshold

Default: 1 hop

Range: 1 to 254 hops

Function: Specifies a time to live (TTL) value for the tunnel. This value is the minimum IP TTL required for a multicast datagram to be forwarded out this tunnel.

Instructions: Use this parameter to control the scope of the datagrams. If the IP TTL is less than the Threshold value you specify, the datagram is dropped by the router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.8

Adding a DVMRP Tunnel

To add a tunnel to an interface, begin at the DVMRP Tunnel Parameters window.

1. Click on the Add button. The DVMRP Tunnel Address window appears (Figure 7-5).
2. Enter a local and remote IP address for the tunnel.
3. Click on OK.

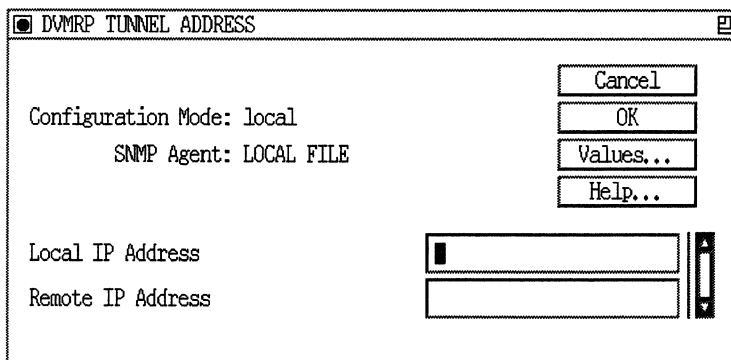


Figure 7-5. DVMRP Tunnel Address Window

Add Tunnel Parameters Descriptions

Use this section as a guide for setting DVMRP Tunnel Address parameters.

Parameter: Local IP Address

Default: Null

Options: The unicast IP address of an interface on a circuit supporting multicasting on the local router

Function: Identifies the local end of the tunnel

Instructions: To identify a unicast tunnel, you must supply the unicast IP address of both ends of the tunnel: the local interface and the remote interface. Use this parameter to enter the local IP address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.4

Parameter: Remote IP address

Default: Null

Options: The unicast IP address of an interface supporting multicasting on a neighboring router

Function: Identifies the remote end of the tunnel

Instructions: To identify a unicast tunnel, you must supply the unicast IP address of both ends of the tunnel: the local interface and the remote interface. Use this parameter to enter the local IP address.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.12.3.1.5

Editing IGMP Global Configuration Parameters

To edit IGMP global parameters:

1. Select IP→Multicast→IGMP→Global. The IGMP Global Configuration window appears (see Figure 7-6).
2. Edit the parameters as described in the following section.
3. Click on Save.

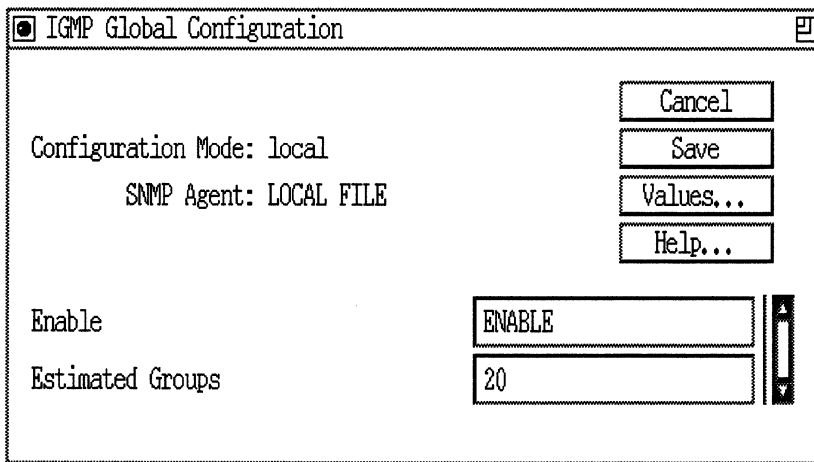


Figure 7-6. IGMP Global Configuration Parameters Window

IGMP Global Configuration Parameters Description

Use this section as a guide for setting IGMP global configuration parameters.

Parameter: **Enable**

Default: Enable

Options: Enable | Disable

Function: Enables or disables this IGMP record.

Instructions: If you have configured IGMP on this router, use this parameter to disable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.1.2

Parameter: **Estimated Groups**

Default: 20 groups

Range: 5 to 65535 groups

Function: Specifies the estimated number of groups that will be simultaneously active for this router.

Instructions: Determine the approximate number of groups and enter the value. This is to allow the router to utilize memory efficiently; exceeding this size during router operation will not cause an error but may cause the router to consume more memory than required.

Note: The following groups are not maintained by IGMP; you do not need to include them in the count: 224.0.0.1, 224.0.0.4, 224.0.0.5, 224.0.0.6,

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.1.4

Editing IGMP Entry Interface Parameters

To edit IGMP Entry Interface parameters:

1. Select IP→Multicast→IGMP→Entry. The IGMP Entry Interface Parameters window appears (see Figure 7-7).
2. Edit the parameters as described in the following section.
3. Click on Save.

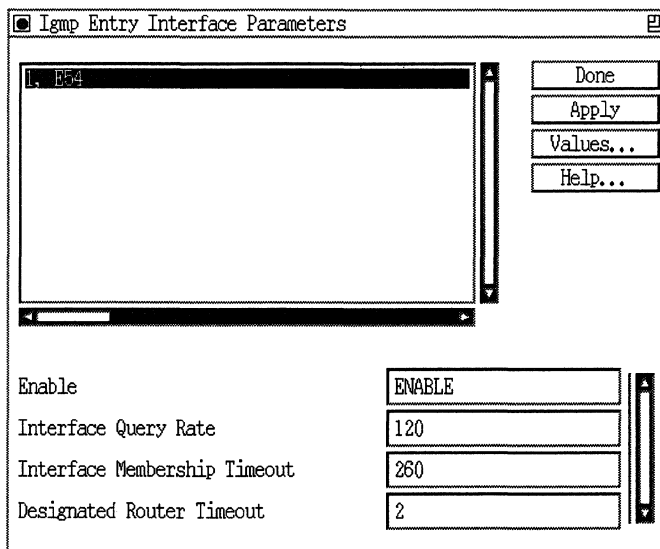


Figure 7-7. IGMP Entry Interface Parameters Window

IGMP Entry Interface Parameters Description

Use this section as a guide for setting IGMP interfaces parameters.

| | |
|-------------------|--|
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Indicates whether this IGMP interface record is to be enabled or disabled. |
| Instructions: | If you have configured IGMP on this interface, use this parameter to disable it. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.13.2.1.2 |

| | |
|-------------------|---|
| Parameter: | Interface Query Rate |
| Default: | 120 seconds |
| Range: | 0 to 4096 seconds |
| Function: | Specifies, in seconds, how often the router sends out group membership queries on the interface. |
| Instructions: | If there are no multicast hosts on this circuit, set the parameter to 0 to disable queries. Specifying 0 affects queries only. The router still forwards multicast datagrams on this circuit. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.13.2.1.5 |

Note: If another IGMP router on this network has taken on the query role, this router will not send out queries unless it has not heard of any queries within the number of seconds specified by the Designated Router Timeout parameter.

Parameter: Interface Membership Timeout

Default: 260 seconds

Range: 30 to 8192 seconds

Function: Specifies, in seconds, the amount of time that a local group membership is valid without the receipt of a subsequent report for that group.

Instructions: The suggested value is $(2 * \text{Query Rate}) + 20$.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.2.1.6

Parameter: Designated Router Timeout

Default: 140 seconds

Range: 10 to 8192 seconds

Function: Specifies, in seconds, the amount of time that can elapse after the last host query message before the IGMP designated router is considered down.

Instructions: The value you specify should be greater than the query rate of all IGMP routers on the network.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.13.2.1.7

Chapter 8

NetBIOS over IP

The Network Basic Input-Output System (NetBIOS) is a session layer communications service used by client and server applications in IBM Token-Ring and PC LAN networks.

NetBIOS provides applications with a programming interface for sharing services and information across a variety of lower-layer network protocols, including IP. Figure 8-1 shows the position of NetBIOS and IP in a simple network architecture.

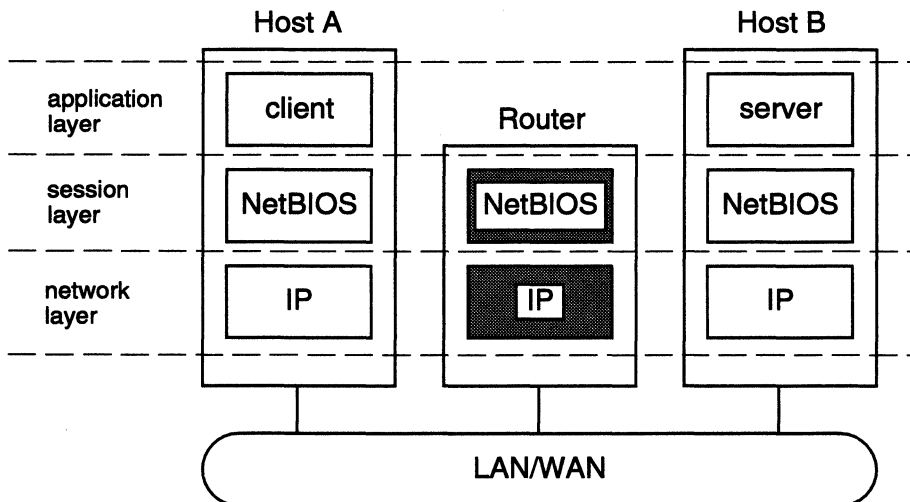


Figure 8-1. NetBIOS over IP

The following sections show you how to configure and customize Wellfleet router software to support NetBIOS in an IP environment:

- “Overview of NetBIOS Services” on page 8-2
- “Customizing IP Support for NetBIOS” on page 8-3
- “Configuring and Customizing a NetBIOS Cache” on page 8-6
- “Editing NetBIOS Parameters” on page 8-9

Overview of NetBIOS Services

There are three categories of NetBIOS services: the name service, the session service, and the datagram service.

The NetBIOS name service allows an application to

- Verify that its own NetBIOS name is unique. The application issues an Add Name Query to NetBIOS. NetBIOS broadcasts the Add Name Query, containing the name. NetBIOS applications that receive the query return an Add Name Response or a Name in Conflict Response. If no response to the query is received after (typically) six broadcasts, the name is considered to be unique.
- Delete a NetBIOS name that the application no longer requires.
- Use a server’s NetBIOS name to determine the server’s network address. The application issues a Name Query Request to NetBIOS containing the target server’s NetBIOS name. NetBIOS broadcasts the Name Query Request. The server that recognizes the name returns a Name Query Response containing its network address.

The NetBIOS session service allows an application to conduct a reliable, sequenced exchange of messages with another application. The messages can be up to 131,071 bytes long.

The NetBIOS datagram service allows an application to exchange datagrams with a specific application or to broadcast datagrams to a group and receive datagrams from the group. Datagrams allow applications to communicate without establishing a session. When a NetBIOS application wants to send information that does not require

acknowledgment from the destination application, the application can transmit a NetBIOS datagram.

This chapter describes Wellfleet IP support for the NetBIOS Name Service, the NetBIOS session service, and the NetBIOS datagram service.

Customizing IP Support for NetBIOS

The NetBIOS name service and datagram service rely on the capability of the underlying network to broadcast Name Query requests to all NetBIOS applications. In a NetBIOS over IP environment, it is the responsibility of the IP router to ensure that the broadcast queries reach all appropriate network segments. To do this, the router

1. Analyzes each NetBIOS packet received on any NetBIOS interface to determine if the packet is a broadcast packet.
2. Rebroadcasts each broadcast packet out all appropriate interfaces except the one on which it was received (readdressing the packet if required).

If alternate paths exist between different network segments, broadcasting loops can occur. To prevent such loops, the router

1. Stamps the data portion of the IP packet with the IP address of the router from which the packet was rebroadcast.
2. Parses the IP addresses included in the data portion of the IP packet to determine if the packet has already been rebroadcast by that router.

In Figure 8-2, for example, client (c) on the network connected to Router B wishes to communicate with server (s), which is located on the network connected to Router A.

1. The client issues a Name Query Request to NetBIOS on the host, specifying the server application by its NetBIOS name. The IP service on the host broadcasts the Name Query Request.

2. Router B receives the Name Query Request, determines that it is a broadcast message, and rebroadcasts it out each of its NetBIOS interfaces (except for the one on which it arrived).
3. Router A receives the broadcast Request and rebroadcasts to its local network.
4. The server on Router A receives the IP broadcast Request and recognizes its own name.

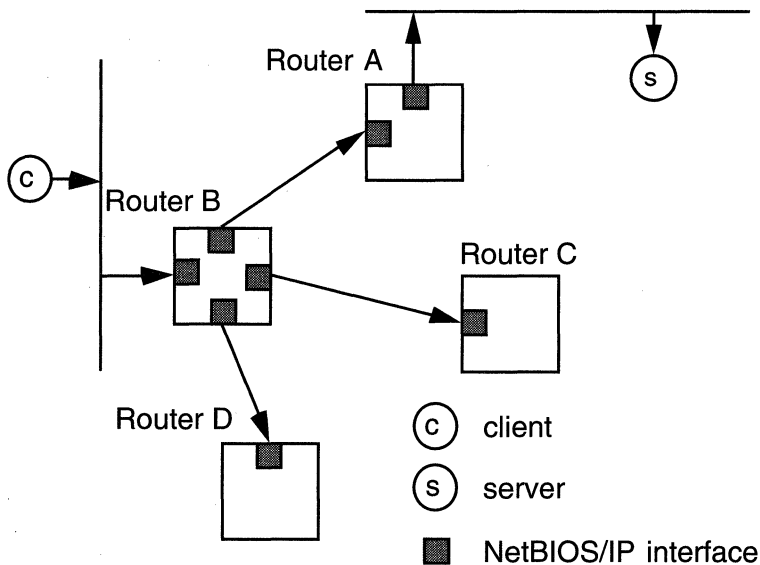


Figure 8-2. Broadcasting a Name Query Request

The server responds to the Name Query Request by issuing a Positive Name Query Response, containing the IP address of the server, to NetBIOS on the host. The following steps occur (see Figure 8-3):

1. NetBIOS sends the Response to Router A as a unicast message.
2. Router A and Router B forward the unicast Response to the awaiting client.

Now that the client has obtained the server's IP address from the Name Query Response, client and server can communicate by exchanging IP messages.

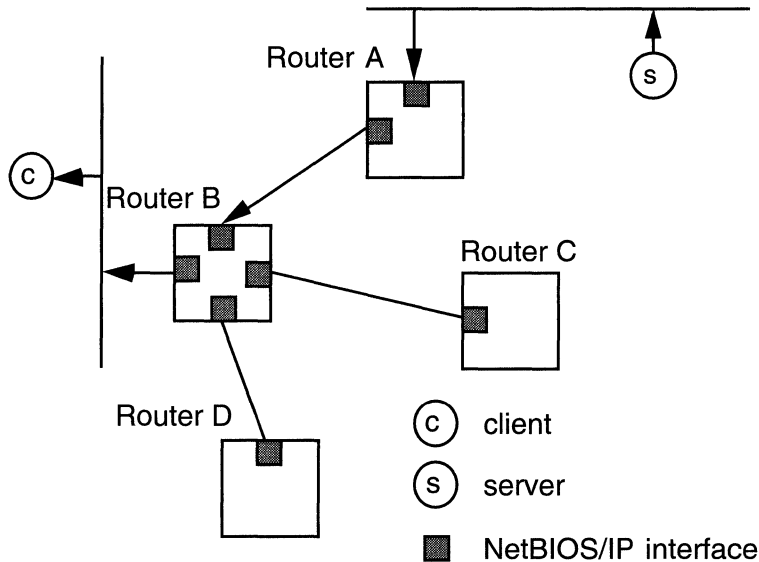


Figure 8-3. Returning a Unicast Name Query Response

You can control the way the IP router rebroadcasts NetBIOS Name Query Requests. For instructions, see the Rebroadcast Packet TTL parameter on page 8-16 and the Rebroadcast Record Route parameter on page 8-16.

Configuring a Static NetBIOS Name

You can add static NetBIOS names into the router. These entries are independent of the name entries learned dynamically in the name cache.

When you configure a static name, you must specify its NetBIOS scope — that is, the area of the network across which the name is known. Each NetBIOS scope has a Scope Identifier, a string of characters meeting the requirements of the Domain Name System. (All NetBIOS names are represented in a manner consistent with the definition for

“compressed name messages” outlined in the Domain Name Service Specification — RFC833.

For a description of NetBIOS static entry parameters and instructions that you can follow to configure name entries, see “Editing NetBIOS/IP Static Entry Table Parameters” on page 1-20.

Configuring and Customizing a NetBIOS Cache

NetBIOS is a broadcast intensive protocol. Much of the broadcast overhead is related to maintaining unique names across the network and providing end users with access to NetBIOS applications. The amount of overhead grows with the number of NetBIOS resources (applications, servers, and clients) on the network.

To keep broadcast traffic to a minimum, each router that runs NetBIOS over IP builds and maintains a cache of NetBIOS name/IP address pairs, using information contained in the Name Query Responses it receives and forwards.

In Figure 8-3, for example:

1. Router A receives a Name Query Response from the server. The router gleans from the Name Query Response the name and IP address of the server.
2. The router stores the name and IP address of the server in its cache.
3. The router forwards the Name Query Response.

Routers that support NetBIOS must analyze each Name Query Request received on a NetBIOS interface to determine if the name of the requested resource (typically, a server) is in the cache. If so, the router replaces the broadcast address on the Request with the unicast IP address of the server. The router then forwards the Name Query Request to the server.

For instructions on customizing a NetBIOS cache, see “Editing NetBIOS/IP Global Parameters” on page 8-9. To enable caching on a NetBIOS interface, see the NetBIOS Name Caching parameter on page 8-18.

Aging a Cache Entry

The router ages cache entries to ensure that cached routes remain consistent with the current network topology. If the cache table lookup mechanism does not access a cache entry within the interval you set in the appropriate Cache Aging Time parameter, the router deletes the entry from the table.

If the router receives a broadcast Name Query Request from a client and finds the name and associated IP address of the requested server in its cache, the router replaces the broadcast address on the Name Query Request with the unicast IP address. The router also assigns the entry a short time to live. If the entry is valid, the router will receive a Positive Name Query Response (which will validate the entry) from the server within the specified time to live. If the entry is invalid, the Name Query Request will not reach the server. In this case, the entry quickly ages out.

For instructions on specifying an age value for cache entries, see the Name Cache Age parameter on page 8-14.

Customizing a Cache Search

The mechanism that NetBIOS uses to search for a name in the cache is based on a fast string hash/search mechanism developed for AppleTalk Zone Name processing. This mechanism uses a hash table that NetBIOS builds and maintains on the router. You can specify the number of entries in the hash table. For instructions, see the Hash Entry Count parameter on page 8-15.

Increasing the number of entries in the hash table

- ❑ Decreases the likely number of names the router must compare before finding a specific cached name
- ❑ Decreases the amount of time it takes the router to find a particular cached name
- ❑ Increases memory usage

Note that increasing the number of entries in the hash table does not increase the number of names the router can cache. This is determined by user configuration and by available memory.

Adding a Traffic Filter to a NetBIOS Interface

If name caching is enabled, a router that receives a Name Query Response (originating from a server and addressed to a client) must be able to deliver the message to the NetBIOS entity on the router (rather than simply forward it out another interface toward its destination).

To enable the router to recognize a unicast IP packet that contains a Name Query Response and pass it to NetBIOS through UDP port 137, you must configure a traffic filter on each NetBIOS interface that receives unicast Name Query Responses.

1. Beginning at the Wellfleet Configuration Manager window, select **Circuits→Edit**. The Circuit List window appears.
2. Click on **Edit**. The Circuit Definition window appears.
3. Select **Protocols→Edit IP→Traffic Filters**. The IP Filters window appears.
4. Click on **Template**. The Filter Template Management window appears.
5. Click on **Create**. The Create IP Template window appears.
6. Select **Criteria→Add→UDP Frame→Destination Port**. The Edit Range screen appears.
7. Enter 137 for the minimum value and the maximum value. Click on **OK**.

8. The Create IP Template window appears.
9. Select Action→Add→Forward to Next Hop. The Next Hop window appears.
10. Enter the IP address of this interface (the interface on which you are configuring the traffic filter). Click on OK.

Editing NetBIOS Parameters

Once you have configured a circuit to support NetBIOS/IP, you invoke NetBIOS windows to edit NetBIOS parameters.

Editing NetBIOS/IP Global Parameters

To access and edit global NetBIOS parameters, complete the following steps:

1. From the Wellfleet Configuration Manager window, select Protocols→IP→NetBIOS→Global.

The Edit NetBIOS/IP Global Parameters window appears (see Figure 8-4).

2. Edit the parameters you want to change.
3. Click on the OK button to save your changes and exit the window

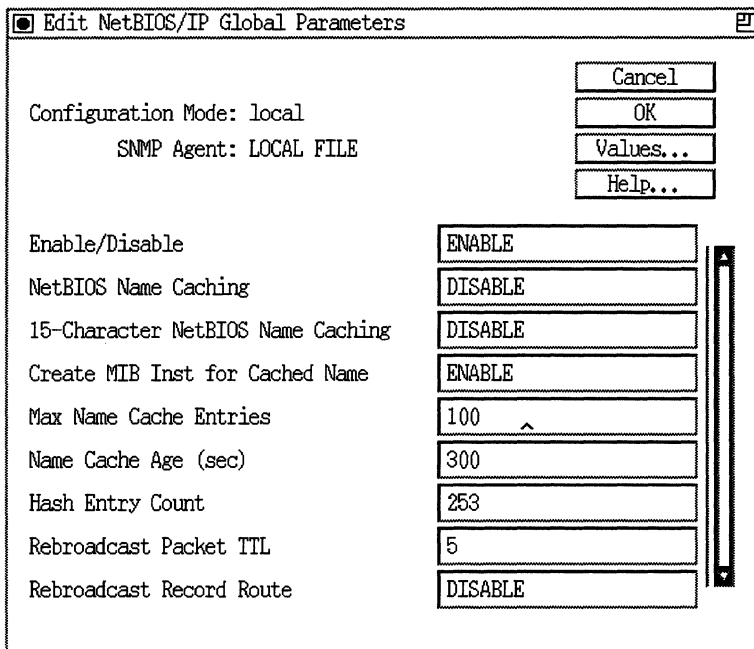


Figure 8-4. Edit NetBIOS/IP Global Parameters Window

NetBIOS Global Parameters

Use the following descriptions as a guide when you configure NetBIOS parameters on the Edit NetBIOS/IP Global Parameters window.

| | |
|-------------------|--|
| Parameter: | Enable/Disable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Enables or disables NetBIOS on this router. |
| Instructions: | If NetBIOS has been configured on this router, use this parameter to disable and re-enable it as required. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.11.1.2 |

Parameter: NetBIOS Name Caching

Default: Disable

Options: Enable | Disable

Function: Globally enables or disables the ability of the router to cache the name associated with each NetBIOS server that is active on the network.

Instructions: Select Enable to activate NetBIOS server name caching at every NetBIOS interface configured on the node.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.4

Parameter: 15-Character NetBIOS Name Caching

Default: Disable

Options: Enable | Disable

Function: Enables or disables the ability of the router to treat a NetBIOS name either as a 15- or 16-character entity.

Instructions: Select Enable to activate 15-character NetBIOS name caching at every NetBIOS interface configured on this router.

Select Disable if you want NetBIOS to treat names as 16-character entities.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.5

Parameter: Create MIB Inst for Cached Name

Default: Enable

Options: Enable | Disable

Function: Enables or disables the ability of the system to

- Create a MIB instance for each name entry stored in the name cache.
- Delete a MIB instance for each NetBIOS name entry that ages out of the name cache.

Instructions: Select Disable if you want to release the system memory and processing resources otherwise dedicated to maintaining cached names in the MIB.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.6

Parameter: Max Name Cache Entries

Default: 100 entries

Range: 1 to 2147483647 entries

Function: Specifies the maximum number of entries you need to provide in the NetBIOS name cache.

Instructions: You can adjust the value of this parameter in direct proportion to the total number of server names expected to be active during intervals of peak traffic load or performance demand on the router. A value of 100 is suitable for networks that include up to 100 NetBIOS names to cache.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.7

| | |
|-------------------|---|
| Parameter: | Name Cache Age |
| Default: | 300 seconds |
| Range: | Any value that can rapidly age infrequently referenced names out of the NetBIOS name cache. |
| Function: | Specifies an age (in seconds) when inactive NetBIOS names expire from the NetBIOS Name Cache. |
| Instructions: | Choose an aging value that allows infrequently referenced or obsolete server names to expire from the name cache. The smaller the value, the less efficient broadcast reduction is, but the more quickly the network recovers topology changes. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.11.1.9 |

Parameter: Hash Entry Count**Default:** 253**Range:** Any integer value**Function:** Specifies the number of entries you want to allow in the cache lookup tables. Each NetBIOS interface has a local table to store and retrieve the names of NetBIOS servers active on the network.**Instructions:** For networks that actively use up to 2500 NetBIOS server names, use the default value (253). To determine a Hash Entry Count for larger networks:

1. Divide the total number of unique NetBIOS server names active in the network by 10.
2. Adjust the quotient to the nearest (higher or lower) prime number. (A prime number can only be divided by itself or 1 and still yield a whole-number quotient.)
3. Replace the default value with the new, calculated number.

Increasing the number of hash table entries does not increase the number of names that a router can cache. With larger networks, increasing the size of the hash tables may, however, reduce internal cache lookup time, thereby improving overall performance.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.10

Parameter: Rebroadcast Packet TTL

Default: 5 seconds

Range: 1 to 255 seconds

Function: Specifies the time-to-live value in seconds to use in rebroadcast packets

Instructions: Use this parameter to restrict the number of routers a rebroadcast packet can traverse. To prevent NetBIOS broadcast packets from traversing the network indefinitely, set the parameter to a minimal value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.13

Parameter: Rebroadcast Record Route

Default: Disable

Options: Enable | Disable

Function: Enables and disables the Insertion of Record Route option in rebroadcast packets.

Instructions: If all IP entities support this option, select Enable to allow the NetBIOS entity in the router to determine whether it has received this packet before on this interface. If so, the router drops it. The Record Route option prevents rebroadcast packets from looping forever.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.1.14

Editing NetBIOS/IP Interface Table Parameters

To edit NetBIOS interface parameters, complete the following steps:

1. From the Wellfleet Configuration Manager window, select **Protocols→IP→NetBIOS→Interface** to display the NetBIOS/IP Interface Table window (see Figure 8-5).

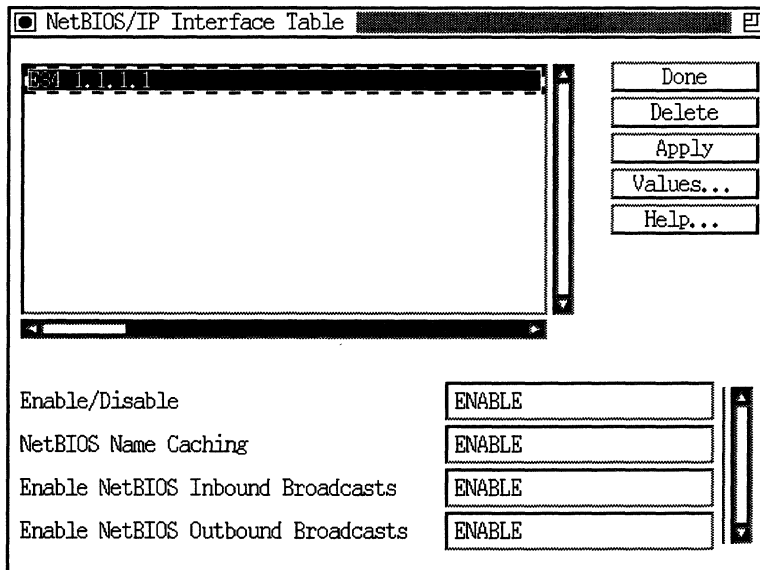


Figure 8-5. NetBIOS/IP Interface Table

2. Select the interface you want to modify.
3. Edit the parameters you want to change.
4. Click on the **Apply** button to save your changes.
5. Click on the **Done** button to exit the IP Interface Table window.

NetBIOS Interface Parameter Descriptions

Use the following descriptions as a guide when you configure parameters on the NetBIOS/IP Interface Table window.

| | |
|-------------------|--|
| Parameter: | Enable/Disable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Enables or disables NETBIOS on this IP interface. |
| Instructions: | If NetBIOS has been configured and enabled on the router, use this parameter to disable and reenable it on this interface as required. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.11.2.1.2 |
| | |
| Parameter: | NetBIOS Name Caching |
| Default: | Enabled |
| Options: | Enable Disable |
| Function: | Enables or disables the ability of this interface to cache the name for each NetBIOS server active in the network. |
| Instructions: | Select Enable if you disabled server name caching previously, and you want now to re-enable that function. Select Disable if you want to release system memory and processing resources otherwise dedicated to server name caching. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.11.2.1.8 |

Parameter: Enable NetBIOS Inbound Broadcasts

Default: Enabled

Options: Enable | Disable

Function: Enables or disables inbound broadcasts on this interface.

Instructions: If NetBIOS is configured and enabled on the router and enabled on this interface, use this parameter to enable and disable inbound broadcasts as required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.2.1.9

Parameter: Enable NetBIOS Outbound Broadcasts

Default: Enable

Options: Enable | Disable

Function: Enables or disables outbound broadcasts on this interface.

Instructions: If NetBIOS is configured and enabled on the router and enabled on this interface, use this parameter to enable and disable outbound broadcasts as required.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.2.1.10

| | |
|-------------------|---|
| Parameter: | Rebroadcast Address |
| Default: | Null |
| Options: | An IP broadcast address |
| Function: | Specifies a broadcast address to use when rebroadcasting NetBIOS packets out this interface. |
| Instructions: | By default, NetBIOS uses the IP broadcast address configured for this interface. Set this parameter if you want to override this broadcast address. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.11.2.1.11 |

Editing NetBIOS/IP Static Entry Table Parameters

The sections that follow describe how to edit, add, and delete statically configured NetBIOS names.

To perform these operations, complete the following steps:

1. From the Wellfleet Configuration Manager window, select Protocols→IP→NetBIOS→Static Name.

The NetBIOS/IP Static Entry Table window appears, showing a list of all statically configured NetBIOS names currently defined (see Figure 8-6).




Figure 8-6. NetBIOS/IP Static Entry Table Window

2. Select the static entry you want to modify.
3. Edit the parameters you want to change, using the descriptions following this procedure as guidelines.
4. Click on Done to exit this screen.

NetBIOS/IP Static Entry Table Parameter Descriptions

Use the following descriptions as guidelines when you configure parameters on the NetBIOS/IP Static Entry Table window.

Parameter: Enable

Default: Enable

Options: Enable | Disable

Function: Enables or disables caching of the NetBIOS name you have selected.

Instructions: Set the parameter to Enabled to activate caching of the name you selected. Set the parameter to Disabled to deactivate caching of the name you selected.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.4.1.2

Parameter: NetBIOS Scope ID

Default: None

Options: A NetBIOS scope identifier

Function: Identifies the area of the network across which the NetBIOS name is known.

Instructions: Enter a name string that meets the requirements of the Domain Name System as described in RFC 833.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.4.1.5

Adding a Statically Configured NetBIOS Name

You may want to statically configure NetBIOS names that are stable elements in your network configuration. Statically configuring a name reduces the use of system memory and processing resources normally required for learning and maintaining NetBIOS names.

To add a statically configured NetBIOS name, complete the following steps:

1. From the NetBIOS/IP Static Entry Table window, click on the Add button.

The NBIP Addresses window appears (see Figure 8-7):

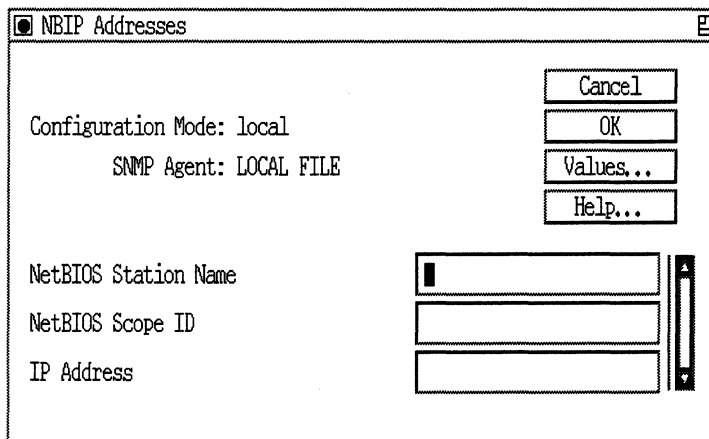


Figure 8-7. NBIP Addresses Window

2. Enter values for the parameters, using the descriptions following this procedure as guidelines.
3. Click on the OK button to save your changes and exit the window.

NetBIOS Addresses Parameter Descriptions

Use the following descriptions as a guide when you configure parameters on the Add NetBIOS Static Name window.

- Parameter:** NetBIOS Station Name
- Default:** None
- Options:** A name string of up to 16 characters.
- Function:** Specifies the name of a NetBIOS station
- Instructions:** Enter the NETBIOS name you want to add. The name must not exceed 16 characters. The system pads names shorter than 16 characters with ASCII space characters. To enter non-ASCII values in the name, use the form `\xbbb`, where *bb* can be any two hexadecimal digits.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.11.4.1.4
-
- Parameter:** NetBIOS Scope ID
- Default:** None
- Options:** A NetBIOS scope identifier
- Function:** Identifies the area of the network across which the NetBIOS name is known
- Instructions:** Enter a name string that meets the requirements of the Domain Name System as described in RFC 833.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.11.4.1.5

Parameter: IP Address

Default: None

Options: The IP address of the NetBIOS station

Function: Specifies an IP address to associate with the statically configured name

Instructions: Enter a valid IP address of a NetBIOS station.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.11.4.1.6



Chapter 9

IP Policies

This chapter provides an overview of IP policies and describes the Site Manager windows you use and the parameters you set to create accept and announce policies for RIP, OSPF, BGP-3, BGP-4, and EGP.

Note: The parameters you set when you construct IP accept and announce policies are a superset of the parameters you set when you construct import and export filters.

We currently support both IP accept and announce policies and import and export filters. In a future release, support for import and export filters will be dropped.

This chapter contains the following sections:

- “IP Routing Table” on page 9-1
- “Configuring Accept Policies” on page 9-5
- “Configuring Announce Policies” on page 9-25

IP Routing Table

Every IP router maintains a table of current routing information. The routing table manager receives routing updates from the network through the Internet protocols running on the router. Periodically, the routing table manager issues routing updates through the protocols. Figure 9-1 shows a router configured with all of the Internet protocols

supported by Wellfleet: OSPF, RIP, BGP-3, BGP-4, and EGP. The arrows indicate the direction of flow of routing information between the network and the protocols running on the router, between the protocols and the routing table manager, and between the routing table manager and the routing table.

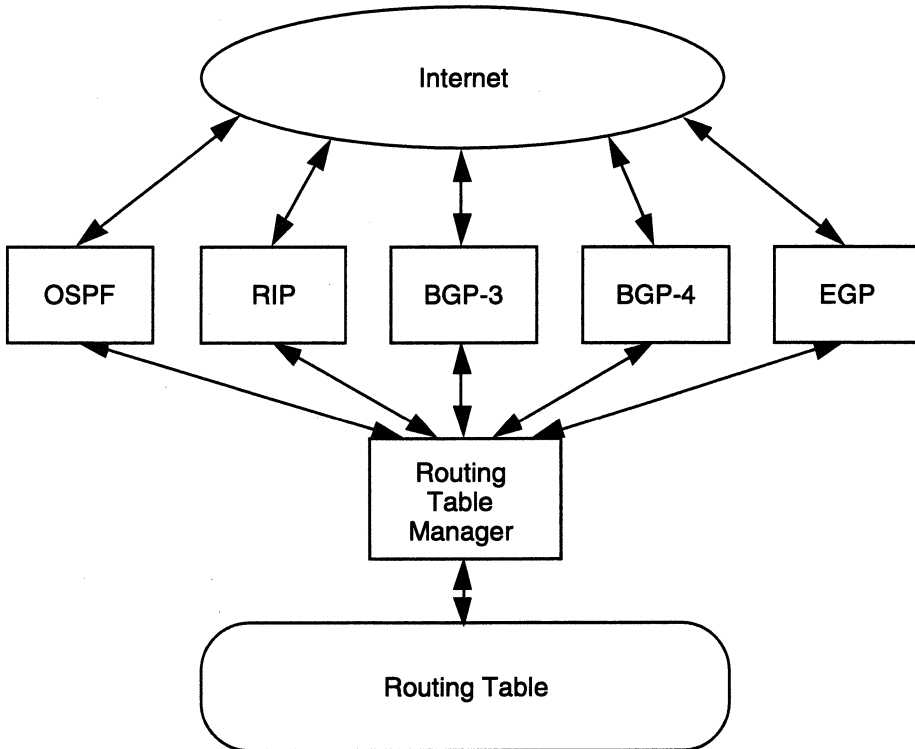


Figure 9-1. IP Routing Table

The flow of routing information between the network, the protocols, and the routing table manager is controlled by *routing information policies*.

Each time a routing update arrives from a remote router, the following steps occur (see Figure 9-2):

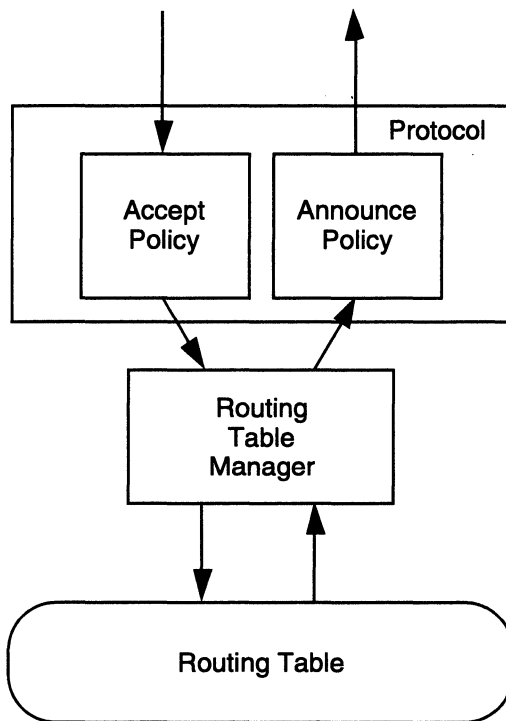


Figure 9-2. Accept and Announce Policies

1. The protocol receiving the route consults an *accept policy* to determine whether to forward the route to the IP routing table manager or drop the route.
2. If the protocol forwards the route, the routing table manager determines whether to inject the route into the routing table.

Periodically, the routing table manager announces routes to other routers in the network.

1. The routing table manager forwards a route for advertisement to the protocol.
2. The protocol consults an *announce policy* to determine whether or not to advertise the route to the network.

Note: The way OSPF applies accept and announce policies to routing information differs in several ways from the procedure shown in Figure 9-2. OSPF link state advertisements (LSAs) are received and placed in the link state database (LSDB) of the router. The information in the LSDB is also propagated to other routers in the OSPF routing domain. According to the OSPF standard, all routers in a given area must maintain a similar database. To maintain database integrity across the network, a router must not manipulate received LSAs before propagating them on to other routers. To accomplish this, OSPF accept and announce policies act in the following manner.

OSPF accept policies control which OSPF non-self-originated external routing information is passed to the routing table manager. The accept policies control only what the local router uses; they do not affect the propagation of OSPF internal and OSPF non-self-originated external information to other routers.

OSPF announce policies control which self-originated external routing updates are placed into the LSDB for distribution according to the OSPF standard. OSPF announce policies affect what other routers learn but only with regard to the local router's self-originated information.

Configuring Accept Policies

To add, edit, or delete an accept policy, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select Protocols→IP→Policy Filters→<protocol>→Accept Policies (*protocol* is a Wellfleet-supported IP protocol: RIP, OSPF, EGP, BGP-3, or BGP-4).

The Accept Policy Filters window for the IP protocol appears. Figure 9-3 shows the Accept Policy Filters window for BGP-3. This window lists all accept policies configured on the router for that protocol and allows you to edit them.

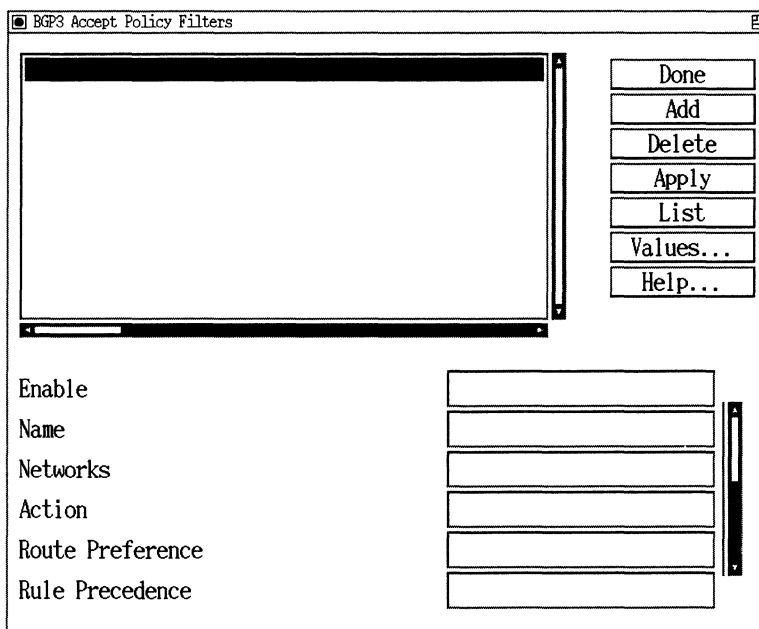


Figure 9-3. BGP-3 Accept Policy Filters Window

2. To add an accept policy, click on the Add button. The Accept Policy Filter Configuration window appears. Figure 9-4 shows the Accept IP Policy Filter Configuration window for BGP-3. Enter the appropriate values and click on the Done button Use the sections that follow as a guide to setting parameter values.

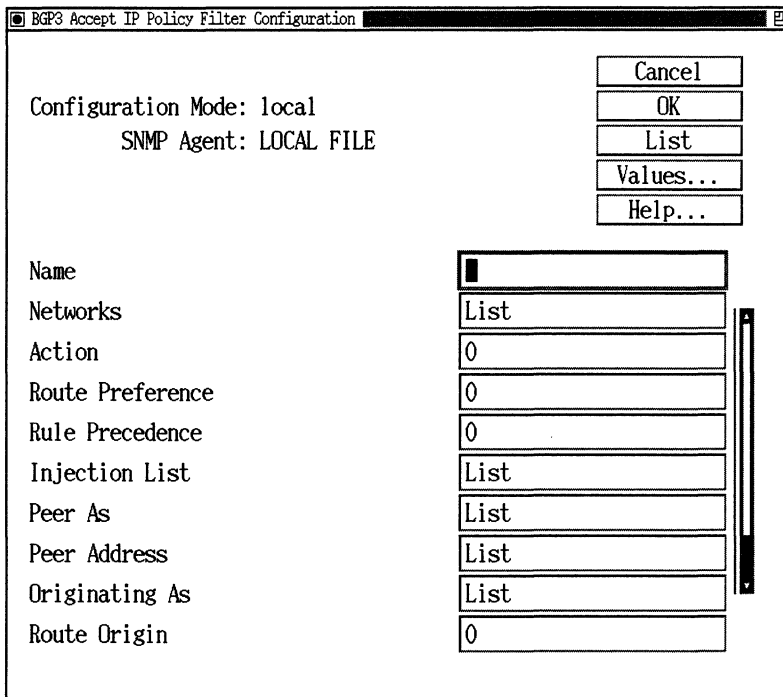


Figure 9-4. BGP-3 Accept IP Policy Filter Configuration Window

IP Accept Policy Parameters Descriptions

IP accept policy parameters fall into two categories: parameters that appear in all IP policies and IP protocol-specific parameters. Accept policy parameters are described in the following sections:

- “Common IP Accept Policy Parameters” on page 9-8
- “RIP-Specific Accept Policy Parameters” on page 9-11
- “OSPF-Specific Accept Policy Parameters” on page 9-12
- “EGP-Specific Accept Policy Parameters” on page 9-13
- “BGP-3-Specific Accept Policy Parameters” on page 9-15
- “BGP-4-Specific Accept Policy Parameters” on page 9-19

Note: Certain accept policy parameters request a list of entries. A list can contain up to 55 entries.

Common IP Accept Policy Parameters

This section describes how to set accept policy parameters common to all IP protocols.

Parameter: **Enable**
Default: Enable
Options: Enable | Disable
Function: Enables or disables this policy.
Instructions: Set to Disable to disable the policy.
MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.2
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.2
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.2
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.2
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.2

Parameter: **Name**
Default: None
Range: Any alphanumeric character string
Function: Identifies this accept policy.
Instructions: Specify a user name for the policy.
MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.4
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.4
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.4
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.4
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.4

| | |
|-------------------|--|
| Parameter: | Networks |
| Default: | An empty list |
| Options: | A list of network identifiers. Each entry consists of a network number, a mask, and a flag to indicate whether the ID refers to a specific network or a range of networks. |
| Function: | Specifies the networks to which this policy applies. |
| Instructions: | Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Use the default empty list to match any route. |
| MIB Object ID: | RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.5 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.5 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.5 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.5 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.5 |
| Parameter: | Action |
| Default: | RIP, OSPF, EGP: Accept BGP-3, BGP-4: Ignore |
| Options: | Accept Ignore |
| Function: | Specifies whether the protocol ignores a route that matches the policy or forwards the route to the routing table manager. |
| Instructions: | Specify Accept to consider the route for insertion in the routing table. To drop the route, specify Ignore. |
| MIB Object ID: | RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.6 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.6 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.6 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.6 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.6 |

Parameter: Route Preference

Default: 1

Range: 1 to 16

Function: Assigns a metric value (the higher the number, the greater the preference) to a route that the protocol forwards to the routing table manager. If confronted with multiple routes to the same destination, the routing table manager may need to use this value to decide which route to insert.

Instructions: Either accept the default value 1, or enter a new value.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.7
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.7
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.7
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.7
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.7

Parameter: Rule Precedence

Default: 0

Range: A metric value

Function: Assigns a metric value to this policy (a policy with a higher value takes precedence over a policy with lower value).

Instructions: Use this value to specify the order of precedence for policies that match the same route.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.8
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.8
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.8
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.8
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.8

RIP-Specific Accept Policy Parameters

This section shows you how to set RIP-specific accept policy parameters.

Parameter: From Gateway

Default: An empty list

Options: A list of IP addresses

Function: Specifies the addresses of one or more routers that could send RIP updates to this router. This policy applies to RIP advertisements from routers on this list.

Instructions: Use the default empty list to indicate that this policy applies to RIP updates from any router.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.10

Parameter: Received on Interface

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP addresses of one or more interfaces on this router. This policy applies to RIP updates received on interfaces that appear on this list.

Instructions: Use the default empty list to indicate that this policy applies to RIP updates received on any interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.11

Parameter: Apply Subnet Mask

Default: Null

Options: Null or IP address mask.

Function: Specifies a mask that will override the interface's subnet mask in the presence of networks with variable length subnet masks.

Instructions: Supply a mask, set the Action parameter to Accept, and use the default Injection List parameter (an empty list).

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.1.1.12

OSPF-Specific Accept Policy Parameters

This section shows you how to set OSPF-specific accept policy parameters.

Parameter: Type

Default: Any

Options: Type 1
Type 2
Any

Function: Describes which types of OSPF ASE routes match this policy.

Instructions: To match either Type 1 or Type 2, use the default, Any.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.3.1.10

| | |
|-------------------|--|
| Parameter: | Tag |
| Default: | An empty list |
| Options: | A list of tag values |
| Function: | Specifies OSPF tag values that could be present in an OSPF ASE advertisement. This policy applies to OSPF ASE advertisements that contain the tag values on this list. |
| Instructions: | Use the default empty list to indicate that this policy applies to OSPF ASE advertisements with any tag value. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.6.3.1.11 |

EGP-Specific Accept Policy Parameters

This section shows you how to set EGP-specific accept policy parameters.

| | |
|-------------------|---|
| Parameter: | Peer List |
| Default: | An empty list |
| Options: | A list of IP addresses |
| Function: | Specifies the IP addresses of one or more EGP peers. This policy applies to EGP advertisements from the peers on this list. |
| Instructions: | Use the default empty list to indicate that this policy applies to EGP advertisements from any EGP peer. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.6.5.1.10 |

Parameter: AS List

Default: An empty list

Options: A list of autonomous system numbers.

Function: Specifies one or more autonomous system numbers. This policy applies to EGP advertisements from peers located in the autonomous systems on this list.

Instructions: Use the default empty list to indicate that this policy applies to EGP advertisements from peers in any AS.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.11

Parameter: Gateway List

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP address of one or more EGP gateways. This policy applies to EGP advertisements that use these gateways as the next hop.

Instructions: Use the default empty list to indicate that this policy applies to EGP advertisements with any gateway address.

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.5.1.12

BGP-3-Specific Accept Policy Parameters

This section shows you how to set BGP-3 specific accept policy parameters.

Parameter: **Injection List**

Default: An empty list

Options: A list of network identifiers

Function: Specifies network IDs to be included in the routing table in place of the network IDs listed in the received advertisement.

Instructions: Specify a non-null value only if the Action parameter is set to Accept. The values you enter in the injection list determine the action taken.

If you supply a list of network IDs, these IDs are injected into the routing table instead of the actual received IDs.

If you use the default (an empty list), the actual received network IDs are injected into the routing table.

If you supply a list that includes the encoding 255.255.255.255/255.255.255.255, the actual received network IDs are injected into the routing table along with the other IDs in the injection list. This allows insertion of an aggregate or default along with the actual networks.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.9

Note: In the current release, the only valid network ID that you can include in an injection list is the default ID, 0.0.0.0/0.0.0.0. The Injection List parameter replaces the received routes with the default route and places the default route in the routing table.

Note: The accept Injection List parameter associates the default route with the attributes of the best route that matches the policy. If you are constructing a BGP-3 or BGP-4 accept policy, keep in mind that the accept Injection List parameter *does not* perform route aggregation as defined in RFC 1654. To aggregate routes in a transit AS, you must construct an announce policy and use the announce Advertise parameter.

Parameter: Peer AS

Default: An empty list

Options: A list of autonomous system numbers, each ranging from 1 to 65536

Function: Specifies one or more autonomous systems. This policy applies to BGP advertisements from peers in those ASs.

Instructions: Use the default empty list to indicate that this policy applies to BGP advertisements from peers in any AS.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.10

Parameter: Peer Address

Default: An empty list

Options: A list of IP addresses

Function: Specifies one or more BGP peers. This policy applies to BGP advertisements from the peers on this list.

Instructions: To indicate that this policy applies to BGP advertisements from any BGP peer, use the default empty list.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.11

Parameter: Originating AS

Default: An empty list

Options: A list of autonomous system numbers

Function: Specifies one or more autonomous systems. This policy applies to BGP advertisements that originate from the ASs on this list.

Instructions: To indicate that the policy applies to BGP advertisements originating from any AS, use the default empty list.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.12

Parameter: Route Origin

Default: Any

Options: Any
IGP
EGP
IGP or EGP
Incomplete
Incomplete or IGP
Incomplete or EGP

Function: Specifies the values of the BGP origin path attribute that apply to this policy.

Instructions: Select the origin values you wish to accept for this policy.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.13

Parameter: BGP-3 Route Preference

Default: 1

Range: 1 to 16

Function: Specifies a value that is used to compare a route that matches this policy with other BGP-3 routes that match the policy. The larger the value, the greater the preference.

Instructions: To specify maximum preference, enter 16. Valid only if the Action parameter is set to Accept.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.14

Parameter: AS Weight Class

Default: Weight Class 1

Options: Weight Class 1 to Weight Class 8

Function: Indicates which weight class value should be used when calculating the AS path weight.

Instructions: Set the Action parameter to Accept and supply a valid BGP-3 weight class.

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.7.1.15

BGP-4-Specific Accept Policy Parameters

This section shows you how to set BGP-4-specific accept policy parameters.

Parameter: **Injection List**

Default: An empty list

Options: A list of network identifiers

Function: Specifies network IDs to be included in the routing table in place of the network IDs listed in the received advertisement.

Instructions: Specify a non-null value only if the Action parameter is Accept. The values you enter in the injection list determine the action taken.

If you supply a list of network IDs, these IDs are injected into the routing table instead of the actual received IDs.

If you use the default (an empty list), the actual received network IDs are injected into the routing table.

If you supply a list that includes the encoding 255.255.255.255/255.255.255.255, the actual received network IDs are injected into the routing table along with the other IDs in the injection list. This allows insertion of an aggregate or default along with the actual network.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.9

Note: In the current release, the only valid network ID that you can include in an injection list is the default ID, 0.0.0.0/0.0.0.0. The Injection List parameter replaces the received routes with the default route and places the default route in the routing table.

Note: The Injection List parameter associates the default route with the attributes of the best route that matches the policy. If you are constructing a BGP-3 or BGP-4 accept policy, keep in mind that the Injection List parameter *does not* perform route aggregation as defined in RFC 1654. To aggregate routes in a transit AS, you must construct an announce policy and use the announce Advertise parameter.

Parameter: Peer AS

Default: An empty list

Options: A list of autonomous system numbers, each ranging from 1 to 65536.

Function: Specifies one or more ASs. This policy applies to BGP advertisements from peers in the autonomous systems on this list.

Instructions: Use the default empty list to indicate that this policy applies to BGP advertisements from peers in any AS.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.10

Parameter: Peer Address

Default: An empty list

Options: A list of IP addresses

Function: Specifies one or more BGP peers. This policy applies to BGP advertisements from the peers on this list.

Instructions: To indicate that this policy applies to BGP advertisements from any BGP peer, use the default empty list.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.9.1.11

Parameter: Originating AS**Default:** An empty list**Options:** A list of autonomous system numbers**Function:** Specifies one or more autonomous systems. This policy applies to BGP advertisements that originate from the ASs on this list.**Instructions:** To indicate that the policy applies to BGP advertisements originating from any AS, use the default empty list.**MIB Object ID:** BGP-4 1.3.6.1.4.1.18.3.5.3.2.6.9.1.12**Parameter: Route Origin****Default:** Any**Options:** Any
IGP
EGP
IGP or EGP
Incomplete
Incomplete or IGP
Incomplete or EGP**Function:** Specifies which values of the BGP origin attribute apply to this policy.**Instructions:** Select the origin values you wish to accept for this policy.**MIB Object ID:** BGP-4 1.3.6.1.4.1.18.3.5.3.2.6.9.1.13

Parameter: Aggregator AS List

Default: An empty list

Options: A list of AS numbers

Function: Specifies one or more autonomous systems. This policy applies to BGP advertisements that contain in their Aggregator path attribute an AS number on this list.

Instructions: To specify that the policy applies to BGP advertisements with any AS number in the Aggregator path attribute, use the default empty list.

MIB Object ID: BGP-4 1.3.6.1.4.1.18.3.5.3.2.6.9.1.14

Parameter: Aggregator Router List

Default: An empty list

Options: A list of IP addresses

Function: Specifies one or more BGP routers. This policy applies to BGP advertisements that contain in their Aggregator path attribute an IP address on this list.

Instructions: To specify that this policy applies to BGP advertisements with any router address in the AGGREGATOR path attribute, use the default empty list.

MIB Object ID: BGP-4 1.3.6.1.4.1.18.3.5.3.2.6.9.1.15

Parameter: Local Preference

Default: 0

Range: 0 to 4294967295

Function: Assigns a local preference value to a route matching this policy. This value overrides the calculated value for EBGp routes or the Local Preference path attribute for IBGP routes.

Instructions: To indicate a preference, enter a value from 1 to 4294967295.

MIB Object ID: BGP-4 1.3.6.1.4.1.18.3.5.3.2.6.9.1.16

Parameter: BGP-4 Preference

Default: 1

Options: 1 to 16

Function: Specifies a value that can be used to compare a route that matches this policy with other BGP-4 routes. The larger the value, the greater the preference.

Instructions: To indicate maximum preference, enter 16. This parameter is valid only if the Action parameter is set for Accept.

MIB Object ID: BGP-4 1.3.6.1.4.1.18.3.5.3.2.6.9.1.17

Parameter: **AS Weight Class**
Default: Weight Class 1
Range: Weight Class 1 to Weight Class 8
Function: Indicates which weight class value should be used when calculating the AS path weight.
Instructions: Enter a valid BGP-4 weight class. Valid only if the Action parameter is set for Accept.
MIB Object ID: BGP-4 1.3.6.1.4.1.18.3.5.3.2.6.9.1.18

Configuring Announce Policies

To add, edit, or delete announce policies, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select Protocols→IP→Policy Filters→*protocol*→Announce Policies Filters (*protocol* is RIP, OSPF, EGP, BGP-3 or BGP-4).

The Announce Policy Filters window appears. Figure 9-5 shows the BGP-3 Announce Policy Filters window. This window lists all announce policies configured on the router for that protocol. You edit announce policies from this window.

2. To add an announce policy, click the Add button. The BGP-3 Announce Policy Filters Configuration window for the protocol appears (see Figure 9-6). Set the parameters and click on the Done button.

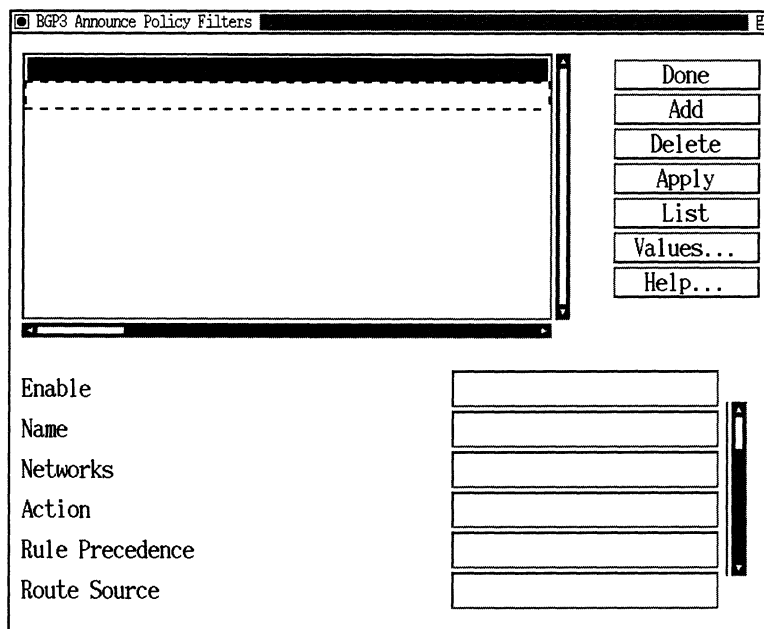


Figure 9-5. BGP-3 Announce Policy Filters Window

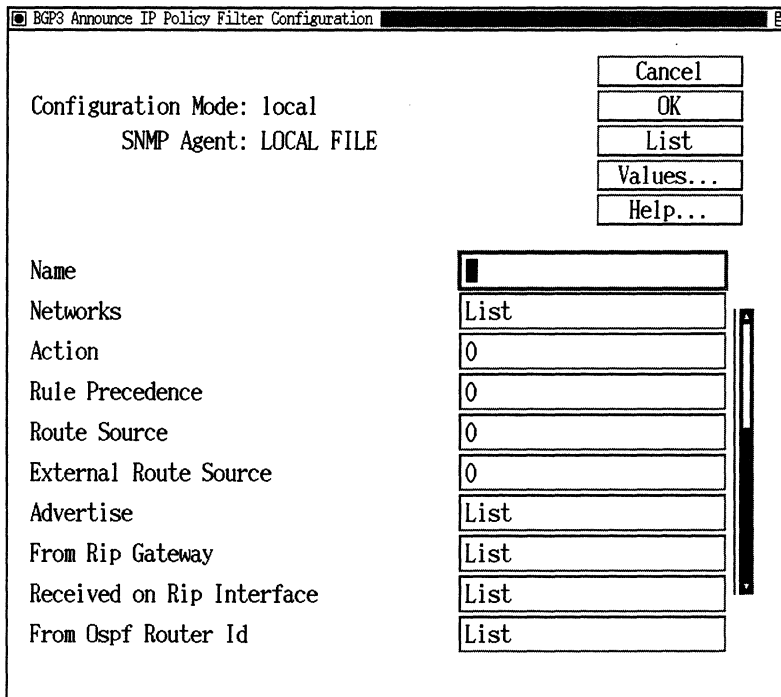


Figure 9-6. BGP-4 Announce IP Policy Filter Configuration Window

IP Announce Policy Parameters

IP announce policy parameters fall into two categories: parameters that appear in all policies and IP protocol-specific parameters. Announce policies are described in the following section:

- ❑ “Common IP Announce Policy Parameters” on page 9-27
- ❑ “RIP-Specific Announce Policy Parameters” on page 9-40
- ❑ “OSPF-Specific Announce Policy Parameters” on page 9-42
- ❑ “EGP-Specific Announce Policy Parameters” on page 9-45
- ❑ “BGP-3-Specific Announce Policy Parameters” on page 9-48
- ❑ “BGP-4-Specific Announce Policy Parameters” on page 9-53

Note: Certain announce policy parameters request a list of entries. A list can contain up to 55 entries.

Common IP Announce Policy Parameters

This section describes how to set common IP announce policy parameters.

Parameter: **Enable**
Default: Enable
Options: Enable | Disable
Function: Enables or disables this policy.
Instructions: Set to Disable to disable the policy.
MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.2
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.2
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.2
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.2
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.2

Parameter: **Name**
Default: None
Options: Any alphanumeric character string
Function: Identifies this policy.
Instructions: Enter a unique name for the policy.
MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.4
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.4
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.4
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.4
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.4

Parameter: Networks

Default: An empty list

Options: A list of network identifiers. Each entry consists of a network number, a mask, and a flag to indicate whether the ID refers to a specific network or a range or networks

Function: Specifies which networks will match this policy.

Instructions: Enter a specific encoding of 0.0.0.0/0.0.0.0 to match the default route. Enter a range encoding of 0.0.0.0/0.0.0.0 to match any route. Enter an empty list to match any route.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.5
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.5
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.5
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.5
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.5

Parameter: Action

Default: RIP, OSPF, EGP: Propagate

Default: BGP-3, BGP-4: Ignore

Options: Propagate | Ignore

Function: Specifies whether or not to advertise a route that matches this policy.

Instructions: To advertise the route, specify Propagate. To drop the route, specify Ignore.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.6
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.6
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.6
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.6
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.6

Parameter: Precedence**Default:** 0**Options:** A metric value**Function:** Specifies a metric value to be used to compare this policy with other policies that a route may match. A policy with a higher metric takes precedence over a policy with a lower metric. In case of a tie, the protocol uses an internal index value assigned to the policy by IP software. (In general, the index value is indicated by the position of the policy in the Site Manager display — the last policy in the display has the highest index value.)**Instructions:** Use this parameter to assign precedence to policies that match the same route.**MIB Object ID:** RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.7
 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.7
 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.7
 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.7
 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.7

Parameter: Route Source

Default: Any

Options: Any
Direct
Static
RIP
OSPF (not valid for OSPF)
EGP
BGP

Function: Specifies one or more route source identifiers. If you select a route source ID, a route from that source that meets the other criteria of this policy matches the policy.

Instructions: To specify any source, use the default.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.8
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.8
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.8
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.8
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.8

| | |
|-----------------------|---|
| Parameter: | Advertise |
| Default: | An empty list |
| Options: | A list of network identifiers |
| Function: | Specifies network IDs to include in place of the network IDs listed in the route to be advertised. |
| Instructions: | <p>Specify a non-null value only if the announce Action parameter is Propagate. The values you enter in the advertise list determine the action taken.</p> <p>If you supply a list of network IDs, these IDs are advertised instead of the actual IDs in the route.</p> <p>If you use the default (an empty list), the actual IDs are advertised. Note that by default, BGP-4 aggregates subnets into their natural network IDs.</p> <p>If you supply a list that includes the encoding 255.255.255.255/255.255.255.255, the actual network IDs are advertised along with the other IDs in the advertise list. This allows advertisement of an aggregate or default along with the actual network. If the actual network is a subnet (and the advertising protocol supports subnet advertisements), the subnet is advertised.</p> |
| MIB Object ID: | RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.10 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.10 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.10 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.10 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.10 |

Parameter: From RIP Gateway

Default: An empty list

Options: A list of IP addresses

Function: Specifies the addresses of one or more routers that could send RIP updates to this router. This policy applies to RIP advertisements from routers on this list. Applicable only for RIP source routes and if RIP is included as a route source.

Instructions: Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to RIP updates from any router.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.11
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.11
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.11
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.11
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.11

Parameter: Received on RIP Interface

Default: An empty list

Options: A list of IP addresses

Function: Specifies the addresses of one or more interfaces on this router. This policy applies to RIP advertisements received on the interfaces in this list. Applicable only for RIP sourced routes and if RIP is included as route source.

Instructions: Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to RIP updates received on any interface.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.12
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.12
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.12
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.12
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.12

| | |
|-------------------|---|
| Parameter: | From OSPF Router ID |
| Default: | An empty list |
| Options: | A list of IP addresses |
| Function: | Specifies the IDs of one or more OSPF routers. This policy applies to OSPF advertisements authored by a router on this list. Applicable only for OSPF sourced routes and if OSPF is included as a route source. |
| Instructions: | Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to OSPF updates from any router. |
| MIB Object ID: | RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.13 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.13 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.13 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.13 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.13 |

Parameter: **Received OSPF Type**

Default: Any

Options: Type 1
Type 2
External
Internal
Any

Function: Specifies which types of OSPF routes match this policy. Applicable only for OSPF sourced routes and if OSPF is included as a route source.

Instructions: To match any route type, enter Type Any. To match any non-ASE route, enter Type Internal. To match any ASE route, enter Type External. To match any external type 1 route, enter Type 1. To match any external type 2 route, enter Type 2.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.14
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.14
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.14
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.14
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.14

| | |
|-------------------|--|
| Parameter: | Received OSPF Tag |
| Default: | An empty list |
| Options: | A list of tag values |
| Function: | Specifies tag values that could be present in an OSPF ASE advertisement. This policy applies to OSPF ASE advertisements that contain tag values in this list. Applicable only for OSPF sourced ASE routes and if OSPF is included as a route source. |
| Instructions: | Specify one or more tag values. Use the default empty list to indicate that this policy applies to OSPF ASs with any tag value. |
| MIB Object ID: | RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.15 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.15 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.15 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.15 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.15 |

Parameter: **From EGP Peer**

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP address of one or more EGP peers. This policy applies to EGP advertisements authored by a router on this list. Applicable only for EGP source routes and if EGP is included as a route source.

Instructions: Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to EGP advertisements from any router.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.16
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.16
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.16
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.16
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.16

Parameter: **From EGP AS**

Default: An empty list

Options: A list of autonomous system numbers

Function: Specifies one or more autonomous system numbers. This policy applies to EGP advertisements received from EGP peers in an AS on this list. Applicable only for EGP sourced routes and if EGP is included as a route source.

Instructions: Specify one or more AS numbers. Use the default empty list to indicate that this policy applies to EGP advertisements from peers in any AS.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.17
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.17
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.17
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.17
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.17

Parameter: Received EGP Gateway

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP address of one or more EGP gateways. This policy applies to EGP advertisements that use a gateway on this list as the next hop. Applicable only for EGP sourced routes and if EGP is included as a route source.

Instructions: Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to EGP advertisements with any gateway address.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.18
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.18
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.18
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.18
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.18

Parameter: From BGP Peer

Default: An empty list

Options: A list of IP addresses

Function: Specifies the IP address of one or more BGP peers. This policy applies to BGP advertisements authored by a router on this list. Applicable only for BGP sourced routes and if BGP is included as a route source.

Instructions: Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to BGP advertisements from any router.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.19
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.19
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.19
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.19
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.19

Parameter: **From BGP AS**

Default: An empty list

Options: A list of autonomous system numbers

Function: Specifies one or more autonomous system numbers. This policy applies to BGP advertisements received from BGP peers in an AS on this list. Applicable only for BGP sourced routes and if BGP is included as a route source.

Instructions: Specify one or more AS numbers. Use the default empty list to indicate that this policy applies to BGP advertisements from peers in any AS

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.20
OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.20
EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.20
BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.20
BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.20

| | |
|-----------------------|---|
| Parameter: | Received BGP Next Hop |
| Default: | An empty list |
| Options: | A list of IP addresses |
| Function: | Specifies one or more IP addresses. This policy applies to BGP advertisements whose Next Hop attribute matches an IP address on this list. Applicable only for BGP sourced routes and if BGP is included as a route source. |
| Instructions: | Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to BGP advertisements with any Next Hop attribute. |
| MIB Object ID: | RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.21 OSPF: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.21 EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.21 BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.21 BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.22 |

RIP-Specific Announce Policy Parameters

This section shows you how to set RIP-specific announce policy parameters.

Parameter: External Route Source

Default: Any

Options: Direct
Static
RIP
OSPF (with type 2 metric)
EGP
BGP
Any

Function: Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy.

Instructions: This parameter applies only to OSPF routes that use the new ASE Type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric. To specify any external route source, use the default.

MIB Object ID: RIP: 1.3.6.1.4.1.18.3.5.3.2.6.2.1.9

Parameter: Outbound Interfaces

- Default:** An empty list
- Options:** A list of IP addresses
- Function:** Specifies a list of outbound RIP interfaces. If an interface appears in this list, the policy applies to RIP advertisements sent via that interface.
- Instructions:** Specify one or more IP addresses. Configure an empty list to indicate that this policy applies to any outbound RIP interface.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.6.2.1.22

Parameter: RIP Metric

- Default:** 0
- Options:** 0 or an export metric
- Function:** Specifies an optional export RIP metric to use when advertising a route that matches this policy.
- Instructions:** Set the Action parameter for Announce. If you use the default, the RIP metric is the routing table metric calculated for RIP plus the interface cost.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.6.2.1.23

OSPF-Specific Announce Policy Parameters

This section shows you how to set OSPF-specific announce policy parameters.

| | |
|-------------------|---|
| Parameter: | Type |
| Default: | 0 |
| Options: | Type 1 Type 2 0 |
| Function: | Specifies an OSPF ASE metric type to use in advertisements for routes that match this policy. |
| Instructions: | Enter a zero value if you want to use the default metric that IP includes in the advertisement, based on the route source. For a BGP, EGP, or RIP route, the default is Type 2. For routes from all other sources, the default is Type 1. |
| | Set the Action parameter for Propagate. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.6.4.1.22 |

Parameter: Tag

Default: Null

Options: Null or a tag value.

Function: Specifies a value for the OSPF external route tag field. If the outgoing route matches this policy, the router places this value in the field.

Instructions: Set the Action parameter to Propagate and set the Automatic Tag parameter to Disable.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.23

Parameter: Automatic Tag

Default: Disable

Options: Enable | Disable

Function: Enables BGP/OSPF automatic tag generation.

Instructions: Select Disable (the default) to use the value you specify with the Tag parameter. Select Enable to generate a tag according to the criteria in RFC 1403 (or any superseding RFC).

This parameter overrides the Tag Generation Method parameter on the OSPF Global Parameters window.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.4.1.24

| | |
|-------------------|---|
| Parameter: | OSPF Metric |
| Default: | 0 |
| Options: | 0 or an export metric |
| Function: | Specifies an optional OSPF metric to use when advertising a route that matches this policy. |
| Instructions: | Set the Action parameter for Announce. If you use the default, the OSPF metric is the routing table metric. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.6.4.1.25 |

EGP-Specific Announce Policy Parameters

This section shows you how to set EGP-specific announce policy parameters.

Parameter: External Route Source

Default: Any

Options: Direct
Static
RIP
OSPF (with type 2 metric)
EGP
BGP
Any

Function: Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy.

Instructions: This parameter applies only to OSPF routes that use the new ASE Type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric. To specify any external route source, use the default.

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.9

Parameter: **EGP Peer List**
Default: An empty list
Options: A list of IP addresses
Function: Specifies a list of IP addresses of EGP peers. If a peer appears in this list, the policy applies to EGP advertisements sent to that peer.
Instructions: Specify one or more IP addresses. Use the default empty list to indicate that the policy applies to any BGP peer.
MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.22

Parameter: **EGP Interface List**
Default: An empty list
Options: A list of IP addresses
Function: Specifies a list of outgoing interfaces. If an interface appears on this list, the policy applies to EGP advertisements sent via that interface.
Instructions: Specify one or more IP addresses. Use the default empty list to indicate that this policy applies to any outbound interface.
MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.23

Parameter: **EGP Metric**

Default: 0

Options: 0 or an export metric value

Function: Specifies an optional export metric to use when advertising a route that matches this policy.

Instructions: Select the default to indicate the routing table metric calculated for EGP is to be used. This parameter is valid only if the Action parameter is set to Propagate.

MIB Object ID: EGP: 1.3.6.1.4.1.18.3.5.3.2.6.6.1.24

BGP-3-Specific Announce Policy Parameters

This section shows you how to set BGP-3-specific announce policy parameters.

Parameter: **External Route Source**

Default: Any

Options: Direct
Static
RIP
OSPF (with type 2 metric)
EGP
BGP
Any

Function: Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy.

Instructions: This parameter applies only to OSPF external routes that use the new ASE Type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric.
To specify any external route source, use the default.

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.9

Parameter: Outbound Peer AS List

- Default:** An empty list
- Options:** A list of AS numbers
- Function:** Specifies a list of autonomous system numbers. If an AS number is included in this list, this policy applies to BGP advertisements being sent to BGP peers in that AS.
- Instructions:** Specify one or more AS numbers. Use the default empty list to indicate that this policy applies to BGP advertisements going to peers in any AS.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.6.8.1.22

Parameter: Outbound Peers

- Default:** An empty list
- Options:** A list of IP numbers
- Function:** Specifies the IP address of one or more BGP peers. If a BGP peer is included in this list, this policy applies to BGP advertisements being sent to that peer.
- Instructions:** Specify one or more IP addresses. Configure an empty list to indicate that this policy applies to BGP advertisements being sent to any peer.
- MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.6.8.1.23

Parameter: Inter-AS Metric Selector

Default: None

Options: None
Specified
Originating

Function: Indicates whether or not an Inter-AS metric is to be advertised for a network matching this policy and, if advertised, what value to use.

Instructions: Select None to indicate that no metric is to be advertised. Select Specified to indicate that the value you specify in the Inter AS Metric parameter is to be used. Select Originating to indicate that the metric from the originating protocol will be used. This parameter is valid only if the Action parameter is set to Propagate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.24

Parameter: Specific Inter-AS Metric

Default: Null

Options: Null or an AS metric

Function: Specifies a value for the Inter-AS metric.

Instructions: Supply a value and set the Use Metric parameter to Specified.

MIB Object ID: BGP-3: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.25

Parameter: **Origin**

Default: As Is

Options: As Is
IGP
EGP
Incomplete

Function: Specifies an Origin attribute override. The Origin attribute of a route matching this policy will be replaced with the indicated value.

Instructions: To allow the existing Origin attribute, use the default.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.26

Parameter: **AS Path Override**

Default: An empty list

Options: A list of AS numbers

Function: Specifies an AS path override.

Instructions: Enter a non-null value to override the AS path attribute of a route matching this policy. Each element of the AS path is an AS number. Valid only if the Action parameter is set to Propagate. Use the default empty list to allow the existing AS path attribute to remain in the route.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.27

Parameter: **Next Hop**
Default: Null
Options: An IP address
Function: Overrides the Next Hop path attribute with the IP address you specify.
Instructions: To allow the existing Next Hop attribute, use the default null value.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.8.1.28

BGP-4-Specific Announce Policy Parameters

This section shows you how to set BGP-4-specific announce policy parameters.

| | |
|-------------------|--|
| Parameter: | External Route Source |
| Default: | Any |
| Options: | Direct Static RIP OSPF (with type 2 metric) EGP BGP Any |
| Function: | Specifies one or more external route source identifiers. If you specify an external route source, a route from that source that meets the other criteria of this policy matches the policy. |
| Instructions: | This parameter applies only to OSPF routes that use the new ASE Type 2 metric. The protocol from which OSPF received the route is encoded in the ASE metric, along with the route's metric. To specify any external route source, use the default. |
| MIB Object ID: | BGP-4: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.9 |

| | |
|-------------------|--|
| Parameter: | Outbound Peer AS |
| Default: | An empty list |
| Options: | A list of AS numbers |
| Function: | Specifies a list of autonomous system numbers. If an AS number is included in this list, this policy applies to BGP advertisements being sent to BGP peers in that AS. |
| Instructions: | Specify one or more AS numbers. Configure an empty list to indicate that this policy applies to BGP advertisements going to peers in any AS. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.6.10.1.22 |

Parameter: Outbound Peers**Default:** An empty list**Options:** A list of IP addresses**Function:** Specifies the IP address of one or more BGP peers. If a BGP peer is included in this list, this policy applies to BGP advertisements being sent to that peer.**Instructions:** Specify one or more IP addresses. Configure an empty list to indicate that this policy applies to BGP advertisements being sent to any peer.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.6.10.1.23**Parameter: Multi-Exit Discriminator****Default:** None**Options:** None
Specified
Originating**Function:** Indicates whether or not a Multi Exit Descriptor metric is to be advertised for a network matching this policy and, if advertised, what value to use.**Instructions:** Select None to indicate that no value is to be advertised. Select Specified to indicate that the value you specify for the Multi-Exit Disc parameter is to be used. Select Originating to indicate that the metric from the originating protocol is to be used. This parameter is valid only if the Action parameter is set for Propagate.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.6.10.1.24

Parameter: Multi-Exit Discriminator Value

Default: Null

Options: Null or a metric value

Function: Specifies a metric for the Multi-Exit Discriminator attribute.

Instructions: To advertise a Multi Exit Discriminator value, set the Action parameter to Propagate and set the Use Multi-Exit Discriminator parameter to Specified.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.25

Parameter: Origin

Default: As Is

Options: As Is
IGP
EGP
Incomplete

Function: Specifies an Origin attribute override. The Origin attribute of a route matching this policy will be replaced with the indicated value.

Instructions: To allow the existing Origin attribute, use the default.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.26

Parameter: AS Path**Default:** Null**Options:** An AS path**Function:** Specifies an AS path that overrides the AS-path attribute of a route matching this policy.**Instructions:** Constructs a BGP-4 AS path composed of AS path segments. Each AS path segment includes a path segment type, a path segment length specifying the number of ASs in the segment, and a path segment value containing one or more AS numbers.

There are two AS path segment types:

Type 1. An unordered set of ASs a route in the UPDATE message has traversed.

Type 2. An ordered set of ASs a route in the UPDATE message has traversed.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.27**Parameter: Local Preference Override****Default:** Local Pref Override False**Options:** False | True**Function:** Indicates whether or not you are supplying an override value for the Local Preference path attribute in the routing Update message. (The Local Pref attribute is valid only in an Update advertised to an IBGP peer.) If you select False, the router uses the IP route weight value to calculate the LOCAL_PREF path attribute.**Instructions:** To override the Local Pref attribute, select True and supply a value for the Local Pref parameter.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.6.10.1.28

Parameter: Local Preference Value

Default: Null

Options: Null or a route weight value

Function: Specifies an override value for the Local Pref attribute.

Instructions: Enter a value and set the Local Pref Override parameter to True.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.29

Parameter: Next Hop

Default: Null

Options: An IP address

Function: Overrides the Next Hop path attribute with the IP address you specify.

Instructions: To allow the existing Next Hop attribute, use the default null value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.30

Parameter: **Atomic**

Default: Automatic

Options: Automatic | Force | Ignore

Function: Allows control over the Atomic path attribute.

Instructions: By default, the router automatically sets this parameter if it knows that certain networks in aggregate range have not been included in an aggregate advertisement.

To include the Atomic attribute even if the router does not think one is required, set the parameter to Force.

To prohibit the Atomic attribute even the router thinks one is required, set the parameter to Ignore.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.6.10.1.31

Chapter 10

Import and Export Route Filters

This chapter describes the procedures you follow to configure import and export route filters for RIP, OSPF, BGP-3, and EGP.

Note: Import and export filters provide a subset of the parameters provided by accept and announce policies. We currently support both IP policies and IP route filters. However, network administrators using import and export filters for routing table management should migrate as quickly as possible to IP policies. In a future release, support for the import and export filters described in this chapter will be dropped.

- “RIP Route Filters” on page 10-2
- “OSPF Route Filters” on page 10-18
- “BGP-3 Route Filters” on page 10-37
- “EGP Route Filters” on page 10-59

RIP Route Filters

The following sections show you how to select RIP route filter windows from the Site Manager and describes all RIP route filter parameters.

Configuring RIP Import Route Filters

To add, edit, or delete RIP Import Route Filters, begin at the Wellfleet Configuration Manager Window and proceed as follows:

1. Select Protocols→IP→Route Filters→RIP→Import Filters.

The RIP Import Route Filters List window appears (Figure 10-1). It lists all RIP import route filters configured on the router. You add, edit and delete import route filters from this window.

2. Add, edit or delete import route filters as described in the following sections.

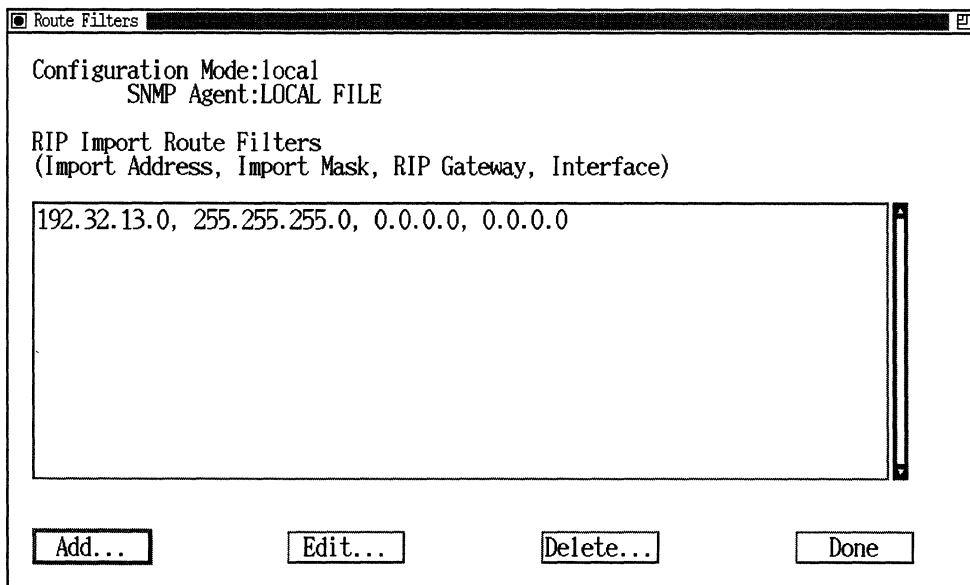


Figure 10-1. RIP Import Route Filters List Window

Adding a RIP Import Route Filter

To add an import route filter, begin at the RIP Import Route Filters window and proceed as follows:

1. Click on Add.

The RIP Import Route Filter Configuration window appears (see Figure 10-2).

2. Specify the RIP import route filter configuration parameters.

All RIP import route filter parameters are described following these instructions.

3. Click on OK.

The RIP Import Route Filter window appears (see Figure 10-3). It displays the default settings for the Enable, Action and Preference parameters (specifically, it enables the filter, sets the Action parameter to Accept, and sets the Preference parameter to 1).

4. Either accept the default settings or specify new settings for the Action and Preference parameters, then click on OK.

The RIP Import Route Filters window now lists the import route filter you added.

5. Click on Done to save your changes and exit the window.

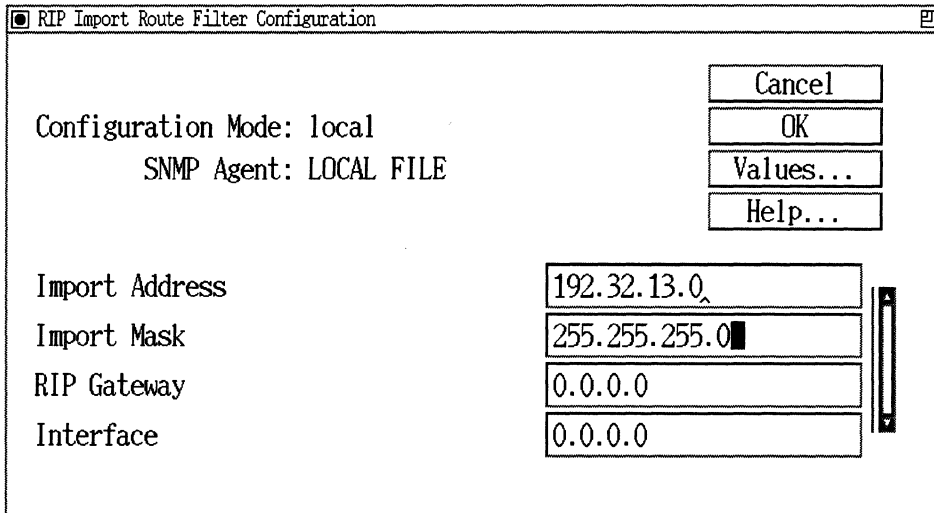


Figure 10-2. RIP Import Route Filter Configuration Window

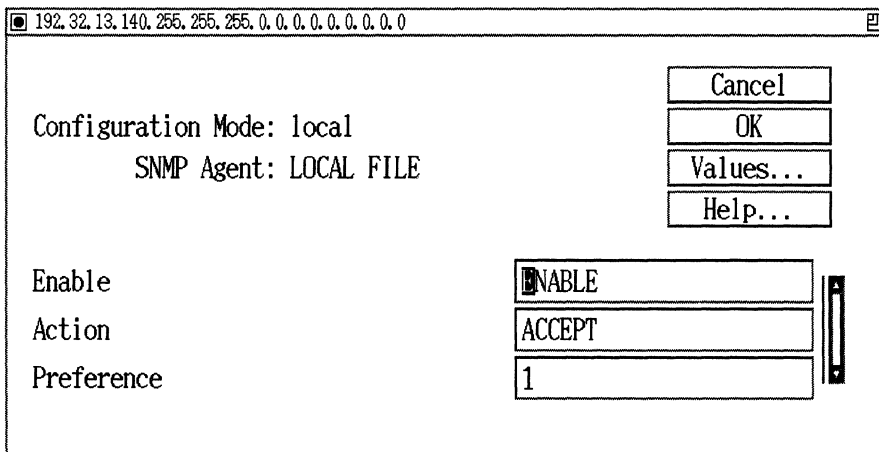


Figure 10-3. RIP Import Route Filter Window

RIP Import Route Filter Parameter Descriptions

This section describes how to set all RIP import route filter parameters.

| | |
|-------------------|---|
| Parameter: | Import Address |
| Default: | 0.0.0.0 |
| Options: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. If this field is set to 0.0.0.0, the filter applies to all networks. |
| Instructions: | Enter the appropriate network address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.8.1.3 |

Parameter: Import Mask

Default: 0.0.0.0

Options: Depends on the address class of the network address.

Function: Specifies the range of addresses this filter acts upon.

For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Import Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Import Mask, the Net ID and Subnet ID portions of the address will be filtered.

If the Import Address field is set to 0.0.0.0, and the Import Mask is set to 0.0.0.0, then the filter applies to *all* routes. If the Import Address field is set to 0.0.0.0, and the Import Mask is set to 255.255.255.255, then the filter applies to the *default* route.

Instructions: Enter the mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.4

Parameter: RIP Gateway

Default: 0.0.0.0

Options: Any IP address

Function: Identifies, by IP address, the router that is sending the updates. This filter will apply to updates from that router.

If this field is set to 0.0.0.0, the filter applies to updates from any router.

Instructions: Enter the appropriate IP address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.7

Parameter: Interface

Default: 0.0.0.0

Options: Any IP address

Function: Specifies the local IP address of the interface that connects this router to the RIP Gateway. This filter will apply only to those updates received on this interface.

If set to 0.0.0.0, this filter applies to all interfaces.

Instructions: Enter the appropriate IP address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.8

Parameter: **Enable**
Default: Enable
Options: Enable | Disable
Function: Enables or disables this import route filter.
Instructions: Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now wish to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.2

Parameter: **Action**
Default: Accept
Options: Accept | Ignore
Function: Specifies whether the route is transferred to the routing tables. If Action is set to Accept (default), the routing information is sent to the routing tables. If Action is set to Ignore, the routing information is dropped.
Instructions: Either accept the default Accept, or select Ignore.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.5

Parameter: Preference

Default: 1

Range: 1 to 16

Function: Assigns a weighted precedence value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, BGP-3, OSPF external, EGP, and RIP.

Instructions: If this hierarchy is acceptable, accept the default value 1 for preference. If you want to grant preference to this RIP-derived route, assign a new preference value in the range of 1 to 16 (the greater the number, the higher the preference).

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.8.1.6

Note: The default preference for static routes is 1, but may be set to any value between 1 and 16 (refer to “Editing Static Route Parameters” for more information). If you want to grant a RIP-derived route preference over a static route, make sure the preference value you assign to the RIP-derived route is greater than the preference value of the static route you want it to override.

Editing a RIP Import Route Filter

You can edit the Enable, Action and Preference parameters for a RIP import route filter.

Note: You cannot reconfigure the Import Address, Import Mask, RIP Gateway or Interface parameters for a RIP import route filter. To change these parameters, you must delete the filter and add a new filter with the proper information. See “Deleting a RIP Import Route Filter” on page 10-10 for instructions.

To edit these parameters, begin at the RIP Import Route Filters window shown in Figure 10-3 and proceed as follows:

1. Click on the import route filter you want to edit.
2. Click on Edit.
3. Edit those parameters that you want to change

All RIP import route filter parameters are described in the section “RIP Import Route Filter Parameter Descriptions.”

4. Click on OK.
5. Click on Done to exit the window and to save your changes.

Deleting a RIP Import Route Filter

To delete a RIP import route filter, begin at the RIP Import Route Filters window shown in Figure 10-3 and proceed as follows:

1. Click on the import route filter you want to delete.
2. Click on Delete.
3. Click on Delete to delete the import route filter.
4. Click on Cancel to exit the window.

Configuring RIP Export Route Filters

To add, edit, or delete RIP Export Route Filters, begin at the Wellfleet Configuration Manager Window and proceed as follows:

1. Select the Protocols→IP→Route Filters→RIP→Export Filters option.

The RIP Export Route Filters window appears (see Figure 10-4). It lists all RIP export route filters configured on the router. You add, edit and delete export route filters from this window.

2. Add, edit or delete export route filters as described in the following sections.

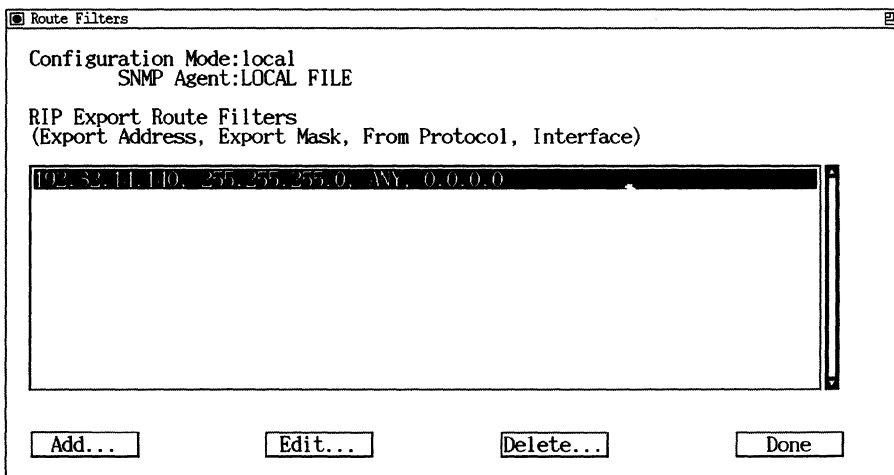


Figure 10-4. RIP Export Route Filters List Window

Adding a RIP Export Route Filter

To add an export route filter, begin at the RIP Export Route Filters window shown in Figure 10-4 and proceed as follows:

1. Click on Add.

The RIP Export Route Filter Configuration Window appears (see Figure 10-5). All parameters on this window display the default settings.

2. Specify the Export Address, Export Mask, From Protocol and Interface parameters.

All RIP Export parameters are described following these instructions.

3. Click on OK .

After you click on the Okay button, the RIP Export Route Filters window appears (see Figure 10-6). It displays the default settings for the Enable, Action, and RIP Metric parameters (specifically, it enables the filter, sets the Action parameter to Propagate, and sets the RIP Metric parameter to 0).

4. Either accept the default settings or specify new settings for the Action and Metric parameters, then click on the Okay button.
5. Click on Done to exit the window.

RIP Export Route Filter Parameter Descriptions

This section describes how to set all RIP export route filter parameters.

| | |
|-------------------|---|
| Parameter: | Export Address |
| Default: | 0.0.0.0 |
| Range: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.9.1.3 |

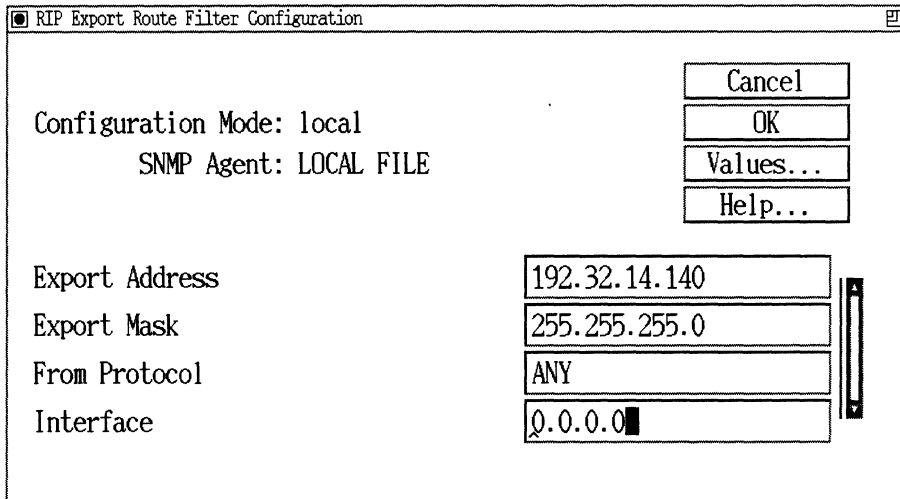


Figure 10-5. RIP Export Route Filter Configuration Window

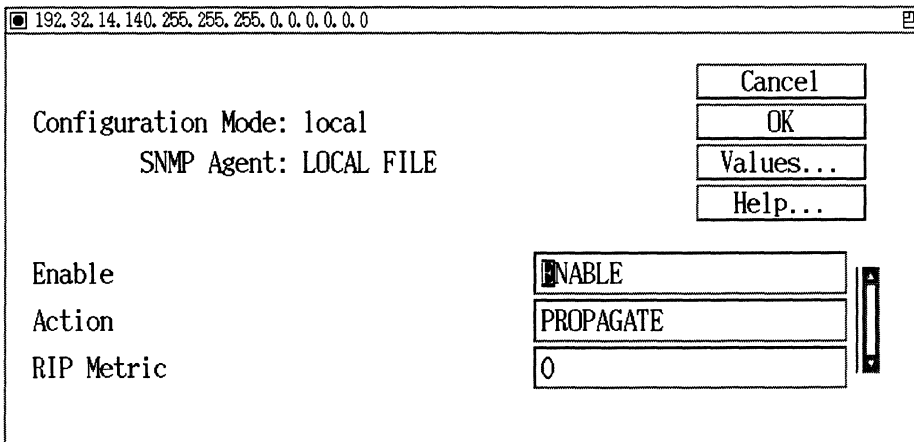


Figure 10-6. RIP Export Route Filters Windows

Parameter: Export Mask

Default: 0.0.0.0

Range: Depends on the address class of the network address.

Function: Specifies the range of addresses upon which this filter acts.

For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Export Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Export Mask, the Net ID and Subnet ID portions of the address will be filtered.

If the Export Address field is set to 0.0.0.0, and the Export Mask is set to 0.0.0.0, then the filter applies to *all* routes. If the Export Address field is set to 0.0.0.0, and the Export Mask is set to 255.255.255.255, then the filter applies to the *default* route.

Instructions: Enter the appropriate mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.4

Parameter: From Protocol

Default: Any

Options: Any | RIP | EGP | OSPF | Direct | Static | BGP-3

Function: Identifies the source of the routing information: direct connection, static route, or RIP, OSPF, EGP or BGP-3-derived route.

Instructions: Select the appropriate option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.5

Parameter: Interface

Default: 0.0.0.0

Range: Any IP address

Function: Identifies the outgoing IP interface for the RIP update. This filter will only apply to this interface. If set to 0.0.0.0, this filter applies to all interfaces.

Instructions: Enter the appropriate IP address in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.7

Parameter: Enable

Default: Enable

Options: Enable | Disable

Function: Enables or disables this export route filter.

Instructions: Set to Disable if you want to disable this export route filter.

Set to Enable if you previously disabled this export route filter and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.2

Parameter: Action

Default: Propagate

Options: Propagate | Ignore | Aggregate

Function: Controls the flow of routing information. If Action is set to Propagate, this route is advertised. If Action is set to Ignore, advertising of this route is suppressed. If Action is set to Aggregate, the network is not explicitly advertised. Instead, the default route (0.0.0.0) is advertised.

Instructions: Either accept the default Propagate, or select Ignore or Aggregate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.6

Parameter: RIP Metric

Default: 0 (0 = the actual route cost as learned)

Range: 0 to 15

Function: Assigns a RIP cost to the propagated route. The value 0 causes the actual route cost (as learned) to be used.

Instructions: Accept the default value 0, or enter a new value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.9.1.8

Note: Do not use a value that exceeds the diameter of the RIP network.

Editing a RIP Export Route Filter

You can edit the Enable, Action and RIP Metric parameters for an export route filter.

Note: You cannot reconfigure the Export Address, Export Mask, Protocol and Interface parameters for a RIP export route filter. To change these parameters, you must delete the filter and add a new filter with the proper information. See “Deleting a RIP Export Route Filter” on page 10-17 for instructions.

To edit these parameters, begin at the RIP Export Route Filters window shown in Figure 10-4 and proceed as follows:

1. Click on the export route filter you wish to edit.
2. Click on Edit.
3. Edit those parameters you want to change.

All RIP export route filter parameters are described in the section “RIP Export Route Filter Parameter Descriptions.”

4. Click on OK.
5. Click on Done to exit the window.

Deleting a RIP Export Route Filter

To delete an export route filter, begin at the RIP Export Route Filters window shown in Figure 10-4 and proceed as follows:

1. Click on the export route filter you wish to delete.
2. Click on Delete.
3. Click on Delete to delete the export route filter.
4. Click on Done to exit the window.

OSPF Route Filters

The following sections show you how to select OSPF route filter windows from the Site Manager and describes all OSPF route filter parameters.

Configuring OSPF Import Route Filters

To add, edit, or delete OSPF Import Route Filters, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols→IP→Route Filters→OSPF→Import Filters option.

The OSPF Import Route Filters List window appears (see Figure 10-7). It lists all OSPF import route filters configured on the router.

2. Add, edit or delete import route filters as described in the following sections.

Note: OSPF route filters pertain only to AS Boundary routers; OSPF import router filters pertain only to external OSPF routes.

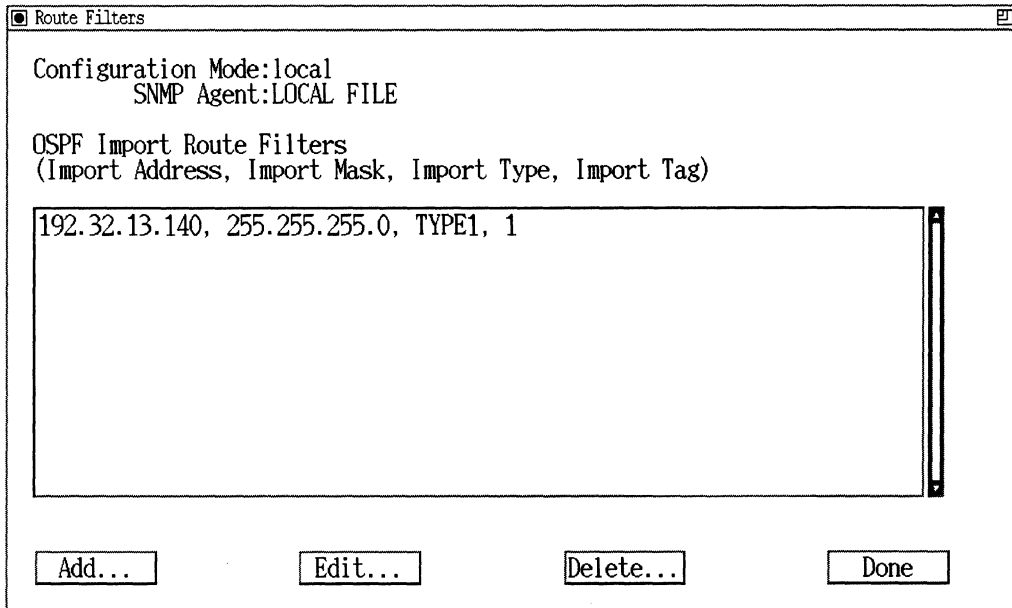


Figure 10-7. OSPF Import Route Filters List Window

Adding an OSPF Import Route Filter

To add an import route filter, begin at the OSPF Import Route Filters List window (see Figure 10-7) and proceed as follows:

1. Click on Add.

The OSPF Import Route Filter Configuration window appears (see Figure 10-8).

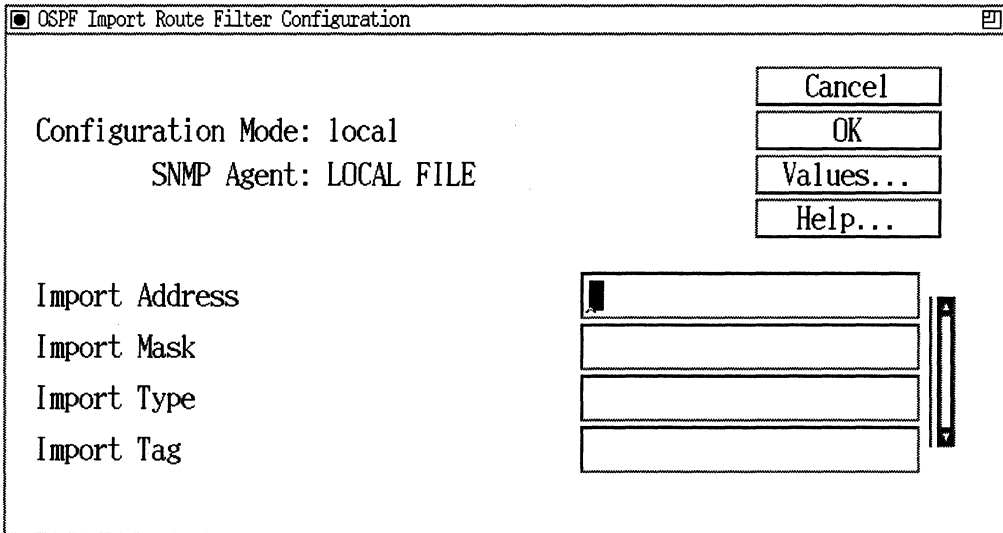


Figure 10-8. OSPF Import Route Filter Configuration Window

2. Specify the Import Address, Import Mask, Import Type and Import Tag parameters.

All OSPF import route filter parameters are described following these instructions.

3. Click on Okay.

The Site Manager then displays the default settings for the Enable, Action and Preference parameters (see Figure 10-9). Specifically, it enables the filter, sets the Action parameter to Accept, and sets the Preference parameter to 1.

4. Either accept the default settings or specify new settings for the Action and Preference parameters, then click on the Okay button.
5. Click on Done to exit the window.

OSPF Import Route Filter Parameter Descriptions

This section describes how to set all OSPF import route filter parameters.

| | |
|-------------------|---|
| Parameter: | Import Address |
| Default: | None |
| Range: | An IP address |
| Function: | Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks. |
| Instructions: | Enter the appropriate network address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.10.1.3 |

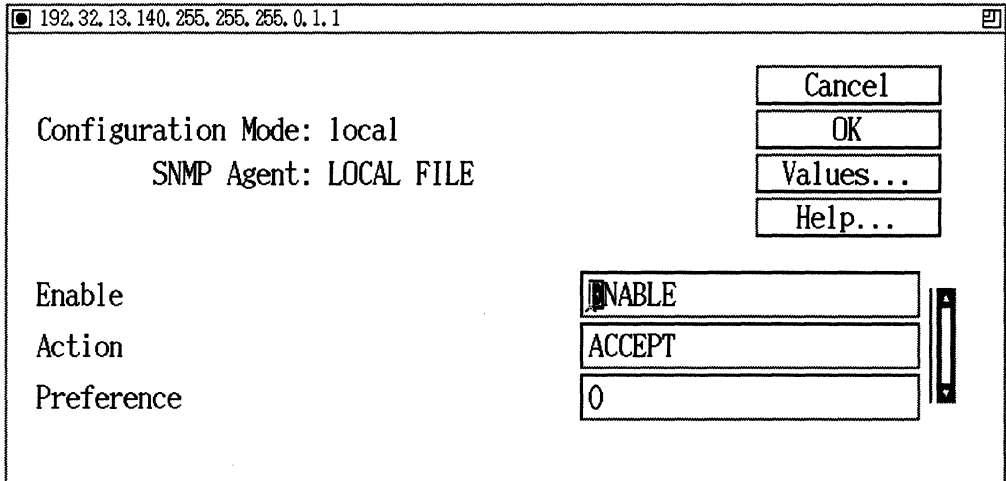


Figure 10-9. OSPF Import Route Filters Window

Parameter: Import Mask**Default:** 0.0.0.0**Range:** Depends on the address class of the network address.**Function:** Specifies the range of addresses upon which this filter acts.

For example, consider Class B Network 172.32.0.0. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Import Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Import Mask, the Net ID and Subnet ID portions of the address will be filtered.

If the Import Address field is set to 0.0.0.0, and the Import Mask is set to 0.0.0.0, then the filter applies to *all* routes. If the Import Address field is set to 0.0.0.0, and the Import Mask is set to 255.255.255.255, then the filter applies to the *default* route.

Instructions: Enter the appropriate mask in dotted decimal notation.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.10.1.4

Parameter: **Import Type**
Default: Type 1
Options: Type1 | Type2
Function: Indicates the type of route to which this filter applies. Type 1 indicates that only AS External Type 1 routes are to be filtered. Type 2 indicates that only AS External Type 2 routes are to be filtered.
Instructions: Select Type 1 or Type 2 as appropriate.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.10.1.7

Parameter: **Import Tag**
Default: 1
Range: 1 to 2147483647
Function: Indicates the tag with which this route filter is concerned. Each AS External Advertisement contains a tag field. If the tag field matches Import Tag, the appropriate action is taken, either the route is accepted or ignored. Import Tag is pertinent to AS External Advertisements only.
Instructions: Enter the appropriate tag number.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.10.1.8

Parameter: **Enable**
Default: Enable
Options: Enable | Disable
Function: Enables or disables this import route filter.
Instructions: Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now wish to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.10.1.2

Parameter: **Action**
Default: Accept
Options: Accept | Ignore
Function: Specifies whether the route is transferred to the routing tables. If Action is set to Accept (default), the routing information is sent to the routing tables. If Action is set to Ignore, the routing information is dropped.
Instructions: Either accept the default Accept, or select Ignore.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.10.1.5

Parameter: Preference

Default: 0

Range: 0 to 16

Function: Assigns a weighted precedence value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, BGP-3, OSPF external, EGP, and RIP.

Instructions: If this hierarchy is acceptable, accept the default value 0 for preference. If you want to grant preference to this OSPF-derived route, assign a new preference value in the range of 1 to 16 (the greater the number, the higher the preference).

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.10.1.6

Note: The default preference for static routes is 0, but may be set to any value between 0 and 16 (refer to “Editing Static Route Parameters” for more information). If you want to grant a OSPF-derived route preference over a static route, make sure the preference value you assign to the OSPF-derived route is greater than the preference value of the static route you want it to override.

Editing an OSPF Import Route Filter

You can edit Enable, Action and Preference parameters for an OSPF import route filter.

Note: You cannot reconfigure the Import Address, Import Mask, Import Type and Import Tag parameters for an OSPF import route filter. To change these parameters, you must delete the filter and add a new filter with the proper information. See “Deleting an OSPF Import Route Filter” on page 10-27 for instructions.

To edit these parameters, begin at the OSPF Import Route Filters List window shown in Figure 10-7 and proceed as follows:

1. Click on the import route filter you want to edit.
2. Click on Edit.
3. Edit those parameters that you want to change.

All OSPF import route filter parameters that you can edit are described in the section “OSPF Import Route Filter Parameter Descriptions.”

4. Click on OK.
5. Click on Done to save your changes and exit the window.

Deleting an OSPF Import Route Filter

To delete an OSPF import route filter, begin at the OSPF Import Route Filters window shown in Figure 10-7 and proceed as follows:

1. Click on the OSPF import route filter you want to delete.
2. Click on Delete.
3. Click on Delete to delete the import route filter.
4. Click on Done to save your changes and exit the window.

Configuring OSPF Export Route Filters

To add, edit, or delete OSPF Export Route Filters, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols→IP→Route Filters→OSPF→Export Filters option.

The OSPF Export Route Filters window appears (see Figure 10-10). It lists all OSPF export route filters configured on the router. You add, edit and delete export route filters from this window.

2. Add, edit or delete OSPF export route filters as described in the following sections.

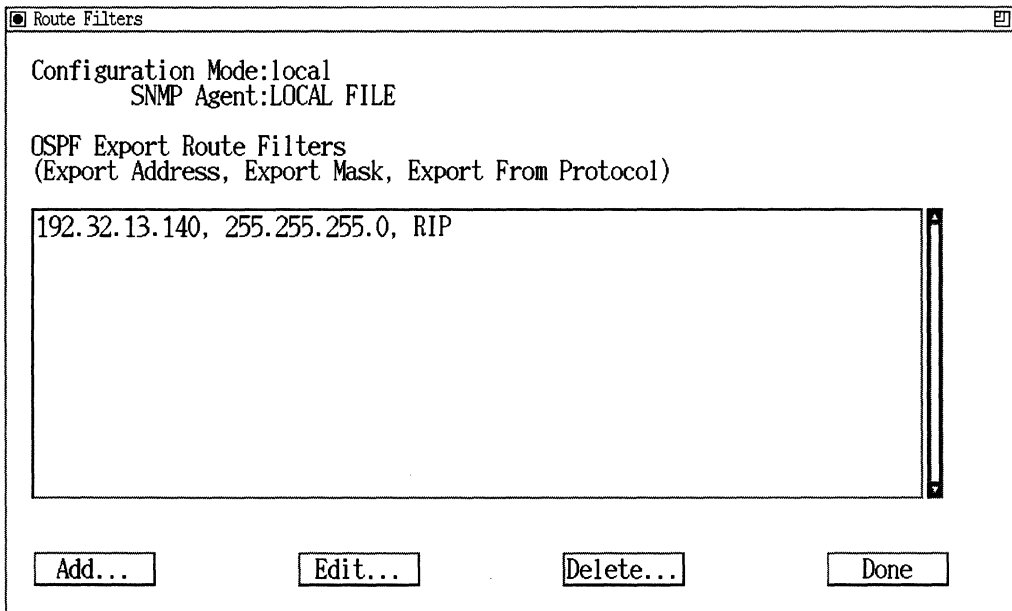


Figure 10-10. OSPF Export Route Filters List Window

Adding an OSPF Export Route Filter

To add an OSPF export route filter, begin at the OSPF Export Route Filters window and proceed as follows:

1. Click on the Add button.

The OSPF Export Route Filter Configuration window appears (see Figure 10-11).

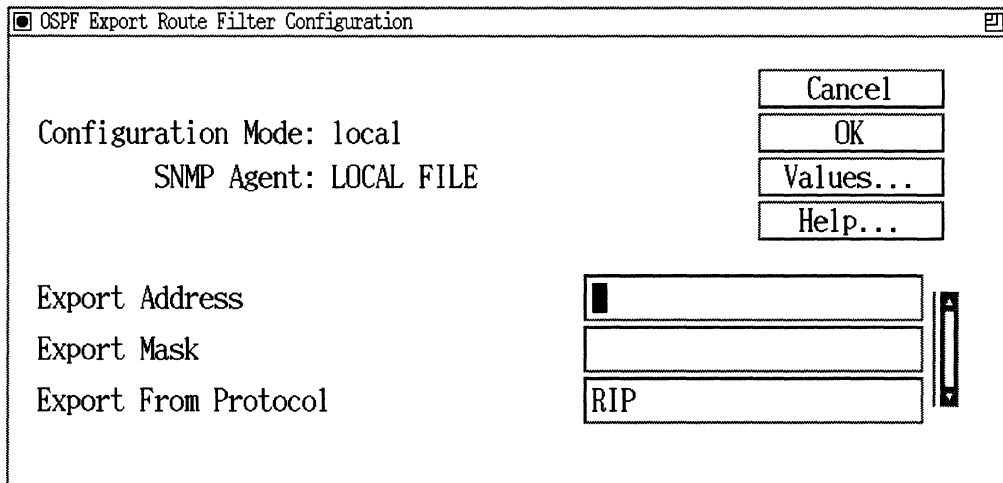


Figure 10-11. OSPF Export Route Filter Configuration Window

2. Specify the Export Address, Export Mask, and Export From Protocol parameters.

All OSPF export route filter parameters are described following these instructions.

3. Click on OK.

The Site Manager then displays the default settings for the Enable, Action, Type, Tag and AutoTag parameters (see Figure 10-12). Specifically, it enables the filter, sets the Action parameter to Propagate, the Type parameter to Type 1, the Tag parameter to 0 and the AutoTag parameter to Disable).

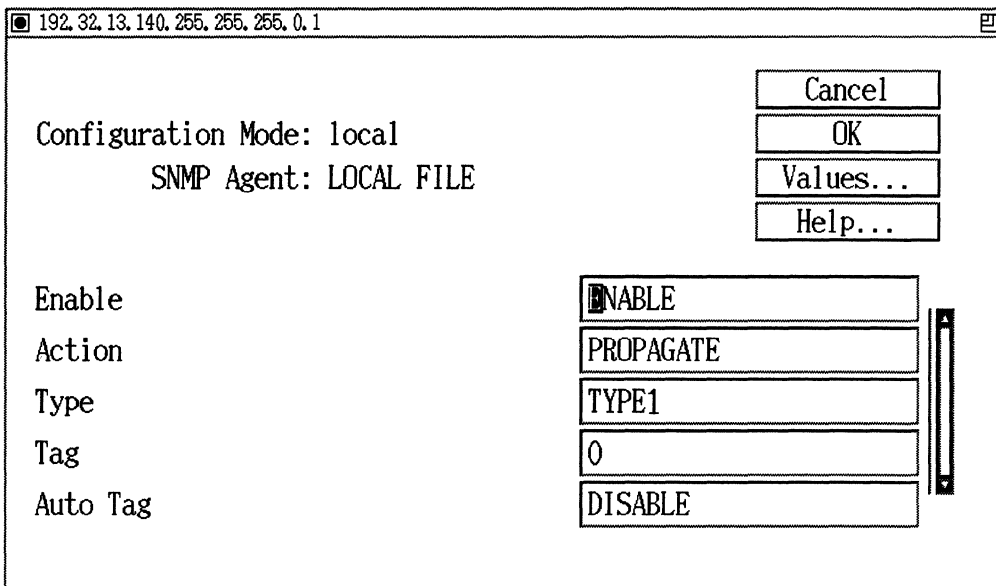


Figure 10-12. OSPF Export Route Filters Window

4. Either accept the default settings or specify new settings, then click on OK.
5. Click on Done to exit the window.

OSPF Export Route Filter Parameter Descriptions

This section describes how to set all OSPF export route filter parameters.

| | |
|-------------------|---|
| Parameter: | Export Address |
| Default: | 0.0.0.0 |
| Range: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.11.1.3 |

Parameter: Export Mask

Default: None

Range: Depends on the address class of the network address.

Function: Specifies the range of addresses upon which this filter acts.

For example, consider Class B Network 172.32.0.0. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Export Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Export Mask, the Net ID and Subnet ID portions of the address will be filtered.

If the Export Address field is set to 0.0.0.0, and the Export Mask is set to 0.0.0.0, then the filter applies to *all* routes. If the Export Address field is set to 0.0.0.0, and the Export Mask is set to 255.255.255.255, then the filter applies to the *default* route.

Instructions: Enter the appropriate mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.4

Parameter: Export From Protocol

Default: RIP

Options: Any, RIP, EGP, OSPF, Direct, static, BGP-3

Function: Identifies the source of the routing information: direct connection, static route, or RIP, EGP, OSPF or BGP-3-derived route.

Instructions: Select the appropriate option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.5

Parameter: Enable

Default: Enable

Options: Enable | Disable

Function: Enables or disables this export route filter.

Instructions: Set to Disable if you want to disable this export route filter.

Set to Enable if you previously disabled this export route filter and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.2

Parameter: Action

Default: Propagate

Options: Propagate | Ignore

Function: Controls the flow of routing information. If Action is set to Propagate, this route is advertised. If Action is set to Ignore, advertising of this route is suppressed.

Instructions: Either accept the default Propagate, or select Ignore.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.6

Parameter: Type

Default: Type1

Options: As Is, Type1, Type2

Function: Specifies an OSPF ASE metric type to use in advertisements for routes that match this policy.

Instructions: Select As Is if you want to use the default metric that IP includes in the advertisement, based on the route source. For a BGP, EGP, or RIP route, the default is Type 2. For routes from all other sources, the default is Type 1.

Set the Action parameter for Propagate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.7

Parameter: Tag

Default: 1

Range: 1 to 2147483647

Function: Sets the tag value for the AS External Advertisement that is generated for this network.

This parameter has meaning only when the Action parameter is set to Propagate.

Instructions: Enter the appropriate tag.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.11.1.8

| | |
|-------------------|--|
| Parameter: | Auto Tag |
| Default: | Disable |
| Options: | Enable Disable |
| Function: | If enabled, the router creates a tag for this route as described in RFC 1364 (BGP/OSPF Interaction). |
| Instructions: | Set to Enable if you are running BGP-3 as your exterior gateway protocol. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.11.1.9 |

Editing an OSPF Export Route Filter

You can edit the Enable, Action, Type, Tag and AutoTag parameters for an OSPF export route filter.

Note: You cannot reconfigure the Export Address, Export Mask, or Export From Protocol parameters for an OSPF export route filter. To change these parameters, you must delete the filter and add a new filter with the proper information. See “Deleting an OSPF Export Route Filter” on page 10-36 for instructions.

To edit these parameters, begin at the OSPF Export Route Filters List window shown in Figure 10-10 and proceed as follows:

1. Click on the export route filter you want to edit.
2. Click on Edit.
3. Edit those parameters you want to change.

All OSPF export route filter parameters are described in the section “OSPF Export Route Filter Parameter Descriptions.”

4. Click on OK.
5. Click on Done to save your changes and exit the window.

Deleting an OSPF Export Route Filter

To delete an OSPF export route filter, begin at the OSPF Export Route Filters List window and proceed as follows:

1. Click on the export route filter you wish to delete.
2. Click on Delete button.
3. Click on Delete to delete the export route filter.
4. Click on Done to save your changes and exit the window.

BGP-3 Route Filters

The following sections show you how to select BGP-3 route filter windows from the Site Manager and describes all BGP-3 route filter parameters.

Configuring BGP-3 Import Route Filters

To add, edit, or delete BGP-3 import route filters, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols→IP→Route Filters→BGP-3→Import Filters option.

The BGP-3 Import Route Filters List window appears. It lists all BGP-3 import route filters configured on the router. You add, edit, and delete BGP-3 import route filters from this window.

2. Add, edit or delete import route filters as described in the following sections.

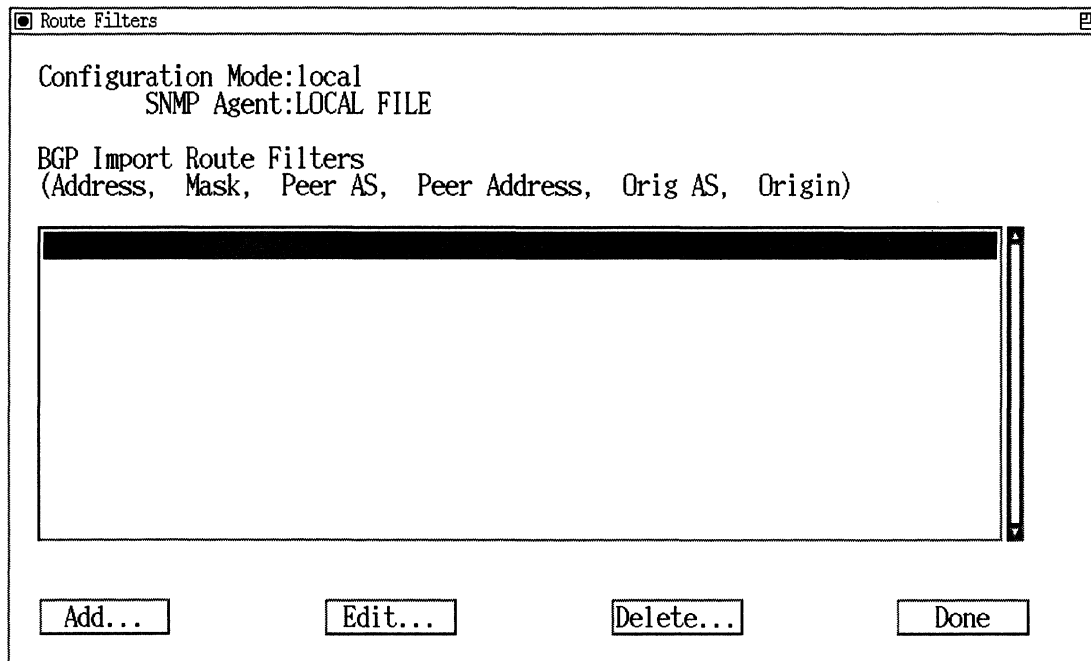


Figure 10-13. BGP-3 Import Route Filters List Window

Adding a BGP-3 Import Route Filter

To add an import route filter, begin at the BGP-3 Import Route Filters window and proceed as follows:

1. Click on Add.

The BGP-3 Import Route Filter Configuration window appears (see Figure 10-14).

BGP Import Route Filter Configuration

Configuration Mode: local
SNMP Agent: LOCAL FILE

Cancel
OK
Values...
Help...

Import Address: 192.32.0.0
Import Mask: 255.255.255.0
Import Peer AS: 0
Import Peer Address: 192.32.14.140
Import Originating AS: 0
Import Route Origin: ANY
Import Action: IGNORE

Figure 10-14. BGP-3 Import Route Filter Configuration Window

2. Specify the BGP-3 import route filter configuration parameters.

All BGP-3 import route filter configuration parameters are described following these instructions.

3. Click on OK.

The BGP-3 Import Route Filter window appears (Figure 10-15). It displays the default settings for the Enable, Action, Preference, and BGP-Preference parameters.

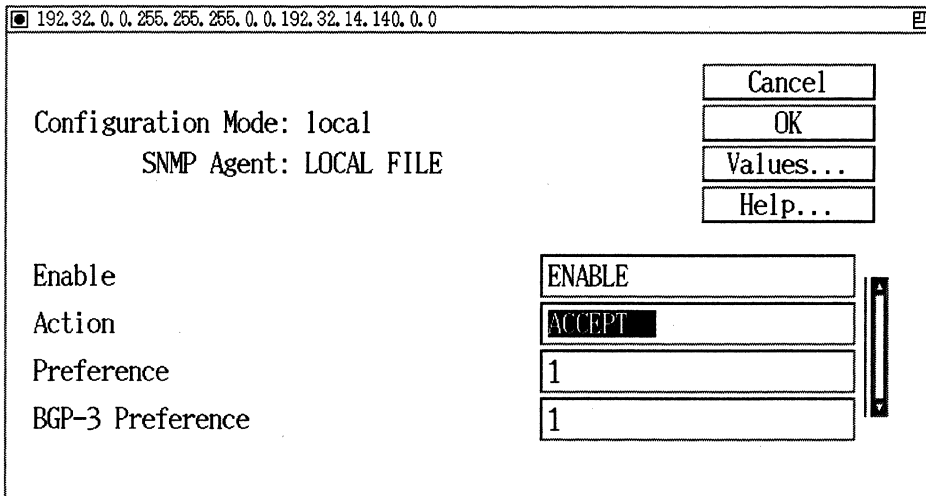


Figure 10-15. BGP-3 Import Route Filter Window

4. Either accept the default settings, or edit these parameters to your network specifications, then click on OK.
5. Click on Cancel to exit the window.

BGP-3 Import Route Filter Parameter Descriptions

This section describes how to set all BGP-3 import route filter configuration parameters.

| | |
|-------------------|--|
| Parameter: | Import Address |
| Default: | 0.0.0.0 |
| Range: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. |
| Instructions: | Enter the appropriate network address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.14.1.3 |

Parameter: Import Mask**Default:** 0.0.0.0**Range:** Depends on the address class of the network address.**Function:** Specifies the range of addresses upon which this filter acts.

For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Import Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Import Mask, the Net ID and Subnet ID portions of the address will be filtered.

If the Import Address field is set to 0.0.0.0, and the Import Mask is set to 0.0.0.0, then the filter applies to *all* routes. If the Import Address field is set to 0.0.0.0, and the Import Mask is set to 255.255.255.255, then the filter applies to the *default* route.

Instructions: Enter the appropriate mask in dotted decimal notation.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.14.1.4

Parameter: Import Peer AS

Default: 0

Range: 0 to 65535

Function: Identifies the Autonomous System to which the BGP router at the remote end of this BGP peer connection belongs. This filter will apply to updates from this router. The value 0 means “any” AS.

Instructions: Enter the appropriate AS number.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.7

Parameter: Import Peer Address

Default: 0.0.0.0

Range: Any IP address

Function: Specifies the IP address of the interface on the remote side of this BGP peer connection. This filter will apply to updates from this router. The value 0 means “any” peer.

Instructions: Enter the IP address in dotted decimal notation. If the peer is in a remote AS, the address must be on the same subnet as the local interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.8

Parameter: Import Peer Original AS

Default: 0

Range: 0 to 65535

Function: Specifies the AS from which the route originated (the last AS in the AS path). The filter will apply to updates created by any routers in this AS. The value 0 means "any" AS.

Instructions: Enter the appropriate AS number.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.9

Parameter: Import Route Origin

Default: Any

Options: Any, IGP, EGP, Incomplete

Function: Specifies the value of the Origin Path Attribute in the Update message received.

Instructions: Set the appropriate Import Route Origin value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.10

Parameter: Action

Default: Ignore

Options: Accept | Ignore

Function: Specifies whether the route is transferred to the routing tables. If Action is set to Accept, the routing information is sent to the routing tables. If Action is set to Ignore, the routing information is dropped.

Instructions: Either accept the default, Ignore, or select Accept.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.5

Parameter: **Enable**
Default: Enable
Options: Enable | Disable
Function: Enables or disable this import route filter.
Instructions: Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now wish to re-enable it.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.2

Parameter: Preference

Default: 1

Range: 1 to 16

Function: Assigns a weighted precedence value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, BGP-3, static, OSPF, external, and RIP. If Intra-AS IBGP routing is used, then any other route source is preferred over a BGP-3 route.

If this hierarchy is acceptable, accept the default value 1 for preference. If you want to grant preference to this BGP-3-derived route, assign a new preference value in the range of 1 to 16 (the greater the number, the higher the preference).

Note: The default preference for static routes is 16, but may be set to any value between 1 and 16. If you want to grant a BGP-3-derived route preference over a static route, make sure the preference you assign to the BGP-3-derived route exceeds the preference value of the static route you want it to override.

Instructions: Either accept the default value 1, or enter a new value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.14.1.11

Parameter: BGP-3 Preference

Default: 1

Range: 1 to 2147483647

Function: Assigns a weighted precedence value to a route included in the routing tables. If confronted with multiple BGP-3 routes to the same destination, the router, by default, grants preference to routes assigned the highest preference value.

Instructions: Either accept the default value 1, or enter a new value.

Editing a BGP-3 Import Route Filter

You can edit the Enable, Action, Preference, and BGP-3 Preference parameters for BGP-3 import route filters.

Note: You *cannot* reconfigure the Import Address, Import Mask, Import Peer AS, Import Peer Address, Import Peer Original AS, and Import Route Origin parameters for a BGP import route filter. To change these parameters, you must delete the filter and add a new filter with the proper information. See “Deleting a BGP-3 Import Route Filter” on page 10-47 for instructions.

To edit a BGP-3 import route filter, begin at the BGP-3 Import Route Filters window and proceed as follows:

1. Click on the import route filter you want to edit.
2. Click on Edit.

The BGP-3 Import Route Filter window appears.

3. Edit those parameters you want to change.

All BGP-3 parameters are described in the section “BGP-3 Import Route Filter Parameter Descriptions.”

4. Click OK to implement your changes.
5. Click on Cancel to exit the window.

Deleting a BGP-3 Import Route Filter

To delete an import route filter, begin at the BGP-3 Import Route Filters window and proceed as follows:

1. Click on the import route filter you wish to delete.
2. Click on Delete.
3. Click on Delete to delete the import route filter.
4. Click on Cancel to exit the window.

Configuring BGP-3 Export Route Filters

To add, edit, or delete BGP-3 export route filters, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols→IP→Route Filters→ BGP-3→Export Filters option.

The BGP-3 Export Route Filters List window appears (see Figure 10-16). It lists all BGP-3 export route filters configured on the router. You add, edit, and delete BGP-3 export route filters from this window.

2. Add, edit or delete export route filters as described in the following sections.

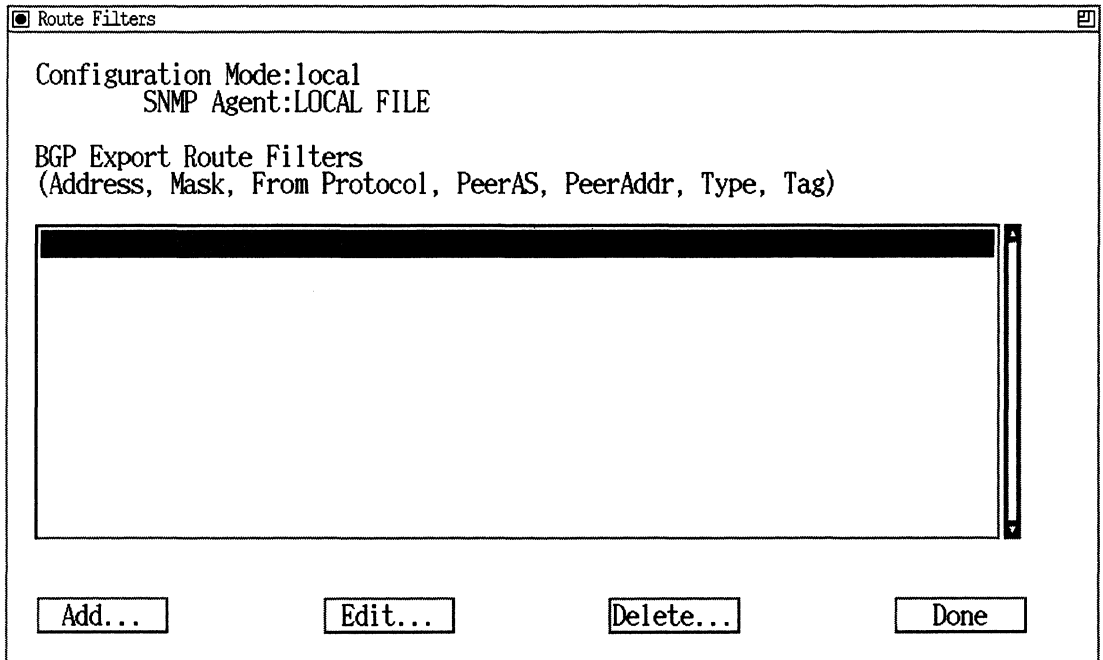


Figure 10-16. BGP-3 Export Route Filters List Window

Adding a BGP-3 Export Route Filter

To add an export route filter, begin at the BGP-3 Export Route Filters window and proceed as follows:

1. Click on Add.

The BGP-3 Export Route Filter Configuration window appears (see Figure 10-17). All parameters on this window display the default settings.

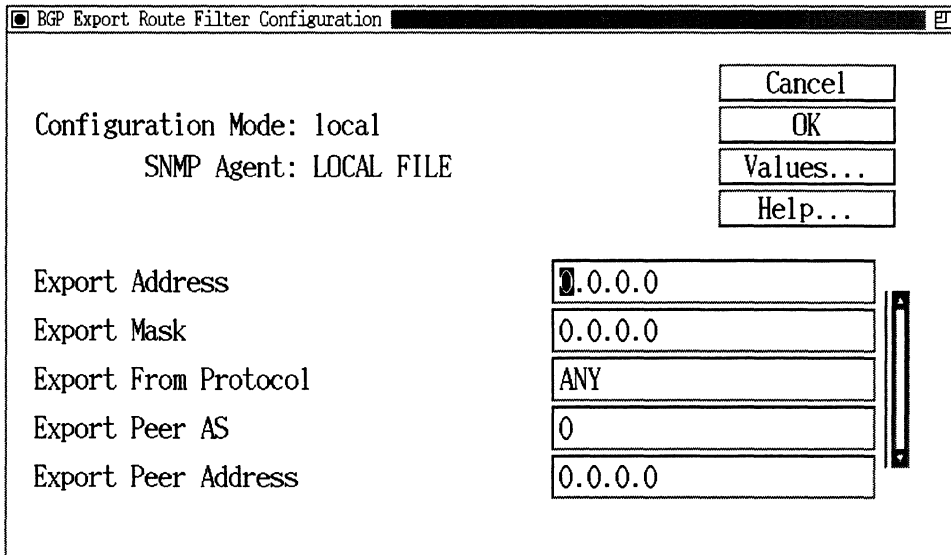


Figure 10-17. BGP-3 Export Route Filter Configuration Window

2. Specify your own settings for the BGP-3 export route filter configuration parameters.

All BGP-3 import route filter parameters are described following these instructions.

3. Click on OK.

After you click on OK, the BGP-3 Export Route Filter window appears (see Figure 10-18). All parameters on this window display the default settings.

BGP-3 Export Route Filter Parameter Descriptions

This section describes how to set all BGP-3 export route filter parameters.

| | |
|-----------------------|---|
| Parameter: | Export Address |
| Default: | 0.0.0.0 |
| Range: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. If this field is left blank, the filter applies to all networks. |
| Instructions: | Enter the appropriate network address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.15.1.3 |

Parameter: Export Mask

Default: 0.0.0.0

Range: Depends on the address class of the network address.

Function: Specifies the range of addresses upon which this filter acts.

For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. In other words, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Export Mask, only the Net ID portion of the address will be filtered. If the mask 255.255.255.0 is entered at Export Mask, the Net ID and Subnet ID portions of the address will be filtered.

If the Export Address field is set to 0.0.0.0, and the Export Mask is set to 0.0.0.0, then the filter applies to *all* routes. If the Export Address field is set to 0.0.0.0, and the Export Mask is set to 255.255.255.255, then the filter applies to the *default* route.

Instructions: Enter the appropriate mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.4

Parameter: **Export from Protocol**
Default: Any
Options: Any, RIP, EGP, OSPF, Direct, Static, BGP-3
Function: Identifies the source of the routing information: Direct connection, static route, or a RIP, EGP, OSPF, or BGP - 3 derived route.
Instructions: Select the appropriate option.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.5

Parameter: **Export Peer AS**
Default: 0
Range: 1 to 65535
Function: Identifies the Autonomous System to which the BGP router at the remote end of this BGP peer connection belongs. This filter will apply to updates sent to any router in this AS. The value 0 means "any" AS.
Instructions: Enter the appropriate AS number.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.7

Parameter: Export Peer Address

Default: 0.0.0.0

Range: Any IP address

Function: Specifies the IP address of the interface on the remote side of this BGP peer connection. This filter will apply to updates sent to this router. The value 0.0.0.0 means “any” peer.

Instructions: Enter the IP address in dotted decimal notation. If the peer is in a remote AS, the address must be on the same subnet as the local interface.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.8

Parameter: Export Enable

Default: Enable

Options: Enable | Disable

Function: Enables or disables this export route filter.

Instructions: Set to Disable if you want to disable this filter. Set to Enable if you want to Enable this filter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.2

Parameter: Export Action

Default: Ignore

Options: Propagate | Ignore | Aggregate

Function: Controls the flow of routing information. If set to Propagate, this route is advertised. If set to Ignore, advertising of this route is suppressed. If set to Aggregate, the network is not explicitly advertised. Instead, the default route (0.0.0.0) is advertised.

Instructions: Select Propagate, Ignore, or Aggregate.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.6

Parameter: Export Use Inter AS Metric

Default: None

Options: None, Specified, Originating

Function: Specifies whether or not an Inter AS metric is advertised for the associated networks. If set to None, then no metric is advertised. If set to Specified, then the value specified for the Export Inter AS Metric parameter is advertised. If set to Originating, then the metric from the originating protocol is advertised. This parameter is only valid if Export Action is set to propagate.

Instructions: Set to the appropriate option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.11

Parameter: Export Inter AS Metric

Default: None

Range: 0 to 65535

Function: If the Export Use Inter AS Metric parameter is set to Specified, then this is the Inter AS Metric value that is advertised.

Instructions: Specify a value within the assigned range.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.12

Parameter: Export Origin

Default: Any

Options: Any, IGP, EGP, Incomplete

Function: If the from protocol is RIP or Static, and Action is propagate, you can use this parameter to change the ORIGIN attribute that is advertised for this network.

Instructions: If you want to change the ORIGIN attribute, select a valid option.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.13

Parameter: Export Neighbor AS

Default: 0

Range: 0 to 65535

Function: If the Export Action is set to Propagate, and the Export Origin is set to EGP, then this parameter must be set to a non-zero value. The value specified here is used as the EGP neighbor AS number when the AS path is constructed.

Instructions: Specify a value within the assigned range.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.15.1.14

Editing a BGP-3 Export Route Filter

You can edit Export Peer Address, Export Enable, Export Action, Export Use Inter AS Metric, Export Inter AS Metric, Export Origin, and Export Neighbor AS parameters for a BGP-3 export route filter.

Note: You *cannot* edit the Export Address, Export Mask, Export From Protocol, or Export Peer AS parameters for a BGP-3 export route filter. To change these parameters, you must delete the filter and add a new filter with the proper information. See “Deleting a BGP-3 Export Route Filter” on page 10-58 for instructions.

To edit the BGP-3 export router filter parameters, begin at the BGP-3 Export Route Filters window and proceed as follows:

1. Select the export route filter you want to edit.
2. Click on Edit.

The BGP-3 Export Route Filter window for that filter appears.

3. Edit those parameters you want to change.

All BGP-3 export route filter parameters are described in the section “Export Route Filter Parameter Descriptions.”

4. Click on OK to implement your changes.
5. Click on Cancel to exit the window.

Deleting a BGP-3 Export Route Filter

To delete a BGP-3 export route filter, begin at the BGP-3 Export Route Filters window and proceed as follows:

1. Click on the BGP-3 export route filter you wish to delete.
2. Click on Delete.
3. Click on Delete to delete the export route filter.

Click on Cancel to exit the window.

EGP Route Filters

The following sections show you how to select EGP route filter windows from the Site Manager and describes all EGP route filter parameters.

Configuring EGP Import Route Filters

To add, edit, or delete EGP import route filters, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols→IP→Route Filters→EGP→Import Filters option.

The EGP Import Route Filters List window appears (see Figure 10-19). It lists all EGP import route filters configured on the router. You add, edit, and delete EGP import route filters from this window.

2. Add, edit or delete import route filters as described in the following sections.

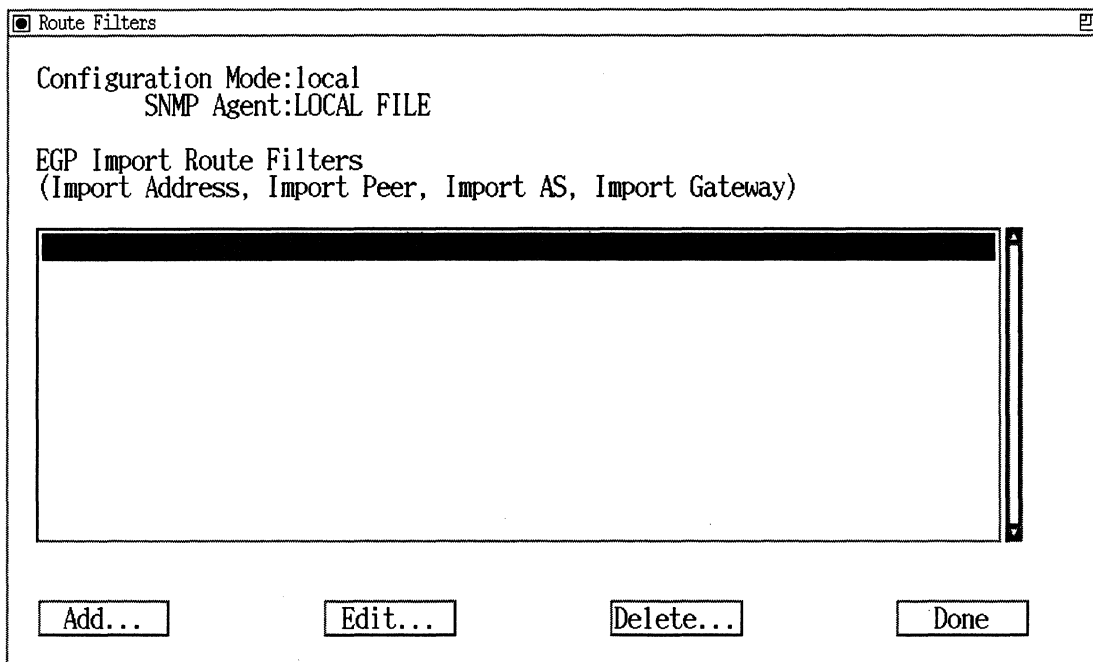


Figure 10-19. EGP Import Route Filters List Window

Adding an EGP Import Route Filter

To add an import route filter, begin at the EGP Import Route Filters window (shown in Figure 10-19) and proceed as follows:

1. Click on Add.
The EGP Import Route Filter Configuration window appears (see Figure 10-20).
2. Specify the EGP import route filter configuration parameters.
All EGP import route filter configuration parameters are described following these instructions.
3. Click on OK.

The EBG Import Route Filters window appears (see Figure 10-21). When you add an import route filter, the Configuration Manager automatically sets the Enable, Action, Preference, and EGP Preference parameters in this window.

4. Either accept the default settings, or edit these parameters to your network specifications, then click on OK.

This section provides information you need to set each parameter.

5. When you are done, click on Cancel to exit the window.

EGP Import Route Filter Configuration

Configuration Mode: local
SNMP Agent: LOCAL FILE

Import Address: 192.32.12.12
Import Peer: 192.32.12.13
Import Autonomous System: 4
Import Gateway: 192.32.12.1

Buttons: Cancel, OK, Values..., Help...

Figure 10-20. EGP Import Route Filter Configuration Window

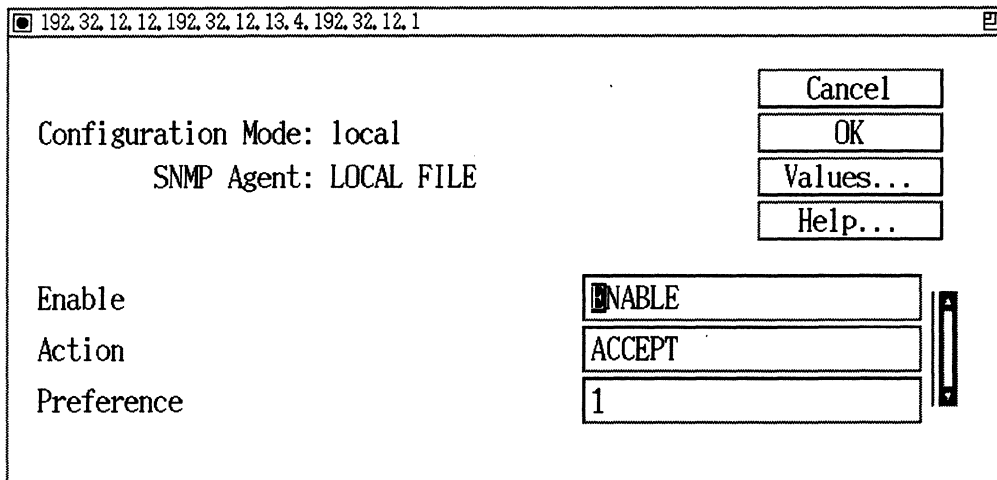


Figure 10-21. EGP Import Route Filter Window

EGP Import Route Filter Parameter Descriptions

This section describes how to set all EGP import route filter configuration parameters.

| | |
|-------------------|---|
| Parameter: | Import Address |
| Default: | 0.0.0.0 |
| Range: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. If this field is set to 0.0.0.0, the filter applies to all networks. |
| Instructions: | Enter the appropriate network address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.12.1.3 |

Parameter: Import Peer**Default:** 0.0.0.0**Range:** Any IP address**Function:** Specifies the IP address of the interface on the remote side of this EGP peer connection. This filter will apply to updates from this router. The default 0.0.0.0 means “any” peer.**Instructions:** Enter the IP address in dotted decimal notation. If the peer is in a remote AS, the address must be on the same subnet as the local interface**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.12.1.7**Parameter: Import Autonomous System****Default:** 0**Range:** 0 - 65536**Function:** Identifies the Autonomous System to which the EGP router at the remote end of this EGP peer connection belongs. This filter will apply to updates from this router. The default 0 means “any” AS.**Instructions:** Enter the appropriate AS number.**MIB Object ID:** 1.3.6.1.4.1.18.3.5.3.2.1.12.1.8

| | |
|-------------------|--|
| Parameter: | Import Gateway |
| Default: | 0.0.0.0 |
| Range: | Any IP address |
| Function: | Specifies the gateway advertised as the next hop for the network. The default value of 0 means “any” gateway. |
| Instructions: | Enter the appropriate gateway number. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.12.1.9 |
| | |
| Parameter: | Enable |
| Default: | Enable |
| Options: | Enable Disable |
| Function: | Enables or disable this import route filter. |
| Instructions: | Set to Disable if you want to disable this filter. Set to Enable if you previously disabled this filter and now wish to re-enable it. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.12.1.2 |
| | |
| Parameter: | Action |
| Default: | Accept |
| Options: | Accept Ignore |
| Function: | Specifies whether the route is transferred to the routing tables. If Action is set to Accept (default), the routing information is sent to the routing tables. If Action is set to Ignore, the routing information is dropped. |
| Instructions: | Either accept the default Accept, or select Ignore. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.12.1.5 |

Parameter: Preference

Default: 1

Range: 1 to 15

Function: Assigns a weighted precedence value to a route included in the routing tables. If confronted with multiple routes to the same destination, the router, by default, grants preference to routes in the following order: direct, OSPF internal, static, BGP-3, OSPF external, and RIP.

If this hierarchy is acceptable, accept the default value 1 for preference. If you want to grant preference to this OSPF-derived route, assign a new preference value in the range of 1 to 15 (the greater the number, the higher the preference).

Instructions: Either accept the default value 1, or enter a new value.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.12.1.6

Editing an EGP Import Route Filter

You can edit the Enable, Propagate, Interface, and Metric parameters for EGP import route filters.

Note: You *cannot* reconfigure the Import Address, Import Mask, Import Peer AS, Import Peer Address, Import Peer Original AS, Import Route Origin, or Import Action parameters for a BGP import route filter. To change these parameters, you must delete the filter and add a new filter with the proper information. See “Deleting an EGP Import Route Filter” on page 10-66 for instructions.

To edit an import route filter, begin at the EGP Import Route Filters List window and proceed as follows:

1. Click on the import route filter you want to edit.

2. Click on Edit.

The EGP Import Route Filter window appears.

3. Edit those parameters you want to change.

All EGP parameters are described in the section “EGP Import Route Filter Parameter Descriptions.”

4. Click on OK to implement your changes.
5. Click on Cancel to exit the window.

Deleting an EGP Import Route Filter

To delete an import route filter, begin at the EGP Import Route Filters List window and proceed as follows:

1. Click on the import route filter you wish to delete.
2. Click on Delete.
3. Click on Delete to delete the import route filter.
4. Click on Cancel to exit the window.

Configuring EGP Export Route Filters

To add, edit, or delete EGP export route filters, begin at the Wellfleet Configuration Manager window and proceed as follows:

1. Select the Protocols→IP→Route Filters→ EGP→Export Filters option.

The EGP Export Route Filters window appears (see Figure 10-22). It lists all EGP export route filters configured on the router. You add, edit, and delete EGP export route filters from this window.

2. Add, edit or delete export route filters as described in the following sections.

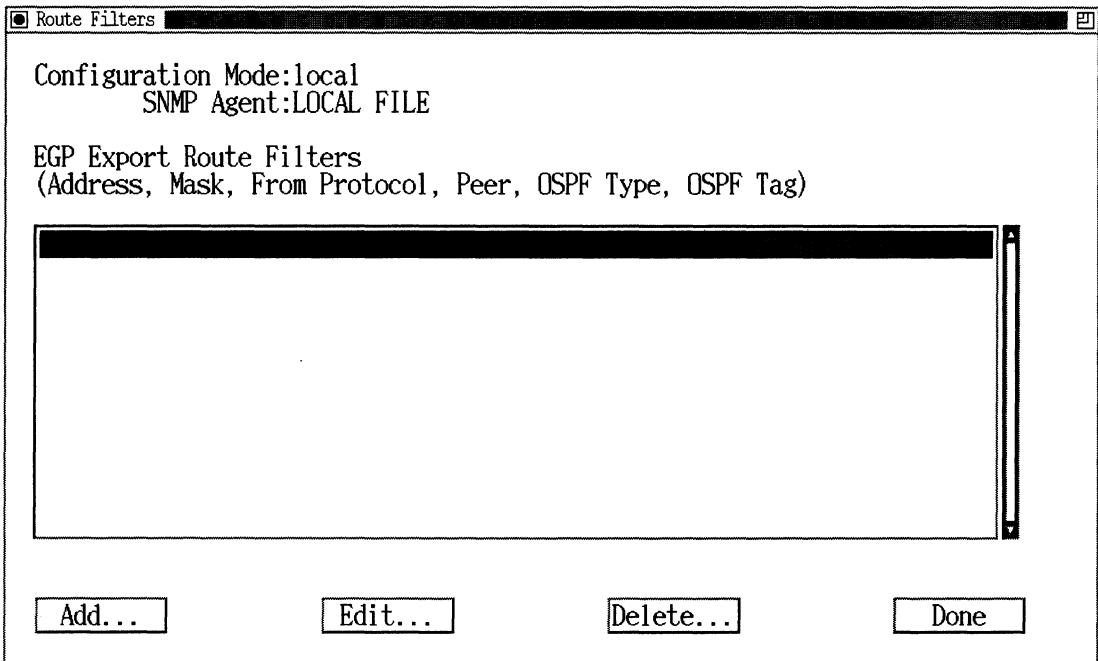


Figure 10-22. EGP Export Route Filters List Window

Adding an EGP Export Route Filter

To add an export route filter, begin at the EGP Export Route Filters List window shown in Figure 10-22 and proceed as follows:

1. Click on Add.

The EGP Export Route Filter Configuration window appears (see Figure 10-23).

2. Specify the EGP export route filter configuration parameters.

All EGP export route filter configuration parameters are described following these instructions.

3. Click on OK.

4. Depending on what you specified at the Export from Protocol parameter, do one of the following tasks:
 - If you specified OSPF, then the EGP OSPF Export Route Filters window appears. Specify the Export OSPF Type and Export OSPF Tag parameters, then click on OK. The EGP Export Route Filters window then appears as described below.
 - If you specified any of the other choices, the EGP Export Route Filters window appears immediately (see Figure 10-24). It displays the default settings for the Enable, Action, Interface, and Metric parameters. Either accept the default settings, or edit these parameters to your network specifications, then click on OK.
5. Click on Cancel to exit the window.

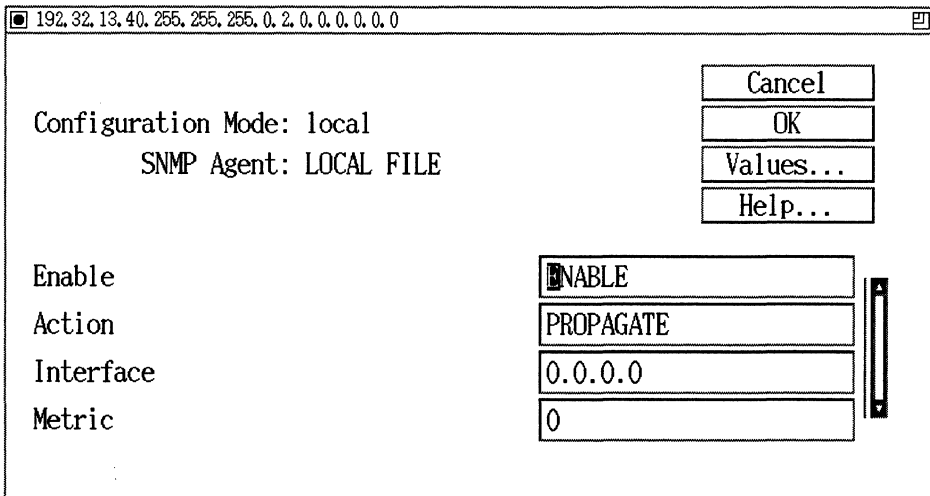


Figure 10-23. EGP Export Route Filter Configuration Window

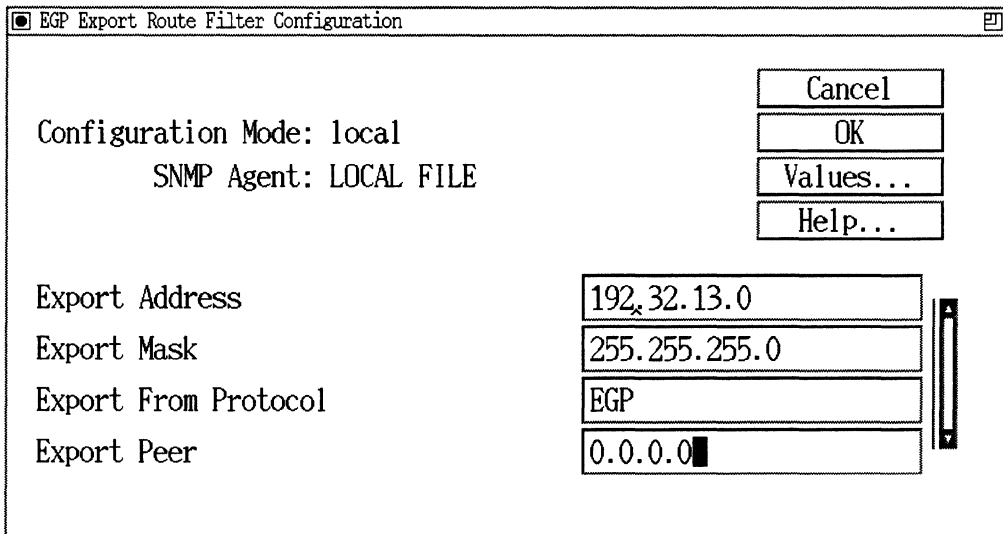


Figure 10-24. EGP Export Route Filter Window

EGP Export Route Filter Parameter Descriptions

This section describes how to set all EGP export route filter parameters.

| | |
|-------------------|---|
| Parameter: | Export Address |
| Default: | 0.0.0.0 |
| Range: | Any IP network address |
| Function: | Identifies, by IP address, the network to which this filter applies. If set to 0.0.0.0, the filter applies to all networks. |
| Instructions: | Enter the appropriate IP address in dotted decimal notation. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.13.1.3 |

Parameter: Export Mask

Default: 0.0.0.0

Range: Depends on the address class of the network address.

Function: Specifies the range of addresses this filter acts upon.

For example, consider Class B Network 172.32.0.0, which allocates the upper 8 bits of the host identification field to the Subnet ID, and the final 8 bits to the Host ID. The address mask directs the filtering process to a specific portion of the IP address. Thus, any IP address that matches the masked portion of 172.32.0.0 is subject to filtering. If 255.255.0.0 is entered at Export Mask, only the Net ID portion of the address is filtered. If the mask 255.255.255.0 is entered at Export Mask, the Net ID and Subnet ID portions of the address is filtered.

If the Export Address field is set to 0.0.0.0, and the Export Mask is set to 0.0.0.0, then the filter applies to *all* routes. If the Export Address field is set to 0.0.0.0 and the Export Mask is set to 255.255.255.255, then the filter applies to the *default* route.

Instructions: Enter the mask in dotted decimal notation.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.4

Parameter: **Export from Protocol**
Default: Any
Options: Any, RIP, EGP, OSPF, Direct, Static, BGP-3
Function: Identifies the source of the routing information: direct connection, static route, or RIP, EGP, OSPF, or BGP-3-derived route.
Instructions: Select the appropriate option.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.5

Parameter: **Export Peer**
Default: 0.0.0.0
Range: Any IP address
Function: Specifies the IP address of the interface on the remote side of this EGP peer connection. This filter will apply to updates from this router. The default value 0.0.0.0 means “any” router.
Instructions: Enter the IP address in dotted decimal notation. The address must be on the same subnet as a local interface.
MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.7

Parameter: Export OSPF Type

Default: None

Options: Type 1, Type 2, Internal

Function: Specifies the type of routes to which this filter applies. If you specify Type 1, then only AS External Type 1 routes are filtered. If you specify Type 2, then only AS External Type 2 routes are filtered.

Note that this parameter is only used if the Export From Protocol parameter is set to OSPF.

Instructions: Depending on the type of routes you want to filter, select Type 1 or Type 2, or Internal.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.8

Parameter: Export OSPF Tag

Default: 0

Range: 0 to 2147483647

Function: Specifies the tag with which this route filter is concerned. Each AS External Advertisement contains a tag field. If the tag field matches Import Tag, the appropriate action is taken; either the route is accepted or ignored.

Note that this parameter is only used if the Export From Protocol parameter is set to OSPF.

Instructions: Enter the appropriate tag number.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.9

Parameter: Enable

Default: Enable

Options: Enable | Disable

Function: Enables or disables this export route filter.

Instructions: Set to Disable if you want to disable this export route filter. Set to Enable if you previously disabled this export route filter and now want to re-enable it.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.2

Parameter: Action

Default: Propagate

Options: Propagate | Ignore

Function: Controls the flow of routing information. If Action is set to Propagate, this route is advertised. If Action is set to Ignore, advertising of this route is suppressed.

Instructions: Either accept the default Propagate, or select Ignore.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.6

Parameter: Interface

Default: 0.0.0.0

Range: Any IP address

Function: Specifies the outbound interface on which to apply this filter

Instructions: Specify the IP address of the interface on which you want to apply this filter.

MIB Object ID: 1.3.6.1.4.1.18.3.5.3.2.1.13.1.10

| | |
|-------------------|--|
| Parameter: | Metric |
| Default: | 0 (0 = the actual route cost as learned) |
| Range: | 0 to 255 |
| Function: | Assigns an EGP cost to the propagated route. The value 0 causes the actual route cost (as learned) to be used. |
| Instructions: | Either accept the default Metric value 0, or enter a new value. |
| MIB Object ID: | 1.3.6.1.4.1.18.3.5.3.2.1.13.1.11 |

Editing an EGP Export Route Filter

You can the edit Enable, Action, Interface, and Metric parameters for an EGP export route filter.

Note: You *cannot* edit the Export Address, Export Mask, Export From Protocol, Export from Peer, Export OSPF Type or Export OSPF Tag parameters for an EGP export route filter. To change these parameters, you must delete the filter and add a new filter with the proper information. See “Deleting an EGP Export Route Filter” on page 10-75 for instructions.

To edit the EGP export router filter parameters, begin at the EGP Export Route Filters window and proceed as follows:

1. Click on the export route filter you want to edit.
2. Click on Edit.

The EGP Export Route Filter window for that filter appears.

3. Edit those parameters you want to change.

All EGP export route filter parameters are described in the section “EGP Export Route Filter Parameter Descriptions.”

4. Click on OK to implement your changes.
5. Click on Cancel to exit the window.

Deleting an EGP Export Route Filter

To delete an EGP export route filter, begin at the EGP Export Route Filters window and proceed as follows:

1. Click on the EGP export route filter you wish to delete.
2. Click on Delete.
3. Click on Delete to delete the export route filter.



Danger Click on Cancel to exit the window.



A

- accept policies, configuring, 9-5
- accept policy parameters, 9-13
 - BGP-3-specific, 9-15
 - BGP-4-specific, 9-19
 - common, 9-8
 - EGP, 9-13
 - OSPF-specific, 9-12
 - RIP-specific, 9-11
- adding
 - a range to an OSPF area, 4-35
 - adjacent hosts, 2-65
 - BGP-3 export route filters, 10-48
 - BGP-3 import route filters, 10-38
 - EGP export route filters, 10-66
 - EGP import route filters, 10-60
 - neighbors to an OSPF interface, 4-51 to 4-53
 - OSPF areas, 4-31
 - OSPF export route filters, 10-29
 - OSPF import route filters, 10-19
 - RIP export route filters, 10-11
 - RIP import route filters, 10-3
 - static routes, 2-59
 - virtual interfaces, 4-56 to 4-58
 - weight value to an AS, 5-46
- Address Resolution Protocol
 - function of, 2-6
 - HP Probe, 2-10
 - Inverse ARP, 2-9
 - proxy ARP, 2-8
 - X.25 DDN and PDN, 2-10

- adjacent hosts
 - adding, 2-65
 - configuring, 2-64
 - definition of, 2-6
 - deleting, 2-70
 - editing, 2-67
- aggregate route, definition of, 1-10
- announce policies, configuring, 9-25
- announce policy parameters
 - BGP-3-specific, 9-48
 - BGP-4-specific, 9-53
 - common, 9-27
 - EGP-specific, 9-45
 - OSPF-specific, 9-42
 - RIP-specific, 9-40
- ARP
 - See* Address Resolution Protocol
- AS weights
 - configuring, 5-45 to 5-52
- autonomous systems
 - definition of, 1-10

B

- BGP
 - AS weight classes, 5-14
 - AS weights, 5-14
 - best route calculation, 5-15, 5-17
 - IBGP intra-AS routing, 5-19
 - IBGP transit AS routing, 5-18
 - interaction with OSPF, 5-17

-
- Keepalive message, 5-8
 - Local Preference Attribute, 5-15
 - message logging, 5-20
 - Notification message, 5-12
 - Open message, 5-7
 - Update message, 5-8
 - BGP parameters
 - BGP
 - AS, 5-47
 - Weight, 5-47, 5-48, 5-49, 5-50, 5-51
 - BGP peers
 - Connect Retry Timer, 5-40
 - Enable, 5-38
 - External Advertisement Timer, 5-39
 - Holdtime, 5-41
 - Keepalive Timer, 5-42
 - Local Address, 5-37
 - Local AS to Advertise to Peer, 5-43
 - Max BGP Version, 5-38
 - Max Update Size, 5-43
 - Min AS Origination Interval, 5-42
 - Min BGP Version, 5-38
 - Peer Address, 5-36
 - Peer AS, 5-36
 - Remote AS, 5-39
 - Route Echo Switch, 5-44
 - event logging
 - Local IP Address, 5-55
 - Message Level, 5-56
 - Message Trace Switch, 5-56
 - Remote Address, 5-55
 - global
 - BGP Enable, 5-24
 - BGP Identifier, 5-24
 - BGP Interval Timer, 5-26
 - BGP Local AS, 5-25
 - Collision Detect, 5-27
 - From Protocols, 5-26
 - IBGP Intra AS Routing, 5-25
 - Multi-hop EBGp Connection, 5-28
 - configuring, 5-32, 5-44
 - BGP-3
 - accept policy parameters, 9-15
 - announce policy parameters, 9-48
 - export route filters
 - adding, 10-48
 - configuring, 10-47 to 10-58
 - deleting, 10-58
 - editing, 10-57
 - import route filter, 10-40
 - import route filters
 - configuring, 10-37
 - deleting, 10-47
 - editing, 10-46
 - BGP-3 parameters
 - BGP-3 Preference, 10-46
 - Enable, 10-44
 - Export Action, 10-55
 - Export Address, 10-51
 - Export Enable, 10-54
 - Export from Protocol, 10-53
 - Export Inter AS Metric, 10-56
 - Export Mask, 10-52
 - Export Neighbor AS, 10-56
 - Export Origin, 10-56
 - Export Peer Address, 10-54
 - Export Peer AS, 10-53
 - Export Use Inter AS Metric, 10-55
 - global
 - Enable, 5-30
 - Import Address, 10-40
 - Import Mask, 10-41
 - Import Peer Address, 10-42
 - Import Peer AS, 10-42
 - Import Peer Original AS, 10-43
 - Import Route Origin, 10-43
 - Preference, 10-45
- BGP-4
 - accept policy parameters, 9-19
 - announce policy parameters, 9-53
- BGP-4 parameters
-

- global
 - Enable, 5-32
- Blacker Front-End support, 2-22, 2-23
 - addressing, 2-24
 - configuring, 2-90
 - required X. 25 network service record parameter settings, 2-94
- broadcast address
 - definition of, 2-4
 - for subnets, 2-5
- broadcast network, 4-3

C

- circuitless IP interfaces, 2-14
 - configuring, 2-57
- Classless Inter-Domain Routing (CIDR), 1-9
- configuring
 - adjacent hosts, 2-64
 - BGP AS weights, 5-45 to 5-52
 - BGP peers, 5-32 to 5-44
 - BGP-3 export route filters, 10-47 to 10-58
 - BGP-3 import route filters, 10-37
 - circuitless IP interfaces, 2-57
 - OSPF import route filters, 10-18
 - OSPF virtual interfaces, 4-55 to 4-63
 - RIP export route filters, 10-11 to 10-17
 - RIPSO support, 2-74
 - static routes, 2-57 to 2-64
- customizing
 - BGP-3 services, 5-1 to 5-57
 - EGP services, 6-1 to 6-24

D

- datagram, 1-2
- DDN X.25 address resolution, 2-10

- deleting
 - a range from an OSPF area, 4-40
 - adjacent hosts, 2-70
 - BGP, 5-57
 - BGP-3, 5-57
 - BGP-3 export route filters, 10-58
 - BGP-3 import route filters, 10-47
 - EGP, 6-24
 - EGP export route filters, 10-75
 - EGP import route filters, 10-66
 - IP from an interface, 2-56
 - OSPF areas, 4-35
 - OSPF export route filters, 10-36
 - OSPF import route filters, 10-27
 - OSPF neighbors, 4-55
 - OSPF virtual interfaces, 4-63
 - RIP export route filters, 10-17
 - RIP import route filters, 10-10
 - static routes, 2-64
 - weight values from an AS
 - BGP, 5-52
- Distance Vector Multicast Routing Protocol (DVMRP), 7-4
- DVMRP (Distance Vector Multicast Routing Protocol), 7-4
- DVMRP parameters
 - circuit
 - Enable, 7-20
 - Metric, 7-22
 - Route Enable, 7-21
 - Threshold, 7-22
 - global
 - Enable, 7-13
 - Estimated Routes, 7-17
 - Full Update Rate, 7-13
 - Garbage Timeout, 7-16
 - Leaf Timeout, 7-14
 - Neighbor Probe Interval, 7-17
 - Neighbor Timeout, 7-15
 - Route Expiration Timeout, 7-15
 - Route Switch Timeout, 7-18

-
- Triggered Update Rate, 7-14
 - tunnel
 - Enable, 7-24
 - Encapsulation Mode, 7-24
 - Local IP Address, 7-27
 - Metric, 7-25
 - Remote IP address, 7-27
 - Threshold, 7-25
- E**
- editing
 - adjacent hosts, 2-67
 - an OSPF area's range, 4-38
 - BGP-3 export route filters, 10-57
 - BGP-3 import route filters, 10-46
 - EGP export route filters, 10-74
 - EGP import route filters, 10-65, 10-66
 - IP parameters, 2-25
 - OSPF area parameters, 4-30 to 4-35
 - OSPF areas, 4-32
 - OSPF export route filters, 10-35
 - OSPF import route filters, 10-27
 - OSPF interface parameters, 4-40 to 4-50
 - OSPF interfaces, 4-42
 - OSPF neighbors, 4-53
 - OSPF parameters, 4-22 to 4-26
 - OSPF virtual interfaces, 4-58 to 4-63
 - RIP export route filters, 10-17
 - RIP import route filters, 10-10
 - RIP parameters, 3-2
 - static routes, 2-61
 - TFTP parameters, 2-71
 - weight value parameters of an AS, 5-51
 - EGP, 9-13
 - announce policy parameters, 9-45
 - deleting neighbors, 6-24
 - export route filter
 - configuring, 10-66 to 10-75
 - deleting, 10-75
 - editing, 10-74
 - import route filter
 - adding, 10-60
 - configuring, 10-59 to 10-66
 - deleting, 10-66
 - editing, 10-65, 10-66
 - modes, 6-4, 6-12
 - Neighbor Acquisition Cease Ack Response, 6-4
 - Neighbor Acquisition Cease Command, 6-4
 - Neighbor Acquisition Confirm Response, 6-3
 - Neighbor Acquisition Phase, 6-3
 - Neighbor Acquisition Refuse Response, 6-4
 - Neighbor Acquisition Request Command, 6-3
 - neighbor reachability phase, 6-7 to 6-10
 - network reachability phase, 6-10 to 6-13
 - overview of, 6-1
 - EGP neighbors
 - configuring, 6-17 to 6-24
 - deleting, 6-24
 - EGP parameters
 - Acquisition Mode, 6-22
 - Action, 10-64, 10-73
 - editing, 6-14 to 6-24
 - Enable, 6-21, 10-64, 10-73
 - Export Address, 10-69
 - Export from Protocol, 10-71
 - Export Mask, 10-70
 - Export OSPF Tag, 10-72
 - Export OSPF Type, 10-72
 - Export Peer, 10-71
 - global, 6-16
 - Hello Timer, 6-23
 - Import Address, 10-62
 - Import AS, 10-63
 - Import Gateway, 10-64
 - Import Peer, 10-63
 - import route filter, 10-62 to 10-65
-

- Interface, 10-73
- Metric, 10-74
- Poll Mode, 6-22
- Poll Timer, 6-23
- Preference, 10-65

Enable Default Route for Subnets, 2-36

H

- host groups, multicasting, 7-1
- HP Probe, definition of, 2-10

I

IBGP

- intra-AS routing, 5-19
- transit AS routing, 5-18

IGMP (Internet Group Management Protocol), 7-2

IGMP parameters

- entry
 - Designated Router Timeout, 7-32
 - Enable, 7-31
 - Interface Membership Timeout, 7-32
- global
 - Enable, 7-29
 - Estimated Groups, 7-29

implementation notes

- EGP, 6-13
- OSPF, 4-20

interface, definition of, 2-2

Interior Gateway Protocol (IGP), 1-10

Internet Group Management Protocol (IGMP), 7-2

Internet Network Information Center (NIC), 1-3

Internet Requests for Comments (RFCs)
IP router compliance, 1-17

Inverse ARP, 2-9

IP address

- definition of, 1-3
- network classes, 1-4
- specifying in dotted decimal notation, 1-5

IP datagram, 1-2

- definition of, 1-2
- Header Checksum field, 1-3
- Options field, 1-3
- Time to Live field, 1-3
- Type of Service field, 1-2

IP parameters

- adjacent host
 - Adjacent Host X.121 Address, 2-70
 - Enable, 2-68
 - Host Encapsulation, 2-70
 - IP Address, 2-66
 - MAC Address, 2-69
 - Next Hop Interface Addr, 2-68
 - Next Hop Interface Mask, 2-69

global

- ARP Forwarding, 2-30
- Default TTL, 2-31
- Enable, 2-28
- Estimated Hosts, 2-36
- Estimated Networks, 2-35
- Forwarding, 2-29
- Maximum Policy Rules, 2-37
- Non Local ARP Source, 2-30
- Nonlocal ARP Destination, 2-31
- RIP Diameter, 2-32
- Route Cache Flush Interval, 2-32
- Routing MIB Table(s), 2-33
- Zero Subnet Enable, 2-34

interface

- Addr Mask Reply, 2-43
- Address Resolution, 2-44
- All Subnet Bcast, 2-43
- Broadcast Address, 2-40
- Checksum, 2-47
- Enable, 2-39

- Enable Security, 2-56
- Enet Arp Encaps, 2-50
- FR Broadcast DLCI, 2-52
- FR Multicast DLCI#1, 2-52
- FR Multicast DLCI#2, 2-53
- Host Cache, 2-46
- Interface Cost, 2-41
- MAC Address, 2-48
- Max Forwarding Table Size, 2-55
- MTU Discovery, 2-42
- Proxy, 2-45
- Redirects, 2-49
- Slot Mask, 2-54
- SMDS Arp Req Address, 2-51
- SMDS Group Address, 2-51
- Subnet Mask, 2-39
- TR Endstation, 2-48

RIPSO

- Default Authority, 2-83
- Default Label, 2-83
- Default Level, 2-84
- Enable Security, 2-75
- Error Authority, 2-85
- Error Label, 2-84
- Implicit Authority, 2-82
- Implicit Label, 2-81
- Implicit Level, 2-82
- Maximum Level, 2-79
- May In Authority, 2-81
- May Out Authority, 2-80
- Minimum Level, 2-78
- Must In Authority, 2-80
- Must Out Authority, 2-79
- Require In Security, 2-78
- Require Out Security, 2-77
- Strip Security, 2-76

static route

- Address Mask, 2-60
- Cost, 2-62
- Destination IP Address, 2-60
- Enable, 2-62
- Next Hop Addr, 2-63

- Next Hop Mask, 2-63
- Preference, 2-64
- TFTP
 - Close Time Out, 2-73
 - Default Volume, 2-72
 - Enable, 2-72
 - Retransmit, 2-73
 - Retry Time Out, 2-73

IP router

- internal routing tables, 1-14
- IP, editing parameters for, 2-25

L

- Local Preference attribute, calculating, 5-15

M

- Multi Exit Discriminator Value (announce), 9-56

multicasting

- aging a route, 7-8
- comparing routes, 7-8
- creating a shortest path tree, 7-8
- DVMRP, 7-4
- host groups, 7-1
- IGMP, 7-2
- leaf network, 7-8
- threshold, 7-9
- tunnel, 7-4

multinet

- definition of, 2-4

N

- NetBIOS over IP, 8-1
 - adding a traffic filter, 8-8
 - aging a cache entry, 8-7

-
- configuring a cache, 8-6
 - configuring a static name, 8-5
 - customizing a cache search, 8-7
 - NetBIOS/IP parameters
 - global
 - 15-Character NetBIOS Name Caching, 8-12
 - Create MIB Inst for Cached Name, 8-13
 - Enable/Disable, 8-11
 - Max Name Cache Entries, 8-13
 - NetBIOS Name Caching, 8-12
 - Rebroadcast Packet TTL, 8-16
 - Rebroadcast Record Route, 8-16
 - interface
 - Enable NetBIOS Inbound Broadcasts, 8-19
 - Enable NetBIOS Outbound Broadcasts, 8-19
 - Enable/Disable, 8-18
 - NetBIOS Name Caching, 8-18
 - Rebroadcast Address, 8-20
 - static entry
 - Enable, 8-22
 - IP Address, 8-25
 - NetBIOS Scope ID, 8-22, 8-24
 - NetBIOS Station Name, 8-24
 - Network Basic Input-Output System (NetBIOS) over IP, 8-1
 - NIC
 - See Internet Network Information Center
 - nonbroadcast multiaccess network, 4-3
 - O**
 - OSPF
 - accept policy parameters, 9-12
 - adding
 - a range to an area, 4-35
 - areas, 4-31
 - neighbors to an interface, 4-51 to 4-53
 - announce policy parameters, 9-42
 - area border routers, 4-9
 - AS boundary routers, 4-9
 - AS external (ASE) route advertisements
 - using route weight in, 4-13
 - backup soloist, 4-18
 - broadcast interface, 4-3
 - configuring virtual parameters, 4-55 to 4-63
 - database synchronization, 4-2
 - deleting
 - a range from an area, 4-40
 - areas, 4-35
 - neighbors, 4-55
 - virtual interfaces, 4-63
 - editing
 - an area's range, 4-38
 - area parameters, 4-30 to 4-35
 - areas, 4-32
 - global parameters, 4-23 to 4-26
 - interface parameters, 4-40 to 4-50
 - neighbors, 4-53
 - virtual interfaces, 4-58 to 4-63
 - editing global parameters, 4-26
 - export route filters
 - adding, 10-29
 - deleting, 10-36
 - editing, 10-35
 - external routes, 4-10
 - features
 - configurable cost metrics, 4-11
 - link state protocol, 4-2
 - routing areas, 4-5
 - stub areas, 4-6
 - virtual links, 4-10, 4-11
 - import route filters
 - adding, 10-19
 - deleting, 10-27
 - editing, 10-27
-

- networks it supports, 4-3
- nonbroadcast multi-access interface, 4-3
- point-to-multipoint interface, 4-3
- point-to-point interface, 4-3
- router types, 4-8
 - area border routers, 4-8
 - AS Boundary routers, 4-9
 - backbone routers, 4-8
 - internal routers, 4-8
- specifying a preferred path, 4-11
- transit area, 4-10, 4-11
- types of routing
 - external routing, 4-9
 - inter-area routing, 4-9
 - intra-area routing, 4-9
- OSPF parameters
 - area
 - Authentication Type, 4-33
 - Enable, 4-33
 - Import AS Extern, 4-34
 - Import Summaries, 4-35
 - Range Mask, 4-38
 - Range Net, 4-37
 - Stub Metric, 4-34
 - area range
 - Enable, 4-39
 - Mask, 4-39
 - export route filters
 - Action, 10-33
 - Auto Tag, 10-35
 - Enable, 10-33
 - Export Address, 10-31
 - Export From Protocol, 10-33
 - Export Mask, 10-32
 - Tag, 10-34
 - Type, 10-34
 - global
 - AS Boundary Router, 4-25
 - ASE Metric Support, 4-26
 - Backup Disable, 4-27
 - Backup Log Mask, 4-28
 - Enable, 4-24
 - Hold Down Timer, 4-25
 - OSPF Slot, 4-26
 - Primary Log Mask, 4-27
 - Router ID, 4-24
 - import route filters
 - Action, 10-25
 - Enable, 10-25
 - Import Address, 10-21
 - Import Mask, 10-23
 - Import Tag, 10-24
 - Import Type, 10-24
 - Preference, 10-26
 - interface
 - Area ID, 4-43
 - Dead Interval, 4-47
 - Enable, 4-42
 - Hello Interval, 4-46
 - Metric Cost, 4-49
 - MTU Size, 4-50
 - Password, 4-50
 - Poll Interval, 4-48
 - Retransmit Interval, 4-45
 - Rtr Priority, 4-44
 - Transit Delay, 4-44
 - Type, 4-43
 - neighbor
 - Enable, 4-54
 - Neighbor Address, 4-53
 - Priority, 4-54
 - virtual interface
 - Dead Interval, 4-62
 - Enable, 4-59
 - Hello Interval, 4-61
 - Neighbors Router ID, 4-58
 - Password, 4-63
 - Retransmit Interval, 4-60
 - Transit Area ID, 4-57
 - Transit Delay, 4-59

P

PDN X.25 address resolution, 2-10

peers

adding for BGP, 5-35

point-to-multipoint network, 4-3

point-to-point network, 4-3

policies, definition of, 1-16

policy parameters

Action (accept), 9-9

Action (announce), 9-28

Advertise (announce), 9-31

Aggregator AS List (accept), 9-22

Aggregator Router List (accept), 9-22

Announce Tag, 9-43

Apply Subnet Mask (accept), 9-12

AS List (Accept), 9-14

AS Path (announce), 9-57

AS Path Override (announce), 9-51

AS Weight Class (accept), 9-18, 9-24

BGP-3 Route Preference (accept), 9-18

BGP-4 Preference (accept), 9-23

EGP Interface List (announce), 9-46

EGP Metric (announce), 9-47

EGP Peer List (announce), 9-46

Enable (accept), 9-8

Enable (announce), 9-27

External Route Source (announce), 9-40,
9-45, 9-48, 9-53

From BGP Peer (announce), 9-37

From BGP Peer AS (announce), 9-38

From EGP Peer (announce), 9-36

From Gateway (accept), 9-11

From OSPF Router ID (announce), 9-33

From RIP Gateway (announce), 9-32

Gateway List (accept), 9-14

Injection List (accept), 9-15, 9-19

Inter-AS Metric Selector (announce),
9-50

Local Preference (accept), 9-23

Local Preference Override (announce),
9-57

Local Preference Value (announce), 9-58

Multi-Exit Discriminator (announce),
9-55

Name (accept), 9-8

Name (announce), 9-27

Networks (accept), 9-9

Networks (announce), 9-28

Origin (announce), 9-51, 9-56

Originating AS (accept), 9-17, 9-21

OSPF Metric (announce), 9-44

Outbound Interface (announce), 9-41

Outbound Peer AS (announce), 9-54

Outbound Peer AS List (announce), 9-49

Outbound Peers (announce), 9-49, 9-55

Peer Address (accept), 9-16, 9-20

Peer AS (accept), 9-16, 9-20

Peer List (accept), 9-13

Precedence (announce), 9-29

Received BGP Next Hop (announce),
9-39

Received EGP Gateway (announce), 9-37

Received on Interface (accept), 9-11

Received on RIP Interface (announce),
9-32

Received OSPF Tag (announce), 9-35

Received OSPF Type (announce), 9-34

Route Origin (accept), 9-17, 9-21

Route Preference (accept), 9-10

Rule Precedence (accept), 9-10

Specific Inter-AS Metric (announce),
9-50

Tag (accept), 9-13

Type (accept), 9-12

Type (announce), 9-42

preference, definition of, 1-14

Proxy ARP, 2-8

R

revised IP security option

RIPSO, 2-15

See RIPSO

RIP

accept policy parameters, 9-11

announce policy parameters, 9-40

export route filters

adding, 10-11

configuring, 10-11 to 10-17

deleting, 10-17

editing, 10-17

import route filters

adding, 10-3

configuring, 10-2 to 10-10

deleting, 10-10

editing, 10-10

overview, 3-1

See Routing Information Protocol

RIP parameters

editing, 3-2 to 3-7

export route filters

Action, 10-16

Enable, 10-15

Export Address, 10-12

Export Mask, 10-14

From Protocol, 10-15

Interface, 10-15

Rip Metric, 10-16

import route filters

Action, 10-8

Enable, 10-8

Import Address, 10-5

Import Mask, 10-6

Interface, 10-7

Preference, 10-9

RIP Gateway, 10-7

interface

Default Route Listen, 3-6

Default Route Supply, 3-6

Enable, 3-4

RIP Listen, 3-5

RIP Supply, 3-4

RIP interface

Poisoned Reverse, 3-7

RIPSO, 2-15

configuring support for, 2-74

example of, 2-19

how it works on the router, 2-17

network example, 2-20

Route Cache Interval, 2-32

Router Discovery

definition of, 2-22

parameters

Broadcast Type, 2-87

Enable, 2-87

Interface Pref, 2-89

Lifetime, 2-89

Maximum Interval, 2-88

Minimum Interval, 2-88

Routing Information Protocol

See RIP

S

security label format, 2-15, 2-16

static black hole routes

configuring, 2-60, 2-63

definition of, 2-21

static routes

adding, 2-59

configuring, 2-57 to 2-64

definition of, 2-21

deleting, 2-64

editing, 2-61

subnet mask

function of, 1-6

specifying, 1-7

subnets, definition of, 1-6

supernets, definition of, 1-9

T

TFTP (Trivial File Transfer Protocol), 2-13

token ring networks, 2-11

Trivial File Transfer Protocol

function of, 2-13

V

virtual interface

adding, 4-56 to 4-58

W

weight parameters

BGP, 5-47, 5-55

weight value

adding to an AS, 5-46

deleting from an AS

BGP, 5-52

weight, definition of, 1-15





Bay Networks

The Merged Company of SynOptics and Wellfleet

8 Federal Street
Billerica, MA 01821



Printed in U.S.A. on Recycled Paper