PortMaster®

Routing Guide

Lucent Technologies
Remote Access Business Unit
4464 Willow Road
Pleasanton, CA 94588
510-737-2100
800-458-9966

December 1997

950-1191A

Copyright and Trademarks

© 1997 Lucent Technologies. All rights reserved.

PortMaster, ComOS, and ChoiceNet are registered trademarks of Lucent Technologies. PMconsole, IRX, True Digital, RAMP, and Total Access. Sure and Simple. are trademarks of Lucent Technologies. ProVision is a service mark of Lucent Technologies. All other marks are the property of their respective owners.

Disclaimer

Lucent Technologies makes no express or implied representations or warranties with respect to the contents or use of this manual, and specifically disclaims any implied warranties of merchantability or fitness for a particular purpose. Lucent Technologies further reserves the right to revise this manual and to make changes to its content at any time, without obligation to notify any person or entity of such revisions or changes.

Contents

About This Guide

This *PortMaster Routing Guide* is designed to aid PortMaster configuration in a routing environment.

This guide provides practical routing guidance for the standard routing protocols in use today. After an introductory chapter summarizing routing concepts, specific protocol chapters discuss implementation of those protocols on PortMaster products, including example configurations for each. A troubleshooting appendix concludes the guide with a series of routing case studies from the Lucent Technical Support team, illustrating common routing configuration problems and their solutions.

Routing protocols supported by the PortMaster ComOS® software include the following:

- Routing Information Protocol (RIPv1)
- Open Shortest Path First (OSPF)
- Border Gateway Protocol (BGP-4)

Support for the routing protocols described is provided by ComOS releases as follows:

Protocol	**Support in ComOS Version**
RIP	All
OSPF	Release 3.5 and later
BGP-4	Release 3.5 and later

Although basic routing concepts and protocols applicable to PortMaster products are described in this guide, this material is not a substitute for a routing textbook, and it does not duplicate routing protocol information contained in the Internet standards documents. The Internet standards are described in Requests for Comments (RFCs). (See "RFCs" on page ix.) To better understand the protocols and their use, Lucent recommends that you read the applicable RFCs. Several books that provide additional background information on TCP/IP and routing issues are listed on page x.

Audience

This guide is designed to be used by qualified system administrators and network managers who need to know how to configure PortMaster products using the Lucent implementation of RIP, OSPF, and BGP-4. Users of this guide should have the following:

- Familiarity with basic PortMaster configuration and ComOS commands
- Understanding of networking concepts and the suite of TCP/IP data communications protocols
- Knowledge of IP addressing, subnetting, and classless interdomain routing (CIDR)
- Knowledge of the Internet Packet Exchange (IPX) protocol and the Service Advertising Protocol (SAP)—required only if you are routing Novell IPX datagrams
- Basic understanding of networking applications and TCP/IP configuration in a PC environment

PortMaster Documentation

The following manuals are available from Lucent Technologies. The hardware installation guides are included with most PortMaster products; other manuals can be ordered through your PortMaster distributor or directly from Lucent.

The manuals are also provided as PDF and PostScript files on the *PortMaster Software CD* shipped with your PortMaster.

In addition, you can download PortMaster information and documentation from **http://www.livingston.com**.

- *ChoiceNet® Administrator's Guide*

  This guide provides complete installation and configuration instructions for ChoiceNet server software.

- *Command Line Administrator's Guide*

  This guide provides the complete description and syntax of each command in the ComOS command set.

- *Configuration Guide for PortMaster Products*

  This guide provides a comprehensive overview of networking and configuration issues related to PortMaster products.

- PortMaster hardware installation guides

  These guides contain complete hardware installation instructions. An installation guide is available for each PortMaster product line—IRX™, Office Router, Communications Server, and Integrated Access Server.

- *PMconsole™ for Windows Administrator's Guide*

  This guide covers PMconsole Administration Software for Microsoft Windows, a graphical tool for configuring the PortMaster. The majority of the material in this guide also applies to the UNIX version of PMconsole.

- *PortMaster Routing Guide*

  This guide describes routing protocols supported by PortMaster products, and how to use them for a wide range of routing applications.

- *RADIUS Administrator's Guide*

  This guide provides complete installation and configuration instructions for Lucent Remote Authentication Dial-In User Service (RADIUS) software.

Additional References

RFCs

Use any World Wide Web browser to find a Request for Comments (RFC) online.

RFC 768, *User Datagram Protocol*
RFC 791, *Internet Protocol*
RFC 792, *Internet Control Message Protocol*
RFC 793, *Transmission Control Protocol*
RFC 854, *Telnet Protocol Specification*
RFC 950, *Internet Standard Subnetting Procedure*
RFC 988, *Host Extensions for IP Multicasting*
RFC 1058, *Routing Information Protocol*
RFC 1144, *Compressing TCP/IP Headers for Low-Speed Serial Links*
RFC 1157, *A Simple Network Management Protocol (SNMP)*

RFC 1166, *Internet Numbers*
RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*
RFC 1256, *ICMP Router Discovery Messages*
RFC 1321, *The MD5 Message-Digest Algorithm*
RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*
RFC 1334, *PPP Authentication Protocols*
RFC 1362, *Novell IPX Over Various WAN Media (IPXWAN)*
RFC 1413, *Identification Protocol*
RFC 1490, *Multiprotocol Interconnect Over Frame Relay*
RFC 1583, *OSPF Version 2*
RFC 1587, *OSPF NSSA Options*
RFC 1597, *Address Allocations for Private Internets*
RFC 1627, *Network 10 Considered Harmful (Some Practices Shouldn't be Codified)*
RFC 1634, *Novell IPX Over Various WAN Media (IPXWAN)*
RFC 1661, *The Point-to-Point Protocol (PPP)*
RFC 1700, *Assigned Numbers*
RFC 1717, *The PPP Multilink Protocol (MP)*
RFC 1771, *A Border Gateway Protocol 4 (BGP-4)*
RFC 1812, *Requirements for IP Version 4 Routers*
RFC 1814, *Unique Addresses are Good*
RFC 1818, *Best Current Practices*
RFC 1824, Requirements for IP Version 4 Routers
RFC 1826, *IP Authentication Header*
RFC 1827, *IP Encapsulating Payload*
RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*
RFC 1878, *Variable Length Subnet Table for IPv4*
RFC 1918, *Address Allocation for Private Internets*
RFC 1965, *Autonomous System Confederations for BGP*
RFC 1966, *BGP Route Reflection, An Alternative to Full Mesh IBGP*
RFC 1997, *BGP Communities Attribute*
RFC 2003, *IP Encapsulating Security Payload*
RFC 2138, *Remote Authentication Dial In User Service (RADIUS)*
RFC 2139, *RADIUS Accounting*

Books

Internetworking with TCP/IP, Volume 1: Principles, Protocols, and Architecture. Douglas Comer. Prentice Hall, Inc. 1995. (ISBN 0-13-216987-8 (v.1))

Routing in the Internet. Christian Huitema. Prentice Hall PTR, 1995. (ISBN 0-13-132192-7)

TCP/IP Network Administration. Craig Hunt. O'Reilly & Associates, Inc. 1994. (ISBN 0-937175-82-X)

TCP/IP Illustrated, Volume 1: The Protocols. W. Richard Stevens. Addison-Wesley Publishing Company. 1994. (ISBN 0-201-63346-9)

Internet Routing Architectures. Bassam Halabi. Cisco Press, 1997.

Document Conventions

The following conventions are used in this guide:

Convention	Use	Examples
Bold font	Indicates a user entry—a command, menu option, button, or key—or the name of a file, directory, or utility, except in code samples.	• Enter **version** to display the version number. • Press **Enter**. • Open the **permit_list** file.
Italic font	Identifies a command-line placeholder. Replace with a real name or value.	• **set** *Ether0* **address** *Ipaddress* • Replace *Area* with the name of the OSPF area.
Square brackets ([ ])	Enclose optional keywords and values in command syntax.	• **set nameserver** [**2**] *Ipaddress* • **set** *S0* **destination** *Ipaddress* [*Ipmask*]

Convention	Use	Examples
Curly braces ({ })	Enclose a required choice between keywords and/or values in command syntax.	**set syslog** *Logtype* {[**disabled**] [*Facility.Priority*]}
Vertical bar (|)	Separates two or more possible options in command syntax.	• **set** *S0*|*W1* **ospf on**|**off** • **set** *S0* **host default**|**prompt**|*Ipaddress*

Document Advisories

Note – means take note. Notes contain information of importance or special interest.

Caution – means be careful. You might do something—or fail to do something—that results in equipment failure or loss of data.

Warning – means danger. You might do something—or fail to do something—that results in personal injury or equipment damage.

Contacting Lucent Technical Support

The PortMaster comes with a 1-year hardware warranty.

For all technical support requests, record your PortMaster ComOS version number and report it to the technical support staff.

New releases and upgrades of PortMaster software are available by anonymous FTP from **ftp://ftp.livingston.com.pub/le/**.

You can schedule a 1-hour software installation appointment in advance by calling the technical support telephone number listed below. Appointments must be scheduled at least one business day in advance.

For Europe, Middle East, and Africa (EMEA)

Contact the Remote Access EMEA Support Center Monday through Friday between the hours of 8 a.m. and 8 p.m. (GMT+1), excluding French public holidays.

- By voice, dial +33-4-92-92-48-88.
- By fax, dial +33-4-92-92-48-40.
- By electronic mail (email) send mail to **emeasupport@livingston-ent.fr**

For North America, Latin America, and the Asia Pacific Region

Contact Lucent, Monday through Friday between the hours of 6 a.m. and 5 p.m. (GMT –8).

- By voice, dial 800-458-9966 within the United States (including Alaska and Hawaii), Canada, and the Caribbean, or +1-510-737-2100 from elsewhere.
- By fax, dial +1-510-737-2110.
- By email, send mail as follows:
  - From North America and Latin America to **support@livingston.com**.
  - From the Asia Pacific Region to **asia-support@livingston.com**.
- Using the World Wide Web, see **http://www.livingston.com/**.

Subscribing to PortMaster Mailing Lists

Lucent maintains the following Internet mailing lists for PortMaster users:

- **portmaster-users**—a discussion of general and specific PortMaster issues, including configuration and troubleshooting suggestions. To subscribe, send email to **majordomo@livingston.com** with **subscribe portmaster-users** in the body of the message.

The mailing list is also available in a daily digest format. To receive the digest, send email to **majordomo@livingston.com** with **subscribe portmaster-users-digest** in the body of the message.

- **portmaster-radius**—a discussion of general and specific RADIUS issues, including configuration and troubleshooting suggestions. To subscribe, send email to **majordomo@livingston.com** with **subscribe portmaster-radius** in the body of the message.

  The mailing list is also available in a daily digest format. To receive the digest, send email to **majordomo@livingston.com** with **subscribe portmaster-radius-digest** in the body of the message.

- **portmaster-announce**—announcements of new PortMaster products and software releases. To subscribe, send email to **majordomo@livingston.com** with **subscribe portmaster-announce** in the body of the message. All announcements to this list also go to the portmaster-users list. You do not need to subscribe to both lists.

Routing Overview 1

This chapter provides a brief overview of routing concepts and describes how the Lucent ComOS implements routing protocols. The chapter is intended as a quick refresher for those already familiar with network routing, and as an introduction for those with little or no knowledge of the standard protocols. See "About This Guide" for pertinent RFCs and some useful books.

You should understand this material before you configure your PortMaster. Refer also to the glossary for unfamiliar terms.

This chapter discusses the following topics:

Understanding Routing

Routing is the process of finding the route to a destination, and routing protocols determine how a router updates its route information. A router is attached to two or more networks, and its primary function is receiving IP packets through one network interface and forwarding them through another. The packets can travel through a number of routers before arriving at their final destination. Each router-to-router transmission is called a hop.

Hosts

When a host sends an IP data packet to a destination host that is on the same network or subnet as the sending host, the packet goes directly to the destination host. If the destination host and sending host are on different networks or subnets, the packet goes

to a router. The destination IP address contained in a data packet is the basic information that makes routing and delivery possible. The major differences between typical hosts and routers are

- Hosts are usually connected to a single network. Routers are connected to two or more networks.
- Hosts have very little routing information, frequently limited to a route for their local network and a default route to a router attached to their local network. Routers can have detailed route information for many routes in their routing tables.

Routes

A router constructs a routing table containing a list of routes. Each route indicates the next place to which the router forwards a packet, to get it closer to a final destination. A route consists of the following information:

- Destination prefix
- Netmask length
- Gateway to the destination
- Metric that indicates how desirable this gateway is for reaching the destination

When a router receives a packet that needs to be forwarded, the router compares the destination address in the packet with addresses in the routing table to find a route for the specified destination. After determining a route, the router forwards the packet to the associated gateway listed in the routing table.

Address Resolution

The IP address and routing table are used to route a data packet to a specific host on a network. In the case of Ethernet, an IP address must be translated to the physical, or media access control (MAC) address of the target host. This translation, or *resolving* of an IP address to a MAC address is carried out by the Address Resolution Protocol (ARP). The Reverse Address Resolution Protocol (RARP) translates in the opposite direction—that is, from a MAC address to an IP address.

In practice, a host broadcasts an ARP request over the Ethernet and listens for a reply. Other hosts on the Ethernet receive the request, and when one of them recognizes its own IP address, it replies with its MAC address. The requester then adds the IP

address and MAC translation to its ARP cache for future use. In the case of a PortMaster, ARP cache entries expire after 15 minutes. Near the end of a 15-minute interval, the PortMaster sends another ARP request to refresh and maintain up-to-date cache entries.

ARP replies do not always have to be sent from the target host but can also be sent by a third party—such as a router—that knows the route information and acts as a proxy for the target. This process—known as proxy ARP—is supported by PortMaster products.

Static Routing

Static routes contain route information that does not respond to changes in network topology or condition. In spite of this inability to respond to change, static routes can be useful for stable networks or segments of networks that are not expected to change frequently.

Other possible uses for static routing are

- For operations that use or connect to routers that do not support the routing protocols supported by the Lucent ComOS
- For operations over WAN links where costs are determined by connect time or number of packets sent, and the use of routing protocols adds to the cost

Static routing tables are constructed manually on a PortMaster by the network administrator, who must update them periodically to reflect any network changes.

Default Gateway

If a network or a segment of a network has only one exit point, you need set only a default gateway to ensure that all packets bound for remote destinations get routed onto the Internet. A default gateway is also useful when you want to set up a gateway for packets that have no alternative route specified in the routing table.

When you specify a default gateway in a router configuration, all packets to destinations not found in the routing table are forwarded to that default gateway.

Dynamic Routing

If a network has more than one possible route to a destination, or if the network is large or complex, or experiences frequent changes, dynamic or learned routing through the use of routing protocols is necessary. Routing protocols collect routing information, broadcast it dynamically, and update routing tables automatically to reflect network changes.

Dynamic routing adds a high degree of flexibility to the routing process by

- Permitting the selection of routes based on cost or congestion criteria
- Allowing redundant routes to compensate for downed lines
- Selecting the best route to a destination when multiple routes exist
- Providing fast response to network changes

Variable-Length Subnet Masks

The commonly used Routing Information Protocol (RIP) requires that a single netmask be used across the entire subnet. This limitation has a major impact on the routing of traffic among many networks. Regardless of the actual need for addresses, each subnet of the network must be assigned the same number of IP addresses.

The use of more than one subnet mask in a subnetted network is described in RFC 1812. A network using this method of subnetting is referred to as a network with variable-length subnet masks (VLSMs) because the network prefixes have different lengths. VLSM support allows more efficient use of TCP/IP addresses by allowing network designers greater freedom in assigning addresses.

On PortMaster products, VLSMs can be used for static routing without a routing protocol, or with a routing protocol that provides VLSM support, such as OSPF and BGP-4. VLSMs overcome the RIP limitation of having to use a single netmask across the network. RIPv1 does not send netmasks in its routing updates.

Advantages of VLSM

The advantages of VLSM over RIP are illustrated in Figure 1-1 and Figure 1-2.

Figure 1-1 and Figure 1-2 both show a network with the following topology:

- Router A is on two subnets, 192.168.1.*x* and 192.168.2.*x*.
- Router B, at address 192.168.3.65, is on a subnet that includes 192.168.3.64 through 192.168.3.127.
- Router C, at address 192.168.3.161, is on a subnet that includes 192.168.3.160 through 192.168.3.191.

Routing without VLSMs

In Figure 1-1, Router A exchanges RIP updates with Routers B and C. Because they are using RIP, Router A is unable to distinguish between the networks of Routers B and C. Router A sees both Router B and Router C as part of network 192.168.3.*x*, where *x* ranges from 0 through 255.

Figure 1-1 RIP Routing—without VLSMs

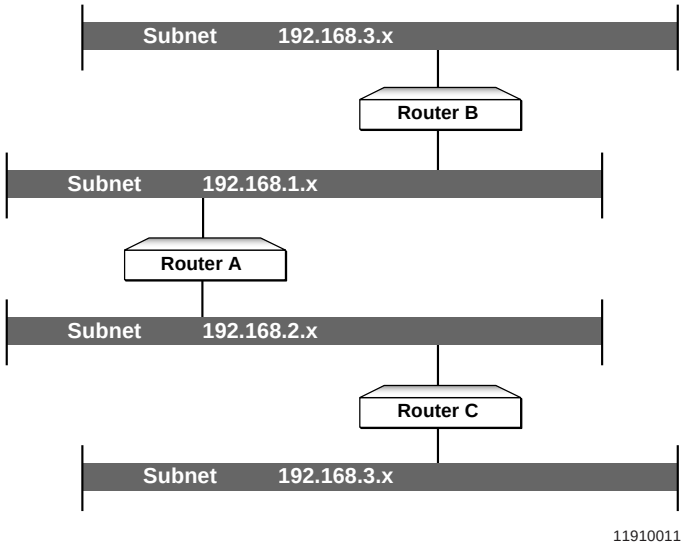

When data is directed from Router A to a 192.168.3.*x* address, Router A sends the packet to either Router B or Router C, depending on which router last provided a RIP update—making that last router appear to be the gateway for all network 192.168.3.*x* addresses. Hence, RIP cannot provide reliable routing in this kind of network.

Routing with VLSMs

Variable length subnet masking provides for more reliable routing. In the Open Shortest Path First (OSPF) network shown in Figure 1-2, Router A has both the IP address and the netmask information required to identify the unique set of addresses associated with the networks of Routers B and C. Router B's OSPF update states that its address is 192.168.3.65 and its netmask is 255.255.255.192. In other words, the subnet associated with Router B is 192.168.3.64/26, where 192.168.3.64 is the network prefix and the /26 represents the number of high-order bits in the netmask.

Figure 1-2 OSPF Routing—with VLSMs

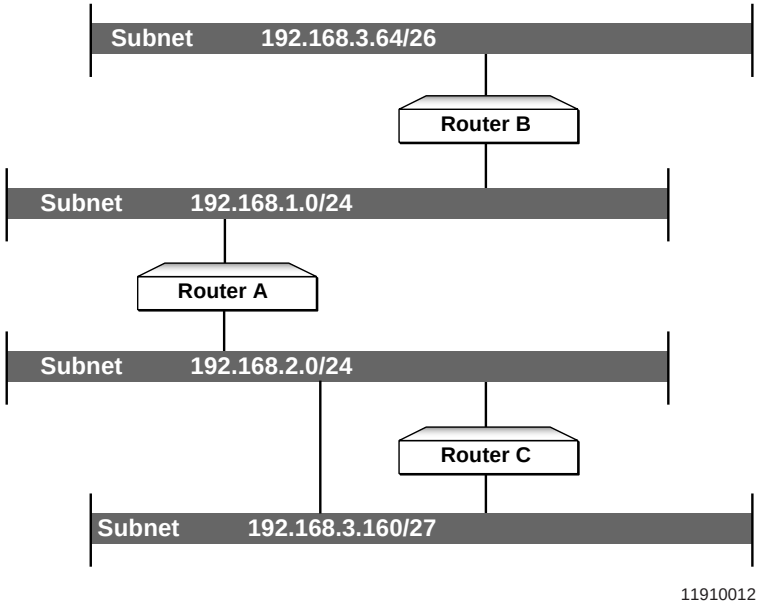

With this information, Router A can forward to Router B any IP traffic sent to addresses 192.168.3.64 through 192.168.3.127. Router C's address, 192.168.3.161, and network mask, 255.255.255.224—192.168.3.160/27—ensure that Router A can properly forward all traffic sent to addresses 192.168.3.160 through 192.168.3.191.

CIDR

Variable-length subnet masking uses classless interdomain routing (CIDR) addressing. Table 1-1 lists the variable-length netmasks from 16 to 32, the CIDR representation form (/*xx*), and the hexadecimal and decimal equivalents. Refer to RFC 1878 for more information.

Table 1-1 CIDR Netmask Table

Mask Value			
CIDR	**Hexadecimal**	**Dotted Decimal**	**Number of Hosts**
/16	0xffff0000	255.255.0.0	65,534
/17	0xffff8000	255.255.128.0	32,766
/18	0xffffc000	255.255.192.0	16,382
/19	0xffffe000	255.255.224.0	8190
/20	0xfffff000	255.255.240.0	4094
/21	0xfffff800	255.255.248.0	2046
/22	0xfffffc00	255.255.252.0	1022
/23	0xfffffe00	255.255.254.0	510
/24	0xffffff00	255.255.255.0	254
/25	0xffffff80	255.255.255.128	126
/26	0xffffffc0	255.255.255.192	62
/27	0xffffffe0	255.255.255.224	30
/28	0xfffffff0	255.255.255.240	14
/29	0xfffffff8	255.255.255.248	6
/30	0xfffffffc	255.255.255.252	2
/32	0xffffffff	255.255.255.255	This is a single-host route.

Routing with a PortMaster

A PortMaster router gathers and maintains information in a routing table that enables the transmission of packets between networks. The routing table contains an entry for each identified route, and can be configured statically, maintained by the router dynamically through a routing protocol, or both.

PortMaster products support host routes and network routes. Variable-length subnet masks (VLSMs) and classless interdomain routing (CIDR) are also supported in ComOS release 3.5 and later, as is the more common use of a single subnet mask per network.

A PortMaster determines whether a packet is local or needs to be routed by examining the destination IP address of packets, and the IP address and netmask of the PortMaster interfaces. For example, if the PortMaster Ethernet address is 192.168.1.1 and the netmask is 255.255.255.0, then all hosts with addresses between 192.168.1.1 and 192.168.1.254 are deemed local hosts, and packets for these hosts do not require routing.

If a packet requires routing, the routing table is consulted, with one of the following results:

- A route is found and the packet is sent to the gateway specified in the route.
- No route is found, the packet is discarded, and an Internet Control Message Protocol (ICMP) Type 3 message is sent to the packet source.

Routing can occur over Ethernet, synchronous, and asynchronous lines. Routing is set on Ethernet interfaces, hardwired network ports, in the user table or RADIUS for dial-in users, and in the location table for dial-out locations.

Routing Table

The PortMaster maintains a routing table, and each entry includes some or all of the following information, depending on the routing protocol used:

- Routing protocol used
- Destination prefix
- Netmask length
- Gateway

- Source of routing data
- Route status flags
- Metrics
- Interface

The longest prefixes in a routing table entry are checked first. If no route is found, the learned gateway (if one exists) is used; otherwise the primary gateway is used. If no route is found, the packet is discarded.

Static routes are never overwritten with learned, or "dynamic," routes. However, if a default gateway with a better metric than the primary gateway is learned and default routes are being listened for, the new gateway is set as the learned gateway. In addition, the learned gateway is replaced if a default gateway with a lower metric is found.

Routes are valid only for active interfaces. If an interface becomes unreachable, all routes that go through the interface are examined. If the specified gateway is available through another interface, the route is moved to the active interface. Otherwise, the routing table entry is marked as obsolete.

Table 1-2 shows an example of a PortMaster routing table.

Table 1-2 PortMaster Routing Table

Destination	Mask	Gateway	Source	Flag	Met	Interface
0.0.0.0	0	198.168.96.2	local	NS	1	ether0
198.168.7.2	32	198.168.96.2	ospf	HD	3	ether0
172.27.195.131	32	198.168.96.2	ospf/E2	HD	4	ether0
198.168.1.9	32	198.168.96.2	ospf	HD	4	ether0
198.168.97.42	32	198.168.96.229	ospf	HD	2	ether0
10.0.0.0	32	198.168.96.2	ospf/E2	HD	3	ether0
198.168.1.245	32	198.168.96.2	ospf	HD	3	ether0

Table 1-3 describes the routing table flags.

Table 1-3 Routing Table Flags

Flag	Description
H	Host route.
N	Network route.
L	Local route—a network or destination is directly attached. If this is a point-to-point link, the PortMaster can reach the network or destination by sending the packet down the link; otherwise, the PortMaster can reach hosts on this network with an ARP request.
S	Static route, either configured and permanent, or temporary via RADIUS.
D	Route learned via a routing protocol.
C	Route that changed recently but is not yet propagated to all interfaces.
O	Obsolete; marked for deletion.

Routing Protocols

Routing protocols supported by the Lucent ComOS include the following:

- Routing Information Protocol Version 1 (RIPv1)
- Open Shortest Path First (OSPF)
- Border Gateway Protocol (BGP-4)

Comparing Routing Protocols

The selection of appropriate routing procedures and protocols for specific network applications depends primarily on the following:

- Size and complexity of the network
- Routing features desired
- Support for routing protocols by other routers in the network

- Availability of technical expertise for network design
- Need for interior and exterior routing protocols

Table 1-4 shows a comparison of significant advantages and limitations of the interior routing protocols currently supported on PortMaster products.

Table 1-5 on page 1-13 shows the major features of the BGP-4 exterior routing protocol.

Table 1-4 Interior Routing Protocols Compared

Protocol	Advantages	Limitations
RIPv1	• Very simple to install and maintain. • Low overhead on system bandwidth in small networks. • In widespread use, supported by all routers. • Suitable for small stable networks.	• Insufficient for large networks. • High overhead on system bandwidth in large networks. • Slow response to network changes. • Loops can occur and cause congestion. • Autonomous systems not supported. • VLSMs not supported. • CIDR not supported. • Multiple paths to a destination not supported. • No support for propagating netmask information with routes.

Table 1-4 Interior Routing Protocols Compared *(Continued)*

Protocol	Advantages	Limitations
OSPF	• Designed for interior routing within autonomous systems. • Responds quickly to network changes. • Moderate overhead on system bandwidth in large networks. • Suitable for small to large dynamic networks. • Fast, loopless convergence of route information. • VLSMs supported. • CIDR supported. • Supports multiple paths to a destination. • Supports message digest algorithm 5 (MD5) authentication. • Well-designed interface with BGP.	• Can be complex to install and maintain. • Network convergence delays can occur in very large networks. • Recommended maximum of 50 routers per OSPF area. • Recommended maximum of 60 neighbors per router. • Not supported by all routers.

Table 1-5 Features of the BGP-4 Exterior Routing Protocol

Protocol	Advantages	Limitations
BGP-4	• Designed for exterior routing between autonomous systems. • For established neighbors, only updates to routing tables are exchanged. • Moderate overhead on system bandwidth in large networks. • Can select a route from competing routes. • Permits route filtering. • Can detect loops at the autonomous system level. • Supports path aggregation to simplify route advertisements. • Supports CIDR.	• Complex to install and maintain. • Not suitable for routing within a single routing domain. • Not supported by all routers. • When peers become neighbors, they exchange complete routing tables, which can be large. • Uses a lot of memory.

RIP

The Routing Information Protocol (RIP) is an interior routing protocol that uses a distance-vector algorithm to optimize routing. It is widely used by routers in TCP/IP networks and Novell IPX networks.

PortMaster RIP updates for each interface can be configured to broadcast routing updates, listen for routing updates, both, or neither. RIP routes are updated in the following ways:

- Broadcast: The PortMaster sends RIP information to the interface every 30 seconds.
- Listen: The PortMaster listens for RIP information on the interface from other routers.

When RIP routing is enabled, a PortMaster broadcasts any default route and static route information as part of a normal RIP message.

The distance-vector algorithm used by RIP for making routing decisions is based on the number of hops to a destination. The route with the least number of hops is considered to have the lowest cost. The number of hops is referred to as the route metric; a route with a metric greater than 15 is considered to be unreachable.

Routing table entries for RIP are aged as follows:

- Updates are expected every 30 seconds.
- After 120 seconds, if information for another route to a given destination is observed the new route is used for the destination.
- After 180 seconds without an update, a route is considered unreachable. RIP makes the route obsolete by setting the route metric to 16.
- The routing timer is suspended if an on-demand interface is suspended.
- Routes are valid only for active interfaces.
- If an interface becomes unreachable, all routes that go through that interface are examined. If the specified gateway is available through another interface, the route is moved to the active interface. Otherwise, the routing table entry is marked as obsolete.

OSPF

The Open Shortest Path First (OSPF) routing protocol is designed for a hierarchical network structure and is generally organized in a form similar to that shown in Figure 1-3.

The community of routers, usually owned and managed by a single group as illustrated in Figure 1-3, is called an autonomous system (AS). OSPF divides the world into two domains—within the autonomous system and external to the autonomous system. The autonomous system itself is divided into a central, or backbone area (area 0) to which all other areas are connected, and other (nonzero) numbered areas

Figure 1-3 OSPF Network

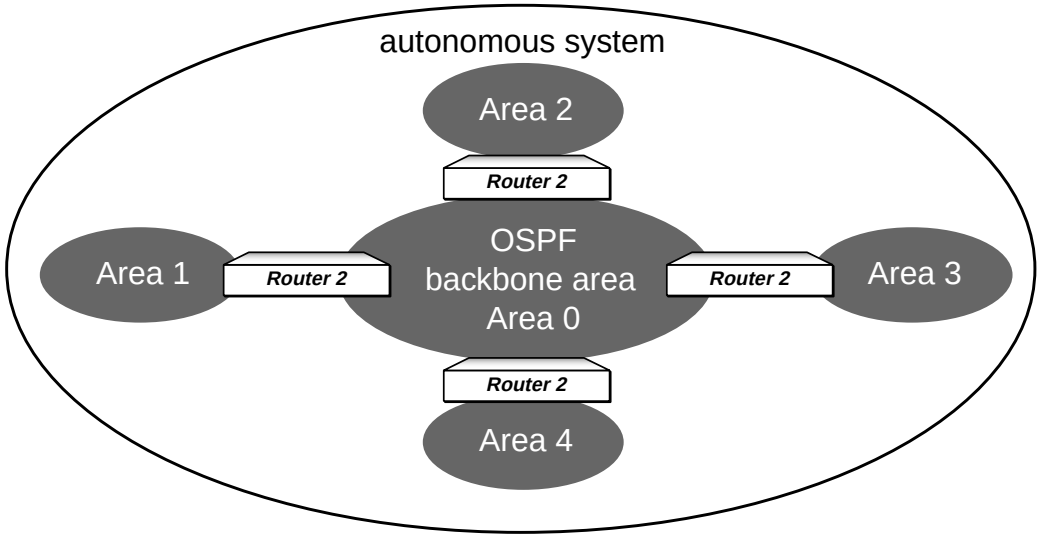

Areas

Each OSPF area has one or more network address/netmask pairs and is identified by a 32-bit area number. An area is a contiguous set of routers sharing network segments between them. An area must be one of the following types, depending upon the routing information that it uses:

- Transit area—An area where routing information for a default route, static routes, intra-area routes, interarea routes, and autonomous system external routes is kept in the area database. The backbone area falls into this category. Route types are described in "OSPF Routes" on page 1-17.
- Stub area—An area where all autonomous system external routes are summarized by a per-area default route in the area database, which also contains intra-area and interarea routes. Summarization reduces the database size and memory requirements for a stub area's internal routers. A stub area frequently has only one area border router (ABR)—described in "Routers" on page 1-16.
- Not-so-stubby area (NSSA)—Similar to a stub area, but it can export limited autonomous system external route information to the backbone. The NSSA is useful for areas that do not need to retain all the autonomous system external route

information flooded into the area, but do need to export some limited external route information not permitted from standard OSPF stub areas. For more information, see RFC 1587.

OSPF Network Relationships

The following terms are used to describe relationships in an OSPF network:

- Neighbors—A set of routers with which a single router can communicate.
- Adjacencies—Neighbors that exchange link state advertisement (LSA) databases are defined as adjacent. Neighbors can communicate over a **multicast** address to enable an OSPF packet to reach many addresses, or over statically configured **unicast** addresses to enable an OSPF packet to reach only one address.
- Hellos—Neighbors exchange hello packets to find each other, and each router's hello packets list the routers it has heard from. When two routers can identify themselves in their neighbor's hello packets, they can communicate with each other. When communication occurs, the following takes place:
  - If the neighbors are active, they can form an adjacency.
  - If an adjacency is formed, then the LSA databases are synchronized.

Routers

Routers in an autonomous system belong to one or more of the following categories:

- Internal router (IR): All router interfaces belong to the same area.
- Area border router (ABR): A router attached to multiple areas.
- Backbone router (BBR): An internal router or area border router that is part of the backbone area.
- Autonomous system border router (ASBR): A router attached to multiple autonomous systems.

On each network interface, a voting process takes place whereby a designated router (DR) is chosen from all the routers. Once a designated router is chosen, it remains the designated router, even when another router comes online that might have been chosen earlier.

A backup designated router (BDR) is also chosen. The backup designated router takes over if the designated router is inoperable for any reason.

The designated router and backup designated router perform the following functions:

- Both the designated router and backup designated router form full adjacencies to all neighbors. Other neighbors do not form adjacencies themselves.
- The designated router alone advertises the network LSA to the area.

When a designated router becomes inoperable, the network node is lost and all routers on the network must regenerate their router LSAs. Although regeneration proceeds quickly, it can cause a temporary routing slowdown.

OSPF Routes

OSPF routes are divided into types described in Table 1-6.

Table 1-6 OSPF Route Types

Type of Route	Properties	Advertised Information
Intra-area (AR)	Stays within a single area.	Throughout the area.
Interarea (IAR)	Stays within the autonomous system, but crosses area border routers with area IDs greater than 0.	Summarized, and flooded throughout the autonomous system by area border routers.
Type 1 external (E1)	Learned outside the autonomous system, with OSPF-like metrics.	Possibly summarized, flooded throughout the autonomous system by autonomous system border routers, with the autonomous system border router itself.
Type 2 external (E2)	Learned outside the autonomous system, with non-OSPF-like metrics.	Same as Type 1 external.

When comparing routes of the same type, the rules given in Table 1-7 show which route is better—has the lowest cost.

Table 1-7 OSPF Route Cost Rules

Type	Rule
Intra-area routes	OSPF chooses the route with the lowest cost.
Interarea routes	OSPF chooses the lowest-cost combination of the following: the route with the lowest cost to the border router plus the interarea route with the lowest cost.
Type 1 external	OSPF chooses the lowest-cost combination of the following: the route with the lowest cost to the autonomous system border router plus the Type 1 external route with the lowest advertised cost.
Type 2 external	OSPF chooses the advertised Type 2 external route with the lowest advertised cost. If a tie occurs, the route with the lowest cost to the autonomous system border router wins.

BGP-4

Note – Currently, the Border Gateway Protocol (BGP) runs only on PortMaster IRX routers and the PortMaster 3. Because a full BGP routing table for the entire Internet requires 7MB of memory, be sure to upgrade your IRX or PortMaster 3 to 16MB of memory.

The concept of a hierarchical network structure was introduced in “OSPF” on page 1-14, along with a description of an autonomous system. In contrast to the OSPF interior routing protocol, which was designed for routing within an autonomous system, the Border Gateway Protocol (BGP) is an exterior routing protocol. BGP version 4, described in RFC 1771, and further defined in version 5 of the BGP-4 Internet Draft RFC of January 1997, was designed for routing between autonomous systems. To use BGP, each autonomous system must have an autonomous system identifier. You obtain autonomous system identifiers from the Internet Network Information Center (InterNIC).

BGP can be thought of as the “glue” that holds the Internet together. Consequently, BGP routing tables can contain tens of thousands of routes. This number is increasingly rapidly as the Internet expands.

The primary purposes of BGP are to allow different autonomous systems to share routes and to connect in redundant ways, controlled by policy and protocol design, so that routing loops are not formed.

Routing with BGP

Figure 1-4 illustrates companies in five autonomous systems that use BGP as an exterior routing protocol.

Figure 1-4 Example of Linked Multiple Autonomous Systems

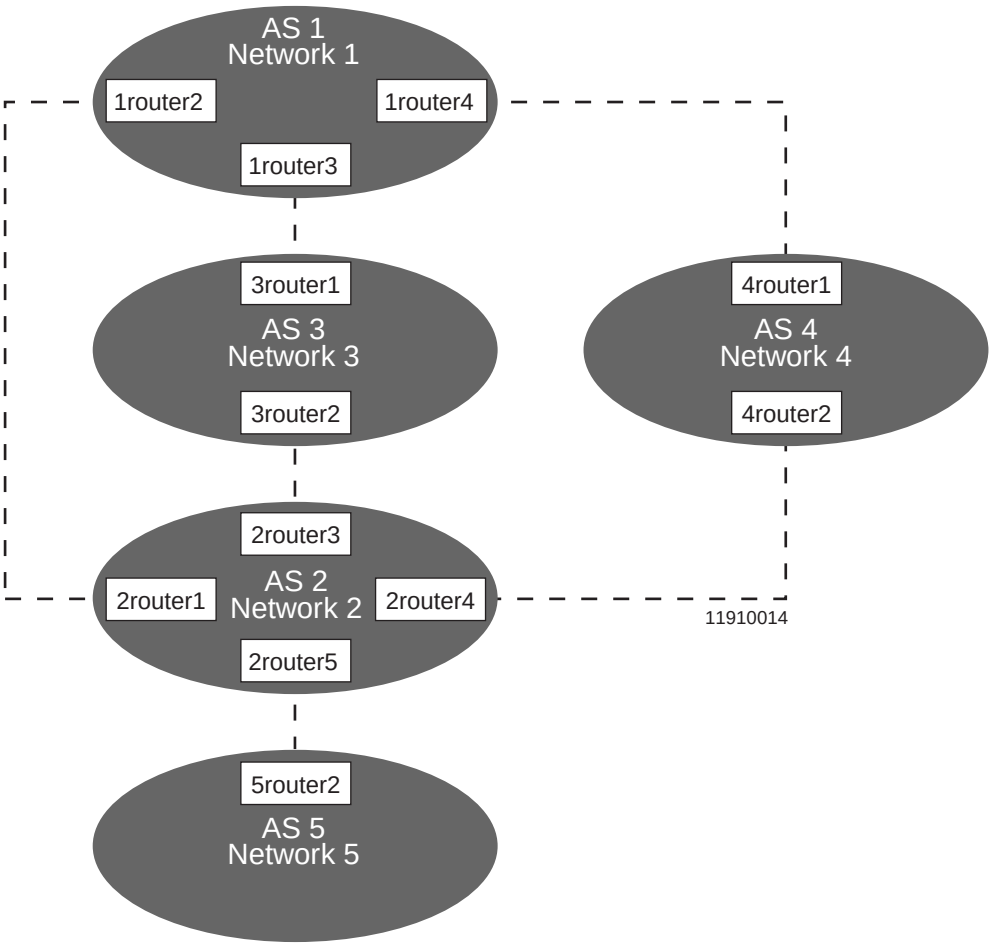

Autonomous system AS 1 advertises Network 1; autonomous system AS 2 advertises Network 2, and so on. A pair of BGP routers is located between each pair of autonomous systems. For example, *1router4* and *4router1* are located between AS1 and AS4.

Policies for Best Path Selection

BGP does not advertise a simple metric representing cumulative link bandwidth costs, as do protocols such as OSPF. Instead, BGP advertises only a path through zero or more autonomous systems, the attributes of the path, and sets of destinations reachable at the end of the path. Sets of destinations are called the network layer reachability information (NLRI). The combined bundle of the path, attributes, and NLRI is called a BGP route. Based on BGP route information, each BGP router executes policy decisions to choose the best route to a final destination.

AS 1 has the following information about paths to Network 5:

- Via *4router1*, along path AS 4 -> AS 2 -> AS 5
- Via *3router1*, along path AS 3 -> AS 2 -> AS 5
- Via *2router1*, along path AS 2 -> AS 5

Although the third path is shortest, it might not be the best path for your purposes. You can create BGP policies to determine the best path according to your own preferences. For example, you can configure policies that implement path preferences such as the following:

- Paths going through the *3router1* peer are preferred over all others.
- Paths going through AS 2 are used only as a last resort.
- Paths going through a single autonomous system are preferred at all costs to other alternatives.

BGP Peers

BGP does not broadcast route information to all listeners as do RIP and OSPF. Instead, each router running BGP must be configured for every other BGP router with which it needs to communicate. Routers that send BGP messages to each other are called BGP speakers, and each pair of BGP speakers that communicate with each other are called peers. Peers are explicitly configured by the PortMaster administrator.

Peers are called internal when they belong to the same autonomous system. When peers belong to other autonomous systems, they are referred to as external peers. BGP treats internal and external peers differently in many details. In particular, unless either route reflection or BGP confederations are configured, all internal peers in a single autonomous system must be fully meshed (directly peered) with each other. For example, if AS 25 has four internal peers (A, B, C, D), then it has six pairs of internal peers (AB, AC, AD, BC, BD, CD). If you are using route reflection or confederations, the routers are partially meshed.

Confederations

BGP requires that all BGP peers within an autonomous system be linked to each other—or fully meshed. As a result, when a BGP peer learns an external route—path attributes and destination—it does not forward this route to the other BGP peers because they already have it. As the number of peers increases in an autonomous system, the number of required links can become large. For example, an autonomous system with 20 peers requires 190 links.

You can reduce the number of BGP peer links by dividing the autonomous system into smaller autonomous systems called confederation member autonomous systems (CMASs). RFC 1965 describes CMASs. If the 20-peer autonomous system is subdivided into a confederation with five CMASs of four peers each, the total number of links is reduced from 190 to 35. This reduction simplifies management of the autonomous system, and reduces message traffic.

Route Reflection

Like BGP autonomous system confederations, route reflection (described in RFC 1966) allows the clustering of peers and reduces the number of links that are otherwise required for a fully meshed autonomous system. Although route reflector clusters are configured differently from CMASs, the functional difference between the two is that route reflectors in each cluster maintain path and attribute information across the entire autonomous system. In this way, the autonomous system still functions like a fully meshed autonomous system.

Route reflection is useful when you want to reduce the traffic and CPU overhead of a fully meshed system. However, confederations allow for policy changes and control across the confederation boundaries within an autonomous system, while route reflection requires the use of identical policies on all internal peers. Therefore, if you want to fine-tune routing within the autonomous system, confederations offer a better solution.

Route Summarization

A BGP speaker can forward to its peers information learned from other peers, as well as originate information into the BGP Internet. BGP originates to its peers only routing information explicitly indicated and supported by the interior routing protocols in use—OSPF, RIP, static routes, or directly attached routes. These special advertisements are called summarizations, and must be explicitly configured.

A summarization is advertised only when an explicit route in that summary is supported through a non-BGP source such as OSPF, RIP, a static route, or a directly attached route—Ethernet, Frame Relay, T1, and so on. The supported route must be more specific than or as specific as the route in the summary. For example, a default route to 0.0.0.0/0 cannot support a summary.

On the PortMaster, all static routes, either configured or learned via RADIUS, can be summarized automatically through propagation. You can define a rule for propagating—translating and advertising—all static routes into BGP.

Rules of Route Precedence

When a PortMaster learns multiple routes to a destination, it uses the following route precedence rules for selecting the preferred path, in order from highest to lowest:

1. Static routes
2. Default routes learned via BGP
3. Routes learned from BGP external peers
4. Routes learned from OSPF and RIP. Routes with the lowest computed hop count are preferred. In the case of a tie, the PortMaster uses the following order of preference, in order from highest to lowest:
    a. OSPF intra-area routes
    b. OSPF interarea routes
    c. OSPF external Type 1 routes
    d. OSPF not-so-stubby-area (NSSA) type 1 routes
    e. OSPF external Type 2 routes

f. OSPF NSSA Type 2 routes

g. RIP routes

5. Routes learned from BGP internal peers

Avoiding Routing Loops

With propagation, you can translate routes from one routing protocol into another—for example, RIP into OSPF. Normally, the route does not include an explicit hop count, and the PortMaster computes a hop count for you. This computed hop count minimizes the chances of routing loops occurring between routing protocol boundaries. However, if you choose to override the automatic hop count computation and specify an explicit hop count, you should consider using propagation rules with route filters to ensure that routing loops are not created.

Figure 1-5 shows four networks connected by four PortMaster IRX routers.

Figure 1-5 Propagation Rule

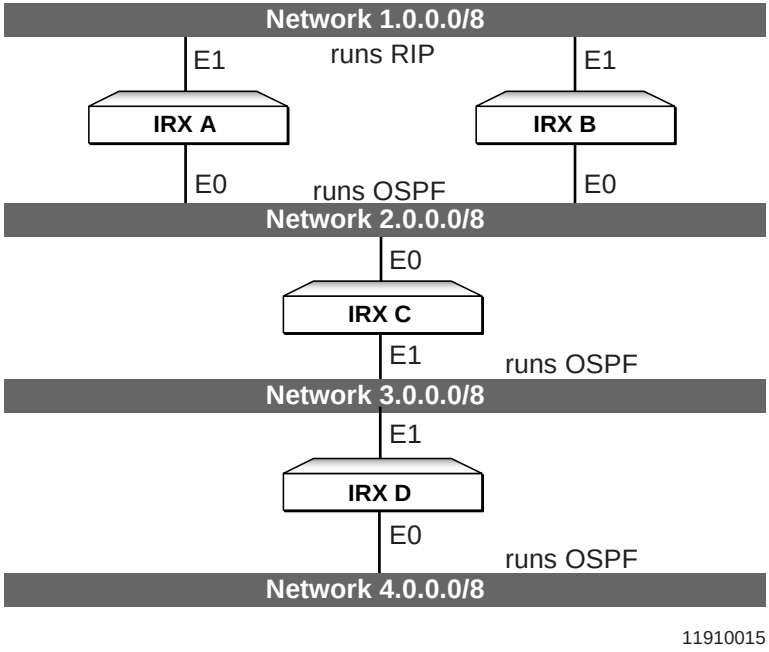

For these networks, assume a propagation rule configured on both IRX A and IRX B to always translate OSPF routes into RIP routes with a hop count of 1. As a result, both IRX A and IRX B advertise network 4.0.0.0 into RIP with a hop count of 1. IRX B has a choice of routes to network 4.0.0.0 as follows:

- The one-hop-count RIP route via IRX A
- The multihop OSPF route via IRX C

The automatic preference rules of the ComOS favor the RIP route for IRX B, and IRX A makes exactly the same choice. As a result, routers IRX A and IRX B point to each other for routes to network 4.0.0.0, thereby creating a routing loop.

This problem is not unique to the way the ComOS determines its route preferences. Mechanisms that translate routing protocols using a fixed metric for all routes, or that allow routes from one protocol—such as OSPF—to be always preferred over another—such as RIP—share the same style of routing loop problems across a routing protocol boundary. These problems are the intrinsic result of using fixed metrics for route advertisement. Consider the example shown in Figure 1-5 with the RIP and OSPF routing protocols reversed. If OSPF routes are always preferred over RIP, IRX A and IRX B have exactly the same problem in choosing a route to network 4.0.0.0. They always choose each other.

To avoid creating routing loops, use a routing protocol input filter. With such a filter, you can indicate on IRX A and IRX B that routes coming from RIP are accepted on a RIP-running Ethernet only and are translated into OSPF. Similarly, on an OSPF-running Ethernet, you can choose to propagate into RIP only those routes from the OSPF side of the network.

If multiple routers in your network perform route propagation using fixed metrics, you can avoid routing loops by applying route filters with your propagation rules.

Configuring RIP Routing 2

This chapter describes the Routing Information Protocol (RIP) operation on PortMaster products and provides sample network configurations that show how to configure each PortMaster product in the network.

This chapter discusses the following topics:

Use this chapter in conjunction with the *PortMaster Command Line Administrator's Guide*. Refer to the glossary and Chapter 1, "Routing Overview," for definitions of terms and for an explanation of how PortMaster products implement RIP.

Understanding RIP

RIP is an interior routing protocol that uses a distance-vector algorithm to optimize routing. The distance-vector algorithm used by RIP for making routing decisions is based on the number of hops to a destination. The route with the least number of hops is considered to have the lowest cost. The number of hops is referred to as the route metric. A route with a metric greater than 15 is considered to be unreachable.

Although the more accurate link state technology of the Open Shortest Path First routing protocol (OSPF) is the most common interior routing protocol, RIP is still widely used on routers in TCP/IP and in Novell IPX networks as the primary method for the exchange of routing information.

When using RIP, routing updates can be turned on or off for individual interfaces on the PortMaster, and each interface can be configured as follows:

- Broadcast routing updates
- Listen for routing updates
- Broadcast and listen for routing updates
- Neither broadcast nor listen for routing updates

When a PortMaster interface is configured to broadcast, it sends RIP information onto the network every 30 seconds. When a PortMaster interface is configured to listen, it receives RIP broadcasts from other routers on the network. Routing table entries are aged as follows for RIP:

- Updates are expected every 30 seconds.
- After 120 seconds if information for another route to a given destination is observed, the new route is used for the destination.
- After 180 seconds without an update, a route is considered unreachable, and is made obsolete by having its metric set to 16.
- The routing timer is suspended if an on-demand interface is suspended.
- Routes are valid for active interfaces only.
- If an interface fails, all routes that go through that interface are examined. If the specified gateway is available on another interface, the route is moved to the active interface. Otherwise the routing table entry is marked as obsolete—also referred to as poisoned.

When routing is enabled, a PortMaster broadcasts any default route and static route information as part of a normal RIP message.

RIP Configuration Examples

The examples in this section illustrate typical scenarios using RIP. Configurations are simplified to focus on the critical elements in each scenario. This section includes the following examples:

- "Routing with Subnets" on page 2-2
- "Using Proxy ARP" on page 2-8
- "Configuring a Simple Dial-In Connection" on page 2-9

Routing with Subnets

Figure 2-1 illustrates an appropriate implementation of subnetting. The example uses two class C addresses, one (192.168.100.0) for the local Ethernet segment and one (192.168.200.0) for dial-in users.

Figure 2-1 Routing with Subnets

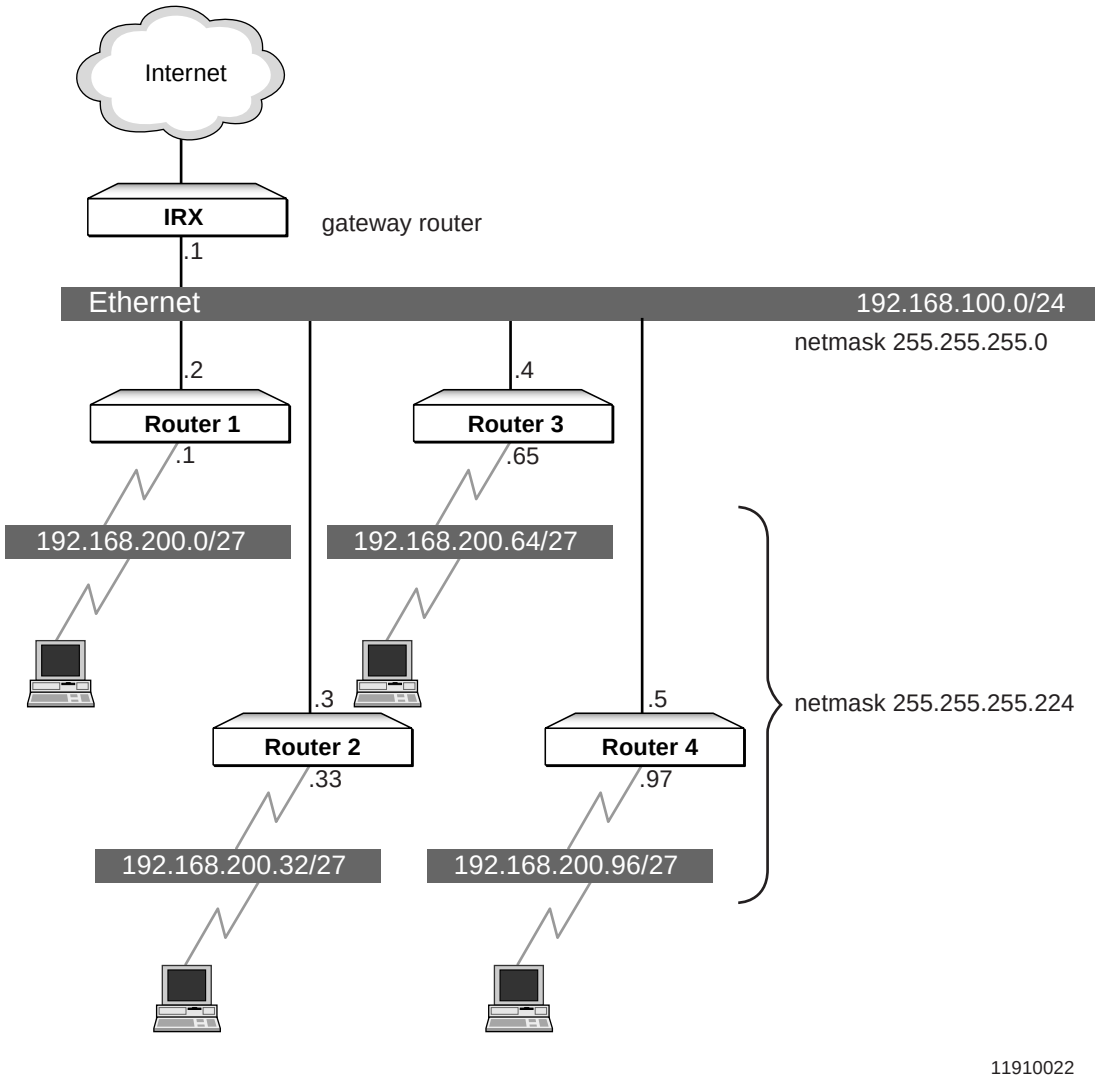

On the local Ethernet segment, an IRX at address 192.168.100.1/24 (netmask 255.255.255.0) is configured as the gateway router serving four PM-2ER-30s: Router 1 at address 192.168.100.2, Router 2 at 100.3, Router 3 at 100.4, and Router 4 at 100.5.

The second class C address, 192.168.200.0/27 (netmask 255.255.255.224) is divided into four subnets of 30 addresses each. Allowing for the use of the .0 subnet, this entire network is used for all the dial-in assigned pools. So the 192.168.200.0 is on 192.168.100.2, the 200.32 is on 100.3, the 200.64 is on 100.4, and the 200.96 is on 100.5. This is a convenient configuration because, with 30 dial-in ports on each PM-2ER-30, no addresses are wasted.

All the ports on the PM-2Es are configured for network dial-in, and routing is on by default—except during user configuration, when routing is set to off to ensure that it is disabled for dial-in clients.

Note – On Ethernet interfaces, routing is off by default in ComOS 3.5 and earlier.

Configuration on the PortMaster Internetwork Router (IRX)

This example shows the configuration for the Ethernet port connecting the PortMaster IRX to the Ethernet network it shares with the PM-2ER-30s. Configuration of the WAN interface (which includes setting the S1 address, netmask, and destination, as well as setting a gateway to the Internet service provider) is not shown.

To configure the IRX, enter the following commands:

1. **Set the Ethernet address:**

```
Command> set ether0 address 192.168.100.1
Local (ether0) address changed from 0.0.0.0 to 192.168.100.1
```

2. **Set the Ethernet netmask**

```
Command> set ether0 netmask 255.255.255.0
Ether0 netmask changed from 0.0.0.0 to 255.255.255.0
```

3. **Set the netmask for the 192.168.200.0 network:**

```
Command> add netmask 192.168.200.0 255.255.255.224
Ether0 netmask changed from 0.0.0.0 to 255.255.255.0224
```

Configuration on Router 1

1. **Set the Ethernet address:**

```
Command> set ether0 192.168.100.2
Local (ether0) changed from 0.0.0.0 to 192.168.100.2
```

2. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.0
Ether0 netmask changed from 0.0.0.0 to 255.255.255.0
```

3. **Set the upstream gateway router:**

```
Command> set gateway 192.168.100.1
Gateway changed from 0.0.0.0 to 192.168.100.1, metric = 1
```

4. **Set the base IP address of the assigned address pool:**

```
Command> set assigned 192.168.200.1
First assigned address changed from 0.0.0.0 to 192.168.200.1
```

5. **Set the assigned address pool size:**

```
Command> set pool 30
Assigned address pool size changed from 0 to 30
```

6. **Add a static netmask to the netmask table:**

```
Command> add netmask 192.168.200.0 255.255.255.224
New netmask successfully added
```

7. **Save the configuration:**

```
Command> save all
```

Configuration on Router 2

1. **Set the Ethernet address:**

```
Command> set ether0 192.100.3
Local (ether0) changed from 0.0.0.0 to 192.168.100.3
```

2. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.0
Ether0 netmask changed from 0.0.0.0 to 255.255.255.0
```

3. **Set the upstream gateway router:**

```
Command> set gateway 192.168.100.1
Gateway changed from 0.0.0.0 to 192.168.100.1, metric = 1
```

4. **Set the base IP address of the assigned address pool:**

```
Command> set assigned 192.168.200.33
First assigned address changed from 0.0.0.0 to 192.168.200.33
```

5. **Set the assigned address pool size:**

```
Command> set pool 30
Assigned address pool size changed from 0 to 30
```

6. **Add a static netmask to the netmask table:**

```
Command> add netmask 192.168.200.0 255.255.255.224
```

7. **Save the configuration:**

```
Command> save all
New netmask successfully added
```

Configuration on Router 3

1. **Set the Ethernet address:**

```
Command> set ether0 192.100.4
Local (ether0) changed from 0.0.0.0 to 192.168.100.4
```

2. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.0
Ether0 netmask changed from 0.0.0.0 to 255.255.255.0
```

3. **Set the upstream gateway router:**

```
Command> set gateway 192.168.100.1
Gateway changed from 0.0.0.0 to 192.168.100.1, metric = 1
```

4. **Set the base IP address of the assigned address pool:**

```
Command> set assigned 192.168.200.65
First assigned address changed from 0.0.0.0 to 192.168.200.65
```

5. **Set the assigned address pool size:**

```
Command> set pool 30
Assigned address pool size changed from 0 to 30
```

6. **Add a static netmask to the netmask table:**

```
Command> add netmask 192.168.1200.0 255.255.255.224
New netmask successfully added
```

7. **Save the configuration:**

```
Command> save all
```

Configuration on Router 4

1. **Set the Ethernet address:**

```
Command> set ether0 192.100.5
Local (ether0) changed from 0.0.0.0 to 192.168.100.5
```

2. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.0
Ether0 netmask changed from 0.0.0.0 to 255.255.255.0
```

3. **Set the upstream gateway router:**

```
Command> set gateway 192.168.100.1
Gateway changed from 0.0.0.0 to 192.168.100.1, metric = 1
```

4. **Set the base IP address of the assigned address pool:**

```
Command> set assigned 192.168.200.97
First assigned address changed from 0.0.0.0 to 192.168.200.97
```

5. **Set the assigned address pool size:**

```
Command> set pool 30
Assigned address pool size changed from 0 to 30
```

6. **Add a static netmask to the netmask table:**

```
Command> add netmask 192.168.200.0 255.255.255.224
New netmask successfully added
```

7. **Save the configuration:**

```
Command> save all
```

Using Proxy ARP

Proxy Address Resolution Protocol (ARP) is ARP for an IP address that is not local on the LAN. Figure 2-2 shows an example in which the PortMaster 3 Integrated Access Server acts as the proxy for dial-in users. The IP address of dial-in users must be in the same range as the Ethernet segment.

Figure 2-2 Proxy ARP

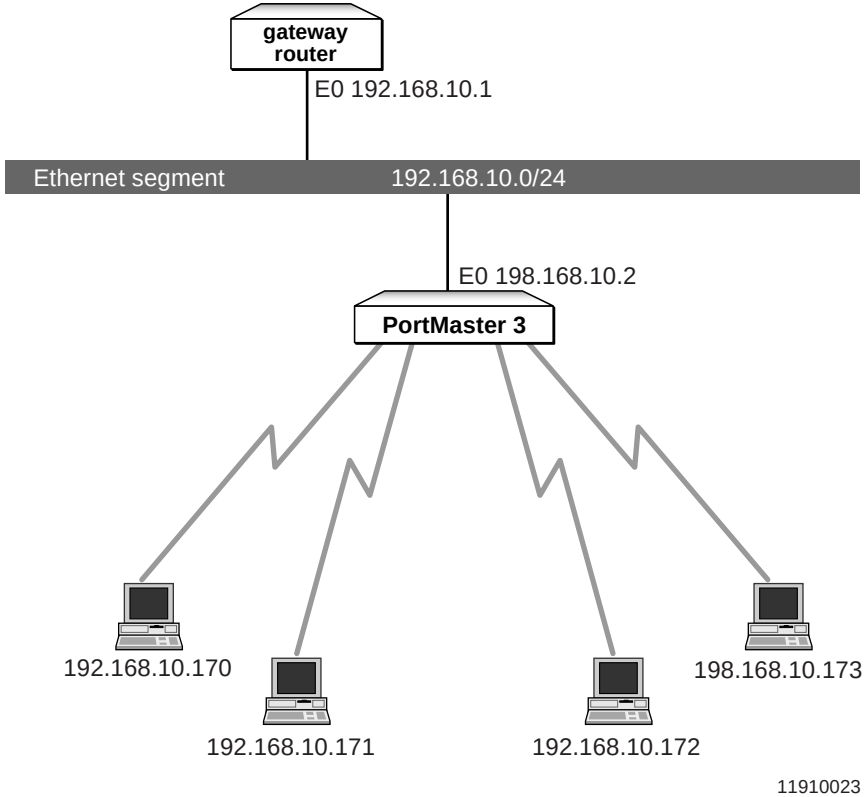

Configuration on the PortMaster 3

In addition to the regular configuration, enter the following commands to enable the PortMaster 3 to use proxy ARP for dial-in users:

1. **1. Set the base address of the assigned address pool:**

   ```
   Command> set assigned_address 198.168.10.65
   First Assigned address changed from 0.0.0.0 to 192.168.10.65
   ```

2. **Set the size of the assigned pool of IP addresses:**

   ```
   Command> set pool 48
   Assigned address pool size changed from 0 to 48
   ```

3. **Save the configuration:**

   ```
   Command> save all
   ```

Configuring a Simple Dial-In Connection

Figure 2-3 shows a simple network structure in which two locations are connected via a serial interface through a dial-in connection. Router 1, which has a connection to the Internet, is the only router connected to Ethernet 1. Router 2 is the only router connected to Ethernet 2. All data traffic between hosts on Ethernet 2 and the Internet must therefore use Router 1 as a gateway.

Figure 2-3 Simple Dial-In Router Connection

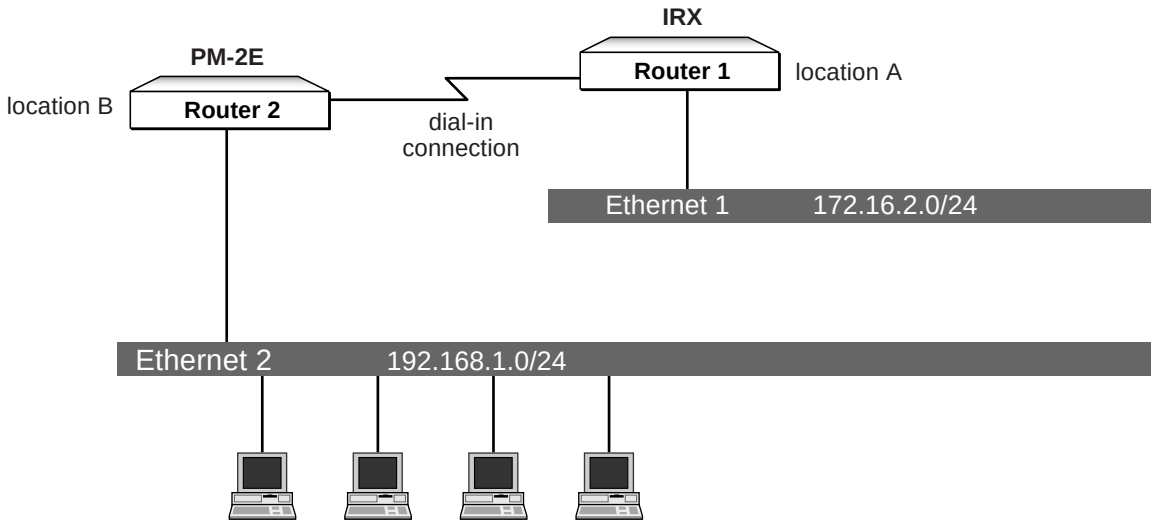

To make this connection possible, Router 2 must be configured so that hosts on Ethernet 2 can access the Internet via a dial-in PPP connection from Router 2 to Router 1, and Router 1 must be configured to accept a dial-in modem connection from Router 2.

Router 2 must broadcast RIP packets over the dial-in connection that contains route information about the Ethernet 2 network. Router 1 must listen for RIP packets over the dial-in connection to acquire the route information broadcast by Router 2.

Configuration on Router 1

In this example, a network user profile is configured for dial-in calls from Router 2. To keep the example focused, configuration of the WAN interface to the Internet is not considered. Router 1 is assumed to be a PortMaster IRX.

For more information, see the instructions for configuring dial-in connections in the *Configuration Guide for PortMaster Products*.

The complete configuration for Router 1 appears in the following box:

```
Command> set ether0 address 172.16.2.1
Command> add netuser prouter2 password ghrimpdezxt1
Command> set user prouter2 protocol ppp
Command> set user prouter2 destination 192.168.1.1
Command> set user prouter2 routing listen
Command> save all
Command> reboot
```

The following procedure shows the separate tasks involved in the configuration of Router 1:

1. **Set the Ethernet interface address:**

   ```
   Command> set ether0 address 172.16.2.1
   Local (ether0) changed from 0.0.0.0 to 172.16.2.1
   ```

2. **Add a network user and password:**

   ```
   Command> add netuser prouter2 password ghrimpdezxt1
   User successfully added
   ```

3. **Set the network protocol to be used:**

   ```
   Command> set user prouter2 protocol ppp
   Username: prouter2 Type: Dial-in Network User
   Address: Negotiated            Netmask: 255.255.255.0
   Protocol: PPP                    Options: Quiet, Compression
   MTU: 1500                Async Map: 00000000
   ```

4. **Set the IP address of the remote router:**

   ```
   Command> set user prouter2 destination 192.168.1.1
   Username: prouter2 Type: Dial-in Network User
   Address: 192.168.1.1 Netmask: 255.255.255.0
   Protocol: PPP                    Options: Quiet, Compression
   MTU: 1500                Async Map: 00000000
   ```

5. **Configure the router to receive RIP updates:**

```
Command> set user prouter2 routing listen
Username: prouter2 Type: Dial-in Network User
Address: 192.168.1.1 Netmask: 255.255.255.0
Protocol: PPP                    Options: Quiet, Listen, Compression
MTU: 1500               Async Map: 00000000
```

6. **Save the configuration and reboot:**

```
Command> save all
Command> reboot
```

The **show routes** command issued on Router 1 now gives the following (partial) response:

```
Command> show routes
Destination         Gateway          Flag     Met     Interface
-----------         -------          ----     ---     ---------
0.0.0.0             10.10.10.1       NS       1       ptp1
172.16.2.0          172.16.2.1       NL       1       ether0
192.168.1.1         192.168.1.1      HL       1       ptp2
192.168.1.0         192.168.1.1      ND       1       ptp2
```

The routes that Router 1 learned by listening to the RIP packets broadcast from Router 2 are shown in this router table.

Configuration on Router 2

This example assumes Router 2 to be a PortMaster 2E-30. The location table commands are used to set up the dial-out profile. Port S1 is used for the dial-out modem. To keep the examples focused on the routing configuration, global and host or user configurations are not shown.

For more information, see the instructions for configuring dial-out connections in the *Configuration Guide for PortMaster Products*.

The complete configuration for Router 2 appears in the following box:

```
Command> set ether0 192.168.1.1
Command> set gateway 172.16.2.1
Command> add location loc_a
Command> set location loc_a on_demand
Command> set location loc_a protocol ppp
Command> set location loc_a destination 172.16.2.1
Command> set location loc_a rip broadcast
Command> set location loc_a group 1
Command> set location loc_a telephone 15105551212
Command> set location loc_a username prouter2
Command> set location loc_a password ghrimpdezxt1
Command> set location loc_a idletime 20 minutes
Command> set location loc_a maxports 1
Command> set s1 group 1
Command> set s1 network dialout
Command> reset s1
Command> save all
Command> reboot
```

The following procedure shows the separate tasks involved in the configuration of Router 2:

1. **Set the Ethernet address on the interface connecting Router 2 to Router 1:**

   ```
   Command> set ether0 192.168.1.1
   Local (ether0) changed from  to 192.168.1.1
   ```

2. **Set the gateway for this router:**

   ```
   Command set gateway 172.16.2.1
   Gateway changed from 0.0.0.0 to 172.16.2.1
   ```

3. **Add the location of Router 1 to the location table:**

   ```
   Command> add location loc_a
   Location loc_a successfully added
   ```

4. **Configure the location to dial on demand:**

```
Command> set location loc_a on_demand
loc_a changed to On-Demand Dial
```

5. **Set the protocol for the location:**

```
Command> set location loc_a protocol ppp
loc_a protocol changed to ppp
```

6. **Set the IP address of the location:**

   Note that the IP address of the location (destination) is the same as the IP address of the gateway (see Step 2).

```
Command> set location loc_a destination 172.16.2.1
loc_a destination changed from 0.0.0.0 to 172.16.2.1
```

7. **Set the interface to send RIP packets to loc_a:**

```
Command> set location loc_a rip broadcast
loc_a rip changed from off to broadcast, no_listen
```

8. **Associate the location with a dial group:**

```
Command> set location loc_a group 1
loc_a group number changed from 0 to 1
```

9. **Set the telephone number used for dialing to the location:**

```
Command> set location loc_a telephone 15105551212
New telephone successfully set for location loc_a
```

10. **Set the username used to authenticate on the remote host:**

```
Command> set location loc_a username prouter2
New username successfully set for location loc_a
```

11. **Set the password used to authenticate on the remote host:**

    The username and password entered here must also be present on the remote host (in the user table, RADIUS, or other authentication mechanism).

```
Command> set location loc_a password ghrimpdezxt1
New password successfully set for location loc_a
```

12. **Set the time period the line can be idle before the connection is dropped:**

```
Command> set location loc_a idletime 20 minutes
loc_a idle timeout changed from 0 to 20 minutes
```

13. **Set the number of dial-out ports to be used for this connection:**

```
Command> set location loc_a maxports 1
loc_a maximumport count changed from 0 to 1
```

14. **Assign the port to the dial-out group:**

```
Command> set s1 group 1
Group number for port S1 changed from 0 to 1
```

15. **Make the port available for dialing to the remote location:**

```
Command> set s1 network dialout
Port type for S1 changed from Login to Network(dialout)
```

16. **Reset the port:**

```
Command> reset s1
Resetting port S1
```

17. **Save changes and reboot:**

```
Command> save all
Command> reboot
```

The **show routes** command issued on Router 2 now gives the following (partial) response:

```
Command> show routes
Destination         Gateway          Flag     Met     Interface
-----------         -------          ----     ---     ---------
192.168.1.0         192.168.1.1      NL       1       ether0
0.0.0.0             172.16.2.1       NS       1       ptp1
172.16.2.1          172.16.2.1       HL       1       ptp2
```

Configuring OSPF Routing 3

This chapter describes how to configure your PortMaster product for routing using the Open Shortest Path First (OSPF) protocol. Lucent implements the OSPF protocol as defined in RFC 1583. Also supported are not-so-stubby-areas (NSSAs), defined in RFC 1587, and message-digest 5 (MD5) algorithm authentication, defined in RFC 1321.

This chapter discusses the following topics:

Use this chapter in conjunction with the *PortMaster Command Line Administrator's Guide*. See the glossary and Chapter 1, "Routing Overview," for definitions of terms and for an explanation of how PortMaster products implement OSPF.

OSPF Configuration Tasks

The order of OSPF configuration is very important. Perform the configuration tasks in the following order:

1. **Enable the use of OSPF on the PortMaster—see page 3-2.**
2. **Set the designated and backup router priority—see page 3-2.**
3. **Set the OSPF areas and ranges—see page 3-2.**
4. **Enable OSPF on the interfaces—see page 3-3.**
5. **Enable configuration changes—see page 3-5.**

Note – You must enter the **save all** and **reboot** commands immediately after enabling or disabling OSPF routing. After you have configured all the desired OSPF settings, enter the **save all** command to save the configuration changes, and then enter the **reset ospf** command to restart OSPF routing.

Large OSPF routing tables might require you to upgrade your PortMaster to 4MB or 16MB of memory. See your hardware installation guide for instructions on upgrading PortMaster memory.

Enabling OSPF on the PortMaster

You must enable OSPF routing on the PortMaster before you can configure OSPF settings. To enable routing, enter the following commands:

```
Command> set ospf enable
Command> save all
Command> reboot
```

Setting Router Priority

You must set an OSPF priority for the router. Priorities range from 0 to 255. A calculation is performed on each interface separately to determine the designated router for that interface. The router with the highest priority on a network segment is the designated router. The router with the second highest priority is chosen as the backup designated router that takes over if the designated router becomes unable to perform its duties. If priorities tie, the router with the lower router ID is selected as the designated router. A router assigned priority 0 can never assume the duties of designated router.

To set an OSPF router priority, enter the following command:

```
Command> set ospf priority Number
```

Adding OSPF Areas

In Lucent's current implementation of OSPF, an area is a contiguous set of routers sharing network segments between them. All routers must have at least one interface in area 0.0.0.0, known as the backbone area. Use area 0.0.0.0 if you have only one OSPF area.

To add an OSPF area, enter the following command:

```
Command> add ospf area Area
```

You can specify *Area* in either dotted decimal notation (also known as dotted quad) or decimal notation. For example, the backbone area can be specified as either 0.0.0.0 or 0.

Note – Lucent does not currently support the use of virtual links either to create a noncontiguous area or to enable an area border router to not be directly attached to the backbone.

Setting OSPF Area Ranges

You must set one or more ranges of network addresses that define each OSPF area. You can set a maximum of eight ranges for each OSPF area. You can also optionally set the type of OSPF route propagation.

To set an OSPF area range and optional route propagation, enter the following command:

```
Command> set ospf area Area range Prefix/NM [advertise|quiet|off]
```

Prefix specifies the destination prefix shared by all IP addresses within the range. *NM* is the netmask that specifies the number of high-order bits in *Prefix* that must match addresses to include those addresses within the area. The netmask is a number from 1 to 30, preceded by a slash (/)—for example, /24. Table 1-1 on page 1-7 maps the CIDR representation of netmasks to the hexadecimal notation and to the more familiar dot-separated representation.

Routes to the networks within the range are summarized and propagated to other areas if you use the **advertise** keyword; this is the default. The routes are not summarized or propagated to other areas if you use the **quiet** keyword. The **off** keyword removes the specified range from the area.

Caution – Make sure that the ranges set with this command include the addresses for all PortMaster interfaces within this OSPF area.

Setting OSPF on the Interfaces

You can enable or disable the OSPF protocol on an Ethernet interface or on an asynchronous or synchronous network hardwired port.

Ethernet Interface

To set OSPF on the interface, enter the following command—all on one line:

```
Command> set Ether0 ospf on|off [cost Number] [hello-interval Seconds]
[dead-time Seconds]
```

The **on** keyword enables OSPF on the specified Ethernet interface. The **off** keyword disables OSPF on that interface.

You can specify the cost of sending a packet on the interface with a link state metric by using the **cost** *Number* keyword and value. The *Number* metric is a 16-bit decimal number between 1 and 65535. The default is 1.

Routers in OSPF networks continually exchange hello packets with their neighbor routers. You can set the interval that elapses between the transmission of hello packets on the interface by using the **hello-interval** *Seconds* keyword and value. *Seconds* can range from 10 to 120 seconds. The default is 10 seconds.

If the PortMaster stops receiving hello packets from a neighbor, it treats that router as inactive, or down. You can specify how long the PortMaster waits for hello packets from neighbors by using the **dead-time** *Seconds* keyword and value. *Seconds* can range from 40 to 1200 seconds. The default is 40 seconds.

Note – You must set the same **cost** value, the same **hello-interval** value, and the same **dead-time** value on all routers attached to a common network.

Asynchronous or Synchronous Port

To set OSPF on an asynchronous or synchronous port, enter the following command—all on one line:

```
Command> set S0|W0 ospf on|off [cost Number] [hello-interval Seconds]
[dead-time Seconds] [nbma|point-to-multipoint|wan-as-stub-ptmp]
```

The **on** keyword enables OSPF on the specified asynchronous (*S0*) or synchronous (*W0*) port. The **off** keyword disables OSPF on that port.

Enabling OSPF Configuration Changes

After you make all the desired configuration changes, enter the following commands:

```
Command> save all
Command> reset ospf
```

The **save all** command writes all configuration changes to the nonvolatile memory of the PortMaster. The **reset ospf** command performs the following steps to recreate OSPF startup conditions:

- Removes old MD5 authentication numbers and secrets
- Resets all active neighbors to use new key numbers and secrets
- Restarts OSPF routing, making all configuration changes effective without requiring the PortMaster to be rebooted

Additional OSPF Settings

This section describes additional commands that you might need for your OSPF configuration.

Setting an Area Password

You can specify an area password or key to use when you communicate to other routers in the area. The password must be an ASCII string of from 1 to 8 characters.

To set the password, enter the following command:

```
Command> set ospf area Area password String
```

Setting MD5 Authentication

You can specify an MD5 secret for an OSPF area. All routers in the area must have the same key number associated with the MD5 secret.

Caution – Overwriting the current key number with the same number causes the secret to be lost immediately—that is, without rebooting.

To set the key number and the secret for MD5 authentication, enter the following command:

```
Command> set ospf area Area md5 Number String
```

The key ID *Number* associated with the MD5 secret must be an integer between 1 and 255. The MD5 secret *String* must be an ASCII string of between 1 and 16 characters.

When an MD5 key number and secret are changed, both the old and new numbers and secrets remain valid until either the **reboot** or **reset ospf** command is entered.

Setting a Router ID

Note – The Ether0 IP address is used by default if you set the OSPF router address to 0.0.0.0 or the ID number to 0. Lucent strongly recommends that you use the default setting.

To set the OSPF router IP address or ID number, enter the following command:

```
Command> set ospf router-id Ipaddress|Number
```

Caution – Be careful when using this feature. When you set a new router ID, the links belonging to an old router ID take as long as 1 hour to expire. Routing instability can result during the expiration period.

Propagating External Routes

The propagation settings enable you to specify how routes coming from one routing protocol are translated and advertised by the PortMaster into another routing protocol.

Because external routes can be propagated into transit areas, but not into stub areas, you can determine external route propagation into OSPF areas by defining an area as either a transit area or a stub area.

The backbone area is always defined as a transit area.

A stub area does not attach to any area except the backbone and has no exit other than to the backbone. Because external routes are not propagated to stub areas, stub areas must be given a default route to reach external destinations. See “Injecting the Default Route into a Stub Area or NSSA” on page 3-11 for more information.

To set the propagation of external routes by defining an area as either a transit area or a stub area, use the following command. The keyword **on** defines the area as a transit area. The keyword **off** defines the area as a stub area.

```
Command> set ospf area Area external on|off
```

Propagating RIP Routing

For routers running both OSPF and RIP on a network, you can enable the propagation of RIP routes learned on a specified Ethernet interface from other routers into OSPF as OSPF Type 2 external routes. See Table 1-6 on page 1-17 for a description of OSPF route types.

To set RIP route propagation, enter the following command:

```
Command> set Ether0 ospf accept-rip on|off
```

If the RIP routes learned on the Ethernet interface originate from routers that are always running both OSPF and RIP, you can avoid duplicating route information by using the default **off** setting for RIP-to-OSPF propagation.

Defining Propagation Filters

The propagation filter is an IP access filter that you create in the filter table on the PortMaster. It uses the source addresses specified in the filter list to indicate routes.

Refer to the instructions for configuring filters in the *Configuration Guide for PortMaster Products* for information on setting filters.

To define a propagation filter, enter the following commands:

```
Command> add filter Filtername
Command> set filter Filtername RuleNumber permit|deny Prefix(src)/NM
Prefix(dest)/NM
```

Add other keywords and values as needed to the **set filter** command.

Defining Propagation Rules

You define a propagation rule to determine how routes coming into the PortMaster in one protocol are translated and advertised in another protocol. You insert the rule into a filter after you create the filter. The filter is stored in the filter table of the PortMaster.

To define a propagation rule, enter the following command—all on one line:

```
Command> add propagation Protocol(src) Protocol(dest) Metric Filtername
Command> reset propagation
```

Use the appropriate keyword—**rip**, **static**, **ospf**, **bgp**—to designate the source and destination protocols. *Metric* is a common metric used to translate from one protocol to another. A 0 metric causes the PortMaster to attempt to build a metric automatically.

Caution – If you plan to use a constant metric instead of the automatically generated metric provided by the ComOS, you run the risk of creating routing loops if you do not provide for filters or policies to screen the route information the PortMaster accepts from each routing protocol.

Modifying or Deleting Propagation Rules

Follow this procedure to change or delete a propagation rule:

1. **Delete the existing propagation rule as follows:**

   ```
   Command> delete propagation Protocol(SRC) Protocol(dest)
   ```

2. **If you are changing a rule, add the revised propagation rule as follows:**

   ```
   Command> add propagation Protocol(SRC) Protocol(dest) Metric Filtername
   ```

3. **Reset the propagation rules system as follows:**

   ```
   Command> reset propagation
   ```

4. **Follow any additional instructions prompted by the PortMaster.**

Applying Interface-Specific Propagation Route Filters

Apply a propagation route filter to a specific interface by entering the following command:

```
Command> set Ether0|S0|W1 route-filter in|out Filtername
Command> reset ospf
```

The route filter is a packet filter that you create in the filter table on the PortMaster. You must specify the Ethernet, asynchronous, or synchronous interface to which the filter is applied. You must also specify whether the filter is applied for inbound or outbound traffic.

Refer to instructions for configuring filters in the *Configuration Guide for PortMaster Products* for information about setting filters.

The effects of route filters depend on the protocol being filtered and on whether the filter is for inbound or outbound routes.

Table 3-1 describes the effects of route filters.

Table 3-1 Effects of Protocol on Route Filters

Protocol	Inbound Route Filter—Route Injection	Outbound Route Filter—Route Advertisement
RIP	The filter permit/deny rule applies and determines which routes are placed into the PortMaster routing table when • The address of the advertiser of the route matches the source address in the filter. • The destination address in the route being advertised matches the destination address in the filter. For RIP, the advertiser is the next-hop (direct) advertiser of the information.	The destination addresses in the filter determine which routes are advertised out of this interface.
OSPF	The filter permit/deny rule applies and determines which routes are placed into the routing table when • The address of the advertiser of the route matches the source address in the filter. • The destination address in the route being advertised matches the destination address in the filter. For OSPF, the advertiser is the ultimate advertiser of the information, not the next-hop OSPF router. Also, the filter specifies only the information that is in the routing table. Because OSPF area flooding rules make filtering inbound or outbound information on a per-interface basis impractical, applying the same inbound filter to all interfaces running OSPF within the same area is generally good practice.	The filter is ignored. OSPF area flooding rules make the definition of outbound route filters impractical on a per-interface basis. Use propagation filters to translate routing information from RIP, static, or BGP routes so that they do not enter OSPF as external Type 2 routes.

Setting an NSSA

Not-so-stubby areas (NSSAs) are similar to stub areas except that a form of Type 1 and Type 2 external routes can be learned from NSSAs. The N1/N2 routes learned from an NSSA are translated into corresponding external routes for the backbone and other transit areas that accept external routes. As for stub areas, default costs can be set for NSSAs. See "Setting a Stub Area or NSSA Default Route" on page 3-59 for more information. External routes are not advertised into NSSAs.

To define an area as an NSSA, enter the following command:

```
Command> set ospf area Area nssa on|off
```

You can specify the cost of sending a packet on the interface with a link state metric by using the **cost** *Number* keyword and value. The *Number* metric is a 16-bit decimal number between 1 and 65535. The default is 1.

Routers in OSPF networks continually exchange hello packets with their neighbor routers. You can set the interval that elapses between the transmission of hello packets on the interface by using the **hello-interval** *Seconds* keyword and value. *Seconds* can range from 10 to 120 seconds. The default is 10 seconds.

If the PortMaster stops receiving hello packets from a neighbor, it treats that router as inactive, or down. You can specify how long the PortMaster waits for hello packets from neighbors by using the **dead-time** *Seconds* keyword and value. *Seconds* can range from 40 to 1200 seconds. The default is 40 seconds.

OSPF Handling in a Frame Relay Network

You can specify how OSPF is handled based on the kind of Frame Relay network being used.

- **Nonbroadcast Multiaccess**—Use the **nbma** keyword for nonbroadcast multiaccess Frame Relay networks that have full mesh connectivity and all routers on the Frame Relay running OSPF. A designated router is elected, and overall OSPF traffic is reduced.
- **Point to Multipoint—**Use the **point-to-multipoint** keyword when the Frame Relay network has partial mesh connectivity or when all OSPF speakers on the network cannot communicate with each other. If you set this keyword, the Frame Relay network is modeled as a series of point-to-point interfaces.

Configure a point-to-multipoint interface when the **show ospf link** command displays a large number of available routes, but the **show routes** command displays no routes learned over the Frame Relay interface.

- **WAN as Stub Point to Multipoint—**Use the **wan-as-stub-ptmp** keyword when interoperating with equipment from other vendors that implements a variant of **point-to-multipoint**. This mode advertises the Frame Relay network as a stub network in the router link state advertisement (LSA) as opposed to the standard host route.

  Use the **show ospf links** command to check the router LSAs of your neighbors on the Frame Relay network. If they show stub network link entries for the Frame Relay network with the netmask for that network, configure the interface as **wan-as-stub-ptmp**. If they show the Frame Relay network as a host route with a mask of 255.255.255.255, then configure the interface as **point-to-multipoint**.

Injecting the Default Route into a Stub Area or NSSA

You can enable an area border router to create and inject the default route (0.0.0.0) into a stub area or NSSA. The default route is advertised to a stub area or NSSA when you specify a cost—an integer from 0 to 15. Lower costs are preferred.

To enable the injection of external routes into a stub area or NSSA, enter the following command:

```
Command> set ospf area Area stub-default-cost Cost
```

To disable the injection of external routes into a stub area, specify a cost of 0.

Displaying OSPF Settings

Several commands enable you to display detailed information about your OSPF configuration.

To display information on the configured OSPF areas, enter the following command:

```
Command> show ospf areas
```

To display a summary of the OSPF database with one line per link state advertisement (LSA), enter the following command—all on one line:

```
Command> show ospf links [router|network|summary|external|nssa]
```

To show information about routers directly accessible through your network interfaces, enter the following command:

```
Command> show ospf neighbor
```

See the *Command Line Administrator's Guide* for detailed explanations and examples of these commands.

OSPF Configuration Examples

This section provides the following examples of OSPF network configurations:

Propagating OSPF over a Single WAN Link

This example shows how to connect two OSPF areas over a single WAN connection using numbered serial ports. The example uses PortMaster Office Routers on Ethernet (see Figure 3-1), although any PortMaster router or network type could be used.

Note – Because OSPF applies an area to each port that falls within the range you set for it, any port or network segment over which you want to propagate OSPF must be in the same range.

For a discussion of typical problems encountered in this type of configuration, see "Propagating OSPF over a WAN Link" on page A-15 in Appendix A.

Figure 3-1 Propagating OSPF over a WAN Link

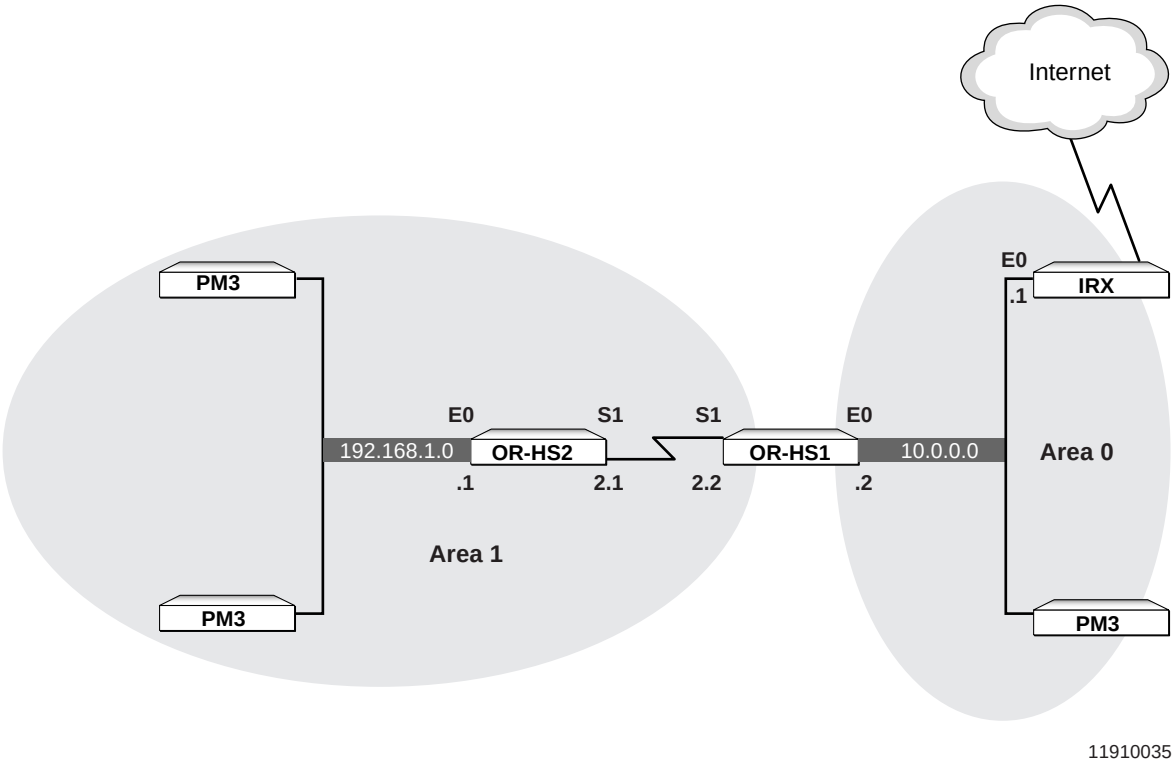

Configuration on OR-HS2

The order of OSPF configuration is important. First set priority (and router ID if desired), then set areas and ranges, and finally enable OSPF for the interfaces.

The complete configuration for OR-HS2 appears in the following box:

```
Command> set S1 address 192.168.2.1
Command> set S1 destination 192.168.2.2
Command> set ospf enable
Command> save all
Command> reboot
Command> add ospf area 1
Command> set ospf area 1 range 192.168.2.0/24
Command> set ospf area 1 range 192.168.1.0/24
Command> set S1 ospf on
Command> set Ether0 ospf on
Command> reset s1
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring OS-HS2:

1. **Set the local IP address of the network hardwired synchronous port to create a numbered interface:**

   ```
   Command> set S1 address 192.168.2.1
   Port S1 local address changed from 0.0.0.0 to 192.168.2.1
   ```

2. **Set the IP address of the remote router for a network hardwired synchronous port:**

   ```
   Command> set S1 destination 192.168.2.2
   Port W1 destination changed from 0.0.0.0 to 192.168.2.2
   ```

3. **Enable the use of OSPF on this router:**

   ```
   Command> set ospf enable
   OSPF will be enabled after next reboot
   ```

4. **Save the configuration and reboot the router:**

   ```
   Command> save all
   Saving global configuration
   Saving ports
   User table successfully saved
   Hosts table successfully saved
   ```

```
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved

Command> reboot
```

5. **Add Area 1 to the area tables of OR-HS2:**

```
Command> add ospf area 1
New Area successfully added
```

6. **Set the ranges of network addresses that define the OSPF area:**

```
Command> set ospf area 1 range 192.168.2.0/24
Area successfully updated
```

7. **Set the Area 1 OSPF range for the Ethernet port:**

```
Command> set ospf area 1 range 192.168.1.0/24
Area successfully updated
```

8. **Enable the OSPF protocol on the synchronous network hardwired port:**

```
Command> set S1 ospf on
S1 ospf state changed from off to on
```

9. **Enable the OSPF protocol on the Ethernet interface:**

```
Command> set Ether0 ospf on
Ether0 ospf state changed from off to on
```

10. **Reset the synchronous network hardwired port:**

```
Command> reset S1
Resetting port S1
```

11. **Save the configuration:**

```
Command> save all
Command> reset ospf
```

Configuration on OR-HS1

The complete configuration for OR-HS1 appears in the following box:

```
Command> set S1 address 192.168.2.2
Command> set S1 destination 192.168.2.1
Command> set ospf enable
Command> save all
Command> reboot
Command> add ospf area 0
Command> set ospf area 0 range 10.0.0.0/8
Command> add ospf area 1
Command> set ospf area 1 range 192.168.2.0/24
Command> set ospf area 1 range 192.168.1.0/24
Command> set S1 ospf on
Command> set Ether0 ospf on
Command> reset s1
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring OS-HS1:

1. **Set the local IP address of the network hardwired synchronous port to create a numbered interface:**

   ```
   Command> set S1 address 192.168.2.2
   Port S1 local address changed from 0.0.0.0 to 192.168.2.2
   ```

2. **Set the IP address of the remote router for a network hardwired synchronous port:**

   ```
   Command> set S1 destination 192.168.2.1
   Port W1 destination changed from 0.0.0.0 to 192.168.2.1
   ```

3. **Enable the use of OSPF on this router:**

   ```
   Command> set ospf enable
   OSPF will be enabled after next reboot
   ```

4. **Save the configuration and reboot the router:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved

Command> reboot
```

5. **Add Area 0 to the area tables of OR-HS1:**

```
Command> add ospf area 0
New Area successfully added
```

6. **Set the ranges of network addresses that define OSPF Area 0:**

```
Command> set ospf area 0 range 10.0.0.0/8
Area successfully updated
```

7. **Add Area 1 to the area tables of OR-HS1:**

```
Command> add ospf area 1
New Area successfully added
```

8. **Set the ranges of network addresses that define OSPF Area 1:**

```
Command> set ospf area 1 range 192.168.2.0/24
Area successfully updated
```

9. **Set the Area 1 OSPF range for the Ethernet port:**

```
Command> set ospf area 1 range 192.168.1.0/24
Area successfully updated
```

10. **Enable the OSPF protocol on the synchronous network hardwired port:**

```
Command> set S1 ospf on
S1 ospf state changed from off to on
```

11. **Enable the OSPF protocol on the Ethernet interface:**

```
Command> set Ether0 ospf on
Ether0 ospf state changed from off to on
```

12. **Reset the synchronous network hardwired port:**

```
Command> reset S1
Resetting port S1
```

13. **Save the configuration:**

```
Command> save all
Command> reset ospf
```

Nonbroadcast Multiaccess

The most basic OSPF scenario in a nonbroadcast multiaccess (NBMA) network is a single network in a single OSPF area. Figure 3-2 shows a more typical scenario, which includes more than one network in a single area. In this configuration, all routers that are members of the network can respond to any ARP packet that goes out on the network.

Although this example uses Ethernet, the example applies as well to any network type.

Note – Because this example is not concerned with configuring the gateway, no gateway address or netmask is included in the illustration or example configuration.

Figure 3-2 NBMA

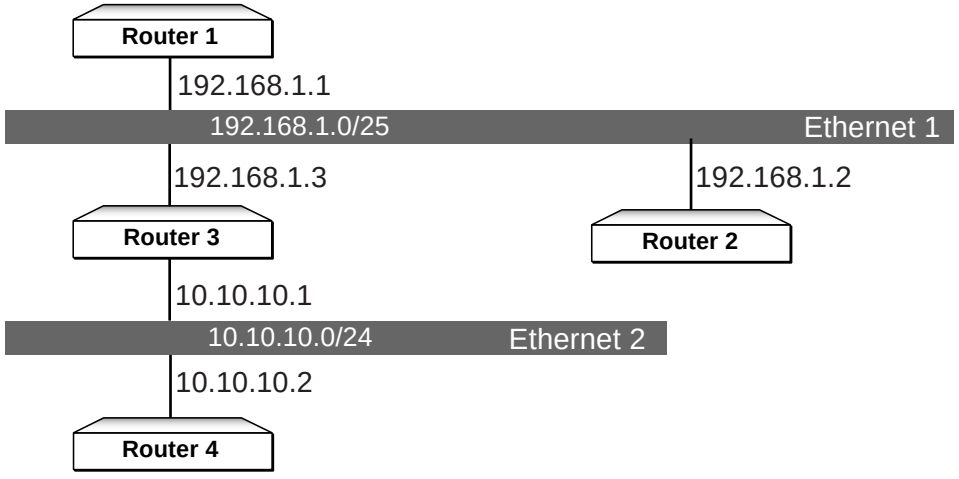

Assumptions:

- Router 1 is the designated router on Ethernet1.
- Router 3 is the backup designated router for Ethernet1.
- Router 3 is the gateway for Ethernet2.
- Router 4 is the designated router on Ethernet2.

Configuration on Router 1

The complete configuration for Router 1 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x
Command> set ether0 address 192.168.1.1
Command> set ether0 netmask 255.255.255.128
Command> set ospf priority 2
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/25
Command> set ospf area 0 range 10.10.10.0/24
Command> set ether0 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 1:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the Ethernet address:**

   ```
   Command> set ether0 address 192.168.1.1
   Local (ether0) address changed from 0.0.0.0 to 192.168.1.1
   ```

4. **Set the Ethernet netmask:**

   ```
   Command> set ether0 netmask 255.255.255.128
   Ether0 netmask changed from 0.0.0.0 to 255.255.255.128
   ```

5. **Set the router priority:**

```
Command> set ospf priority 2
OSPF priority changed from 0 to 2
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Ethernet 1:**

```
Command> set ospf area 0 range 192.168.1.0/25
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set the OSPF area range for Ethernet 2:**

```
Command> set ospf area 0 range 10.10.10.0/24
Range 10.10.10.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

9. **Set Ethernet OSPF on:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

10. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

11. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 2

The complete configuration for Router 2 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway 192.168.1.1
Command> set ether0 address 192.168.1.2
Command> set ether0 netmask 255.255.255.128
Command> set ospf priority 0
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/25
Command> set ospf area 0 range 10.10.10.0/24
Command> set ether0 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 2:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the gateway router:**

   ```
   Command> set gateway 192.168.1.1
   Gateway changed from 0.0.0.0 to 192.168.1.1, metric = 1
   ```

3. **Set the Ethernet address:**

   ```
   Command> set ether0 address 192.168.1.2
   Local (ether0) address changed from 0.0.0.0 to 192.168.1.2
   ```

4. **Set the Ethernet netmask:**

   ```
   Command> set ether0 netmask 255.255.255.128
   Ether0 netmask changed from 255.255.255.128 to 255.255.255.128
   ```

5. **Set the router priority:**

```
Command> set ospf priority 0
OSPF priority changed from 5 to 0
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range:**

```
Command> set ospf area 0 range 192.168.1.0/25
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set the OSPF area range:**

```
Command> set ospf area 0 range 10.10.10.0/24
Range 10.10.10.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

9. **Set Ethernet OSPF on:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

10. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

11. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 3

The complete configuration for Router 3 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway 192.168.1.1
Command> set ether0 address 192.168.1.3
Command> set ether0 netmask 255.255.255.128
Command> set ether1 address 10.10.10.1
Command> set ether1 netmask 255.255.255.0
Command> set priority 1
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/25
Command> set ospf area 0 range 10.10.10.0/24
Command> set ether0 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 3:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the gateway router:**

   ```
   Command> set gateway 192.168.1.1
   Gateway changed from 0.0.0.0 to 192.168.1.1, metric = 1
   ```

3. **Set the Ethernet 0 address:**

   ```
   Command> set ether0 address 192.168.1.3
   Local (ether0) address changed from 0.0.0.0 to 192.168.1.3
   ```

4. **Set the Ethernet 0 netmask:**

   ```
   Command> set ether0 netmask 255.255.255.128
   Ether0 netmask changed from 0.0.0.0 to 255.255.255.128
   ```

5. **Set the Ethernet 1 address:**

```
Command> set ether1 address 10.10.10.1
Local (ether0) address changed from 0.0.0.0 to 10.10.10.1
```

6. **Set the Ethernet 1 netmask:**

```
Command> set ether1 netmask 255.255.255.0
Ether0 netmask changed from 0.0.0.0 to 255.255.255.0
```

7. **Set the router priority:**

```
Command> set ospf priority 1
OSPF priority changed from 5 to 1
```

8. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

9. **Set the OSPF area range:**

```
Command> set ospf area 0 range 192.168.1.0/25
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

10. **Set the OSPF area range:**

```
Command> set ospf area 0 range 10.10.10.0/24
Range 10.10.10.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

11. **Set Ethernet OSPF on:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

12. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
```

```
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

13. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 4

The complete configuration for Router 4 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway 10.10.10.1
Command> set ether0 address 10.10.10.2
Command> set ether0 netmask 255.255.255.0
Command> set priority 0
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/25
Command> set ospf area 0 range 10.10.10.0/24
Command> set ether0 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 4:

1. **Enable OSPF on the router:**

```
Command> set ospf enable
Command> save all
Command> reboot
```

2. **Set the gateway router:**

```
Command> set gateway 10.10.10.1
Gateway changed from 0.0.0.0 to 10.10.10.1, metric = 1
```

3. **Set the Ethernet 0 address:**

```
Command> set ether0 address 10.10.10.2
Local (ether0) address changed from 0.0.0.0 to 10.10.10.2
```

4. **Set the Ethernet 0 netmask:**

```
Command> set ether0 netmask 255.255.255.0
Ether0 netmask changed from 0.0.0.0 to 255.255.255.0
```

5. **Set the router priority:**

```
Command> set ospf priority 0
OSPF priority changed from 5 to 0
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range:**

```
Command> set ospf area 0 range 192.168.1.0/25
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set the OSPF area range:**

```
Command> set ospf area 0 range 10.10.10.0/24
Range 10.10.10.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

9. **Set Ethernet OSPF on:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

10. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
```

```
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

11. Make OSPF configuration effective:

```
Command> reset ospf
Resetting OSPF
```

Nonbroadcast Multiaccess Multiple Areas

Figure 3-3 shows a slightly more complex OSPF network configuration than that shown in Figure 3-2, involving multiple areas.

Note – Because this example is not concerned with configuring the gateway, no gateway address or netmask is included in the illustration or example configuration.

Figure 3-3 NBMA Multiple Areas

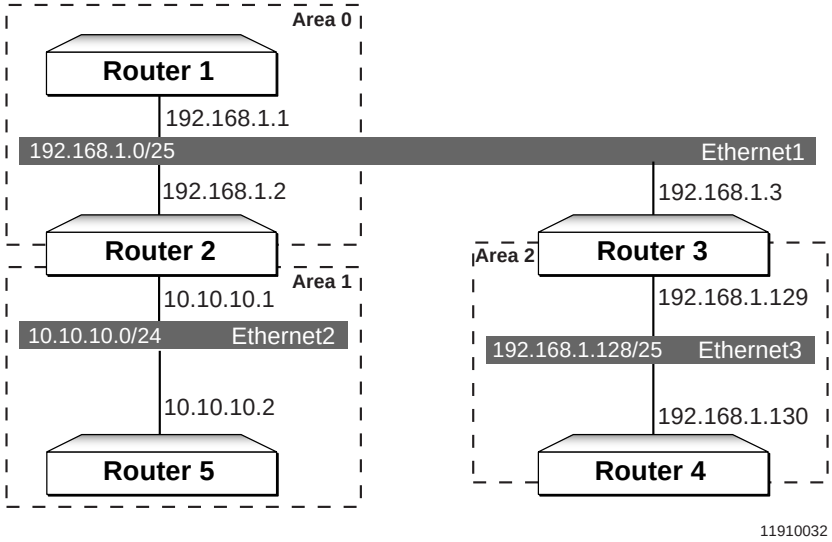

Assumptions:

- Area 0 consists of Router1, Router 2, and Router 3 on Ethernet1.
- Area 1 consists of Router 2 and Router 5 on Ethernet2.
- Area 2 consists of Router 3 and Router 4 on Ethernet3.

Configuration on Router 1

The complete configuration for Router 1 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x
Command> set ether0 address 192.168.1.1
Command> set ether0 netmask 255.255.255.128
Command> set ospf priority 2
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/25
Command> set ether0 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 1:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the upstream gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the Ethernet address:**

   ```
   Command> set ether0 address 192.168.1.1
   Local (ether0) address changed from 0.0.0.0 to 192.168.1.1
   ```

4. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.128
Ether0 netmask changed from 0.0.0.0 to 255.255.255.128
```

5. **Set the router priority:**

```
Command> set ospf priority 2
OSPF priority changed from 4 to 2
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Ethernet 1:**

```
Command> set ospf area 0 range 192.168.1.0/25
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set Ethernet OSPF on:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

10. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 2

The complete configuration for Router 2 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway 192.168.1.1
Command> set ether0 address 192.168.1.2
Command> set ether0 netmask 255.255.255.128
Command> set ether1 address 10.10.10.1
Command> set ether1 netmask 255.255.255.0
Command> set ospf priority 0
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/25
Command> add ospf area 1
Command> set ospf area 1 range 10.10.10.0/24
Command> set ether0 ospf on
Command> set ether1 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 2:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the gateway router:**

   ```
   Command> set gateway 192.168.1.1
   Gateway changed from 0.0.0.0 to 192.168.1.1, metric = 1
   ```

3. **Set the Ethernet 0 address:**

   ```
   Command> set ether0 address 192.168.1.2
   Local (ether0) address changed from 0.0.0.0 to 192.168.1.2
   ```

4. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.128
Ether0 netmask changed from 0.0.0.0 to 255.255.255.128
```

5. **Set the Ethernet 1 address:**

```
Command> set ether1 address 10.10.10.1
Local (ether0) address changed from 0.0.0.0 to 10.10.10.1
```

6. **Set the Ethernet netmask:**

```
Command> set ether1 netmask 255.255.255.0
Ether0 netmask changed from 0.0.0.0 to 255.255.255.128
```

7. **Set the router priority:**

```
Command> set ospf priority 0
OSPF priority changed from 2 to 0
```

8. **Define the OSPF Area 0:**

```
Command> add ospf Area 0
New Area successfully added
```

9. **Set the OSPF Area 0 range:**

```
Command> set ospf area 0 range 192.168.1.0/25
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

10. **Define the OSPF Area 1:**

```
Command> add ospf area 1
New Area successfully added
```

11. **Set the OSPF Area 1 range:**

```
Command> set ospf area 1 range 10.10.10.0/24
Range 10.10.10.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

12. **Set Ethernet 0 OSPF on:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

13. Set Ethernet 1 OSPF on:

```
Command> set ether1 ospf on
Ether1 ospf state changed from off to on
```

14. Save changes:

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

15. Make OSPF configuration effective:

```
Command> reset ospf
Resetting OSPF
```

Configuration for Router 3

The complete configuration for Router 3 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway 192.168.1.1
Command> set ether0 address 192.168.1.3
Command> set ether0 netmask 255.255.255.128
Command> set ether1 address 192.168.1.129
Command> set ether1 netmask 255.255.255.128
Command> set ospf priority 1
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/25
Command> add ospf area 2
Command> set ospf area 2 range 192.168.1.128/15
Command> set ether0 ospf on
Command> set ether1 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 3:

1. **Enable OSPF on the router:**

```
Command> set ospf enable
Command> save all
Command> reboot
```

2. **Set the gateway router:**

```
Command> set gateway 192.168.1.1
Gateway changed from 0.0.0.0 to 192.168.1.1, metric = 1
```

3. **Set the Ethernet 0 address:**

```
Command> set ether0 address 192.168.1.3
Local (ether0) address changed from 0.0.0.0 to 192.168.1.3
```

4. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.128
Ether0 netmask changed from 0.0.0.0 to 255.255.255.128
```

5. **Set the Ethernet 1 address:**

```
Command> set ether1 address 192.168.1.129
Local (ether0) address changed from 0.0.0.0 to 192.168.1.129
```

6. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.128
Ether0 netmask changed from 0.0.0.0 to 255.255.255.128
```

7. **Set the router priority:**

```
Command> set ospf priority 1
OSPF priority changed from 5 to 1
```

8. **Define the OSPF Area 0:**

```
Command> add ospf area 0
New Area successfully added
```

9. **Set the OSPF Area 0 range:**

```
Command> set ospf area 0 range 192.168.1.0/25
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

10. **Define the OSPF Area 2:**

```
Command> add ospf area 2
New Area successfully added
```

11. **Set the OSPF Area 2 range:**

```
Command> set ospf area 2 range 192.168.1.128/15
Range 192.168.1.128 for area 0.0.0.0 successfully updated
Area successfully updated
```

12. **Set Ethernet 0 OSPF on:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

13. **Set Ethernet 1 OSPF on:**

```
Command> set ether1 ospf on
Ether1 ospf state changed from off to on
```

14. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

15. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration for Router 4

The complete configuration for Router 4 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway 192.168.1.129
Command> set ether0 address 192.168.1.130
Command> set ether0 netmask 255.255.255.128
Command> set ospf priority 0
Command> add ospf area 1
Command> set ospf area 1 range 10.10.10.0/24
Command> set ether0 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 4:

1. **Enable OSPF on the router:**

```
Command> set ospf enable
Command> save all
Command> reboot
```

2. **Set the gateway router:**

```
Command> set gateway 192.168.1.129
Gateway changed from 0.0.0.0 to 192.168.1.129, metric = 1
```

3. **Set the Ethernet 0 address:**

```
Command> set ether0 address 192.168.1.130
Local (ether0) address changed from 0.0.0.0 to 192.168.1.130
```

4. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.128
Ether0 netmask changed from 0.0.0.0 to 255.255.255.128
```

5. **Set the router priority:**

```
Command> set ospf priority 0
OSPF priority changed from 5 to 0
```

6. **Define the OSPF Area 1:**

```
Command> add ospf area 1
New Area successfully added
```

7. **Set the OSPF Area 1 range:**

```
Command> set ospf area 1 range 10.10.10.0/24
Range 10.10.10.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set Ethernet 0 OSPF on:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

10. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration for Router 5

The complete configuration for Router 5 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway 10.10.10.1
Command> set ether0 address 10.10.10.2
Command> set ether0 netmask 255.255.255.0
Command> set ospf priority 1
Command> add ospf area 1
Command> set ospf area 1 range 10.10.10.0/24
Command> set ether0 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 5:

1. **Enable OSPF on the router:**

```
Command> set ospf enable
Command> save all
Command> reboot
```

2. **Set the gateway router:**

```
Command> set gateway 10.10.10.1
Gateway changed from 0.0.0.0 to 10.10.10.1, metric = 1
```

3. **Set the Ethernet 0 address:**

```
Command> set ether0 address 10.10.10.2
Local (ether0) address changed from 0.0.0.0 to 10.10.10.2
```

4. **Set the Ethernet netmask:**

```
Command> set ether0 netmask 255.255.255.0
Ether0 netmask changed from 0.0.0.0 to 255.255.255.0
```

5. **Set the router priority:**

```
Command> set ospf priority 1
OSPF priority changed from 5 to 1
```

6. **Define the OSPF Area 1:**

```
Command> add ospf area 1
New Area successfully added
```

7. **Set the OSPF Area 1 range:**

```
Command> set ospf area 1 range 10.10.10.0/24
Range 10.10.10.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set Ethernet 0 OSPF on:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. Save changes:

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

10. Make OSPF configuration effective:

```
Command> reset ospf
Resetting OSPF
```

Fully Meshed Frame Relay

The fully meshed Frame Relay network in Figure 3-4 is another example of a nonbroadcast multiaccess configuration. The significant difference between this configuration and the configuration in Figure 3-2 is that the network medium type in Figure 3-4 is Frame Relay.

Note – Because this example is not concerned with configuring the gateway, no gateway address or netmask is included in the illustration or example configuration.

Figure 3-4 NBMA Fully Meshed Frame Relay

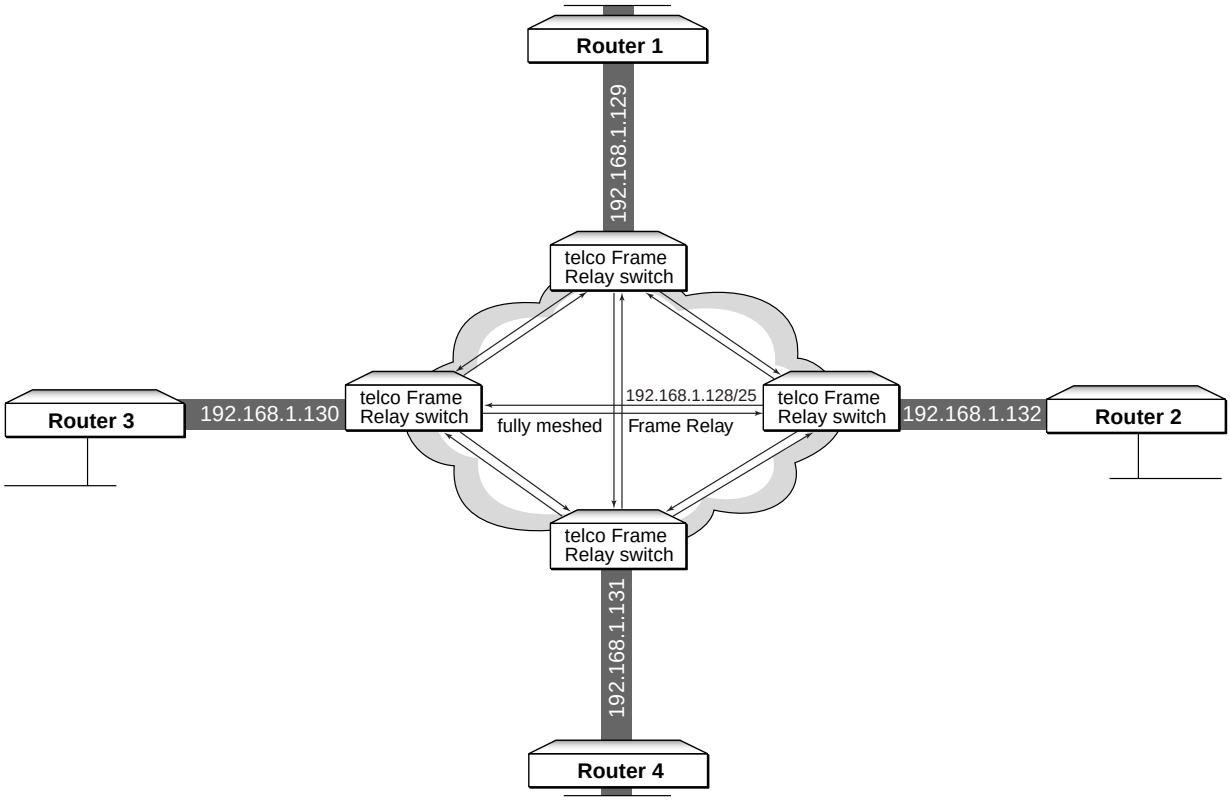

Configuration on Router 1

The complete configuration for Router 1 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x
Command> set ether0 address y.y.y.y
Command> set ether0 netmask 255.255.255.z
Command> set ospf priority 1
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/24
Command> set ether0 ospf on
Command> set w1 network hardwired
Command> set w1 protocol frame
Command> set w1 lmi 10
Command> set w1 address 192.168.1.129
Command> set w1 netmask 255.255.255.128
Command> set w1 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 1:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the IP address of the upstream gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the address of the Ethernet interface:**

   ```
   Command> set ether0 address y.y.y.y
   Local (ether0) address changed from 0.0.0.0 to y.y.y.y
   ```

4. **Set the netmask of the Ethernet subnet:**

```
Command> set ether0 netmask 255.255.255.z
Ether0 netmask changed from 0.0.0.0 to 255.255.255.z
```

5. **Set the router priority:**

```
Command> set ospf priority 1
OSPF priority changed from 3 to 1
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Area 0:**

```
Command> set ospf area 0 range 192.168.1.0/24
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set OSPF on for ether0:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Set the network type for the synchronous port:**

```
Command> set w1 network hardwired
Port type for port W1 changed from Network to Network(hardwired)
```

10. **Set the transport protocol for the hardwired synchronous port:**

```
Command> set w1 protocol frame
Port type for port W1 changed from ppp to frame relay
```

11. **Set the local management interface polling interval:**

```
Command> set w1 lmi 10
LMI keepalive timer for W1 changed from 0 to 10
```

12. **Set the local IP address for the hardwired synchronous port to create a numbered interface:**

```
Command> set w1 address 192.168.1.129
Port W1 local address changed from 0.0.0.0 to 192.168.1.129
```

13. **Set the IP netmask of the remote router for the network hardwired synchronous port:**

```
Command> set w1 netmask 255.255.255.128
W1 netmask changed from 0.0.0.0 to 255.255.255.128
```

14. **Enable the OSPF protocol on the hardwired synchronous port:**

```
Command> set w1 ospf on
W1 ospf state changed from off to on
```

15. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

16. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 2

The complete configuration for Router 2 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x
Command> set ether0 address y.y.y.y
Command> set ether0 netmask 255.255.255.z
Command> set ospf priority 2
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/24
Command> set ether0 ospf on
Command> set w1 network hardwired
Command> set w1 protocol frame
Command> set w1 lmi 10
Command> set w1 address 192.168.1.132
Command> set w1 netmask 255.255.255.128
Command> set w1 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 2:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the IP address of the upstream gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the address of the Ethernet interface:**

   ```
   Command> set ether0 address y.y.y.y
   Local (ether0) address changed from 0.0.0.0 to y.y.y.y
   ```

4. **Set the netmask of the Ethernet subnet:**

```
Command> set ether0 netmask 255.255.255.z
Ether0 netmask changed from 0.0.0.0 to 255.255.255.z
```

5. **Set the router priority:**

```
Command> set ospf priority 2
OSPF priority changed from 4 to 2
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Area 0:**

```
Command> set ospf area 0 range 192.168.1.0/24
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set OSPF on for ether0:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Set the network type for the synchronous port:**

```
Command> set w1 network hardwired
Port type for port W1 changed from Network to Network(hardwired)
```

10. **Set the transport protocol for the hardwired synchronous port:**

```
Command> set w1 protocol frame
Port type for port W1 changed from ppp to frame relay
```

11. **Set the local management interface polling interval:**

```
Command> set w1 lmi 10
LMI keepalive timer for W1 changed from 0 to 10
```

12. **Set the local IP address for the hardwired synchronous port to create a numbered interface:**

```
Command> set w1 address 192.168.1.132
Port W1 local address changed from 0.0.0.0 to 192.168.1.132
```

13. **Set the IP netmask of the remote router for the network hardwired synchronous port:**

```
Command> set w1 netmask 255.255.255.128
W1 netmask changed from 0.0.0.0 to 255.255.255.128
```

14. **Enable the OSPF protocol on the hardwired synchronous port:**

```
Command> set w1 ospf on
W1 ospf state changed from off to on
```

15. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

16. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 3

The complete configuration for Router 3 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x
Command> set ether0 address y.y.y.y
Command> set ether0 netmask 255.255.255.z
Command> set ospf priority 0
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/24
Command> set ether0 ospf on
Command> set w1 network hardwired
Command> set w1 protocol frame
Command> set w1 lmi 10
Command> set w1 address 192.168.1.131
Command> set w1 netmask 255.255.255.128
Command> set w1 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 3:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the IP address of the upstream gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the address of the Ethernet interface:**

   ```
   Command> set ether0 address y.y.y.y
   Local (ether0) address changed from 0.0.0.0 to y.y.y.y
   ```

4. **Set the netmask of the Ethernet subnet:**

```
Command> set ether0 netmask 255.255.255.z
Ether0 netmask changed from 0.0.0.0 to 255.255.255.z
```

5. **Set the router priority:**

```
Command> set ospf priority 0
OSPF priority changed from 3 to 0
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Area 0:**

```
Command> set ospf area 0 range 192.168.1.0/24
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set OSPF on for ether0:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Set the network type for the synchronous port:**

```
Command> set w1 network hardwired
Port type for port W1 changed from Network to Network(hardwired)
```

10. **Set the transport protocol for the hardwired synchronous port:**

```
Command> set w1 protocol frame
Port type for port W1 changed from ppp to frame relay
```

11. **Set the local management interface polling interval:**

```
Command> set w1 lmi 10
LMI keepalive timer for W1 changed from 0 to 10
```

12. **Set the local IP address for the hardwired synchronous port to create a numbered interface:**

```
Command> set w1 address 192.168.1.131
Port W1 local address changed from 0.0.0.0 to 192.168.1.131
```

13. **Set the IP netmask of the remote router for the network hardwired synchronous port:**

```
Command> set w1 netmask 255.255.255.128
W1 netmask changed from 0.0.0.0 to 255.255.255.128
```

14. **Enable the OSPF protocol on the hardwired synchronous port:**

```
Command> set w1 ospf on
W1 ospf state changed from off to on
```

15. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

16. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 4

The complete configuration for Router 4 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x
Command> set ether0 address y.y.y.y
Command> set ether0 netmask 255.255.255.z
Command> set ospf priority 0
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/24
Command> set ether0 ospf on
Command> set w1 network hardwired
Command> set w1 protocol frame
Command> set w1 lmi 10
Command> set w1 address 192.168.1.130
Command> set w1 netmask 255.255.255.128
Command> set w1 ospf on
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 4:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the IP address of the upstream gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the address of the Ethernet interface:**

   ```
   Command> set ether0 address y.y.y.y
   Local (ether0) address changed from 0.0.0.0 to y.y.y.y
   ```

4. **Set the netmask of the Ethernet subnet:**

```
Command> set ether0 netmask 255.255.255.z
Ether0 netmask changed from 0.0.0.0 to 255.255.255.z
```

5. **Set the router priority:**

```
Command> set ospf priority 0
OSPF priority changed from 4 to 0
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Area 0:**

```
Command> set ospf area 0 range 192.168.1.0/24
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set OSPF on for ether0:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Set the network type for the synchronous port:**

```
Command> set w1 network hardwired
Port type for port W1 changed from Network to Network(hardwired)
```

10. **Set the transport protocol for the hardwired synchronous port:**

```
Command> set w1 protocol frame
Port type for port W1 changed from ppp to frame relay
```

11. **Set the local management interface polling interval:**

```
Command> set w1 lmi 10
LMI keepalive timer for W1 changed from 0 to 10
```

12. **Set the local IP address for the hardwired synchronous port to create a numbered interface:**

```
Command> set w1 address 192.168.1.130
Port W1 local address changed from 0.0.0.0 to 192.168.1.130
```

13. **Set the IP netmask of the remote router for the network hardwired synchronous port:**

```
Command> set w1 netmask 255.255.255.128
W1 netmask changed from 0.0.0.0 to 255.255.255.128
```

14. **Enable the OSPF protocol on the hardwired synchronous port:**

```
Command> set w1 ospf on
W1 ospf state changed from off to on
```

15. **Save changes:**

```
Command> save all
```

16. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Point-to-Multipoint Partially Meshed Frame Relay

Because permanent virtual circuits (PVCs) are provided by the access carrier and can be expensive, a fully meshed configuration might not be cost-effective. Also, if your organization has multiple remote sites with virtual circuits built into concentration routers at the central site, a hub and spoke configuration might be a better solution.

In Figure 3-5, although Router 4 is not directly linked to Route 3 and Router 5, it can communicate with them (is partially meshed) via its PVC connection with Router 2.

Note – Because this example is not concerned with configuring the gateway, no gateway address or netmask is included in the illustration or example configuration.

Figure 3-5 Point-to-Multipoint Partially Meshed Frame Relay

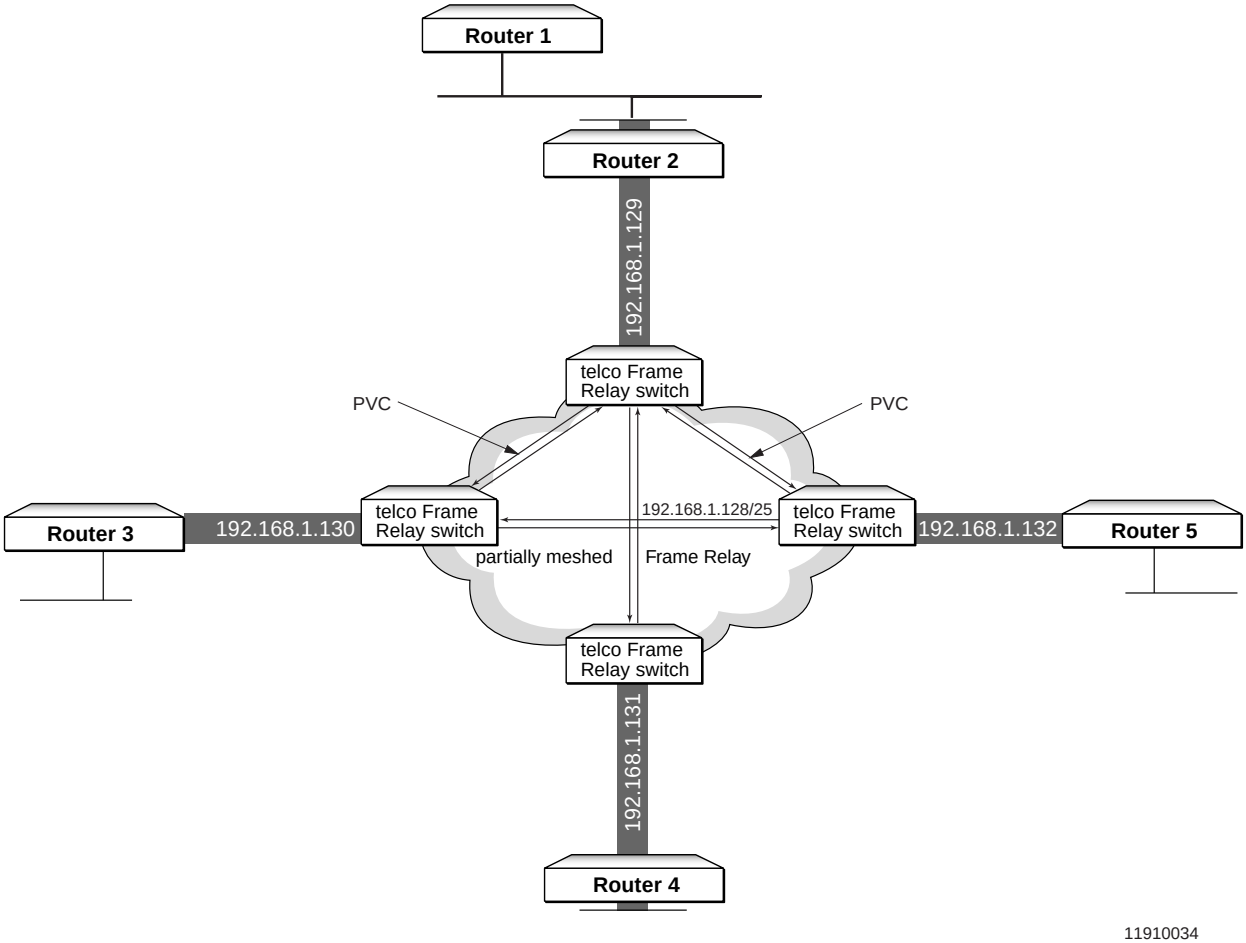

Configuration on Router 2

The complete configuration for Router 2 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x ==> IP address of gateway
Command> set ether0 address y.y.y.y ==> IP address of Ethernet interface
Command> set ether0 netmask 255.255.255.z ==> Netmask of Ethernet subnet
Command> set ospf priority 1
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/24
Command> set ether0 ospf on
Command> set w1 network hardwired
Command> set w1 protocol frame
Command> set w1 lmi 10
Command> set w1 address 192.168.1.129
Command> set w1 netmask 255.255.255.128
Command> set w1 ospf on point-to-multipoint
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 2:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the IP address of the upstream gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the address of the Ethernet interface:**

   ```
   Command> set ether0 address y.y.y.y
   Local (ether0) address changed from 0.0.0.0 to y.y.y.y
   ```

4. **Set the netmask of the Ethernet subnet:**

```
Command> set ether0 netmask 255.255.255.z
Ether0 netmask changed from 0.0.0.0 to 255.255.255.z
```

5. **Set the router priority:**

```
Command> set ospf priority 1
OSPF priority changed from 5 to 1
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Area 0:**

```
Command> set ospf area 0 range 192.168.1.0/24
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set OSPF on for ether0:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Set the network type for the synchronous port:**

```
Command> set w1 network hardwired
Port type for port W1 changed from Network to Network(hardwired)
```

10. **Set the transport protocol for the hardwired synchronous port:**

```
Command> set w1 protocol frame
Port type for port W1 changed from ppp to frame relay
```

11. **Set the local management interface polling interval:**

```
Command> set w1 lmi 10
LMI keepalive timer for W1 changed from 0 to 10
```

12. **Set the local IP address for the hardwired synchronous port to create a numbered interface:**

```
Command> set w1 address 192.168.1.129
Port W1 local address changed from 0.0.0.0 to 192.168.1.129
```

13. **Set the IP netmask of the remote router for the network hardwired synchronous port:**

```
Command> set w1 netmask 255.255.255.128
W1 netmask changed from 0.0.0.0 to 255.255.255.128
```

14. **Enable the OSPF protocol on the hardwired synchronous port and set it as the interface to a point-to-multipoint Frame Relay network:**

```
Command> set w1 ospf on point-to-multipoint
W1 ospf state changed from off to on
```

15. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

16. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 3

The complete configuration for Router 3 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x ==> IP address of gateway
Command> set ether0 address y.y.y.y ==> IP address of Ethernet interface
Command> set ether0 netmask 255.255.255.z ==> Netmask of Ethernet subnet
Command> set ospf priority 0
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/24
Command> set ether0 ospf on
Command> set w1 network hardwired
Command> set w1 protocol frame
Command> set w1 lmi 10
Command> set w1 address 192.168.1.130
Command> set w1 netmask 255.255.255.128
Command> set w1 ospf on point-to-multipoint
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 3:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the IP address of the upstream gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the address of the Ethernet interface:**

   ```
   Command> set ether0 address y.y.y.y
   Local (ether0) address changed from 0.0.0.0 to y.y.y.y
   ```

4. **Set the netmask of the Ethernet subnet:**

```
Command> set ether0 netmask 255.255.255.z
Ether0 netmask changed from 0.0.0.0 to 255.255.255.z
```

5. **Set the router priority:**

```
Command> set ospf priority 0
OSPF priority changed from 5 to 0
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Area 0:**

```
Command> set ospf area 0 range 192.168.1.0/24
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set OSPF on for ether0:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Set the network type for the synchronous port:**

```
Command> set w1 network hardwired
Port type for port W1 changed from Network to Network(hardwired)
```

10. **Set the transport protocol for the hardwired synchronous port:**

```
Command> set w1 protocol frame
Port type for port W1 changed from ppp to frame relay
```

11. **Set the local management interface polling interval:**

```
Command> set w1 lmi 10
LMI keepalive timer for W1 changed from 0 to 10
```

12. **Set the local IP address for the hardwired synchronous port to create a numbered interface:**

```
Command> set w1 address 192.168.1.130
Port W1 local address changed from 0.0.0.0 to 192.168.1.130
```

13. **Set the IP netmask of the remote router for the network hardwired synchronous port:**

```
Command> set w1 netmask 255.255.255.128
W1 netmask changed from 0.0.0.0 to 255.255.255.128
```

14. **Enable the OSPF protocol on the hardwired synchronous port and set it as the interface to a point-to-multipoint Frame Relay network:**

```
Command> set w1 ospf on point-to-multipoint
W1 ospf state changed from off to on
```

15. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

16. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 4

The complete configuration for Router 4 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x ==> IP address of gateway
Command> set ether0 address y.y.y.y ==> IP address of Ethernet interface
Command> set ether0 netmask 255.255.255.z ==> Netmask of Ethernet subnet
Command> set ospf priority 0
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/24
Command> set ether0 ospf on
Command> set w1 network hardwired
Command> set w1 protocol frame
Command> set w1 lmi 10
Command> set w1 address 192.168.1.131
Command> set w1 netmask 255.255.255.128
Command> set w1 ospf on point-to-multipoint
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 4:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the IP address of the upstream gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the address of the Ethernet interface:**

   ```
   Command> set ether0 address y.y.y.y
   Local (ether0) address changed from 0.0.0.0 to y.y.y.y
   ```

4. **Set the netmask of the Ethernet subnet:**

```
Command> set ether0 netmask 255.255.255.z
Ether0 netmask changed from 0.0.0.0 to 255.255.255.z
```

5. **Set the router priority:**

```
Command> set ospf priority 0
OSPF priority changed from 5 to 0
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Area 0:**

```
Command> set ospf area 0 range 192.168.1.0/24
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set OSPF on for ether0:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Set the network type for the synchronous port:**

```
Command> set w1 network hardwired
Port type for port W1 changed from Network to Network(hardwired)
```

10. **Set the transport protocol for the hardwired synchronous port:**

```
Command> set w1 protocol frame
Port type for port W1 changed from ppp to frame relay
```

11. **Set the local management interface polling interval:**

```
Command> set w1 lmi 10
LMI keepalive timer for W1 changed from 0 to 10
```

12. **Set the local IP address for the hardwired synchronous port to create a numbered interface:**

```
Command> set w1 address 192.168.1.131
Port W1 local address changed from 0.0.0.0 to 192.168.1.131
```

13. **Set the IP netmask of the remote router for the network hardwired synchronous port:**

```
Command> set w1 netmask 255.255.255.128
W1 netmask changed from 0.0.0.0 to 255.255.255.128
```

14. **Enable the OSPF protocol on the hardwired synchronous port and set it as the interface to a point-to-multipoint Frame Relay network:**

```
Command> set w1 ospf on point-to-multipoint
W1 ospf state changed from off to on
```

15. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

16. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuration on Router 5

The complete configuration for Router 5 appears in the following box:

```
Command> set ospf enable
Command> save all
Command> reboot
Command> set gateway x.x.x.x ==> IP address of gateway
Command> set ether0 address y.y.y.y ==> IP address of Ethernet interface
Command> set ether0 netmask 255.255.255.z ==> Netmask of Ethernet subnet
Command> set ospf priority 0
Command> add ospf area 0
Command> set ospf area 0 range 192.168.1.0/24
Command> set ether0 ospf on
Command> set w1 network hardwired
Command> set w1 protocol frame
Command> set w1 lmi 10
Command> set w1 address 192.168.1.132
Command> set w1 netmask 255.255.255.128
Command> set w1 ospf on point-to-multipoint
Command> save all
Command> reset ospf
```

The following procedure shows the separate tasks for configuring Router 5:

1. **Enable OSPF on the router:**

   ```
   Command> set ospf enable
   Command> save all
   Command> reboot
   ```

2. **Set the IP address of the upstream gateway router:**

   ```
   Command> set gateway x.x.x.x
   Gateway changed from 0.0.0.0 to x.x.x.x, metric = 1
   ```

3. **Set the address of the Ethernet interface:**

   ```
   Command> set ether0 address y.y.y.y
   Local (ether0) address changed from 0.0.0.0 to y.y.y.y
   ```

4. **Set the netmask of the Ethernet subnet:**

```
Command> set ether0 netmask 255.255.255.z
Ether0 netmask changed from 0.0.0.0 to 255.255.255.z
```

5. **Set the router priority:**

```
Command> set ospf priority 0
OSPF priority changed from 5 to 0
```

6. **Define the OSPF area:**

```
Command> add ospf area 0
New Area successfully added
```

7. **Set the OSPF area range for Area 0:**

```
Command> set ospf area 0 range 192.168.1.0/24
Range 192.168.1.0 for area 0.0.0.0 successfully updated
Area successfully updated
```

8. **Set OSPF on for ether0:**

```
Command> set ether0 ospf on
Ether0 ospf state changed from off to on
```

9. **Set the network type for the synchronous port:**

```
Command> set w1 network hardwired
Port type for port W1 changed from Network to Network(hardwired)
```

10. **Set the transport protocol for the hardwired synchronous port:**

```
Command> set w1 protocol frame
Port type for port W1 changed from ppp to frame relay
```

11. **Set the local management interface polling interval:**

```
Command> set w1 lmi 10
LMI keepalive timer for W1 changed from 0 to 10
```

12. **Set the local IP address for the hardwired synchronous port to create a numbered interface:**

```
Command> set w1 address 192.168.1.132
Port W1 local address changed from 0.0.0.0 to 192.168.1.132
```

13. **Set the IP netmask of the remote router for the network hardwired synchronous port:**

```
Command> set w1 netmask 255.255.255.128
W1 netmask changed from 0.0.0.0 to 255.255.255.128
```

14. **Enable the OSPF protocol on the hardwired synchronous port and set it as the interface to a point-to-multipoint Frame Relay network:**

```
Command> set w1 ospf on point-to-multipoint
W1 ospf state changed from off to on
```

15. **Save changes:**

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
New configurations successfully saved
```

16. **Make OSPF configuration effective:**

```
Command> reset ospf
Resetting OSPF
```

Configuring BGP Routing 4

This chapter describes how to configure PortMaster products to route using the Border Gateway Protocol (BGP). Lucent implements version 4 of the protocol, as defined in RFC 1771. Also supported are the BGP communities attribute, defined in RFC 1997; BGP autonomous system confederation, defined in RFC 1965; and BGP route reflection, defined in RFC 1966.

Use this chapter in conjunction with the *PortMaster Command Line Administrator's Guide*. Refer to the glossary and Chapter 1, "Routing Overview," for definitions of terms and for an explanation of how PortMaster products implement BGP.

Note – BGP runs only on PortMaster IRX Routers and the PortMaster 3. Because a full BGP routing table for the entire Internet requires about 7MB of memory, be sure to upgrade your IRX or PortMaster 3 memory to 16MB.

This chapter discusses the following topics:

BGP routing has many configuration options. See the *Command Line Administrator's Guide* for detailed command descriptions and instructions.

BGP Configuration Tasks

This section lists the tasks you must perform to configure BGP on your PortMaster for both a simple and a more complex configuration.

The order of BGP configuration is important. You first enable BGP on the PortMaster, and then complete the configuration steps in order for either a simple or advanced configuration.

Simple BGP Configuration

A simple **multihome** configuration for an autonomous system with multiple exit points to the Internet requires only the following key steps:

1. **Enable BGP routing—see page 4-4.**
2. **Set the BGP identifier—see page 4-4.**
3. **Set the autonomous system identifier—see page 4-4.**
4. **Save the settings by entering the following command:**

   ```
   Command> save all
   ```

5. **Add the peers and apply a routing method—see page 4-14.**
6. **Save the settings and reset BGP—see page 4-23.**

Advanced BGP Configuration

If you want more control over your BGP routing, you can create and apply your own propagation filters and routing policies through additional configuration steps:

1. **Enable BGP routing—see page 4-4.**
2. **Set the BGP identifier—see page 4-4.**
3. **Set the autonomous system identifier—see page 4-4.**
4. **Save the settings by entering the following command:**

   ```
   Command> save all
   ```

5. **(Optional) define propagation filters—see page 4-6.**
6. **(Optional) define propagation rules—see page 4-7.**
7. **Define BGP policies—see page 4-8.**
8. **Add peers and apply policies as needed—see page 4-14 and page 4-15.**

9. **Configure other options as needed—see page 4-16 through page 4-22.**
10. **Save the settings and reset BGP—see page 4-23.**

Configuring BGP on the PortMaster

This section describes how to configure the PortMaster 3 and the PortMaster IRX for BGP routing. Topics include the following:

Enabling BGP Routing

You must enable BGP routing on the PortMaster before you can configure BGP settings. To enable routing, enter the following commands:

```
Command> set bgp enable
Command> save all
Command> reboot
```

The **set bgp enable** command causes the PortMaster to load BGP software into memory when it is next rebooted. You must save the setting and reboot the PortMaster for the setting to take effect.

You must save all BGP configuration changes into nonvolatile RAM on the PortMaster. You can save after each setting change, or after a series of changes.

To save the settings you have just configured, enter the following command:

```
Command> save all
```

Setting the BGP Identifier

The BGP identifier is an IP address on the PortMaster that identifies the PortMaster as a BGP router to other routers. It is usually the IP address of one of the interfaces on the PortMaster.

To set the BGP identifier, enter the following command:

```
Command> set bgp id Ipaddress
```

Ipaddress must be specified in dotted decimal notation.

Setting the Autonomous System Identifier

The autonomous system (AS) identifier is the number used to identify the autonomous system to which the router belongs. The autonomous system identifier is supplied by InterNIC.

To set the autonomous system identifier, enter the following command:

```
Command> set bgp as ASN
```

If autonomous system confederations are in use, this identifier functions as your confederation's autonomous system number as it appears to peers outside the confederation.

Defining Confederations

You can avoid the overhead of having all peers within an autonomous system fully meshed by dividing an autonomous system into multiple smaller autonomous systems called confederation member autonomous systems (CMAS). These CMASs are grouped into a single confederation. You specify an identifier for the CMAS. The router advertises this autonomous system identifier to peers that are marked as confederation members in its configuration so that the confederation can be recognized by other confederation members. A confederation appears like a single autonomous system to external autonomous systems.

Subdividing an autonomous system into a confederation changes the peer relationships of confederation members in different CMASs from internal to external. If you confederate an autonomous system, you must ensure that all routers in the autonomous system belong to a CMAS; however, the policies used by BGP peers can change across confederation boundaries.

To set the autonomous system identifier for the BGP confederation member, enter the following command:

```
Command> set bgp cma ASN
```

Choosing a value of 0 (zero) disables use of confederations in this router. By default, this parameter is not set.

Confederations are one method for avoiding the overhead of having all peers within an autonomous system fully communicate to—be fully meshed with—each other. Route reflection clusters provide an alternative method, but require the use of identical policies on all peers within the autonomous system.

Defining the Route Reflector Cluster ID

A cluster is a set of BGP internal peers within an autonomous system. An autonomous system can be divided into many clusters. Each cluster must have one or more internal peers configured as route reflectors. The remaining peers in the cluster are called route reflector clients. Peers configured as route reflectors in an autonomous system are fully meshed with each other and with all other nonreflector clients in the autonomous system. Clients are configured as peers only with route reflectors in their cluster.

Dividing an autonomous system into route reflector clusters results in less network traffic and CPU overhead than a fully meshed system. Route reflector clusters are simpler to configure than confederations but do not allow the degree of policy control that is possible across confederation boundaries. The primary advantage of route reflectors is that they allow the PortMaster to interoperate with BGP peers that cannot be configured into confederations.

All route reflectors within a cluster must be configured with the same cluster ID. The cluster ID is not configured on the route reflector clients in the cluster.

To set the BGP route reflector cluster ID, enter the following command:

```
Command> set bgp cluster-id Ipaddress
```

The *Ipaddress* is in dotted decimal format. It can be any value but is typically the router ID of one of the route reflectors. Setting the value to 0.0.0.0 removes the cluster ID and prevents this router from being a route reflector.

Propagating Routing Protocols

Propagation settings let you specify how routes coming from one routing protocol are translated and advertised by the PortMaster into another routing protocol. You control route propagation in the following ways:

- "Defining Propagation Filters" on page 4-6
- "Defining Propagation Rules" on page 4-7
- "Modifying or Deleting Propagation Rules" on page 4-7

Defining Propagation Filters

The propagation filter is an IP access filter that you create in the filter table on the PortMaster. It uses source and destination IP prefixes and netmasks to specify protocol translation by route.

To define a propagation filter and specify the route you want to apply protocol translation to, enter the following commands:

```
Command> add filter Filtername
Command> set filter Filtername RuleNumber permit|deny Prefix(src)/NM
Prefix(dest)/NM
```

Add other keywords and values as needed to the **set filter** command. For information about setting filters, refer to the *Configuration Guide for PortMaster Products*.

Defining Propagation Rules

You define a propagation rule to determine how routes coming into the PortMaster in one protocol are translated and advertised in another protocol. A filter must first be created in the filter table of the PortMaster.

To define a propagation rule, enter the following command:

```
Command> add propagation Protocol(src) Protocol(dest) Metric Filtername
```

Use the appropriate keyword—**rip**, **static**, **ospf**, **bgp**—to designate the source and destination protocols. *Metric* is a common metric used to translate from one protocol to another. A metric of 0 causes the PortMaster to attempt to build a metric automatically.

Caution – If you plan to use a constant metric instead of the automatically generated metric provided by the ComOS, you run the risk of creating routing loops if you do not provide for filters or policies to screen the route information the PortMaster accepts from each routing protocol.

Each time the propagation rules are changed, you must reset the propagation rules system using the following command:

```
Command> reset propagation
```

To delete an existing propagation rule, enter the following command:

```
Command> delete propagation Protocol(SRC) Protocol(dest)
```

Modifying or Deleting Propagation Rules

Follow this procedure to change or delete a propagation rule:

1. **Delete the existing propagation rule as follows:**

   ```
   Command> delete propagation Protocol(SRC) Protocol(dest)
   ```

2. **If you are changing a rule, add the revised propagation rule as follows:**

   ```
   Command> add propagation Protocol(SRC) Protocol(dest) Metric Filtername
   ```

3. **Reset the propagation rules system as follows:**

   ```
   Command> reset propagation
   ```

4. **Follow any additional instructions prompted by the PortMaster.**

Working with BGP Policies

A BGP policy is a list of rules that constrain the BGP routes your PortMaster accepts from its peers, uses, and advertises to its peers. You can use the **easy-multihome** alternative to policies when you add each BGP peer to your peer group, or you can define your own policies.

A PortMaster uses an **acceptance policy** to determine whether to admit an IP prefix (received in an update from a BGP peer) into its BGP database for further consideration as a route. If the PortMaster accepts the IP prefix, it uses an **injection policy** to determine whether to use the route to forward packets, and an **advertisement policy** to determine whether to advertise the route to its BGP peers.

This section discusses the three types of policies and describes how to create them. Topics include the following:

Understanding Acceptance Policies

When the PortMaster learns a route from a peer, it scans the policy list. If a policy is applied to the peer and the route does not match any of the rules, the PortMaster does not use the route. If the list contains a rule whose **if** criteria match information

associated with the route, and the rule says **deny**, the route is also dropped. If the rule says **permit**, the route is accepted. If a **degree-of-preference** metric is specified in the **then** portion of the rule, the metric is used as the degree of preference for this route.

Summarization reduces the number of advertised routes (see "Advertising with Summarization" on page 4-16). Some BGP routes received by your PortMaster might not be summarized. Unsummarized routes can include IP prefixes containing as many as 32 high-order bits—many specific addresses rather than fewer route summaries. If your BGP policy rules accept such routes into your BGP database, you can propagate extremely large numbers of routes to your BGP peers and possibly overwhelm them. To avoid this problem, use the **prefix-longer-than** keyword in a BGP acceptance policy to deny IP prefixes with a netmask longer than a particular *NM* value. You might specify, for example, that **prefix-longer-than 16** not be advertised.

You can use policy statements to permit or deny certain routes from being reflected.

Understanding Injection Policies

For each route for which the PortMaster has determined the best possible route, it scans the policy list. If the route does not appear in a rule, the PortMaster denies the route, and it is not passed to the routing table—displayed with the **show route** command. If the list contains a rule whose **if** criteria match the route's IP address prefix and path attribute information, and the rule says **deny**, the route is also dropped. If the rule says **permit**, the route is accepted and passed to the routing table.

However, before any BGP route information is propagated into another routing protocol, such as OSPF or RIP, the route information must also pass any propagation filtering rules imposed by the **add propagation** command. See "Propagating Routing Protocols" on page 4-6 for more information.

An injection policy allows the PortMaster to receive and forward BGP routing information, but to forward packets based on simpler criteria. For example, you might want to forward packets only on routes received from OSPF or on a configured default route.

Understanding Advertisement Policies

For each route the PortMaster considers for advertisement, it scans the policy list. If the route does not appear in a rule, the PortMaster denies the route, and it is dropped. If the list contains a rule whose **if** criteria match the route's IP address prefix and path attribute information, and the rule says **deny**, the route is also dropped. If the rule says

permit, the route is advertised to the peer to which this policy applies. If the rule has any **then** attributes, these override any path attribute values the PortMaster would otherwise send to the other BGP peer.

You can use policy statements to permit or deny reflection of selected routes.

Effects of Route Reflection on BGP Policies

When a route reflector reflects an internal route (learned from internal peers) either from or to a reflector client, the BGP policies for the cluster change as follows:

- For advertisement policies, the route reflector ignores **then** portions and forwards every permitted route as learned. As a result, no modifications are made to the community, next-hop, multiexit discriminator, or local preference values.
- For acceptance policies, any multiexit discriminator is advertised as it was originally received and is not modified upon acceptance.

This modified behavior applies only to reflected internal routes learned from other internal peers, and not to routes originating from the route reflector itself. The route reflector can generate routes from locally configured summarizations or from routing information learned via external peers attached to the route reflector.

You can use policy statements to permit or deny certain routes from being reflected. For more information on policies, see "Working with BGP Policies" on page 4-8.

Creating a BGP Policy

To create a BGP policy, enter the following command:

```
Command> add bgp policy Policyname
```

You can create any number of policies.

To delete a BGP policy, enter the following command:

```
Command> delete bgp policy Policyname
```

The reserved policy name **all** is a predefined policy you can use as an acceptance, injection, and/or advertisement policy to permit or deny all routes.

Defining BGP Policies

You can create any number of acceptance policies, injection policies, or advertisement policies. You can create a single policy that includes all three functions, or you can create separate policies for each function. By avoiding the use of **if** or **then** clauses when defining rules, you can create rules that permit or deny all prefixes, with no modification.

The **set bgp policy** command allows you to define rules that determine how BGP routing information is handled. Table 4-1 describes the three parts of a policy rule—also called a policy statement:

Table 4-1 Parts of a Policy Rule

Header	The portion of the policy rule that defines the policy list. The header contains the following: • The name of the policy—*Policyname*. • Sequential location in the list—*RuleNumber*. • A keyword that indicates whether this rule will **permit** or **deny** routing information based on the rule's criteria. • A keyword that allows you to **include,** which creates nested policies (a policy within a policy). Policy inclusions can be nested to a maximum depth of 10 levels. Any inclusions beyond the 10th level are ignored.
If portion	The path attribute and IP prefix criteria that must be matched to make the current rule applicable.
Then portion	One or more attributes that the PortMaster applies (instead of computed or default parameters) when it accepts or advertises routing information passed by this rule. For example, a rule can specify that a local preference be advertised instead of letting the router compute that preference automatically.

Each **then** attribute applies to either acceptance or advertisement of the routing information, not to both. No attribute applies to injection of BGP routing information into the IP routing table of the PortMaster. As a result, you can combine acceptance, injection, and advertisement policy rules into one rule if they have the same path attribute and IP prefix criteria and are all either permit or deny rules.

Caution – The creation of long, complex lists of policy rules can adversely affect PortMaster CPU performance.

To define a BGP acceptance policy, use the following command:

```
Command> set bgp policy PolicyName [before] RuleNumber
permit|deny|include Policyname
[if
      [prefix [exactly] Prefix/NM]
      [prefix-longer-than NM]
      [as-path String|empty]
      [community Tag]]
[then
      [input-multi-exit-disc Number|strip]
      [degree-of-preference Number]]
```

To define a BGP injection policy, use the following command:

```
Command> set bgp policy PolicyName [before] RuleNumber
permit|deny|include Policyname
[if
      [prefix [exactly] Prefix/NM]
      [as-path String|empty]
      [community Tag]]
```

To define a BGP advertisement policy, use the following command:

```
Command> set bgp policy PolicyName [before] RuleNumber
permit|deny|include Policyname
[if
      [prefix [exactly] Prefix/NM]
      [as-path String|empty]
      [community Tag]]
[then
      [local-pref Number]
      [output-multi-exit-disc Number]
      [next-hop Ipaddress]
      [community add|replace|strip Tag]
      [ignore community restrictions]]
```

Applying a Policy

After adding or changing a rule in a BGP policy, use one of the following commands to apply and save the modified policy:

- Enter **reset bgp peer *Ipaddress(dest)*** to reset only those peers that use a policy.
- Enter **reset bgp** to reset all peers.

Removing a Policy Rule

To remove a rule, specify only the rule number *RuleNumber* in the command. For example, the command **set bgp policy policyname 1**, removes rule number 1 from the BGP policy.

Creating a Common Policy

You might want to use a common set of policies to define network routing plans and administrative needs. To create a common BGP policy for inclusion in other BGP policies, follow this procedure:

1. **Create and define a common BGP policy with the following commands:**

   ```
   Command> add bgp policy permit1011
   Command> set bgp policy permit1011 1 permit if prefix 10.0.0.0/8
   Command> set bgp policy permit1011 2 permit if prefix 11.0.0.0/8
   ```

2. **Include this policy by reference in another policy.**

   For example, to insert the **permit1011** policy at line 5 of the policy **otherone**, enter the following command:

   ```
   Command> set bgp policy otherone 5 include permit1011
   ```

3. **Apply and save the modified policy with the following command:**

   ```
   Command> reset bgp
   ```

See the *Command Line Administrator's Guide* for a details about the **set bgp policy** command.

Defining BGP Peers

BGP peers are pairs of routers that send BGP messages to each other. To pass BGP routing information, you must define the peers for your PortMaster router. You define the peer relationship by setting the IP address of the PortMaster (which will be placed in outgoing packets), the destination address of the peer, and the autonomous system number of the peer.

For example, to add a peer to your PortMaster, you enter the following command:

```
Command> add bgp peer Ipaddress(src) Ipaddress(dest) ASN
```

The IP addresses of the peers are specified in dotted decimal notation. The autonomous system number of the peer (ASN) is a 16-bit decimal number ranging from 1 to 65535.

You must define the relationship on the PortMaster for each of its peers. The relationship must also be configured on each peer router.

You must include a routing policy for each peer you define. You can create and apply specific routing policies to the peers you create or, alternatively, specify the **easy-multihome** routing method for a peer. If you do not specify a routing policy or method for a peer, all routes from the peer are denied.

Applying the Easy-Multihome Routing Method to a Peer

The **easy-multihome** method—the default—is a built-in routing policy that combines the functions of acceptance, injection, and advertisement. It restricts the BGP routing table to accept only paths through the remote autonomous system and, optionally, through one additional autonomous system.

To add or modify (set) a peer and specify the default routing method **easy-multihome**, enter the following command:

```
Command> add|set bgp peer Ipaddress(src) Ipaddress(dest) ASN
easy-multihome [assume-default [Number]] [confederation-
member] [normal] [route-reflector-client] [always-next-hop]
```

When modifying a peer with the **set** version of this command, you must re-enter all the keywords and values you want to associate with the peer. See the *Command Line Administrator's Guide* for a complete description of the **add | set bgp peer** command.

Applying Policies to Peers

To apply specific routing policies to BGP peers, use the **add bgp peer** command. Use the optional keywords and values to control how BGP policies are implemented for route selection.

To modify the characteristics of peers that have already been added to the PortMaster, use the **set** version of the command. When using this version, you must re-enter keywords and values you want to associate with the peer.

You can specify some combination of acceptance, injection, and advertisement policies. See "Working with BGP Policies" on page 4-8 for descriptions of BGP policies and instructions for defining policies.

To add or modify (set) a peer and specify the routing policy as some combination of acceptance, injection, and advertisement, enter the following command:

```
Command> add|set bgp peer Ipaddress(src) Ipaddress(dest) ASN
[accept-policy Policyname] [inject-policy Policyname]
[advertise-policy Policyname] [assume-default [Number]]
[confederation-member] [normal] [route-reflector-client]
[always-next-hop]
```

Note – If you do not specify a policy, all routes are denied for that policy.

See the *Command Line Administrator's Guide* for a complete description of the **add | set bgp peer** command.

Advertising with Summarization

BGP advertises to peers only routing information that is explicitly specified. These special advertisements are known as summarizations. BGP summarization entries control how Interior Gateway Protocol (IGP) routing information from OSPF, RIP, or static routing is forwarded into BGP for advertisement to BGP peers.

To add or modify a summarization entry, enter the following command:

```
Command> add|set bgp summarization Prefix/NM [as ASN [off]]
[cma ASN [off]]
[multi-exit-disc Number] [local-pref Number]
[community Tag]
```

Include your local autonomous system number (**as** *ASN*) in this list to enable the summarization to go to internal peers. You can list up to 14 autonomous systems. Include your CMAS number (**cma** *ASN*) in this list to enable the summarization to go to internal peers in your CMAS.

See the *Command Line Administrator's Guide* for a complete description of the **add | set bgp summarization** command.

Multiexit Discriminator

You can assign an arbitrary rating (**multi-exit-disc** *Number)* to an external route for advertisement to external or confederation-member peers only. *Number* is a 32-bit integer. Lower numbers indicate an increased preference for a specific route.

Use this metric to discriminate among multiple exit or entry points between the same pair of neighboring autonomous systems. If you do not assign a multiexit discriminator, the value 1 is assigned by default.

To explicitly prevent advertisement of a multiexit discriminator for IP prefixes matching this rule, set this keyword to 0 (zero). The PortMaster never forwards a 0 value of this metric to any peer, even if 0 was explicitly received from a peer.

A multiexit discriminator configured in a policy takes precedence over one configured in a route summarization. A decision tree showing the rules for applying a multiexit discriminator is shown in Figure 4-1.

For details about multiexit discriminators (MEDs), input multiexit discriminators (IMEDs), and output multiexit discriminators (OMEDs) refer to the BGP configuration chapter in the *Command Line Administrator's Guide*.

Figure 4-1 Decision Tree for Multiexit Discriminator (MED) Rules

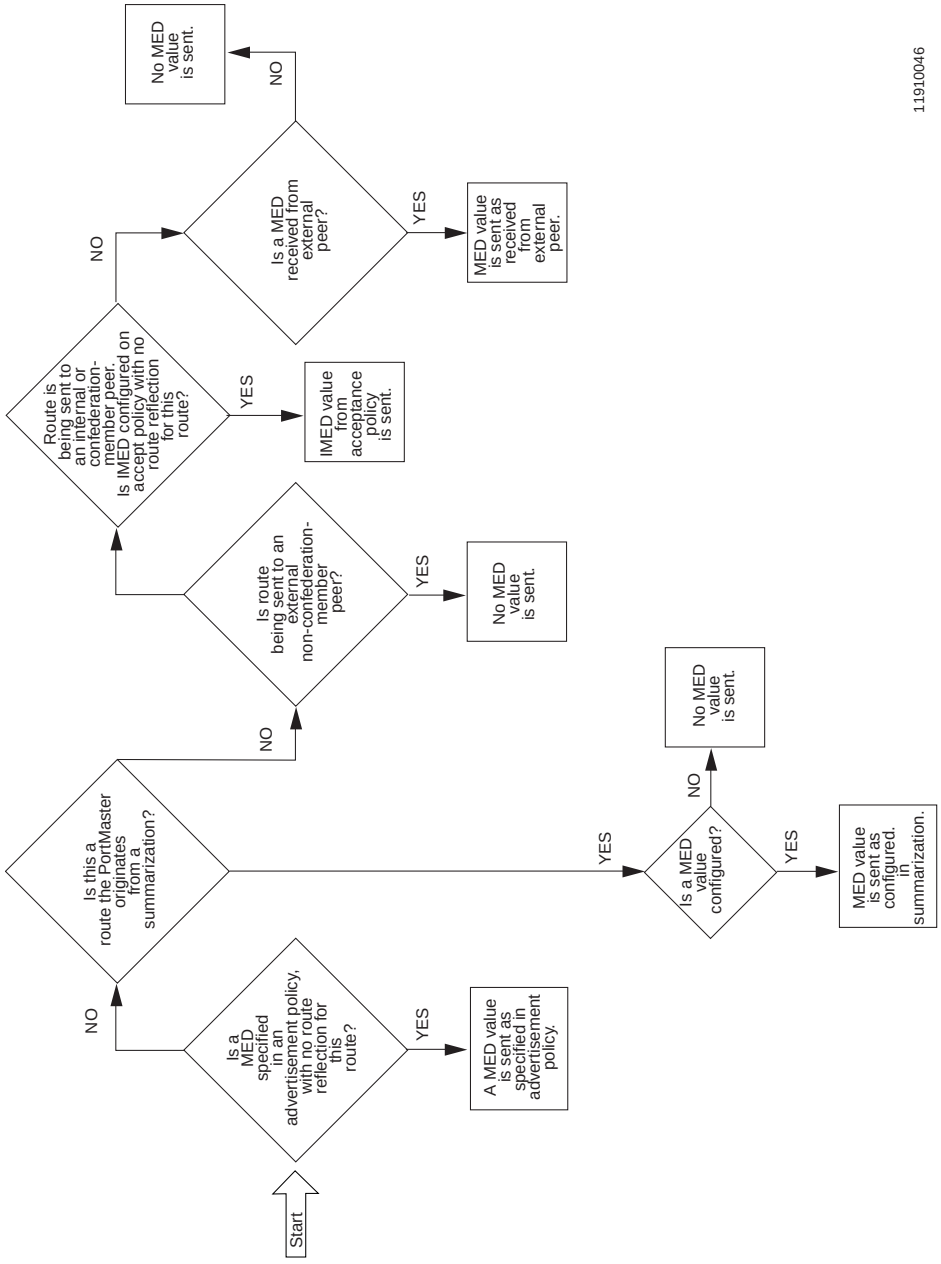

Local Preference

You can assign an arbitrary rating (**local-pref** *Number*) to an external route for advertisement to internal or confederation-member peers only. *Number* is a 32-bit integer. Higher numbers indicate an increased preference for a specific route when more than one route exists. Use this metric to screen a particular autonomous system from your map of routes, for example.

If you do not assign a local preference rating to the IP prefix, a value is assigned as follows:

- If the route comes from an internal peer, the learned local preference number is assigned.
- If the route comes from an external peer, *Number* is based on the autonomous system path length, with a shorter path being preferred.

The rules depend on the source of the information, as shown in Figure 4-2.

Figure 4-2 Decision Tree for Degree of Preference (DOP) Rules

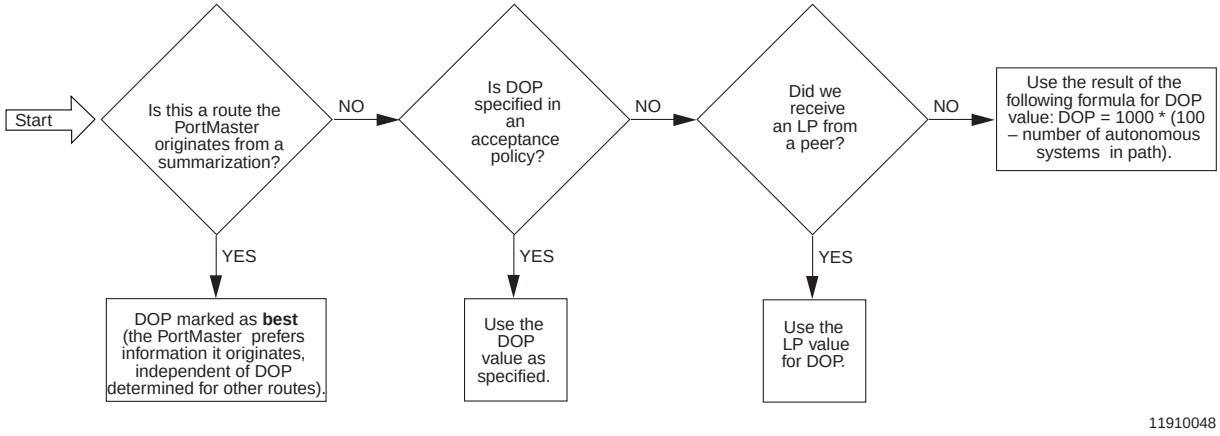

The rules for determining local preference (LP) values to be used for route selection depend on the source of the information, as show in Figure 4-3.

Figure 4-3 Decision Tree for Local Preference (LP) Rules

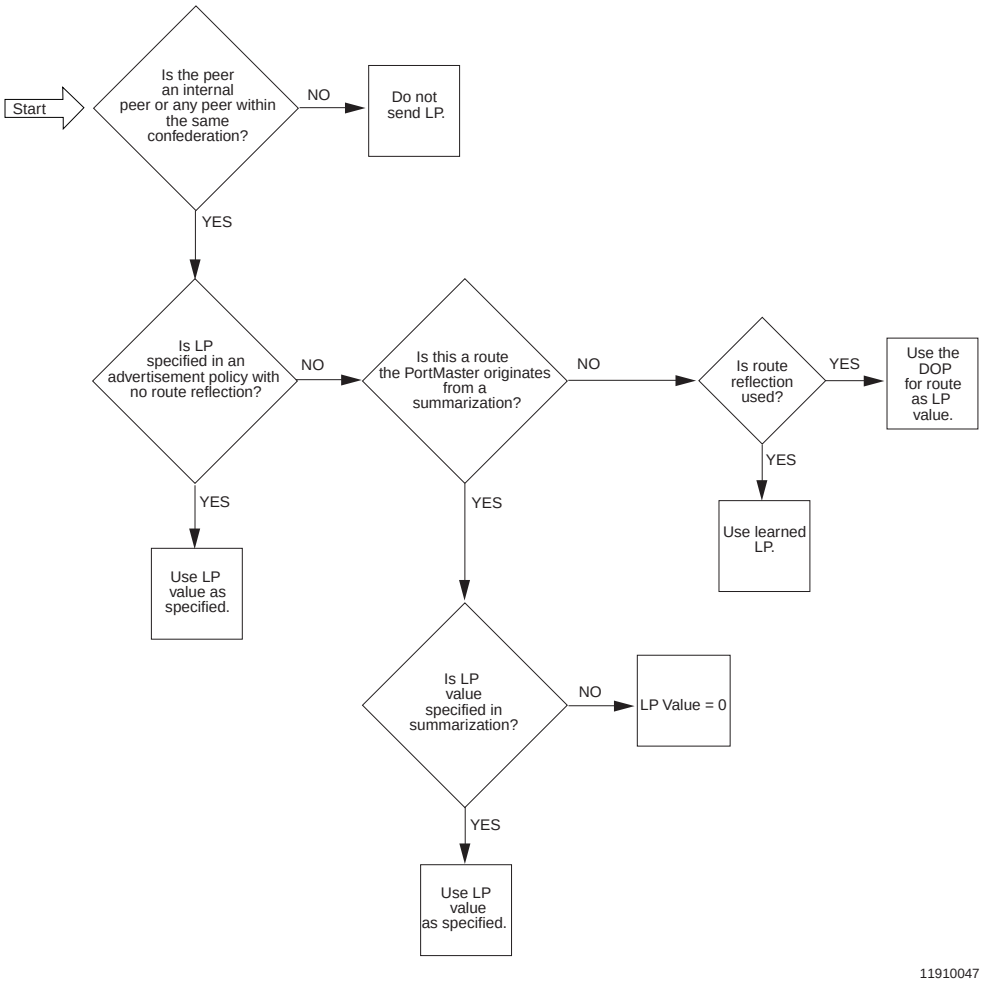

For automatic summarizations from static routes, LP is calculated from the formula LP = (16 – propagation metric of static route); if the propagation method of static routes is the automatic zero (0), LP is calculated from the formula LP = (16 – hop count metric of static route).

Creating BGP Communities

The communities attribute (**community** ***Tag***) variable lets you group a number of BGP destinations under a single name. By assembling destinations into identifiable communities, BGP peers can base policy decisions on the identity of the group rather than on individual destinations. This attribute simplifies the distribution of routing information by grouping a larger number of individual destinations into a smaller number of communities. The community identifier, which can be either one 32-bit value or two 16-bit values, is advertised in update messages between BGP peers.

The settings for community, local preference, and multiexit discriminator in the summarization command interact with those in advertisement policy definitions in the following ways:

- The advertisement policy settings override values for local preference and multi-exit discriminator
- If the advertisement policy definition includes an **add** community, that information is added to the community information specified in the summarization.
- If the advertisement policy definition includes a **replace** community, that information replaces any community information specified in the summarization.

To help provide stability in the Internet, summarizations are advertised only when supported by one or more specific routes that must exist for at least 30 seconds before the summarization is advertised.

Setting IGP Lockstep

When the source router learns a route from internal peers, it forwards the information to any external peers as soon as possible. Enabling the **igp-lockstep** feature forces the source router to wait until it finds a suitable IGP route (an OSPF, RIP, or static route, or a static route via RADIUS) that supports the route before advertising it. An IGP route supports a BGP route if it has the same address and prefix as the BGP route.

You enable **igp-lockstep** only when providing a transit service between two autonomous systems.

To enable or disable IGP lockstep, enter the following command:

```
Command> set bgp igp-lockstep on|off
```

Setting the Connection Retry Interval

You can set the interval after which the source router attempts to open sessions to peers that are not fully established.

To set the connection retry interval, enter the following command:

```
Command> set bgp connect-retry-interval Seconds
```

The valid range is from 30 to 1000 seconds. The default is 120 seconds.

Note – You must set the same value on all peers.

Setting the Keepalive Timer Interval

You can set the interval after which the source router sends keepalive messages to its peers to let them know it is still reachable.

To set the keepalive timer interval, enter the following command:

```
Command> set bgp keepalive-timer Seconds
```

The valid range is from 30 to 1000 seconds. The default is 30 seconds.

Note – You must set the same value on all peers.

Setting the Hold Time Interval

You can set the interval the source router waits between keepalive, update, or notification messages from a peer. When the peer is identified as no longer operational, all information learned from that peer is dropped.

To set the hold time interval, enter the following command:

```
Command> set bgp hold-time Seconds
```

The valid range is from 30 to 1000 seconds. The default is 90 seconds.

Note – You must set the same value on all peers.

Saving and Resetting BGP Routing

You must save all the peer configuration changes you make and reset BGP routing.

To save configuration settings and reset BGP routing, enter the following commands:

```
Command> save all
Command> reset bgp [peer Ipaddress]
```

The **reset bgp** command causes the PortMaster to delete all currently known BGP information. Configuration information is not deleted. The PortMaster rereads BGP configuration information and reestablishes sessions with peers. This process can take some time. If you enter the command and specify **peer** *Ipaddress*, the PortMaster deletes information for that peer only, and resets the configuration with that peer only.

Displaying BGP Settings

You can display the BGP settings configured on your PortMaster by using the commands described in this section. Table 4-2 lists BGP **show** commands and definitions

Table 4-2 BGP **show** Commands

Command	Description
show bgp memory	Provides information on BGP memory usage—an important issue when running BGP because of the large number of routes that are stored in the BGP database.
show bgp next-hop	Shows the known BGP next hop addresses and gateways. Use this command to determine where packets go when forwarded. The information displayed is based on entries in the routing table that are used to forward BGP packets to their destinations.
show bgp paths [*Prefix/NM* [**verbose**]]	Shows learned BGP path information.
show bgp peers [**verbose**]	Shows a list of BGP peers.

Table 4-2 BGP **show** Commands

show bgp policy [*Policyname*]	Shows BGP policy names and policy definitions.
show bgp summarization [**all**]	Shows route summaries advertised to BGP peers.
show routes [*String*|*Prefix*/*NM*	Shows the IP routing table.

See the *Command Line Administrator's Guide* for more information on these commands.

Debugging BGP

Use the **set debug** command to troubleshoot BGP. To track debug command output, enable the option with the appropriate **set debug** *Option* **on** command. This command sends output to the system console, which you set with the **set console** command.

After completing the debugging process, disable the debug commands by using the appropriate **set debug** *Option* **off** command.

To set BGP debug options, enter the following command:

```
Command> set debug bgp-fsm|bgp-decision-process|bgp-opens|
bgp-keepalives|bgp-updates|bgp-notifications|bgp-errors|
bgp-packets|bgp-max on|off
```

Caution – You should use the **set debug bgp-max** command in limited environments —for example, when investigating problems of peer interaction. Executing the **set debug bgp-max** command on a connection where large routing tables are exchanged between peers can create a flood of output that is useless for debugging.

Table 4-3 describes debug keywords.

Table 4-3 Debug Keywords

Keyword	Description
bgp-fsm	Displays events that change the state of the BGP session with any peer.
bgp-decision-process	Displays decisions among routes about the best path to a destination.
bgp-opens	Displays open messages sent and received between any peers.
bgp-keepalives	Displays keepalive messages sent and received between any peers.
bgp-updates	Displays update messages sent and received between any peers.
bgp-notifications	Displays notification messages sent and received between any peers.
bgp-errors	Displays protocol errors occurring between BGP peers.
bgp-packets	Displays bgp-opens, bgp-keepalives, bgp-updates, and bgp-notifications outputs.
bgp-max	Displays all BGP debugging output.

BGP Configuration Examples

This section provides the following examples of BGP routing configurations:

- "Easy-Multihome—Example 1" on page 4-26
- "Easy-Multihome—Example 2" on page 4-30
- "Easy-Multihome—Example 3" on page 4-39
- "Route Reflectors Example" on page 4-47
- "Confederations Example" on page 4-54

These examples provide only the BGP component of a complete routing configuration.

Easy-Multihome—Example 1

Figure 4-4 illustrates a basic multihome policy routing configuration using the default **all** policy.

Figure 4-4 Simple Multihome Configuration

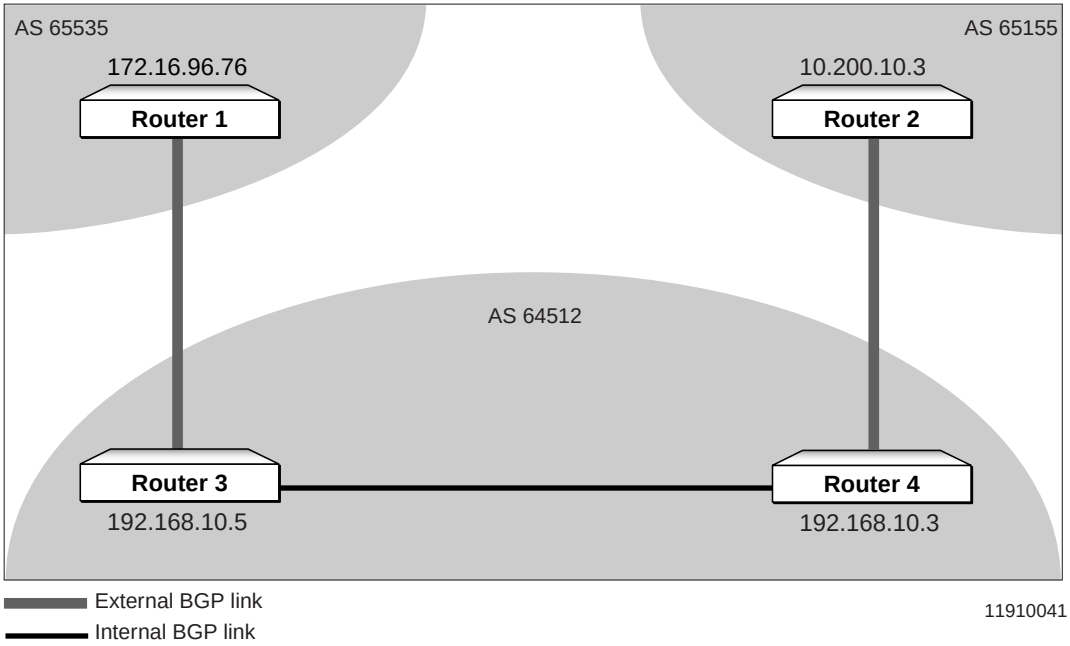

Assumptions:

- Router 1—an external BGP peer—in autonomous system 65535 advertises routes to Router 3 in autonomous system 64512.
- Router 2—an external BGP peer—in autonomous system 65155 advertises routes to Router 4 in autonomous system 64512.
- Routers 3 and 4 are fully meshed PortMaster routers in autonomous system 64512.

Goals:

- Establish internal BGP peer relationships between Routers 3 and 4.
- Each router is always the next hop in the BGP path. All routes are permitted to be accepted, injected, and advertised.

Configuration of Router 3

The complete configuration for Router 3 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.5
Command> set bgp as 64512
Command> add bgp peer 192.168.10.5 172.16.96.76 65535
easy-multihome
Command> add bgp peer 192.168.10.5 192.168.10.3 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp summarization 192.168.10.0/24 as 65535
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 3:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

   ```
   Command> set bgp enable
   BGP will be enabled on the next reboot
   Command> save all
   Command> reboot
   ```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

   ```
   Command> set bgp id 192.168.10.5
   BGP ID changed from 0.0.0.0 to 192.168.10.5
   Command> set bgp as 64512
   BGP AS number changed from 0 to 64512
   ```

3. **Identify Router 1 as an external BGP peer and specify the built-in easy-multihome routing policy:**

```
Command> add bgp peer 192.168.10.5 172.16.96.76 65535 easy-multihome
New BGP peer successfully added
```

4. **Identify Router 4 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.5 192.168.10.3 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

5. **Add summarization for internal routes to be advertised to external peers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65535
BGP summarization successfully added
```

6. **Save the configuration, clear all currently known BGP information, and restart this router with the new configuration:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 4

The complete configuration for Router 4 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.3
Command> set bgp as 64512
Command> add bgp peer 192.168.10.3 10.200.10.3 65155
easy-multihome
Command> add bgp peer 192.168.10.3 192.168.10.5 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp summarization 192.168.10.0/24 as 65535
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 4:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.3
BGP ID changed from 0.0.0.0 to 192.168.10.3
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Identify Router 2 as an external BGP peer and specify the built-in easy-multihome routing policy:**

```
Command> add bgp peer 192.168.10.3 10.200.10.3 65155 easy-multihome
New BGP peer successfully added
```

4. **Identify Router 3 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.3 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

5. **Add summarization for internal routes to be advertised to external peers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65155
BGP summarization successfully added
```

6. **Save the configuration, clear all currently known BGP information, and restart this router with the new configuration:**

```
Command> save all
Command> reset bgp
```

Easy-Multihome—Example 2

Figure 4-5 illustrates a multihome policy routing configuration using the predefined **all** policy.

Figure 4-5 Multihome Configuration with Simple Policies

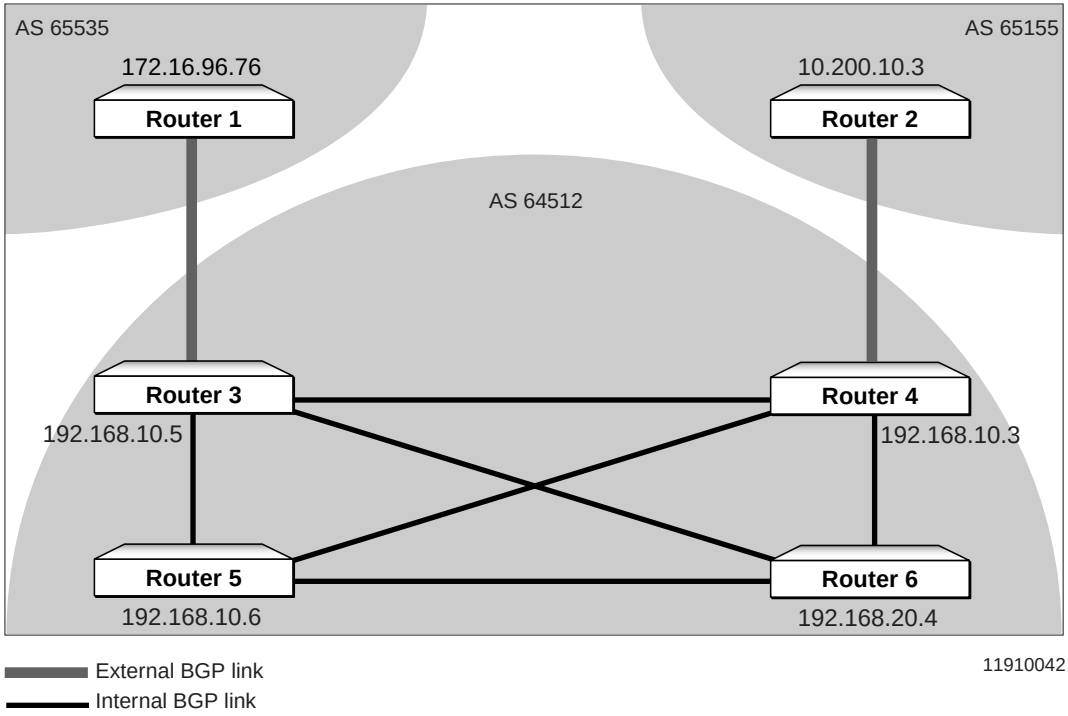

Assumptions:

- Router 1—an external BGP peer—in autonomous system 65535 advertises routes to Router 3 in autonomous system 64512.
- Router 2—an external BGP peer—in autonomous system 65155 advertises routes to Router 4 in autonomous system 64512.
- Routers 3, 4, 5, and 6 are fully meshed PortMaster routers in autonomous system 64512.

Goals:

- Establish internal BGP peer relationships between Routers 3 and 4, Routers 3 and 5, Routers 3 and 6, Routers 4 and 5, Routers 4 and 6, and Routers 5 and 6.
- Each router is always the next hop in the BGP path. All routes are permitted to be accepted, injected, and advertised.

Configuration of Router 3

The complete configuration for Router 3 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.5
Command> set bgp as 64512
Command> add bgp peer 192.168.10.5 172.16.96.76 65535 easy-multihome
Command> add bgp peer 192.168.10.5 192.168.10.3 64512 accept-policy all
inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.5 192.168.10.6 64512 accept-policy all
inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.5 192.168.20.4 64512 accept-policy all
inject-policy all advertise-policy all
Command> add bgp summarization 192.168.10.0/24 as 65535
Command> add bgp summarization 192.168.20.0/24 as 65535
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 3:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.5
BGP ID changed from 0.0.0.0 to 192.168.10.5
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Identify Router 2 as an external BGP peer and specify the built-in easy-multihome routing policy:**

```
Command> add bgp peer 192.168.10.5 172.16.96.76 65535 easy-multihome
New BGP peer successfully added
```

4. **Identify Router 4 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.5 192.168.10.3 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

5. **Identify Router 5 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.5 192.168.10.6 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

6. **Identify Router 6 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.5 192.168.20.4 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

7. **Add summarization for internal routes to be advertised to external peers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65535
BGP summarization successfully added
Command> add bgp summarization 192.168.20.0/24 as 65535
BGP summarization successfully added
```

8. **Save the configuration, clear all currently known BGP information, and restart this router with the new configuration:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 4

The complete configuration for Router 4 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.3
Command> set bgp as 64512
Command> add bgp peer 192.168.10.3 10.200.10.3 65155
easy-multihome
Command> add bgp peer 192.168.10.3 192.168.10.5 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.3 192.168.10.6 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.3 192.168.20.4 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp summarization 192.168.10.0/24 as 65155
Command> add bgp summarization 192.168.20.0/24 as 65155
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 4:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the router as a BGP device and set the autonomous system number:**

```
Command> set bgp id 192.168.10.3
BGP ID changed from 0.0.0.0 to 192.168.10.3
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Identify Router 2 as an external BGP peer and specify the built-in easy-multihome routing policy:**

```
Command> add bgp peer 192.168.10.3 10.200.10.3 65155 easy-multihome
New BGP peer successfully added
```

4. **Identify Router 3 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.3 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

5. **Identify Router 5 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.3 192.168.10.6 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

6. **Identify Router 6 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.3 192.168.20.4 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

7. **Add summarization for internal routes to be advertised to external peers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65155
BGP summarization successfully added
Command> add bgp summarization 192.168.20.0/24 as 65155
BGP summarization successfully added
```

8. **Save the configuration, clear all currently known BGP information, and restart this router with the new configuration:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 5

The complete configuration for Router 5 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.6
Command> set bgp as 64512
Command> add bgp peer 192.168.10.6 192.168.10.5 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.6 192.168.10.3 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.6 192.168.20.4 64512 accept-
policy all inject-policy all advertise-policy all
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 5:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the router as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.6
BGP ID changed from 0.0.0.0 to 192.168.10.6
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Identify Router 3 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.6 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

4. **Identify Router 4 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.6 192.168.10.3 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

5. **Identify Router 6 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.6 192.168.20.4 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

6. **Save the configuration, clear all currently known BGP information, and restart this router with the new configuration:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 6

The complete configuration for Router 6 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.20.4
Command> set bgp as 64512
Command> add bgp peer 192.168.20.4 192.168.10.5 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp peer 192.168.20.4 192.168.10.3 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp peer 192.168.20.4 192.168.10.6 64512 accept-
policy all inject-policy all advertise-policy all
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 6:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

   ```
   Command> set bgp enable
   BGP will be enabled on the next reboot
   Command> save all
   Command> reboot
   ```

2. **Identify the router as a BGP router and set the autonomous system number:**

   ```
   Command> set bgp id 192.168.20.4
   BGP ID changed from 0.0.0.0 to 192.168.20.4
   Command> set bgp as 64512
   BGP AS number changed from 0 to 64512
   ```

3. **Identify Router 3 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.20.4 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

4. **Identify Router 4 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.20.4 192.168.10.3 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

5. **Identify Router 5 as an internal BGP peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.20.4 192.168.10.6 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

6. **Save the configuration, clear all currently known BGP information, and restart this router with the new configuration:**

```
Command> save all
Command> reset bgp
```

Easy-Multihome—Example 3

Figure 4-6 illustrates a fairly complex multihome configuration that establishes advertising policies with community attributes.

Figure 4-6 Multihome Configuration Using Nondefault Policy

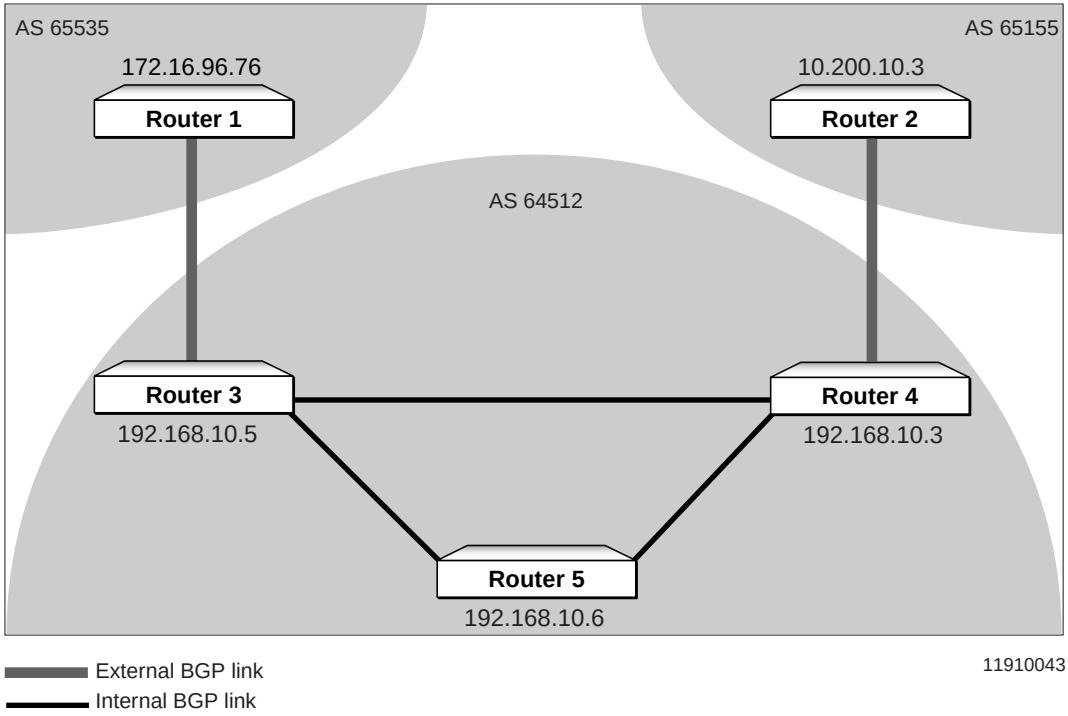

Assumptions:

- Router 1—an external BGP peer to Router 3—in autonomous system 65535 is advertising five routes to Router 3 in autonomous system 64512.
- Router 2—an external BGP peer to Router 4—in autonomous system 65155 is advertising the same five routes to Router 4 in autonomous system 64512.
- Routers 3, 4, and 5 are PortMaster routers in autonomous system 64512.

Goals:

- Ensure that the five routes advertised by Routers 1 and 2 are accepted by Routers 3 and 4.
- Ensure that if one external peer fails, all traffic is routed via the remaining peer. When both external peers are operating again, route preferences resume.
- Establish internal BGP peer relationships among Routers 3, 4, and 5.
- Determine policies to set the routing preferences shown in Table 4-4.

Table 4-4 Routing Preferences

Route Advertised by External Peers—Prefix/NM	**Route Advertised within AS 64512—Prefix/NM**	**Preferred Routing from within AS 64512**
172.32.64.0/24	172.32.64.0/24	via AS 65535
172.80.32.0/24	172.80.32.0/24	via AS 65535
172.112.200.0/24	172.112.200.0/24	via AS 65535
10.108.48.0/24	10.108.0.0/16	via AS 65155
10.108.80.0/24	10.108.0.0/16	via AS 65155

Configuration of Router 3

The complete configuration for Router 3 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.5
Command> set bgp as 64512
Command> add bgp policy adtag105
Command> set bgp policy adtag105 1 permit then community add
105
Command> add bgp policy acdeg10
Command> set bgp policy acdeg10 1 permit then dop 10
Command> add bgp peer 192.168.10.5 172.16.96.76 65535 accept-
policy acdeg10 inject-policy all
Command> add bgp peer 192.168.10.5 192.168.10.3 64512 accept-
policy all inject-policy all advertise-policy adtag105
Command> add bgp peer 192.168.10.5 192.168.10.6 64512
advertise-policy adtag105
Command> add bgp summarization 192.168.10.0/24 as 65535
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 3:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

   ```
   Command> set bgp enable
   BGP will be enabled on the next reboot
   Command> save all
   Command> reboot
   ```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

   ```
   Command> set bgp id 192.168.10.5
   BGP ID changed from 0.0.0.0 to 192.168.10.5
   Command> set bgp as 64512
   BGP AS number changed from 0 to 64512
   ```

3. **Add and set an advertisement policy adtag 105 to label the advertised BGP paths with the community attribute 105:**

```
Command> add bgp policy adtag105
New BGP policy adtag105 successfully added
Command> set bgp policy adtag105 1 permit then community add 105
```

4. **Add and set acceptance policy acdeg10 to assign the degree-of-preference 10 to all accepted BGP paths:**

```
Command> add bgp policy acdeg10
New BGP policy acdeg10 successfully added
Command> set bgp policy acdeg10 1 permit then dop 10
BGP policy acdeg10 updated
```

5. **Set Router 1 as an external peer with appropriate options to control policy implementation:**

```
Command> add bgp peer 192.168.10.5 172.16.96.76 65535 accept-policy
acdeg10 inject-policy all
New BGP peer successfully created
```

6. **Set Router 4 and Router 5 as internal peers with appropriate options to control policy implementation:**

```
Command> add bgp peer 192.168.10.5 192.168.10.3 64512 accept-policy all
inject-policy all advertise-policy adtag105
New BGP peer successfully created
Command> add bgp peer 192.168.10.5 192.168.10.6 64512 advertise-policy
adtag105
New BGP peer successfully created
```

7. **Add summarization for internal routes to be advertised to external peers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65535
BGP summarization successfully created
```

8. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 4

The complete configuration for Router 4 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.3
Command> set bgp as 64512
Command> add bgp policy adtag108
Command> set bgp policy adtag108 1 permit if prefix 10.108.0.0/16 then
community add 108
Command> set bgp policy adtag108 2 permit then lp 5 community add 108
Command> add bgp policy acdeg29
Command> set bgp policy acdeg29 1 permit then dop 29
Command> add bgp peer 192.168.10.3 10.200.10.3 65155 accept-policy
acdeg29 inject-policy all
Command> add bgp peer 192.168.10.3 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy adtag108
Command> add bgp peer 192.168.10.3 192.168.10.6 64512 advertise-policy
adtag108
Command> add bgp summarization 192.168.10.0/24 as 65155
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 4:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.3
BGP ID changed from 0.0.0.0 to 192.168.10.3
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Add and set an advertisement policy adtag108 to label the advertised BGP paths with a community attribute of 108:**

```
Command> add bgp policy adtag108
New BGP policy acdeg108 successfully added
Command> set bgp policy adtag108 1 permit if prefix 10.108.0.0/16 then
community add 108
BGP polity adtag108 updated
```

4. **Add a rule to advertisement policy adtag108 to lower the local preference to 5 on all routes except those included in the IP prefix 10.108.0.0/16:**

```
Command> set bgp policy adtag108 2 permit then lp 5 community add 108
BGP polity adtag108 updated
```

5. **Add and set an acceptance policy acdeg29 to assign a degree of preference of 29 to all accepted BGP paths:**

```
Command> add bgp policy acdeg29
New BGP policy acdeg29 successfully added
Command> set bgp policy acdeg29 1 permit then dop 29
BGP policy acdeg29 updated
```

6. **Set Router 2 as an external peer with appropriate options to control policy implementation:**

```
Command> add bgp peer 192.168.10.3 10.200.10.3 65155 accept-policy acdeg29
inject-policy all
New peer successfully created
```

7. **Set Router 3 and Router 5 as internal peers with appropriate options to control policy implementation:**

```
Command> add bgp peer 192.168.10.3 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy adtag108
New BGP peer successfully created
Command> add bgp peer 192.168.10.3 192.168.10.6 64512 advertise-policy
adtag108
New BGP peer successfully created
```

8. **Add summarization for internal routes to be advertised to external providers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65155
BGP summarization successfully created
```

9. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 5

The complete configuration for Router 5 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.6
Command> set bgp as 64512
Command> add bgp peer 192.168.10.6 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.6 192.168.10.3 64512 accept-policy all
inject-policy all advertise-policy all
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 5:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.6
BGP ID changed from 0.0.0.0 to 192.168.10.6
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Set Router 3 and Router 4 as internal peers and specify that they use the predefined** all **policy to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.6 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully created
Command> add bgp peer 192.168.10.6 192.168.10.3 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully created
```

4. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Route Reflectors Example

Figure 4-7 illustrates a BGP configuration using route reflectors.

Figure 4-7 Route Reflector Configuration

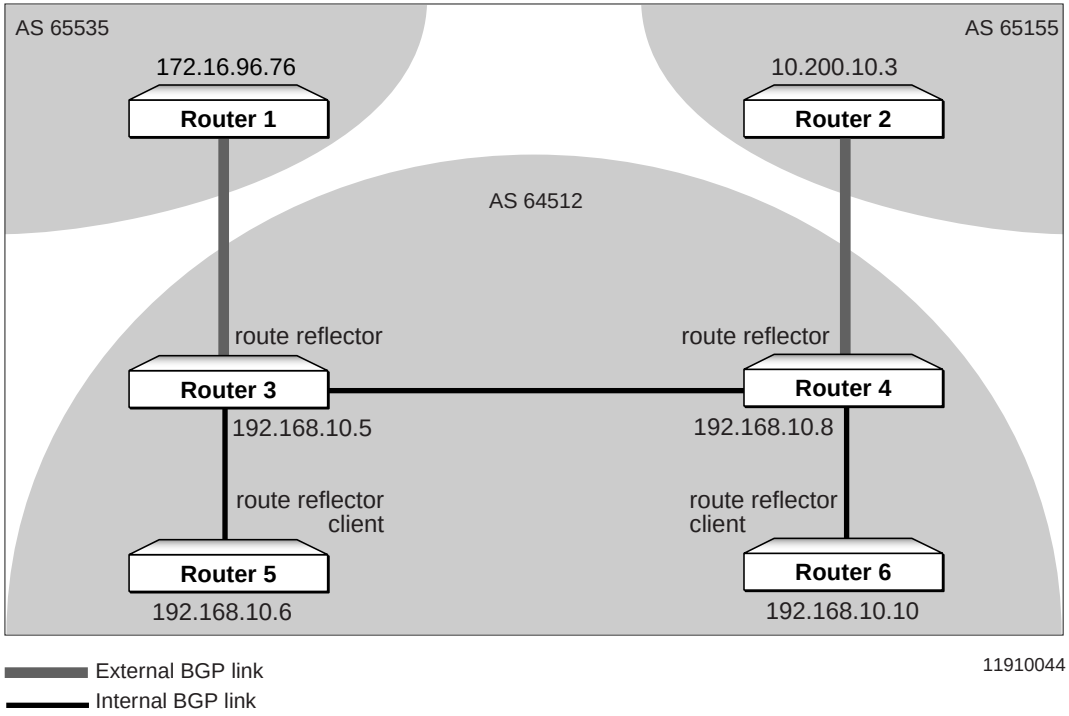

Assumptions:

- Router 1—an external BGP peer—in autonomous system 65535 advertises routes to Router 3 in autonomous system 64512.
- Router 2—an external BGP peer—in autonomous system 65155 advertises routes to Router 4 in autonomous system 64512.
- Routers 3, 4, 5, and 6 are PortMaster routers in autonomous system 64512.

Goals:

- Establish Routers 3 and 4 as route reflectors sharing an internal BGP peer relationship.
- Establish Router 5 as a client of route reflector Router 3.
- Establish Router 6 as a client of route reflector Router 4.
- All routes are permitted to be accepted, injected, and advertised.

Configuration of Router 3

The complete configuration for Router 3 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.5
Command> set bgp as 64512
Command> add bgp peer 192.168.10.5 172.16.96.76 65535 accept-
policy all inject-policy all advertise-policy all
Command> set bgp cluster-id 192.168.10.5
Command> reset bgp
Command> add bgp peer 192.168.10.5 192.168.10.8 64512 accept-
policy all inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.5 192.168.10.6 64512 accept-
policy all inject-policy all advertise-policy all route-
reflector-client
Command> add bgp summarization 192.168.10.0/24 as 65535
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 3:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.5
BGP ID changed from 0 to 192.168.10.5
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Identify Router 1 as an external BGP peer and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.5 172.16.96.76 65535 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

4. **Set the route cluster identifier:**

```
Command> set bgp cluster-id 192.168.10.5
BGP Cluster ID changed from 0.0.0.0 to 192.168.10.5
```

5. **Establish a peer relationship between the two route reflectors:**

```
Command> add bgp peer 192.168.10.5 192.168.10.8 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully created
```

6. **Set the internal peer that is a client to the route reflector:**

```
Command> add bgp peer 192.168.10.5 192.168.10.6 64512 accept-policy all
inject-policy all advertise-policy all route-reflector-client
New BGP peer successfully added
```

7. **Add summarization for internal routes to be advertised to external providers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65535
BGP summarization successfully added
```

8. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 4

The complete configuration for Router 4 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.8
Command> set bgp as 64512
Command> add bgp peer 192.168.10.8 10.200.10.3 65155 accept-policy all
inject-policy all advertise-policy all
Command> set bgp cluster-id 192.168.10.8
Command> add bgp peer 192.168.10.8 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.8 192.168.10.10 64512 accept-policy all
inject-policy all advertise-policy all route-reflector-client
Command> add bgp summarization 192.168.10.0/24 as 65155
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 4:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.8
BGP ID changed from 0 to 192.168.10.8
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Identify Router 2 as an external BGP peer and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.8 10.200.10.3 65155 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

4. **Set the route cluster identifier:**

```
Command> set bgp cluster-id 192.168.10.8
BGP Cluster ID changed from 0.0.0.0 to 192.168.10.8
```

5. **Establish a peer relationship between the two route reflectors:**

```
Command> add bgp peer 192.168.10.8 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

6. **Set the internal peer that is a client to the route reflector:**

```
Command> add bgp peer 192.168.10.8 192.168.10.10 64512 accept-policy all
inject-policy all advertise-policy all route-reflector-client
New BGP peer successfully created
```

7. **Add summarization for internal routes to be advertised to external providers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65155
BGP summarization successfully added
```

8. **Reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 5

The complete configuration for Router 5 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.6
Command> set bgp as 64512
Command> add bgp peer 192.168.10.6 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 5:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.6
BGP ID changed from 0.0.0.0 to 192.168.10.6
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Set the internal peer that is the route reflector for this client router:**

```
Command> add bgp peer 192.168.10.6 192.168.10.5 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

4. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 6

The complete configuration for Router 6 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.10
Command> set bgp as 64512
Command> add bgp peer 192.168.10.10 192.168.10.8 64512 accept-policy all
inject-policy all advertise-policy all
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 6:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.10
BGP ID changed from 0.0.0.0 to 192.168.10.64
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Set the internal peer that is the route reflector for this client router:**

```
Command> add bgp peer 192.168.10.10 192.168.10.8 64512 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

4. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Confederations Example

Figure 4-8 illustrates a BGP configuration using confederation member autonomous systems (CMAS).

Figure 4-8 Confederation Configuration

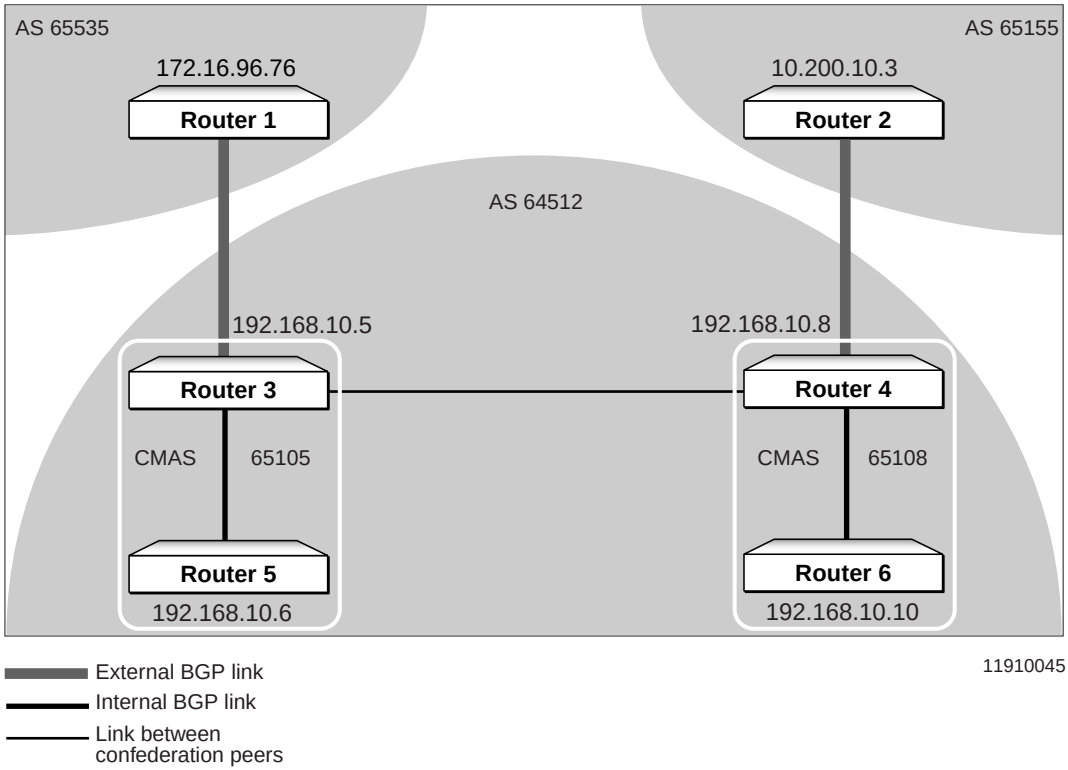

Assumptions:

- Router 1—an external BGP peer—in autonomous system 65535 advertises routes to Router 3 in autonomous system 64512.
- Router 2—an external BGP peer—in autonomous system 65155 advertises routes to Router 4 in autonomous system 64512.
- Routers 3, 4, 5, and 6 are PortMaster routers in autonomous system 64512.

Goals:

- Establish CMAS 65105 with Routers 3 and 5 as members.
- Establish CMAS 65108 with Routers 4 and 6 as members.
- Establish a peer relationship between the two CMASs.
- All routes are permitted to be accepted, injected, and advertised.

Configuration of Router 3

The complete configuration for Router 3 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.5
Command> set bgp as 64512
Command> add bgp peer 192.168.10.5 172.16.96.76 65535 accept-policy all
inject-policy all advertise-policy all
Command> set bgp cma 65105
Command> add bgp peer 192.168.10.5 192.168.10.8 65108 accept-policy all
inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.5 192.168.10.6 65105 accept-policy all
inject-policy all advertise-policy all confederation-member
Command> add bgp summarization 192.168.10.0/24 as 65535
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 3:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.5
BGP ID changed from 0.0.0.0 to 192.168.10.5
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Set Router 1 as an external peer, and specify that it use the predefined policy with the name** all **to permit all routes to be accepted, injected, and advertised:**

```
Command> add bgp peer 192.168.10.5 172.16.96.76 65535 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

4. **Set membership in CMAS 65105:**

```
Command> set bgp cma 65105
BGP Confederation member AS number changed from 0 to 65105
```

5. **Establish a peer relationship between CMAS 65105 and CMAS 65108:**

```
Command> add bgp peer 192.168.10.5 192.168.10.8 65108 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

6. **Establish the internal peer relationship within CMAS 65105:**

```
Command> add bgp peer 192.168.10.5 192.168.10.6 65105 accept-policy all
inject-policy all advertise-policy all confederation-member
New BGP peer successfully added
```

7. **Add summarization for internal routes to be advertised to external providers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65535
BGP summarization successfully added
```

8. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 4

The complete configuration for Router 4 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.8
Command> set bgp as 64512
Command> add bgp peer 192.168.10.8 10.200.10.3 65155 accept-policy all
inject-policy all advertise-policy all
Command> set bgp cma 65108
Command> add bgp peer 192.168.10.8 192.168.10.5 65105 accept-policy all
inject-policy all advertise-policy all
Command> add bgp peer 192.168.10.8 192.168.10.10 65108 accept-policy all
inject-policy all advertise-policy all confederation-member
Command> add bgp summarization 192.168.10.0/24 as 65155
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 4:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.8
BGP ID changed from 0.0.0.0 to 192.168.10.5
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Set Router 2 as an external peer:**

```
Command> add bgp peer 192.168.10.8 10.200.10.3 65155 accept-policy all
inject-policy all advertise-policy all
BGP peer successfully added
```

4. **Set membership in CMAS 65108:**

```
Command> set bgp cma 65108
BGP Confederation member AS number changed from 0 to 65108
```

5. **Establish a peer relationship between CMAS 65108 and CMAS 65105:**

```
Command> add bgp peer 192.168.10.8 192.168.10.5 65105 accept-policy all
inject-policy all advertise-policy all
New BGP peer successfully added
```

6. **Establish the internal peer relationship within CMAS 65108:**

```
Command> add bgp peer 192.168.10.8 192.168.10.10 65108 accept-policy all
inject-policy all advertise-policy all confederation-member
New BGP peer successfully added
```

7. **Add summarization for internal routes to be advertised to external providers:**

```
Command> add bgp summarization 192.168.10.0/24 as 65155
BGP summarization successfully added
```

8. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 5

The complete configuration for Router 5 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.6
Command> set bgp as 64512
Command> set bgp cma 65105
Command> add bgp peer 192.168.10.6 192.168.10.5 65105 accept-
policy all inject-policy all advertise-policy all
confederation-member
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 5:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

   ```
   Command> set bgp enable
   BGP will be enabled on the next reboot
   Command> save all
   Command> reboot
   ```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

   ```
   Command> set bgp id 192.168.10.6
   BGP ID changed from 0.0.0.0 to 192.168.10.5
   Command> set bgp as 64512
   BGP AS number changed from 0 to 64512
   ```

3. **Set membership in CMAS 65105:**

   ```
   Command> set bgp cma 65105
   BGP Confederation member AS number changed from 0 to 65105
   ```

4. **Establish the internal peer relationship within CMAS 65105:**

```
Command> add bgp peer 192.168.10.6 192.168.10.5 65105 accept-policy all
inject-policy all advertise-policy all confederation-member
New BGP peer successfully added
```

5. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Configuration of Router 6

The complete configuration for Router 6 appears in the following box:

```
Command> set bgp enable
Command> save all
Command> reboot
Command> set bgp id 192.168.10.10
Command> set bgp as 64512
Command> set bgp cma 65108
Command> add bgp peer 192.168.10.10 192.168.10.8 65108 accept-
policy all inject-policy all advertise-policy all
confederation-member
Command> save all
Command> reset bgp
```

The following procedure shows the separate tasks for configuring Router 6:

1. **Enable the use of BGP on the PortMaster, save the configuration, and reboot the router for changes to take effect:**

```
Command> set bgp enable
BGP will be enabled on the next reboot
Command> save all
Command> reboot
```

2. **Identify the PortMaster as a BGP router and set the autonomous system number:**

```
Command> set bgp id 192.168.10.10
BGP ID changed from 0.0.0.0 to 192.168.10.10
Command> set bgp as 64512
BGP AS number changed from 0 to 64512
```

3. **Set membership in CMAS 65108:**

```
Command> set bgp cma 65108
BGP Confederation member AS number changed from 0 to 65108
```

4. **Establish the internal peer relationship within CMAS 65108:**

```
Command> add bgp peer 192.168.10.10 192.168.10.8 65108 accept-policy all
inject-policy all advertise-policy all confederation-member
New BGP peer successfully added
```

5. **Save the configuration and reset BGP routing on the PortMaster:**

```
Command> save all
Command> reset bgp
```

Troubleshooting A

This appendix offers case studies of some typical routing problems experienced by Lucent customers and the solutions provided by the Lucent technical support team. The chapter begins with a discussion of troubleshooting programs you can use to help resolve routing problem.

This chapter discusses the following topics:

Troubleshooting Tools

This section describes some of the most useful tools for troubleshooting a routing problem or a suspected routing problem. For more information on these and other network troubleshooting or debugging tools, refer to the *Configuration Guide for PortMaster Products,* the *Command Line Administrator's Guide,* TCP/IP textbooks, and your host's system administration manuals.

Troubleshooting tools discussed in this section include the following:

Ping

The **ping** command, available as part of the PortMaster command set, sends 10 Internet Control Message Protocol (ICMP) echo request packets to the target and listens for an ICMP echo reply. The **ping** command, also part of the UNIX operating system (and other operating systems), is fairly common on networked hosts and on routers. Because it is simple to use and checks end-to-end connectivity, **ping** is generally the best first choice to use for isolating and resolving a routing problem.

The PortMaster **ping** command is used as shown (with the system response) in the following example:

```
Command> ping 192.168.16.0
192.168.16.0 is alive
```

If the **ping** sent to the target is not returned, the response is

```
no answer from 192.168.16.0
```

When used with the **-s** option, some versions of **ping** return statistics on packets sent and the time elapsed for the **ping** round trip, as in the following abbreviated example:

```
Command> ping -s 192.168.16.0
PING 192.168.16.0: 56 data bytes
64 bytes from 192.168.16.0: icmp_seq=0 ttl=255 time=12 ms
64 bytes from 192.168.16.0: icmp_seq=1 ttl=255 time=12 ms
64 bytes from 192.168.16.0: icmp_seq=2 ttl=255 time=12 ms
```

Ptrace

This PortMaster command allows you to see packet information as it passes through a PortMaster product. You use filters to define which packets you want to display. **Ptrace** does not display ICMP or User Datagram Protocol (UDP) packets that originate on the PortMaster itself.

The **ptrace** command is frequently used in conjunction with **ping**. Typically, you open a Telnet session into a PortMaster, activate the **ptrace** command with an appropriate protocol filter, such as TCP, ICMP, UDP, or IPX, and trace the **ping** packets as they arrive at or pass through the PortMaster. The following example shows how to configure a PortMaster product to use **ptrace** with an ICMP packet filter:

```
Command> add filter test
New Filter successfully added

Command> set filter test 1 permit icmp
Filter test updated

Command> set console
Setting CONSOLE to admin session

Command> ptrace test
Packet Tracing Enabled
```

Ptrace output of pings arriving at a PortMaster appear as follows:

```
icmp from 192.168.148.1 to 10.0.0.3 type Echo Request
icmp from 192.168.148.1 to 10.0.0.3 type Echo Request
icmp from 192.168.148.1 to 10.0.0.3 type Echo Request
```

Ptrace output of pings passing through a PortMaster appear as follows:

```
icmp from 192.168.148.1 to 10.0.0.15 type Echo Request
icmp from 10.0.0.15 to 192.168.148.1 type Echo Reply
```

To stop the **ptrace**, issue the following commands:

```
Command> ptrace
Packet Tracing Disabled

Command> reset console
Console RESET
```

Traceroute

Traceroute is another useful troubleshooting tool included in the PortMaster command set. It is also widely available with UNIX and other operating systems. While **ping** checks only the end-to-end connectivity, **traceroute** tries to identify each hop in the path to a destination by sending IP packets with the *time-to-live* (TTL) set incrementally from 1 to a maximum of 30. It prints the addresses that return ICMP *TTL expired* packets.

Use the **traceroute** command as shown, with response, in the following example:

```
Command> traceroute 172.16.1.2
traceroute to (172.16.1.2), 30 hops max
1 192.168.96.2
2 192.168.1.3
3 172.16.1.2
```

Traceroute (UDP) takes its source address from the interface through which it exits, while **ping** (ICMP) takes its source address from the Ether0 interface. When you use **traceroute** and **ping** together, if **traceroute** reaches its target and **ping** does not (or vice versa), check routing to see if the two addresses are being routed differently. This is a common routing problem over Frame Relay connections

Case Studies

This section includes the following case studies:

Host Routing versus Network Routing

This case illustrates some common problems that arise when you configure host routes and network routes. The discussion assumes you are running ComOS 3.5 or later and that you have set the user-netmask to **on** (see **set user-netmask** in the *Command Line Administrator's Guide*).

As shown in Figure A-1, an Internet service provider (ISP) using a PortMaster 3 Integrated Access Server (labeled Router1) receives dial-up connections from a remote user on a PC and from another PortMaster (labeled Router2) serving a corporate LAN. Typical problems in this situation include the following:

- Incorrect netmask for the router (Router2) serving the corporate LAN
- Incorrect entries on the RADIUS server for the remote PC dial-up user or the router (Router2) serving the corporate LAN
- PortMaster not set to support user netmasks

Note – This discussion assumes the router serving the corporate LAN and the remote PC user are correctly configured.

Figure A-1 Host Routing versus Network Routing

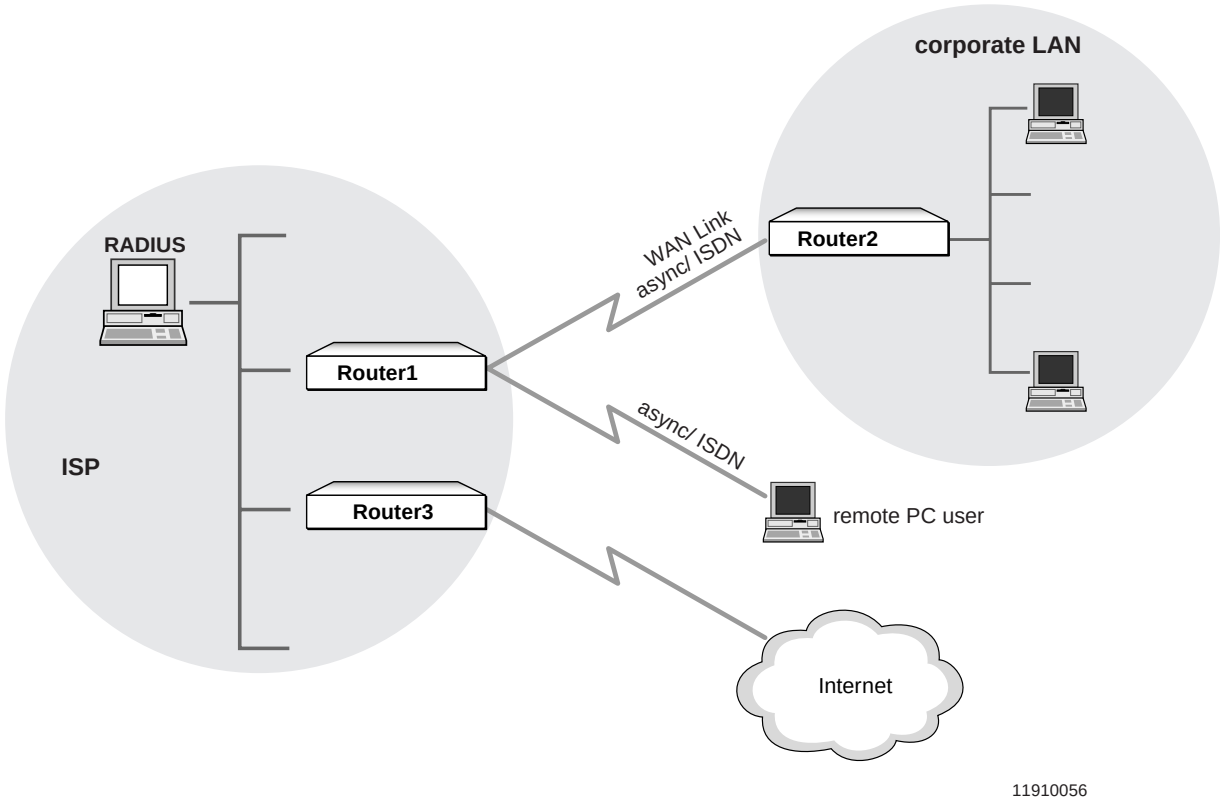

Symptoms

Symptoms experienced by Router2 on the corporate LAN or the remote PC user due to configurations problems described previously include the following:

- Corporate LAN administrator has verified connection to the ISP, but machines on the corporate LAN have no access beyond Router2
- Remote PC user connects to Router1 at the ISP but the connection is very slow or drops, or the PC fails
- ISP administrator can reach Router2 but nothing beyond it

Diagnosis

To verify that the remote PC user or Router2 is in fact connected, the ISP administrator can send a **ping** or **traceroute** command to the device. A **show all**, **show ports**, or **show sessions** command can also reveal whether these devices are connected.

If the administrator can verify that the connection is active, a **show routes** command or a **show table user** command entered on Router1 (if users are added locally) provides information about the nature of the session and the netmask of the device on the other end. Information about users maintained in the RADIUS database can be viewed on the RADIUS server.

For the remote PC user, the netmask in the user entry (locally or on the RADIUS server) must be /32 (255.255.255.255) to indicate that it is a single host IP address. When the PC establishes a connection with Router1, a **show routes** command on Router1 must display a host local (HL) flag associated with the route to the PC.

For Router2, the netmask in the user entry (locally or on the RADIUS server) must reflect the size of the network behind it. When Router2 establishes a connection with Router1, a **show routes** command on Router1 must display a network local (NL) flag associated with the route to the Router2.

If the netmask for Router2 is /32 netmask (255.255.255.255), Router1 might treat Router2 as a host. If it does, Router1 will not route to any of the hosts on Router2's network. Router2, also, will not route for any of the hosts on its network; it will route only for itself. Conversely, anything other than a /32 netmask for the remote user on the PC can cause Router1 to treat the PC as a router with a network behind it.

Solution

If users are maintained locally, set the appropriate netmask for Router2 on Router1, and enter the command **set user-netmask on** to help Router1 distinguish a host from a router. For the remote PC user, set the user netmask to /32 (255.255.255.255). Because the netmask table defaults to a full class C, you must set the user netmask when adding users locally. (See **set user-netmask** in the *Command Line Administrator's Guide*.)

If users are maintained in a RADIUS database, use the Framed-IP-Netmask reply item to configure the appropriate netmask associated with Router2 (RADIUS defaults to a /32 netmask). Refer to the *RADIUS Administrator's Guide* for instructions about setting Framed-IP-Netmask. Because the RADIUS server defaults to a 255.255.255.255 netmask, you need not set a netmask in the user entry for the dial-up PC.

If the netmask for the PC is not 255.255.255.255, and user-netmask is **on**, the ISP will experience severe Proxy ARP and routing problems on Router1. For effective network operation, the remote PC must be configured with a /32 netmask, and the netmask for Router2 must be correctly configured on the RADIUS server.

Note – set user netmask and **set user-netmask** are separate commands.

Configuring the Gateway

This case illustrates the kinds of problems experienced when the gateway router is improperly configured.

Figure A-2 shows the network structure for this case.

Figure A-2 Configuring the Gateway

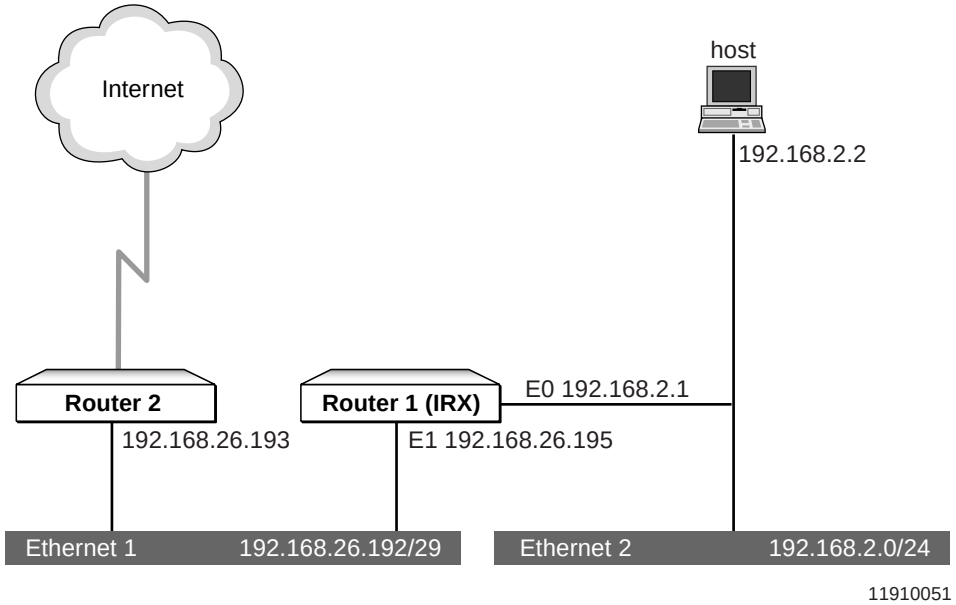

Configuration Information

The specifics of this configuration are as follows:

Ethernet 2	Address 192.168.2.0 Netmask 255.255.255.0
Ethernet 1	Address 192.168.26.192 Netmask 255.255.255.248
Router 1 (IRX 211)	Ether0 address 192.168.2.1 Ether1 address 192.168.26.195
	RIP set on (broadcast and listen) at Ether0 interface
Router 2	Ether0 address 192.168.26.193
	RIP set off on Ether0 interface Static route to Ethernet 2 via gateway on Router 1

Symptom

The host on Ethernet 2 is not able to communicate with the Internet

Diagnosis

From Host, a **ping** to Router 1 shows that Router 1 is active. A **ping** to Router 2 generates no response. A **ptrace** with an ICMP packet filter set up via Telnet to Router 1 shows the **ping** from Host to Router 1 and the **ping** from Host to Router 2, but no return **ping** from Router 2 to Host. This behavior points to a problem with the routing setup on Router 2.

When the source of the problem is isolated as shown here, a **show routes** command entered on Router 2, provides the following (partial) response:

```
Command> show routes
Destination         Gateway          Flag     Met     Interface
-----------------   --------------   -----    ----    ------------
192.168.2.0         192.168.20.195   NS       1       ether0
                    ^
```

The output clearly shows that in the static route the destination (192.168.2.0) is correct, but the gateway route was entered incorrectly as 192.168.20.195 instead of 192.168.26.195, the Ether1 interface address on Router 1.

Solution

The solution to this problem is to issue the following commands on Router 2:

```
Command> delete route 192.168.2.0
Route successfully deleted

Command> add route 192.168.2.0 192.168.26.195 1
New route entry successfully added

Command> save all
```

After completing the configuration change, repeating the **ping** and **ptrace** testing confirmed that routing was functioning correctly.

Configuring Subnets

This case illustrates the kinds of problems experienced because of invalid use of subnetting.

Figure A-3 shows the network structure for this case.

Figure A-3 Configuring Subnets

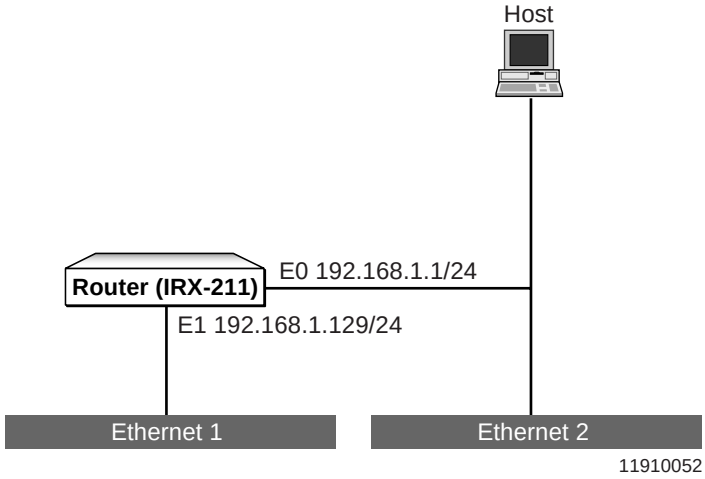

Configuration Information

The configuration information in this case is as follows:

Router (IRX 211)	Ether0 address 192.168.1.1 Netmask 255.255.255.0
	Ether1 address 192.168.1.129 Netmask 255.255.255.0

Symptoms

The PortMaster IRX-211 cannot not communicate with Host, and Host cannot communicate with Ethernet 1.

Solution

An examination of the PortMaster IRX-211 configuration shows that the two Ethernet interfaces are on the same subnet of the same network. This is an invalid configuration; the two Ethernet interfaces must be on different networks or subnets.

One solution is to use a 26-bit netmask to divide the Class C IP address into four subnets of 62 host addresses each. In this way, two of the available subnets are used, one for Ethernet 2 and one for Ethernet 1, leaving an additional two subnets for future use. The IP addresses originally assigned to the Ether0 and Ether1 interfaces are then in different subnets and need not be renumbered.

Correct the netmasks for the interfaces with the following commands:

```
Command> set ether0 netmask 255.255.255.192
ether0 netmask changed from 255.255.255.0 to 255.255.255.192

Command> set ether1 netmask 255.255.255.192
ether1 netmask changed from 255.255.255.0 to 255.255.255.192

Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
Netmasks table successfully saved
Modem table successfully saved
New configurations successfully saved.

Command> reboot
```

Note – You must reboot the PortMaster router to rebuild the routing table.

Configuring Unnumbered Interfaces

This case involves some confusion about IP numbered and unnumbered interfaces. Figure A-4 shows the network structure for this case.

Figure A-4 Configuring Unnumbered Interfaces

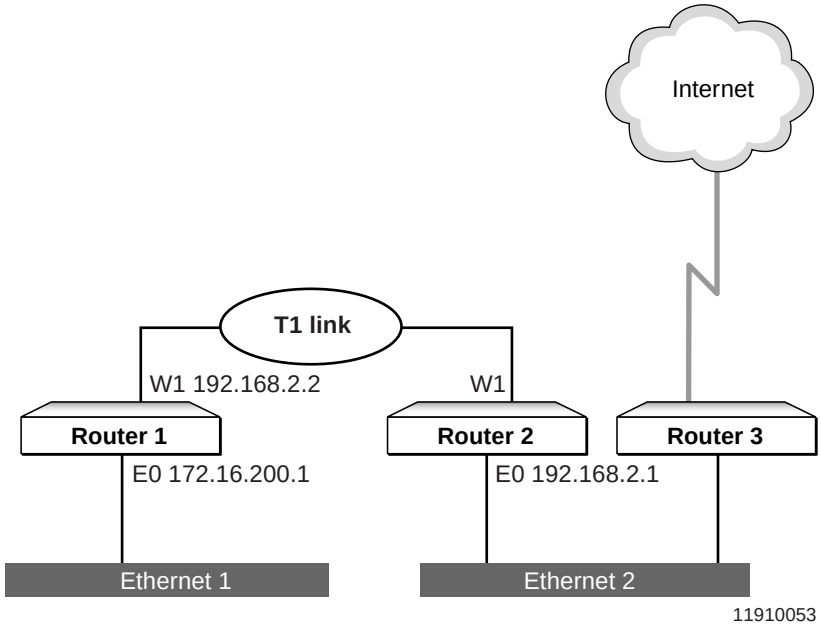

Configuration Information

Device configurations in this case are as follows:

Router 1 (PortMaster)	Ether0 address 172.16.200.1 Netmask 255.255.255.0
	Port W1 address 192.168.2.2
Router 2 (PortMaster)	Ether0 address 192.168.2.1 Netmask 255.255.255.0
T1 link	Hardwired connection between Routers 1 and 2

Symptoms

Router 1 and Router 2 cannot communicate.

Solution

An examination of the configurations shows that on Router 1 the W1 port is set to create a numbered interface, while the W1 port on Router 2 is not configured to create an interface.

In this situation, the solution is to configure unnumbered interfaces for the W1 ports on Router 1 and Router 2 by entering a series of commands on both routers.

Use the following commands on Router 1:

```
Command> set w1 address
Port W1 address changed from 192.168.2.2 to 0.0.0.0

Command> set w1 netmask 0.0.0.0
W1 netmask changed from 0.0.0.0 to 0.0.0.0

Command> set w1 destination 192.168.2.1 255.255.255.0
Port destination changed from 0.0.0.0 to 192.168.2.1
```

```
Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
Netmasks table successfully saved
Modem table successfully saved
New configurations successfully saved.

Command> reboot
```

Use the following commands on Router 2:

```
Command> set w1 address
Port W1 address changed from 0.0.0.0 to 0.0.0.0

Command> set w1 netmask 0.0.0.0
W1 netmask changed from 0.0.0.0 to 0.0.0.0

Command> set w1 destination 172.16.200.1 255.255.255.0
Port destination changed from 0.0.0.0 to 172.16.200.1

Command> save all
Saving global configuration
Saving ports
User table successfully saved
Hosts table successfully saved
Static route table successfully saved
Location table successfully saved
SNMP table successfully saved
Filter table successfully saved
Netmasks table successfully saved
Modem table successfully saved
New configurations successfully saved.

Command> reboot
```

Propagating OSPF over a WAN Link

In this case study, two Ethernet networks in two OSPF areas are incorrectly connected over a WAN link.

To conserve IP addresses, administrators typically do not assign IP addresses to the device serial interfaces on the WAN link—in this case, between OR-HS2 and OR-HS1. Instead, the destination address for the serial interface on each router is set to the Ether0 address of the router on the opposite end of the link. For example, in Figure A-5 the Ethernet port 10.0.0.2 is set as the destination for S1 on OR-HS2 and the Ethernet port 192.168.1.1 is set as the destination for OR-HS1.

OSPF is enabled on the S1 and Ethernet interfaces on both devices, and the OSPF area ranges are set as shown in "Configuration Information."

Note – Although this example uses Lucent OR-HS devices on Ethernet, this problem can occur on any network type or router in a similar configuration.

Figure A-5 Propagating OSPF over a WAN Link (Incorrect)

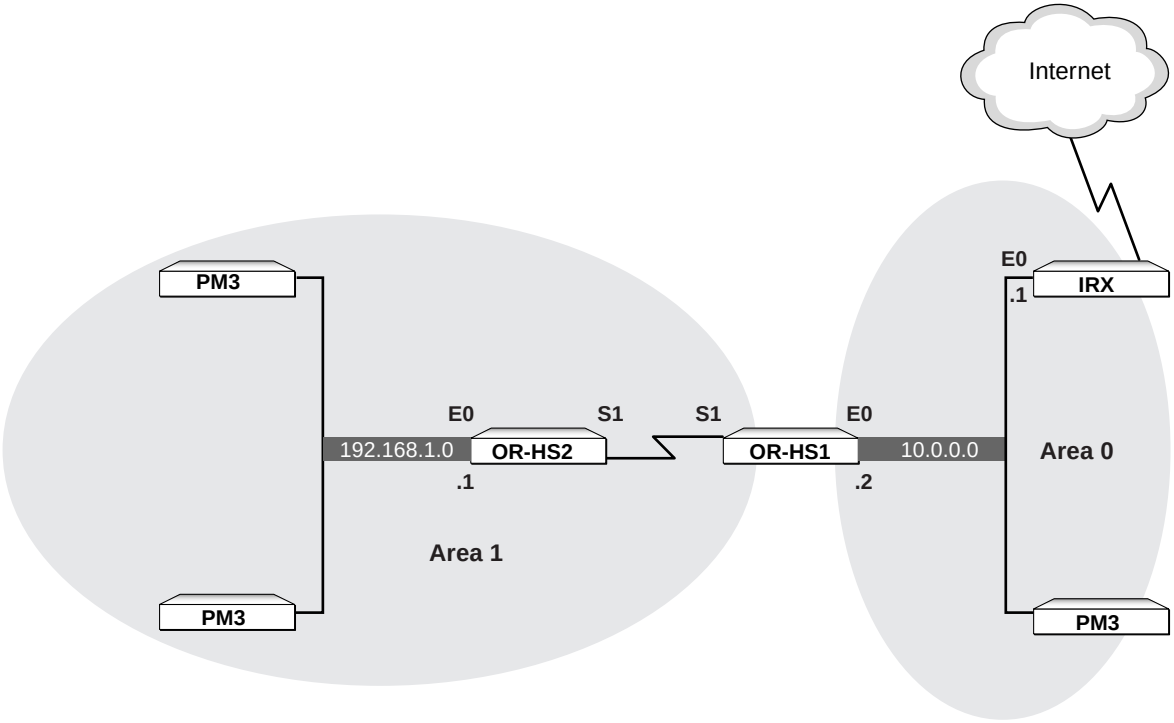

Configuration Information

The incorrect configuration settings for OR-HS2 appear in the following box:

```
Command> set S1 destination 10.0.0.2
Command> set ospf enable
Command> save all
Command> reboot
Command> add ospf area 1
Command> set ospf area 1 range 192.168.1.0/24
Command> set S1 ospf on
Command> set ether0 ospf on
Command> reset S1
```

The incorrect configuration settings for OR-HS1 appear in the following box:

```
Command> set S1 destination 192.168.1.1
Command> set ospf enable
Command> save all
Command> reboot
Command> add ospf area 0
Command> set ospf area 0 range 10.0.0.0/8
Command> add area 1
Command> set ospf area 1 range 192.168.1.0/24
Command> set S1 ospf on
Command> set ether0 ospf on
Command> reset S1
```

Symptoms

OSPF is not being propagated between Area 1 and Area 0. Output of the **show ospf neighbo**r command on either of the OR-HS devices reveals that neither is the OSPF neighbor of the other, and the **if config** command shows that OSPF is not enabled on serial ports.

Solution

When you enable OSPF, it applies an area to each port that falls within the range you set for it. Because the S1 port on OR-HS2 is directly connected to the 10.0.0.0 network (which is its destination), it is outside of the range set for Area 1, and OSPF cannot assign an area to the S1 port.

To correct this problem, you assign an IP address to the S1 port on OR-HS2 (192.168.2.1), and you assign an IP address to the S1 port on OR-HS1 (192.168.2.2). Now the address of the S1 port on OR-HS2 falls within the range of Area 1, so it is assigned to the area; and the address of the S1 port on OR-HS1 also falls within the range of Area 1, so it, too, is assigned to the area. These two routers can now negotiate OSPF back and forth across the WAN link.

Figure A-6 illustrates this address assignment scheme and is followed by the correct configuration for each router in this example. See also “Propagating OSPF over a Single WAN Link” on page 3-12 for a full description of the configuration.

Figure A-6 Propagating OSPF over a WAN Link (Correct)

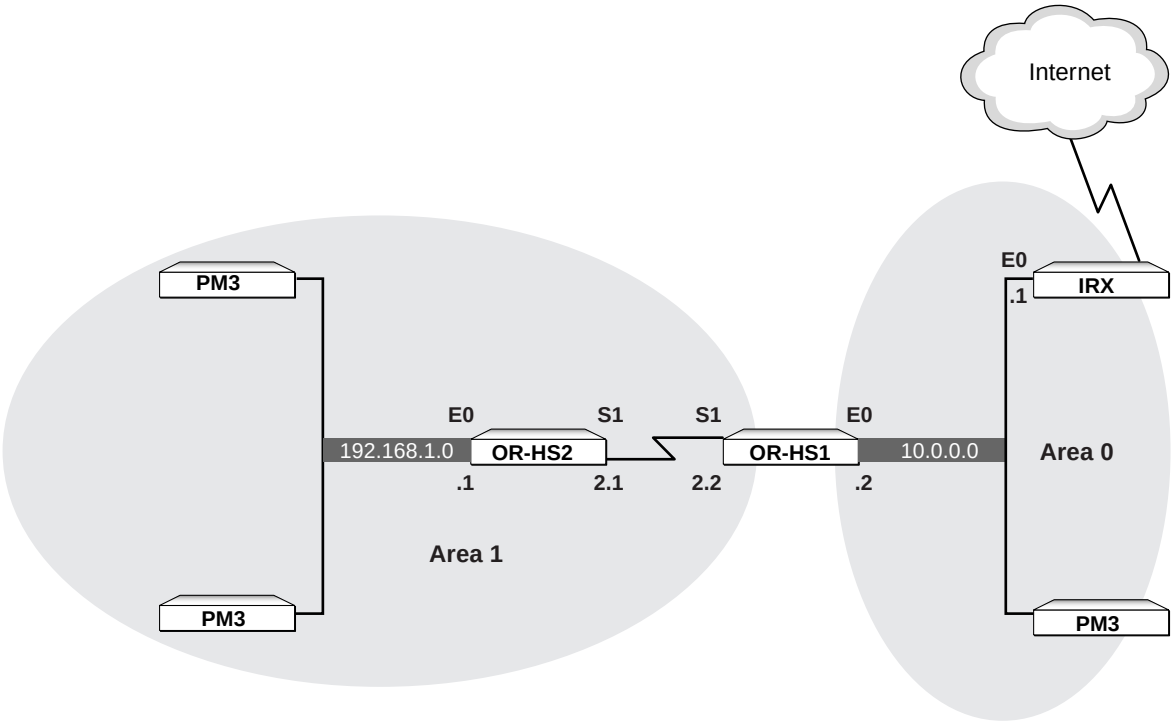

The correct settings for OR-HS2 appear in the following box:

```
Command> set S1 address 192.168.2.1
Command> set S1 destination 192.168.2.2
Command> set ospf enable
Command> save all
Command> reboot
Command> add ospf area 1
Command> set ospf area 1 range 192.168.2.0/24
Command> set ospf area 1 range 192.168.1.0/24
Command> set S1 ospf on
Command> set ether0 ospf on
Command> reset S1
```

The correct settings for OR-HS1 appears in the following box:

```
Command> set S1 address 192.168.2.2
Command> set S1 destination 192.168.2.1
Command> set ospf enable
Command> save all
Command> reboot
Command> add ospf area 0
Command> set ospf area 0 range 10.0.0.0/8
Command> add ospf area 1
Command> set ospf area 1 range 192.168.2.0/24
Command> set ospf area 1 range 192.168.1.0/24
Command> set S1 ospf on
Command> set ether0 ospf on
Command> reset S1
```

Glossary

A

abort error

An error indicating an attempted and failed connection.

acceptance policy

A set of rules that determine the path and route information the PortMaster accepts from a BGP peer for further processing. See also **policy**.

address

A number used to identify a computer or other device on a network or internetwork. See also **IP address**; **MAC address**.

address resolution

A method for translating one type of address into another—for example, an IP address into a media access control (MAC) address.

Address Resolution Protocol

See **ARP**.

adjacency

A relationship between two routers on the same physical network or between the endpoints of a virtual link that controls the distribution of routing protocol packets by limiting their exchange to those routers or endpoints.

advertisement policy

A set of rules that determine the path and route information the PortMaster advertises to a BGP peer. See also **policy**.

agent

A software program installed in a managed network device. An agent stores management information and responds to the manager's request for this information.

aggregation

The process of combining multiple prefixes from one or several routes so that a single prefix and route can be advertised. Route aggregation reduces the amount of information that a device running BGP must store and exchange with its BGP peers. See also **summarization**.

Annex-D

The ANSI T1.617 Frame Relay Annex-D version of the Local Management Interface (LMI) protocol. The Annex-D protocol has a more robust feature set than the proprietary Cisco/Stratacom LMI, but was developed later. Recent versions of the PortMaster software support either type of LMI. Earlier versions supported only the Cisco/Stratacom version. See also **LMI**.

area

In OSPF, a contiguous collection of networks and hosts. Each area runs a separate copy of the shortest-path-first (SPF) algorithm and has its own topological database.

area border router

In OSPF, a router that attaches to the backbone and one other area. An area border router runs separate copies of the shortest-path-first (SPF) algorithm for each area it attaches to. Area border routers condense the topological information of their attached areas and distribute it over the backbone to the other areas.

ARP

Address Resolution Protocol. A protocol that discovers the unique physical hardware address of a node or a LAN from its IP address. When an ARP request is sent to the network, naming the IP address, the machine with that IP address returns its physical address so that it can receive the transmission.

ASCII

American Standard Code for Information Interchange. A standard 8-bit code commonly used by computers and communications equipment.

autonomous system

A collection of routers under the control of a single technical administration, using one or more Interior Gateway Protocols (IGPs)—such as OSPF—to route packets within itself, and an Exterior Gateway Protocol (EGP)—such as BGP—to route packets to other autonomous systems. An autonomous system typically uses a common BGP policy and always presents a consistent view of network reachability to other autonomous systems.

autonomous system border router

In OSPF, a router that exchanges information with routers from other autonomous systems. Autonomous system border routers are also used to import routing information about RIP, direct, or static routes from non-OSPF attached interfaces.

autonomous system path list

In BGP, the list of autonomous systems that a packet must traverse to reach a given set of IP address destinations located within a single autonomous system destination. The list can consist of **sequences**, which are series of autonomous systems, that must be traversed in the order specified, and **sets**, which are collections of autonomous systems one of more of which must be traversed in any order to the destination.

For example, an autonomous system path list might consist of *Sequence 1, 2, 3, Set 4, 5, Sequence 6, 7*. This list indicates that a packet traverses autonomous systems 1, 2, and 3 in order, then one or both of autonomous systems 4 and 5 in any order, and finally autonomous systems 6 and 7 in order. Autonomous system 7 is the destination autonomous system.

B

backbone

A network topology consisting of a single length of cable with multiple network connection points.

backbone area

In OSPF, an area consisting of networks and routers not contained in any area and autonomous system border routers. The backbone area is responsible for distributing routing information between areas. This backbone area must be contiguous either physically or through a virtual link. The number reserved for the backbone area is 0.0.0.0.

backbone router

In OSPF, a router that has an interface into the backbone area by a direct attachment or a virtual link.

Basic Rate Interface

See **BRI**.

baud

The number of discrete signal events per second occurring on a communications channel. Although not technically accurate, baud is commonly used to mean bit rate.

B channel

Bearer channel. A 64Kbps synchronous channel that is part of an ISDN Basic Rate Interface (BRI).

BGP

Border Gateway Protocol. A routing protocol for exchanging network reachability information among autonomous systems. A routing device can use this information to construct a "map" of autonomous system connectivity. Version 4 of this protocol (BGP-4), which supports classless interdomain routing (CIDR) and route aggregation, is the predominant routing protocol used to propagate routes between autonomous systems on the Internet. BGP uses TCP as its transport protocol.

BGP-4

Version 4 of BGP. See also **BGP**.

BONDING

Bandwidth on Demand Interoperability Group. A method for combining two B channels into a single 128Kbps channel.

booting

The process in which a device obtains information and begins to process it to attain a state of normal operation.

Border Gateway Protocol

See **BGP**.

bps

Bits per second. A unit for measuring the data rate.

BRI

Basic Rate Interface. An ISDN interface that consists of two 64Kbps
B channels for voice or data and one 16Kbps D channel for signaling. Compare **PRI**.

broadcast address

A special address reserved for sending a message to all stations. Generally, a broadcast address is a media access control (MAC) destination address of all 1s (ones).

broadcast packets

Packets that are sent to all network nodes.

C

callback

A port configuration allowing the PortMaster to call back dial-in users before providing access. Callback provides an extra layer of security and can simplify telephone charges.

CCITT

Consultative Committee for International Telegraph and Telephone. International organization formerly responsible for the development of communications standards. Now called the ITU-T. See also **ITU-T**.

CD

Carrier Detect. A signal that indicates whether an interface is active. Also, a signal generated by a modem indicating that a call has been connected.

Challenge Handshake Authentication Protocol

See **CHAP**.

channelized T1

An access link operating at 1.544Mbps that is subdivided into 24 channels of 56Kbps each for dial-in use.

channel service unit

See **CSU**.

CHAP

Challenge Handshake Authentication Protocol. A Point-to-Point Protocol (PPP) authentication method for identifying a dial-in user. CHAP does not itself prevent unauthorized access, it merely identifies the remote end. See also **PAP**.

CIDR

Classless interdomain routing. A technique supported by BGP-4 that eliminates the necessity for network address classes by explicitly advertising the length (netmask) associated with each prefix.

CIR

Committed information rate. The minimum bandwidth guaranteed to be available if required on a virtual circuit. This value is also known as *guaranteed bandwidth*.

classless interdomain routing

See **CIDR**.

client/server environment

An environment where a computer system or process requests a service from another computer system. For example, a workstation can request services from a file server across a network.

cluster

A group of internal BGP peers that share a common set of route reflectors. See also **cluster ID**; **route reflection**; **route reflector**. Compare **confederation**.

cluster ID

An identifier, in dotted decimal format, that uniquely identifies a BGP route reflection cluster within an autonomous system. All route reflectors within the cluster must be configured with the same cluster ID. Internal peers that are not reflectors within the cluster must not be configured with a cluster ID. The cluster ID is typically set to the BGP router ID of one of the route reflectors within the cluster. See also **cluster**; **route reflection**; **route reflector**.

CMAS

Confederation member autonomous system. A subdivision of an autonomous system that is recognized only by other peers within the confederation and not by peers external to the confederation. Within the confederation, each BGP peer treats only the peers in its own CMAS as internal peers. Peers in different CMASs are treated as external peers.

committed information rate

See **CIR**.

community

A label that identifies a group of BGP destinations for the purpose of policy enforcement. Assembling destinations into identifiable "communities" lets BGP peers base policy decisions on the identity of the group rather than on individual destinations. The community identifier, which consists either of one 32-bit value or two 16-bit values, is advertised in update messages between BGP peers.

community string

A character string assigned to a Simple Network Management Protocol (SNMP) agent to restrict read and write access to the SNMP variables.

ComOS

The operating system for Lucent communications servers, routers, and access servers.

confederation

In BGP, an autonomous system that has been subdivided into smaller autonomous systems called *confederation member autonomous systems*. (CMASs). A confederation appears like a single autonomous system to other autonomous systems and is recognized only by other confederation members. Subdivision of an autonomous system into a confederation changes the peer relationships of confederation members in different CMASs from internal to external. Use of confederations in an autonomous system requires that all routers in the autonomous system belong to a CMAS; however, the policies used by BGP peers can change across confederation boundaries.

Confederations are one method for avoiding the overhead of having all peers within an autonomous system fully communicate to—be fully meshed with—each other. Route reflection clusters provide an easier method, but require the use of identical policies on all peers within the autonomous system. See also **route reflection**.

confederation member

Any router running BGP and recognizing that its autonomous system is subdivided into smaller autonomous systems called *confederation member autonomous systems*. (CMASs). The CMASs are recognized only by confederation members and not by peers external to the confederation. Subdivision of an autonomous system into a confederation changes the peer relationships of confederation members in different CMASs from internal to external.

confederation member autonomous system

See **CMAS**.

console port

A serial port on a PortMaster attached to a terminal or PC through which you enter commands to communicate with ComOS.

CRC error

Cyclic redundancy check error. These errors can indicate problems with source station hardware, receivers, retiming modules and/or repeaters, bridges, cabling, or transceivers.

CSU

Channel service unit. An ancillary device needed to adapt the V.35 interface to a port on a telephone carrier switch. The CSU is placed between the data terminal equipment (DTE) and the switch.

cyclic redundancy check

See **CRC error**.

D

data communications equipment

See **DCE**.

data link connection identifier

See **DLCI**.

data service unit

See **DSU**.

Data Set Ready

See **DSR**.

data terminal equipment

See **DTE**.

Data Terminal Ready

See **DTR**.

DCE

Data communications equipment. Devices and connections of a communications network that make up the network end of the interface between the network and the user. The DCE provides a physical connection to the network, forwards traffic, and provides a clocking signal to synchronize data transmission between DCE and DTE devices. Modems and interface cards are DCEs.

DDE

Dynamic data exchange. A form of interprocess communication that uses shared memory to exchange data between applications. Applications can use a one-time data transfer or ongoing exchanges.

degree of preference

In BGP, an arbitrary rating number that the PortMaster assigns to every route it receives from a BGP peer. A higher numbers indicates a greater preference for a route when more than one exists to a destination. A route from an internal peer is assigned the local preference number that the PortMaster learned with the route. For a route learned from an external peer, the PortMaster calculates a number based on the autonomous system path length; the shortest path is preferred. You can use a routing policy rule to override the calculated or learned value and assign your own degree of preference to a route. See also **local preference**.

destination

In BGP, the final autonomous system in the autonomous system path whose IP address prefixes and associated netmasks are reported in the network layer reachability information (NLRI) field of an update message. A destination and its path comprise a BGP route. See also **path**; **route**.

dialback

See **callback**.

dial group

A number that is used to associate dial-out locations with ports.

digital service unit

See **DSU**.

direct memory access

See **DMA**.

DLCI

Data link connection identifier. A unique number that represents a particular permanent virtual circuit (PVC) on a particular physical segment of the Frame Relay network. As the frame is passed through each switch, the DLCI is remapped automatically by the switch as necessary.

DMA

Direct memory access. Transfer of data from a peripheral device, such as a hard disk drive, into a computer memory without mediation by a microprocessor.

DNS

Domain Name System. The system used on the Internet for translating the names of network hosts into IP addresses.

DRAM

Dynamic random access memory. A type of semiconductor random access memory (RAM) that stores information in integrated circuits containing capacitors.

DSR

Data Set Ready. The circuit that is activated when data communications equipment (DCE) is powered up and ready for use. See also **DCE**.

DSU

Digital service unit or data service unit. An ancillary device needed to adapt the physical interface on a data terminal equipment (DTE) device—such as a V.35 interface on a port—to a transmission facility—such as leased line or a Frame Relay switch. If the DTE lacks complete digital line interface capability, the DSU can be located with the channel service unit (CSU) on the customer's site and known as a CSU/DSU. See also **CSU**.

DTE

Data terminal equipment. A device at the user end of the interface between the network and the user. The DTE connects to a data network through a data communications equipment (DCE)—such as a modem or an interface card. DTEs convert user information into data signals for transmission, and reconvert received data signals into user information. Compare **DCE**.

DTR

Data Terminal Ready. The circuit that is activated to inform the data communications equipment (DCE) when the data terminal equipment (DTE) is ready to send and receive data. See also **DCE**; **DTE**.

dynamic data exchange

See **DDE**.

dynamic random access memory

See **DRAM**.

E

E1

Digital WAN carrier facility used predominantly in Europe that carries data at a rate of 2.048Mbps. E1 lines can be leased for private use from common carriers. Compare **T1**.

easy-multihome

A specialized, predefined BGP policy that simplifies the use of PortMaster routers in straightforward multihomed environments. When you define easy-multihome for a peer, you restrict what the PortMaster handles from the peer to information that is no more than two autonomous system hops away from the PortMaster. Only information that meets this criterion is accepted from the peer, put into the routing table used to forward packets to their destinations, and advertised to other peers. If you define easy-multihome for a peer, you must also define a default route on each router in your autonomous system to point them to destinations more distant than two hops. See also **multihome routing**; **policy**.

EBGP

Exterior BGP. The BGP used between peers in different autonomous systems, or, when confederations are in use, between peers in different confederation member autonomous systems (CMASs). Unlike internal BGP peers, EBGP peers need not have full connectivity with one another.

echo test

A diagnostic test used to check network reachability in which an Internet Control Message Protocol (ICMP) Echo Request packet or Simple Network Management Protocol (SNMP) test packet is sent to elicit a standard response.

Ethernet

A network communications system developed and standardized by Digital Equipment Corporation, Intel, and Xerox using baseband transmission, carrier sense multiple access/carrier detect (CSMA/CD) access, logical bus topology, and coaxial cable. The successor IEEE 802.3 standard provides for integration of Ethernet into the Open System Interconnection (OSI) model and extends the physical layer and media with repeaters and implementations that operate on fiber optic cable, broadband, and unshielded twisted pair.

external peer

A peer that resides in a different autonomous system—or, when confederations are in use, in a different confederation member autonomous system (CMAS)—from the current PortMaster.

Exterior BGP

See **EBGP**.

F

File Transfer Protocol

See **FTP**.

filter

Generally, a process or device that screens network traffic for certain characteristics, such as source address, destination address, or protocol, and determines whether to forward or discard that traffic based on the established criteria.

filter table

A database used to store filters.

Flash RAM

See **nonvolatile RAM**.

flow control

A technique for ensuring that a transmitting entity, such as a modem, does not overwhelm a receiving entity with data. When the buffers on the receiving device are full, a message is sent to the sending device to suspend the transmission until the data in the buffers has been processed. Flow control can be software-based, or hardware-based.

FRAD

Frame Relay access device. A network device that links any non-Frame Relay connection to a Frame Relay WAN.

frame

A packaging structure for network data and control information. A frame consists of a destination address, source address, length field, data, padding, and frame check sequence. The 802.3 standard for Ethernet specifies that the minimum size data frame is 64 bytes and the maximum size data frame is 1518 bytes.

Frame Relay

An industry-standard switched data link layer protocol that handles multiple virtual circuits using high-level data link layer control (HDLC) encapsulation between connected devices. It is used across the interface between user devices (for example, hosts and routers) and network equipment (for example, switching nodes). Frame Relay is more efficient than X.25, the protocol it replaced.

Frame Relay Access Device

See **FRAD**.

FTP

File Transfer Protocol. A TCP/IP protocol used to transfer files between network hosts.

G

gateway

A device that connects two or more networks that use different protocols. Gateways provide address translation services, but do not translate data. Gateways must be used in conjunction with special software packages that allow computers to use networking protocols not originally designed for them.

graphical user interface

See **GUI**.

GUI

Graphical user interface. A software interface based on pictorial representations and menus of operations and files.

H

hardwired

A continuous connection between two sites. A port on a PortMaster that is configured for hardwired use cannot be simultaneously used for any other type of connection.

hello

Protocol used by OSPF routers to acquire neighbors and to synchronize their topological databases.

high-water mark

The number of bytes of queued network traffic required to open an additional dial-out line to a remote location.

hop

The transmission of a data packet between two network nodes—for example, between two routers.

hop count

Measurement of the distance between a source and destination that is used as a metric to compare routes. If a packet traverses six routers between source and destination nodes, the hop count for the packet will be 6 when it arrives at its destination node.

host

A single, addressable device on a network. Computers, networked printers, and routers are hosts.

hunt group

A group of multiple telephone circuits that allows telephone calls to find an idle circuit to establish a link.

I

IBGP

Interior BGP. The BGP used between peers in the same autonomous system, or, when confederations are in use, between peers in the same confederation member autonomous system (CMAS). All IBGP peers must maintain direct BGP connections to—be **fully meshed** with—every other internal peer, but need not be physically attached to one another.

ICMP

Internet Control Message Protocol. The part of the Internet Protocol (IP) that allows for generation of error messages, test packets, and informational messages related to IP. This protocol is used by the ping function to send an ICMP Echo Request to a network host, which replies with an ICMP Echo Reply.

in-band signaling

Signaling over a network.

injection policy

A set of rules that determine the path and route information the PortMaster takes from BGP and places into its routing table used to forward packets to their destinations. The PortMaster uses the information to determine how packets it receives are forwarded to their ultimate destinations. See also **policy**.

interface

Connection and interaction between hardware, software, and the user. The interface between components in a network is called a protocol. On the PortMaster, the virtual connection between a PortMaster port and the network to which it is connected is called an interface. The connection can be permanent as with the Ethernet interface or network hardwired ports, or it can be temporary, as with ports used for dial-in or dial-out connections.

Integrated Services Digital Network

See **ISDN**.

Interior BGP

See **IBGP**.

internal peer

A peer that resides in the same autonomous system—or, when confederations are in use, in the same confederation member autonomous system (CMAS)—as the current PortMaster.

internal router

In OSPF, a router with all of its directly connected interfaces or physical networks belonging to the same area and containing no virtual connections to the backbone area.

International Organization for Standards

See **ISO**.

internetwork

A network of networks.

Internet

The world-wide internetwork consisting of several large national backbone networks and several regional and campus networks.

Internet Control Message Protocol

See **ICMP**.

Internet Protocol

See **IP**.

Internet Network Information Center

See **InterNIC**.

InterNIC

Internet Network Information Center. An organization that provides information and services related to networking technologies.

IP

Internet Protocol. The protocol defined in RFC 791.

IP address

A 32-bit number assigned by the system administrator, usually written in the form of four decimal fields separated by periods—for example, 192.9.200.1. Any computing device that uses IP must be assigned an Internet or IP address. Part of the Internet address is the IP network number (IP network address), and part is the host address (IP host address). All machines on a given IP network use the same IP network number, and each machine has a unique IP host address. The system administrator sets the subnet mask to specify how much of the address is network number and how much is host address.

IP address prefix

An IP address number that, when paired with a netmask length, represents a range of addresses rather than a single IP network. For example, the prefix and netmask length 128.0.0.0/8 describe all networks whose IP addresses begin with 128. See also **netmask length**.

IP Control Protocol

See **IPCP**.

IPCP

IP Control Protocol. A protocol used by the Point-to-Point Protocol (PPP) for establishing and configuring an IP link over PPP.

IPX

Internet Packet Exchange. Internet protocol defined by Novell, Inc.

IPXWAN

IPX Wide Area Network protocol. The protocol used to establish and configure an IPX link over the Point-to-Point Protocol (PPP), as described in RFC 1634.

IPX Wide Area Network

See **IPXWAN**.

ISDN

Integrated Services Digital Network. A digital communications standard designed to allow the transmission of voice, data, images, and video over existing copper phone lines.

ISO

International Organization for Standards. The international organization that sets standards for network communication protocols.

ITU-T

International Telecommunication Union Telecommunication Standardization Sector. International organization that develops worldwide standards for telecommunications technologies. The ITU-T carries out the functions of the former CCITT. See also **CCITT**.

K

KB

Kilobyte(s). 1024 bytes.

Kb

Kilobit(s). 1024 bits.

Kbps

Kilobits per second.

keepalive message

A periodic message sent between BGP peers to keep their BGP sessions open. If a preset amount of time elapses between keepalive messages from a peer, the PortMaster identifies the peer as no longer operational and drops the session—and any information learned from that peer.

L

LAN

Local area network. A local collection, usually within a single building or several buildings, of personal computers and other devices connected by cabling to a common transmission medium, allowing users to share resources and exchange files. Compare **WAN**.

latency

1) The delay between the time a device requests access to a network and the time it is granted permission to transmit. 2) The delay between the time when a device receives a frame and the time that frame is forwarded out the destination port.

LCP

Link Control Protocol. The protocol used by the Point-to-Point Protocol (PPP) for establishing, configuring, and testing the data link connection.

LED

Light-emitting diode.

line speed

The speed of the physical wire attached to the interface or interface hardware. The line speed is 10Mbps for Ethernet and 1.544Mbps for T1. Fractional T1 is often implemented with a wire speed of T1 (1.544Mbps) and a lower port speed. Upgrading line speed is generally a hardware change. See also **port speed**.

Link Control Protocol

See **LCP**.

link state advertisement

See **LSA**.

LMI

Local Management Interface. A protocol used to communicate link status and permanent virtual circuit (PVC) status in Frame Relay. Two types of LMI are available on Frame Relay: the original proprietary Cisco/Stratacom LMI, and the ANSI T1.617 Annex-D LMI. Although the PortMaster supports both, LMI on the PortMaster refers to the Cisco/Stratacom implementation. See also **Annex-D**.

local area network

See **LAN**.

Local Management Interface

See **LMI**.

local preference

In BGP, the degree-of-preference number that the PortMaster assigns to every external route it advertises to an internal or confederation-member BGP peer. A higher number indicates a greater preference for a route when more than one exists to a destination. Internal and confederation-member peers receiving this route use this local preference rather than calculating their own degree of preference for a route. You can use a routing policy rule to override this value and assign your own local preference to a route you advertise. See also **degree of preference**.

location

A dial-out destination.

location table

A database on the PortMaster where location settings are stored. See **location**.

lockstep

A feature of BGP on the PortMaster that ensures consistency of routing information between the BGP and non-BGP routers within its autonomous system. Lockstep forces the PortMaster to advertise a route learned from an internal BGP peer only when it has learned the same route via an Interior Gateway Protocol (IGP)—OSPF or RIP—or a static route. See also **transit service**.

LSA

Link state advertisement. The state of the router links (interfaces), networks, summaries, or autonomous system external links of an OSPF router that it periodically advertises. Link states are also advertised when a link state changes.

M

MAC address

Media access control address. A unique 48-bit binary number—usually represented as a 12-digit hexadecimal number—encoded in the circuitry of a device to identify it on a LAN.

Management Information Base

See **MIB**.

management station

A workstation or PC capable of retrieving and analyzing statistical information from networked Simple Network Management Protocol (SNMP) agents.

master

In Multichassis PPP, the PortMaster through which an initial connection for a given user is made. Every master also has a corresponding slave. Masters are for a given connection only, and a PortMaster that functions as a master for one user's connection can be a slave for a different user's connection. See also **slave**.

maximum transmission unit

See **MTU**.

MB

Megabyte(s). 1,048,576 bytes.

Mbps

Megabits per second. A unit for measuring data rates.

MD5

Message digest algorithm 5. The algorithm used for message authentication in Simple Network Management Protocol (SNMP) v.2. MD5 verifies the integrity of the communication, authenticates the origin, and checks for timeliness.

media access control address

See **MAC address**.

message digest algorithm 5

See **MD5**.

MIB

Management Information Base. A set of variables that a Simple Network Management Protocol (SNMP)-based management station can query from the SNMP agent of a network device.

modem

Modulator-demodulator. A device that converts the digital signals used by computers to analog signals that can be transmitted over telephone lines.

modem table

A database resident on the PortMaster containing configuration information for commonly used modems.

MTU

Maximum transmission unit. The largest frame or packet that can be sent through a port on a PortMaster without fragmentation.

Multichassis PPP

Multilink PPP over two or more chassis.

Multilink PPP

A protocol defined in RFC 1717 that allows a PortMaster to automatically bring up additional ISDN B channels as bandwidth needs increase. As soon as traffic decreases, the PortMaster disconnects unneeded B channels. See also **Multichassis PPP**.

multiexit discriminator

In BGP, an arbitrary rating number that the PortMaster can use to enforce the use of preferred exit and entry points when multiple connections exist between its autonomous system and another. The PortMaster assigns the multiexit discriminator to any route that it advertises to its external peers, and forwards any multiexit discriminator it learns from its external peers on to its internal peers. A lower number indicates a greater preference for a route when more than one exists to a destination through multiple peers within the same neighboring autonomous system. You can use a routing policy rule to override this value and assign your own multiexit discriminator to a route that you learn or advertise.

multihome routing

In BGP, the process of choosing among multiple exit points to route packets out of a single autonomous system, typically to the Internet. Routers in a multihomed autonomous system usually store large amounts of network reachability information to help them select the best exit point. See also **easy-multihome**.

multiline load balancing

The ability of a PortMaster to add additional lines when network traffic is heavy. If more than one line to a remote location is established, the PortMaster balances the traffic among the lines. Distinct from Multilink PPP.

N

name server

A server connected to a network that resolves hostnames into network addresses.

name service

The software system that provides a database of authorized users for a computer, subnet, or network. The system can reside on one device, or be distributed across several devices in a network.

neighbor

(1) In OSPF, two routers that have interfaces to a common network are neighbors. On multiaccess networks, neighbors are dynamically discovered by the OSPF Hello protocol.

(2) In Multichassis PPP, PortMasters in the same Multichassis PPP domain.

netmask

A 32-bit number that distinguishes the portion of an IP address referring to the network or subnet from the portion referring to the host. Compare **subnet mask**.

netmask length

A number between 0 and 32 preceded by a slash (/) and following an IP address prefix. The netmask length indicates the number of high-order bits in the prefix that an IP address must match to fall within the range indicated by the prefix. For example, the prefix and netmask length 128.0.0.0/8 describe all networks whose IP addresses begin with 128. See also **IP address prefix**.

network

A collection of computers, terminals, and other devices and the hardware and software that enable them to exchange data and share resources over short or long distances.

network handle

A number assigned to an active socket that can be used to close the socket manually, rather than by a request from the client.

network interface card

See **NIC**.

Network Information Service

See **NIS**.

network layer reachability information

See **NLRI**.

network management

In the Open System Interconnection (OSI) model, the five functional application areas of accounting management, configuration management, fault management, performance management, and security management.

NIC

Network interface card. A card that provides network communication capabilities to and from a computer system. A NIC is also known as an *adapter*.

NIS

Network Information Service. A protocol developed by Sun Microsystems for the administration of network-wide databases.

NLRI

Network layer reachability information. The part of a BGP route containing the IP address prefixes and associated netmask lengths that are reachable via the path described in the route. The networks indicated by these prefixes and netmasks reside in the destination autonomous system—the final one listed in the path.

node

A device, such as a PC, server, switching point, bridge, or gateway, connected to a network at a single location. A node can also be called a *station*. See **host**.

nonvolatile RAM

Nonvolatile random access memory. Nonvolatile storage that can be erased and reprogrammed electronically, allowing software images to be stored, booted, and rewritten as necessary.

notification message

A message sent between BGP peers to inform the receiving peer that the sending peer must terminate the BGP session because an error occurred. The message contains information that explains the error. See also **keepalive message**; **open messag**e; **update message**.

not-so-stubby-area

See **NSSA**.

NSSA

Not-so-stubby-area. In OSPF, an area similar to a stub area except that Type 1 and Type 2 external routes can be learned from it. Any external routes learned from an NSSA are translated into Type 1 and Type 2 external routes for the backbone area or other areas that accept external routes. Like stub areas, NSSAs can have default costs set for them but cannot have external routes advertised into them.

NT1

Network termination 1 device. The device that provides an interface between the ISDN Basic Rate Interface (BRI) line used by the telephone company and a customer's terminal equipment. The NT1 also provides power for the terminal equipment, if necessary. In North America, where ISDN BRI is a U loop, the customer must supply the NT1 device; in Japan and the European countries where BRI is an S/T bus, the telephone company supplies the NT1. The PortMaster integrates the NT1 device into its ISDN BRI ports that are U interfaces.

NVRAM

See **nonvolatile RAM**.

O

ODI

Open Datalink Interface. A Novell specification that isolates the protocol stack from the network adapter drivers to provide hardware independence for network connectivity.

Open Datalink Interface

See **ODI**.

open message

A message sent between BGP peers to establish communication. See also **keepalive message**; **notification message**; **update message**.

Open Shortest Path First

See **OSPF**.

OSPF

Open Shortest Path First. A link-state interior gateway routing protocol designed for a hierarchical routing structure. OSPF chooses routes on a best-path, least-cost basis and supports variable-length subnet masks (VLSMs) for "classless" networking, allows up to 255 hops between routers, and provides packet authentication. See also **RIP**.

out-of-band connection

A remote connection, or a connection outside connected networks, established over a modem. This type of connection is useful when network communications are not available.

P

packet

A unit of data sent across a network.

PAP

Password Authentication Protocol. An authentication protocol that allows PPP peers to authenticate one another. The remote router attempting to connect to the local router is required to send an authentication request. Unlike the Challenge Handshake Authentication Protocol (CHAP), PAP passes unencrypted passwords. PAP does not itself prevent unauthorized access, but merely identifies the remote end. The router or access server then determines if that user is allowed access. See also **CHAP**.

parity check

A process for checking the integrity of a character. A parity check appends a bit to a character or word to make the total number of binary 1 digits in the character or word (excluding the parity bit) either odd (for odd parity) or even (for even parity).

partition

Electronic isolation of an Ethernet device from network communications.

Password Authentication Protocol

See **PAP**.

path

In BGP, a autonomous system path list and a collection of attributes that provide descriptions of and explain how to reach a given collection of IP address destinations in a single autonomous system. A path and its destination comprise a BGP route. See also **destination**; **autonomous system path list**; **route**.

peer

(1) In BGP, a router with which a BGP speaker exchanges open messages, notification messages, update messages, and keepalive messages. A PortMaster can have both internal and external peers. See also **internal peer**; **external peer**.

(2) In Multichassis PPP, the relationship between a master and slave. A peer is distinct from a *neighbor*.

permanent virtual circuit

See **PVC**.

physical circuit

A physical connection between two devices.

ping

Packet Internet Groper. A program that is useful for testing and debugging networks. Ping sends an ICMP echo packet to the specified host and waits for a reply. Ping reports success or failure and sometimes statistics about its operation.

Point-to-Point Protocol

See **PPP**.

policy

In BGP, the rule or set of rules the PortMaster follows for accepting, injecting, and/or advertising BGP routes to its BGP internal and external peers. You assign policies to a peer when you add it to the PortMaster during configuration. You can use the default policy **easy-multihome**, or create and assign your own policies. One policy can handle all three functions, or you can create separate policies for acceptance, injection, and advertisement. See also **acceptance policy**; **advertisement policy**; **injection policy**.

port

The physical channel or connection through which data flows.

port speed

The rate at which data is accepted by the port at the end of the wire. For example, when a T1 line exists between a site and a telecommunications provider, the telecommunications provider accepts only the number of bits per second ordered by the customer into the port on its equipment. Upgrading port speed is generally a software change.

PPP

Point-to-Point Protocol. A protocol that provides connections between routers and between hosts and networks over synchronous and asynchronous circuits. See also **SLIP**.

PRI

Primary Rate Interface. The ISDN interface to primary rate access. Primary rate access consists of a single 64Kbps D channel plus 23 (T1) or 30 (E1) 64Kbps B channels for voice or data. Compare **BRI**.

Primary Rate Interface

See **PRI**.

propagation

The process of translating and forwarding routes from one routing protocol into another. Route propagation is also known as route *redistribution*. Lucent recommends using route filters in propagation rules to ensure that you redistribute information without creating routing loops. Compare **summarization**.

proxy Address Resolution Protocol

See **proxy ARP**.

proxy ARP

Proxy Address Resolution Protocol. A variation of the ARP protocol in which a router or other device sends an ARP response to the requesting host on behalf of another node. Proxy ARP can reduce the use of bandwidth on slow-speed WAN links. See also **ARP**.

PVC

Permanent virtual circuit. A circuit that defines a permanent connection in a switched digital service such as Frame Relay. Frame Relay is the only switched digital service that uses PVCs supported by PortMaster products.

R

RADIUS

Remote Authentication Dial-In User Service. A client/server security protocol created by Lucent.

RARP

Reverse Address Resolution Protocol. A protocol used in network routers that provides a method for finding IP addresses based on media access control (MAC) addresses. Compare **ARP**.

Remote Authentication Dial-In User Service

See **RADIUS**.

Request for Comments

See **RFC**.

Reverse Address Resolution Protocol

See **RARP**.

RFC

Request for Comments. One of a series of documents that communicate information about the Internet. Most RFCs document protocol specifications, such as those for IP and BGP. Some RFCs are designated as standards.

RIP

Routing Information Protocol. A protocol used for the transmission of IP or IPX routing information.

rlogin

Remote login. A terminal emulation program, similar to Telnet, offered in most UNIX implementations.

route

A way for a packet to reach its target via the Internet. A BGP route provides a path of autonomous systems—plus any path attributes—to a single destination autonomous system that contains particular IP address prefixes and associated netmasks. Packets whose targets fall within the networks identified by these prefixes and netmasks can use this BGP route. BGP peers advertise routes to each other in **update messages**.

router

A device that connects two or more networks and can direct traffic based on addresses.

route reflection

In BGP, a method for maintaining path and attribute information across an autonomous system, while avoiding the overhead of having all peers within an autonomous system fully communicate to—be fully meshed with—each other. To reduce the number of links, all internal peers are divided into clusters, each of which has one or more route reflectors. A route received by a route reflector from an internal peer is transmitted to its clients, which are the other peers in the cluster that are not route reflectors. Route reflection requires that all internal peers use identical policies.

Confederations are another way to avoid configuring a fully meshed set of peers in a single autonomous system. In contrast to route reflection clusters, confederations require all routers in the autonomous system to operate as confederation members. However, confederations provide a finer control of routing within the autonomous system by allowing for policy changes across confederation boundaries. See also **cluster**; **cluster ID**; **confederation**; **route reflector**.

route reflector

A router configured to transmit routes received from internal BGP peers to one or more other internal peers within its same cluster. These peers are called the route reflector's *clients*. See also **cluster**; **cluster ID**; **route reflection**.

router ID

One of the interface addresses configured on a BGP speaker. The router ID is chosen as the address that uniquely identifies the BGP speaker on the Internet.

Routing Information Protocol

See **RIP**.

routing table

A database of routes to particular network destinations, stored on a router or other device. The routing table stored on the PortMaster contains the following information for each route: IP address and netmask length of the destination, IP address of the gateway, source of the route (if any), type of route, hop-count metric, and PortMaster interface used to forward packets along the route.

RS-232 interface

A standard for data communication using serial data and control signals.

runt packet

A packet with a frame size between 8 and 63 bytes with frame check sequence (FCS) or alignment errors. The runt packet is presumed to be a fragment resulting from a collision.

S

SAP

Service Advertisement Protocol. An IPX protocol that provides a means of informing network clients, via routers and servers, of available network resources and services. See also **IPX**.

Serial Line Internet Protocol

See **SLIP**.

serial port

A bidirectional channel through which data flows one bit as a time. Asynchronous serial ports most often use 10 bits for a character of data including 1 start bit, 8 data bits, and 1 stop bit.

server

A computer or a specialized device that provides and manages access to shared network resources, such as hard disks and printers.

Service Advertisement Protocol

See **SAP**.

service profile identifier

See **SPID**.

Simple Network Management Protocol

See **SNMP**.

slave

In Multichassis PPP, a PortMaster through which a subsequent connection for a particular user is made. (The port through which the connection is made is called the **slave port**.) Every slave has a corresponding master. Slaves are for a given connection only, and a PortMaster that functions as a slave for one user's connection can be a master for a different user's connection. See also **master**.

SLIP

Serial Line Internet Protocol. The protocol, obsoleted by the Point-to-Point Protocol (PPP), for point-to-point serial connections using TCP/IP. See also **PPP**.

SNMP

Simple Network Management Protocol. A protocol defined in RFC 1157, used for communication between management consoles and network devices.

speaker

A single BGP router that is able to communicate with other routers that run BGP. When two BGP speakers communicate with each other, they are called BGP *peers*.

SPID

Service profile identifier. A number used by some service providers to define the services to which an ISDN device subscribes. The ISDN device uses the SPID when accessing the switch that initializes the connection to a service provider.

station

See **host**.

stub area

In OSPF, an area into which no external routes are imported. A stub area cannot contain autonomous system border routers and cannot be a transit area for virtual links. Summary advertisements external to the area are by default imported into the stub area but might be squelched to further reduce area database size. In this case, the default route advertisement by the autonomous system border routers handle all routes external to the area.

subnet mask

A 32-bit netmask used to indicate the bits of an IP address that are being used for the subnet address. Compare **netmask**.

summarization

The process of combining routing information from one routing protocol into another for advertisement. For example, the PortMaster summarizes non-BGP route information it receives internally via the Interior Gateway Protocol (IGP) OSPF or RIP, or via a static route, into BGP for advertisement to BGP internal and external peers. Summarized routing information must comply with BGP advertisement policy rules before advertisement. Compare **propagation**.

SVC

Switched virtual circuit. A connection established between two physical circuits, such as an ordinary telephone call. The call creates a virtual circuit between the originator and the party called.

switched virtual circuit

See **SVC**.

T

T1

Digital WAN carrier facility used to transmit data formatted for digital signal level 1 (DS-1) at 1.544Mbps through the telephone-switching network, using alternate mask inversion (AMI) or binary 8-zero substitution (B8ZS) coding. Compare **E1**.

TCP/IP

An open network standard that defines how devices from different manufacturers communicate with each other over interconnected networks. TCP/IP protocols are the foundation of the Internet.

Telnet

The Internet standard protocol, described in RFC 854, for remote terminal connection service.

terminal adapter

A device that provides ISDN compatibility to non-ISDN devices. An asynchronous terminal adapter turns an asynchronous bit stream into ISDN and is treated by the PortMaster as if it were a modem. A synchronous terminal adapter takes a synchronous bit stream and turns it into ISDN, typically supports V.25bis dialing, and connects to a PortMaster synchronous port. Some terminal adapters can be configured for either synchronous or asynchronous operation.

terminal emulator

A program that makes a PC screen and keyboard act like the video display terminal of another computer.

TFTP

Trivial File Transfer Protocol. A simplified version of the File Transfer Protocol (FTP) that transfers files but does not provide password protection or user directory capability. TFTP can be used by diskless devices that keep software in ROM and use it to boot themselves. The PortMaster can be booted from the network by means of Reverse Address Resolution Protocol (RARP) and TFTP.

transit service

In BGP, the function provided by an autonomous system that is in the path of a route but not the origination or destination. To provide reliable transit service, an autonomous system must ensure that its BGP and non-BGP routers agree on the interior routes and exit and entry points for each transit route through the autonomous system. The PortMaster synchronizes routing information between the BGP and non-BGP routers within its autonomous system by means of the **lockstep** feature. See also **lockstep**.

Trivial File Transfer Protocol

See **TFTP**.

two-way

Relating to a port configuration that allows both incoming and outgoing calls.

U

UDP

User Datagram Protocol. A connectionless protocol defined in RFC 768. UDP exchanges datagrams but does not provide guaranteed delivery.

U interface

The ISDN interface defined as the connection between the network termination 1 device (NT1) and the telephone company local loop. The U interface standard is set by each country. The U interface described in Lucent documentation refers to the U.S. definition.

UNIX

A multiuser, multitasking operating system originally developed by AT&T that runs on a wide variety of computer systems.

UNIX-to-UNIX Copy Program

See **UUCP**.

update message

A message sent between BGP peers to convey network reachability information in two parts. The first part lists the IP address prefixes and associated netmasks for one or more routes that the PortMaster is withdrawing from service because it can no longer reach them. The second part of an update message consists of a single BGP route. See also **route**.

User Datagram Protocol

See **UDP**.

UUCP

UNIX-to-UNIX Copy Program. Interactive communication system for connecting two UNIX computers to send and receive data.

V

V.120

An ITU-T standard for performing asynchronous rate adaptation into ISDN.

V.25bis

An ITU-T standard defining how to dial on synchronous devices such as ISDN or switched 56Kbps.

V.32bis

An ITU-T standard that extends the V.32 connection range from 4800bps to 14.4Kbps. V.32bis modems fall back to the next lower speed when line quality is impaired, and fall back further as necessary. They fall forward to the next higher speed when line quality improves.

V.34

An ITU-T standard that allows data rates as high as 28.8Kbps.

variable-length subnet mask

See **VLSM**.

virtual circuit

A logical connection between two endpoints on a switched digital network. Virtual circuits can be switched or permanent. A switched virtual circuit (SVC) is used when you make an ordinary telephone call, an ISDN connection, or a V.25 switched 56Kbps connection. A permanent virtual circuit (PVC) is used in Frame Relay. See also **PVC**; **SVC**.

virtual connection

In Multichassis PPP, a connection made when a slave forwards all the packets it receives for a particular connection to its corresponding master for processing.

virtual port

In Multichassis PPP, a port corresponding to the physical port of the slave.

VLSM

Variable-length subnet mask. A means of specifying a different subnet mask for the same network number on different subnets. VLSM often allows addresses to be assigned more efficiently. OSPF and BGP support “classless” or VLSM routes.

W

WAN

Wide area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers. Frame Relay is an example of a WAN. Compare **LAN**.

Index

C

K

L

M

N

O

P

R

S

T

U

V

W