Lucent Technologies
Bell Labs Innovations

NavisCore IP Navigator Configuration Guide

For B-STDX 6.5.2.x, CBX 3.5.2.x, and GX 1.5.2.x

Product Code: 80114
Revision 03
January 2002

Notice. Every effort was made to ensure that the information in this document was complete and accurate at the time of printing. However, information is subject to change.

Limited Warranty. Lucent Technologies provides a limited warranty to this product. For more information, see the software license agreement in this document.

Ordering Information. To order copies of this document, contact your Lucent Technologies account representative.

Support Telephone Numbers. For technical support and other services, see the customer support contact information in the “About This Guide” section of this document.

LUCENT TECHNOLOGIES END-USER LICENSE AGREEMENT

LUCENT TECHNOLOGIES IS WILLING TO LICENSE THE ENCLOSED SOFTWARE AND ACCOMPANYING USER DOCUMENTATION (COLLECTIVELY, THE "PROGRAM") TO YOU ONLY UPON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT. PLEASE READ THE TERMS AND CONDITIONS OF THIS LICENSE AGREEMENT CAREFULLY BEFORE OPENING THE PACKAGE(S) OR USING THE LUCENT SWITCH(ES) CONTAINING THE SOFTWARE, AND BEFORE USING THE ACCOMPANYING USER DOCUMENTATION. OPENING THE PACKAGE(S) OR USING THE LUCENT SWITCH(ES) CONTAINING THE PROGRAM WILL INDICATE YOUR ACCEPTANCE OF THE TERMS OF THIS LICENSE AGREEMENT. IF YOU ARE NOT WILLING TO BE BOUND BY THE TERMS OF THIS LICENSE AGREEMENT, LUCENT IS UNWILLING TO LICENSE THE PROGRAM TO YOU, IN WHICH EVENT YOU SHOULD RETURN THE PROGRAM WITHIN TEN (10) DAYS FROM SHIPMENT TO THE PLACE FROM WHICH IT WAS ACQUIRED, AND YOUR LICENSE FEE WILL BE REFUNDED. THIS LICENSE AGREEMENT REPRESENTS THE ENTIRE AGREEMENT CONCERNING THE PROGRAM BETWEEN YOU AND LUCENT, AND IT SUPERSEDES ANY PRIOR PROPOSAL, REPRESENTATION OR UNDERSTANDING BETWEEN THE PARTIES.

1. License Grant. Lucent hereby grants to you, and you accept, a non-exclusive, non-transferable license to use the computer software, including all patches, error corrections, updates and revisions thereto in machine-readable, object code form only (the "Software"), and the accompanying User Documentation, only as authorized in this License Agreement. The Software may be used only on a single computer owned, leased, or otherwise controlled by you; or in the event of inoperability of that computer, on a backup computer selected by you. You agree that you will not pledge, lease, rent, or share your rights under this License Agreement, and that you will not, without Lucent's prior written consent, assign or transfer your rights hereunder. You agree that you may not modify, reverse assemble, reverse compile, or otherwise translate the Software or permit a third party to do so. You may make one copy of the Software and User Documentation for backup purposes. Any such copies of the Software or the User Documentation shall include Lucent's copyright and other proprietary notices. Except as authorized under this paragraph, no copies of the Program or any portions thereof may be made by you or any person under your authority or control.

2. Lucent's Rights. You agree that the Software and the User Documentation are proprietary, confidential products of Lucent or Lucent's licensor protected under US copyright law and you will use your best efforts to maintain their confidentiality. You further acknowledge and agree that all right, title and interest in and to the Program, including associated intellectual property rights, are and shall remain with Lucent or Lucent's licensor. This License Agreement does not convey to you an interest in or to the Program, but only a limited right of use revocable in accordance with the terms of this License Agreement.

3. License Fees. The license fees paid by you are paid in consideration of the license granted under this License Agreement.

4. Term. This License Agreement is effective upon your opening of the package(s) or use of the switch(es) containing Software and shall continue until terminated. You may terminate this License Agreement at any time by returning the Program and all copies or portions thereof to Lucent. Lucent may terminate this License Agreement upon the breach by you of any term hereof. Upon such termination by Lucent, you agree to return to Lucent the Program and all copies or portions thereof. Termination of this License Agreement shall not prejudice Lucent's rights to damages or any other available remedy.

5. Limited Warranty. Lucent warrants, for your benefit alone, for a period of 90 days from the date of shipment of the Program by Lucent (the “Warranty Period”) that the program diskettes in which the Software is contained are free from defects in material and workmanship. Lucent further warrants, for your benefit alone, that during the Warranty Period the Program shall operate substantially in accordance with the User Documentation. If during the Warranty Period, a defect in the Program appears, you may return the Program to the party from which the Program was acquired for either replacement or, if so elected by such party, refund of amounts paid by you under this License Agreement. You agree that the foregoing constitutes your sole and exclusive remedy for breach by Lucent of any warranties made under this Agreement. EXCEPT FOR THE WARRANTIES SET FORTH ABOVE, THE PROGRAM IS LICENSED “AS IS”, AND LUCENT DISCLAIMS ANY AND ALL OTHER WARRANTIES, WHETHER EXPRESS, IMPLIED OR STATUTORY, INCLUDING, WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTIES OF NONINFRINGEMENT.

6. Limitation of Liability. Lucent’s cumulative liability to you or any other party for any loss or damages resulting from any claims, demands, or actions arising out of or relating to this License Agreement shall not exceed the greater of: (i) ten thousand US dollars ($10,000) or (ii) the total license fee paid to Lucent for the use of the Program. In no event shall Lucent be liable for any indirect, incidental, consequential, special, punitive or exemplary damages or lost profits, even if Lucent has been advised of the possibility of such damages.

7. Proprietary Rights Indemnification. Lucent shall at its expense defend you against and, subject to the limitations set forth elsewhere herein, pay all costs and damages made in settlement or awarded against you resulting from a claim that the Program as supplied by Lucent infringes a United States copyright or a United States patent, or misappropriates a United States trade secret, provided that you: (a) provide prompt written notice of any such claim, (b) allow Lucent to direct the defense and settlement of the claim, and (c) provide Lucent with the authority, information, and assistance that Lucent deems reasonably necessary for the defense and settlement of the claim. You shall not consent to any judgment or decree or do any other act in compromise of any such claim without first obtaining Lucent's written consent. In any action based on such a claim, Lucent may, at its sole option, either: (1) obtain for you the right to continue using the Program, (2) replace or modify the Program to avoid the claim, or (3) if neither (1) nor (2) can reasonably be effected by Lucent, terminate the license granted hereunder and give you a prorata refund of the license fee paid for such Program, calculated on the basis of straight-line depreciation over a five-year useful life. Notwithstanding the preceding sentence, Lucent will have no liability for any infringement or misappropriation claim of any kind if such claim is based on: (i) the use of other than the current unaltered release of the Program and Lucent has provided or offers to provide such release to you for its then current license fee, or (ii) use or combination of the Program with programs or data not supplied or approved by Lucent to the extent such use or combination caused the claim.

8. Export Control. You agree not to export or disclose to anyone except a United States national any portion of the Program supplied by Lucent without first obtaining the required permits or licenses to do so from the US Office of Export Administration, and any other appropriate government agency.

9. Governing Law. This License Agreement shall be construed and governed in accordance with the laws and under the jurisdiction of the Commonwealth of Massachusetts, USA. Any dispute arising out of this Agreement shall be referred to an arbitration proceeding in Boston, Massachusetts, USA by the American Arbitration Association.

10. Miscellaneous. If any action is brought by either party to this License Agreement against the other party regarding the subject matter hereof, the prevailing party shall be entitled to recover, in addition to any other relief granted, reasonable attorneys' fees and expenses of arbitration. Should any term of this License Agreement be declared void or unenforceable by any court of competent jurisdiction, such declaration shall have no effect on the remaining terms hereof. The failure of either party to enforce any rights granted hereunder or to take action against the other party in the event of any breach hereunder shall not be deemed a waiver by that party as to subsequent enforcement of rights or subsequent actions in the event of future breaches.

Contents

About This Guide

Chapter 1 Overview

Chapter 2 Configuring Ethernet Logical Ports

Chapter 3 Configuring IP Logical Ports and IP Servers

Chapter 4 Configuring IP Packet Filters

Chapter 5 Configuring Policy-Based Forwarding

Chapter 9 Configuring IP OSPF and VNN OSPF

Chapter 10 Configuring Static Routes

Chapter 11 Configuring Route Policies

Chapter 12 Configuring Label Switched Paths

Chapter 13 About Next Hop Resolution Protocol

Chapter 14 Configuring Next Hop Resolution Protocol

Chapter 15 Configuring IP Multicast Routing

Chapter 16 Configuring IP Virtual Private Networks

Appendix A PRAM Upload

Appendix B Troubleshooting IP Navigator Problems

Index

List of Figures

List of Tables

About This Guide

The *NavisCore IP Navigator Configuration Guide* is a task-oriented guide that describes, step-by-step, the process for configuring an IP interface. This guide is intended for users who will be accessing the NavisCore NMS to configure IP interfaces in a network.

This guide supports the following Network Management Station (NMS) and switch software releases:

- NavisCore, Release 05.01.00.00
- B-STDX, Release 06.05.02.00
- CBX 500, Release 03.05.02.00

What You Need to Know

As a reader of this guide, you should be familiar with UNIX and HP OpenView. You should also know about relational databases to properly maintain Sybase, which is used by NavisCore.

This guide assumes that you have already installed the Lucent switch hardware and the NMS and switch software. See the "Related Documents" section for a list of documents that describe these and other tasks.

Be sure to read the Software Release Notice (SRN) that accompanies each product. The SRN contains the most current feature information and requirements.

Reading Path

This section describes all of the documents that support the NavisCore NMS and switch software.

NMS Documentation

Read the following documents to install and operate NavisCore Release 5.0.

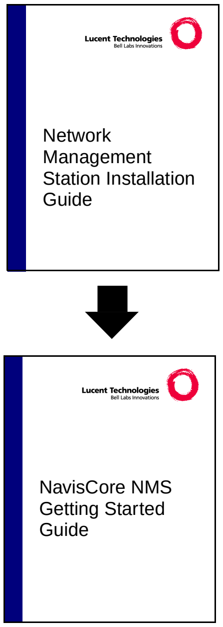
Lucent Technologies
Bell Labs Innovations
Network Management Station Installation Guide
Lucent Technologies
Bell Labs Innovations
NavisCore NMS Getting Started Guide

This guide describes prerequisite tasks, hardware and software requirements, and instructions for installing Solaris, HP OpenView, and NavisCore on the NMS.

This guide describes how to configure and manage NavisCore, network maps, and Lucent switches. It also describes how to add third-party objects to the map and access them through NavisCore.

Switch Software Documentation

Read the following documents to configure switch software for B-STDX Release 06.05.xx.xx, CBX Release 03.05.xx.xx, and GX Release 01.05.xx.xx.

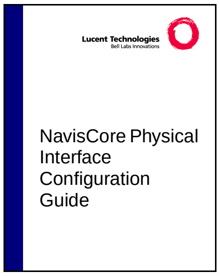

Lucent Technologies
Bell Labs Innovations

NavisCore Physical Interface Configuration Guide

This guide describes the processor and I/O modules on each switch platform, and how to configure physical ports, timing, and other attributes through NavisCore.

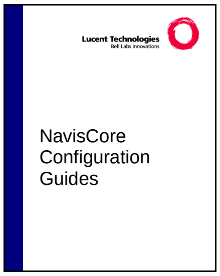

Lucent Technologies
Bell Labs Innovations

NavisCore Configuration Guides

These guides describe how to configure WAN services on the B-STDX, CBX, and GX switch platforms:

- *NavisCore Frame Relay Configuration Guide*
- *NavisCore ATM Configuration Guide*
- *NavisCore IP Navigator Configuration Guide*

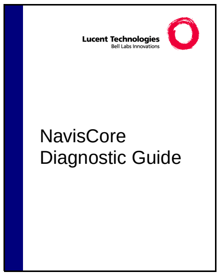

Lucent Technologies
Bell Labs Innovations

NavisCore Diagnostic Guide

This guide describes how to monitor and diagnose problems in your NavisCore switch network.

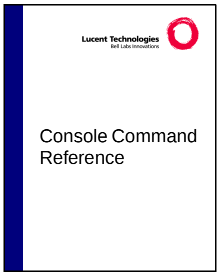

Lucent Technologies
Bell Labs Innovations

Console Command Reference

This reference lists and describes the NavisCore switch console commands.

How to Use This Guide

This guide contains the following information:

Read	To Learn About
Chapter 1	An overview of the product.
Chapter 2	How to configure logical ports on the following cards: • 4-Port Ethernet Card on the CBX 500 switch • 2-Port Ethernet Card on the B-STDX 8000/9000 switch
Chapter 3	How to configure: • IP logical ports on B-STDX 8000/9000 and CBX 500 switches. IP logical ports are ports that support IP routing. • IP server logical ports on the CBX 500 switch. IP server logical ports provide a method of accepting or transmitting IP traffic on a cell-based card.
Chapter 4	How to configure and assign IP packet filters.
Chapter 5	How to provision IP policy-based forwarding.
Chapter 6	How to define static ARP entries.
Chapter 7	How to configure Routing Information Protocol (RIP) parameters on an IP logical port and how to configure equal-cost multipath routing for RIP.
Chapter 8	How to configure Border Gateway Protocol (BGP) parameters including: • BGP switch parameters • BGP neighbors • BGP aggregates • BGP peer groups • BGP route dampening This chapter also describes how to configure IP loopback addresses.
Chapter 9	How to configure Open Shortest Path First (OSPF) parameters for IP Navigator and Virtual Network Navigator (VNN).
Chapter 10	How to configure static routes.
Chapter 11	How to create a route map. The purpose of a route map is to control and modify routing information and to define the parameters that your system uses to redistribute routes between routing domains.

Read	To Learn About
Chapter 12	How IP Navigator uses label switched paths (LSPs) as a means of forwarding IP traffic over switched paths through the Lucent network.
Chapter 13	An overview of the Next Hop Resolution Protocol (NHRP), which Lucent uses to implement absolute QoS.
Chapter 14	How to configure NHRP to implement absolute QoS.
Chapter 15	How to configure IP multicast routing protocols, including: • Internet Group Management Protocol (IGMP) • Distance Vector Multicast Routing Protocol (DVMRP) • Multicast Open Shortest Path First (MOSPF)
Chapter 16	How to configure IP virtual private networks (VPNs).
Appendix A	How to use the Upload PRAM function.
Appendix B	How to troubleshoot IP Navigator problems on Lucent switches.

What’s New in This Release?

This guide describes the following new product features:

Feature or Enhancement	Enables You To	See
New Features in This Release		
Absolute QoS	Set up an SVC “shortcut” to a destination using the Non-Broadcast Multiple Access (NBMA) address of the next hop toward a destination. These shortcuts provide QoS guarantees for traffic that passes over them. Absolute QoS is implemented through the Next Hop Resolution Protocol (NHRP).	Chapter 13 and Chapter 14
IP Virtual Private Networks (VPNs)	Reserve IP network resources for use by private networks, such as corporate intranets and extranets. Examples of the kinds of resources you can reserve for an IP VPN include ARP entries, static routes, route maps, and point-to-point label switched paths (LSPs).	Chapter 16
Equal-Cost Multipath Routing	Load balance IP traffic across multiple routes of equal cost. These routes can be learned through BGP, OSPF, and RIP. These routes can also be manually configured static routes.	Chapter 7, Chapter 8, Chapter 9, and Chapter 10
Policy-Based Forwarding	Configure policy-based forwarding, which is a technique for forwarding IP packets based on criteria defined in forwarding policies. Policy-based forwarding allows switches to forward packets based on policies rather than on destination IP addresses.	Chapter 5
IP Multicast Routing Protocols	Configure the Internet Group Management Protocol (IGMP), Distance Vector Multicast Routing Protocol (DVMRP), and Multicast Open Shortest Path First (MOSPF) on switches that run IP Navigator.	Chapter 15
Separate Instances of OSPF	Configure separate instances of OSPF on B-STDX 8000/9000, CBX 500, and GX 550 switches. Whereas in prior releases Virtual Network Navigator (VNN) and IP Navigator shared a common OSPF component, VNN and IP Navigator now have their own OSPF components or *instances:* the VNN instance of OSPF and the IP instance of OSPF. This feature reduces the number of OSPF routes that non-switch equipment such as routers must track, and allows network managers to conceal network resources from nodes outside the Lucent network.	Chapter 9

Feature or Enhancement	Enables You To	See
Multiple IP OSPF Authentication Keys	Configure multiple authentication keys for an IP OSPF interface. You can then activate these keys at different times, making it extremely difficult for network intruders to compromise your security.	Chapter 9
IP OSPF MD5 Authentication	Configure MD5 authentication for IP OSPF interfaces.	Chapter 9
QoS Guarantees for Point-to-Point LSPs	Assign QoS classes and traffic descriptors to point-to-point LSPs, which are circuits used to forward IP traffic through a Lucent network in both directions. ***Note***: *LSPs were formerly known as MPTs.*	Chapter 12
VBR Traffic Descriptor Support for Point-to-Point LSPs	Assign VBR traffic descriptors for point-to-point LSP connections.	Chapter 12
Point-to-Point LSP Support for IP VPNs	Configure one point-to-point LSP connection between two switches per private IP VPN. For example, suppose that you have three IP VPNs that require point-to-point LSP connections between the same two switches. You can configure one point-to-point LSP connection for each VPN. See the *NavisCore IP Navigator Configuration Guide* (Product Code 80114) for details.	Chapter 12
Multicast LSPs	Configure multicast LSPs, which are point-to-multipoint circuits used to forward IP multicast traffic through a Lucent network.	Chapter 12
Multiple Roots for MPT LSPs at Area Border Routers	Support multiple roots at an area border router (the boundary between OSPF areas) for a multipoint-to-point label switched path (MPT LSP). This feature eliminates the 2048-leaf limit for the aggregate of spanned areas. ***Note***: *MPT LSPs were formerly referred to as "reverse MPTs."*	Chapter 12
BGP Route Dampening	Configure BGP route dampening, which suppresses routes that have become unstable. Route dampening reduces bandwidth utilization by reducing the amount of BGP UPDATE and WITHDRAWN messages that are transmitted as a result of route flapping. Routes that flap frequently are suppressed (that is, are not advertised) until a user-defined parameter expires.	Chapter 8
BGP Peer Groups	Configure BGP peer groups, which are groups of BGP neighbors that share the same route maps. Peer groups provide simplified route map management and efficient route map updates.	Chapter 8

Feature or Enhancement	Enables You To	See
BGP MD5 Authentication	Configure a BGP peer to protect its BGP sessions, using MD5 authentication, against the introduction of spoofed TCP segments into a TCP connection stream.	Chapter 8
BGP Regular Expressions	Define UNIX-style regular expressions in route maps for filtering BGP traffic.	Chapter 11
IP Server PVCs through an ATM Network	Create an IP server PVC between two IP Server switch endpoints. In effect, you can create an IP Server PVC between two CBX 500 switches separated by an intermediate ATM network. You can use this type of PVC as an alternative to ATM OPTimum trunks.	Chapter 3
Additional QoS Support for IP Server PVCs	Assign traffic descriptors other than UBR to IP Server PVCs, as well as traffic policing and traffic shaping.	Chapter 3
DTE Automatic DLCI Recognition	Configure the automatic recognition of DLCIs by DTEs. For Frame Relay UNI DTE logical ports that are configured as IP logical ports, IP Navigator will automatically recognize DLCIs received from the DCE in link management interface (LMI) status messages. The received DLCIs become the equivalent of pre-configured DLCIs. This eliminates the need to pre-configure DLCIs.	Chapter 3
New Route Map Sources	Configure route map sources for the VNN instance of OSPF, DVMRP, and MOSPF.	Chapter 11
B-STDX 1-Port Channelized DS3/1/0	Configure IP logical ports on the B-STDX 1-port channelized DS3/1/0 module.	Chapter 1 and Chapter 3
General Enhancements		
Documentation references	Updated documentation references as needed.	Chapter 14 and Chapter 15

Conventions

This guide uses the following conventions, when applicable:

Convention	Indicates	Example
`Courier Regular`	System messages and output, prompts, pathnames, filenames, and command names.	`Please wait...`
`Courier Italics`	Variable text for which you supply a value.	*`cdrompath`*`/docs/atmcfg.pdf`
`Courier Bold`	User input.	> **`show mpt all`**
Menu => Option	A selection from a menu.	NavisCore => Logon
Italics	Book titles, new terms, and emphasized text.	*Network Management Station Installation Guide*
Boxes around text	Notes, warnings, cautions.	See examples below.

▶ Notes provide additional information or helpful suggestions that may apply to the subject text.

⚠ Cautions notify the reader to proceed carefully to avoid possible equipment damage or data loss.

⚠ Warnings notify the reader to proceed carefully to avoid possible personal injury.

Related Documents

This section lists the related Lucent documentation that may be helpful to read.

- *B-STDX 8000/9000 Hardware Installation Guide* (Product Code: 80005)
- *CBX 500 Hardware Installation Guide* (Product Code: 80011)
- *GX 550 Multiservice WAN Switch Hardware Installation Guide* (Product Code: 80077)
- *Network Management Station Installation Guide* (Product Code: 80097)
- *Network Management Station Upgrade Guide* (P/N 104-00099-00)
- *NavisCore NMS Getting Started Guide* (Product Code: 80106)
- *NavisCore Physical Interface Configuration Guide* (Product Code: 80099)
- *NavisCore Frame Relay Configuration Guide* (Product Code: 80100)
- *NavisCore ATM Configuration Guide* (Product Code: 80101)
- *NavisCore Diagnostics Guide* (Product Code: 80105)
- *NavisCore Troubleshooting Guide* (Product Code: 80104)
- *Console Command Reference (for B-STDX 6.5.2.x, CBX 3.5.2.x, and GX 1.5.2.x)* (Product Code: 80125)

All manuals for Core Systems and the *Master Glossary* are available on the Core Systems Documentation Library CD-ROM (Product Code: 80025).

Customer Comments

Customer comments are welcome. Please respond in one of the following ways:

- Fill out the Customer Comment Form located at the back of this guide and return it to us.
- E-mail your comments to cspubs@lucent.com
- FAX your comments to 978-692-1510, attention Technical Publications.

Technical Support

The Lucent Technical Assistance Center (TAC) is available to assist you with any problems encountered while using this Lucent product. Log on to our Customer Support web site to obtain telephone numbers for the Lucent TAC in your region:

`http://www.lucent.com/support`

Acronyms

This guide uses the following acronyms:

Acronym	Description
ABR	area border router
ABS	area border switch
AESA	ATM End System Address
AFI	authority and format identifier
ARP	Address Resolution Protocol
AS	autonomous system
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
CBT	Core Based Trees
CDV	cell delay variation
CIDR	classless inter-domain routing
CIE	client information element
CLR	cell loss ratio
CTD	cell transfer delay
DCC	data country code
DVMRP	Distance Vector Multicast Routing Protocol
ECMP	equal-cost multipath
ESI	end system identifier
FTP	File Transfer Protocol
HO-DSP	high-order domain specific part
IANA	Internet Assigned Numbers Authority
ICD	international country designator
ICMP	Internet Control Message Protocol
IDI	initial domain identifier
IDP	initial domain part

Acronym	Description
IFMP	Ipsilon Flow Management Protocol (RFC 1953)
IGMP	Internet Group Management Protocol
InARP	Inverse Address Resolution Protocol
IP	Internet Protocol
LIS	logical internet subnet
LSA	link state advertisement
LSP	label switched path
MBONE	(Internet) Multicast Backbone
MOSPF	Multicast Open Shortest Path First
MPOA	Multi Protocol Over ATM
MPT	Multipoint-to-Point Tunnel
MTU	maximum transfer unit
NBMA	Non-Broadcast Multiple-Access
NHC	next hop client
NHRP	Next Hop Resolution Protocol
NHS	next hop server
NSAP	Network Service Access Point
NSSA	Not So Stubby Area
OSPF	Open Shortest Path First
PIM-DM	Protocol Independent Multicast – Dense Mode
PIM-SM	Protocol Independent Multicast – Sparse Mode
PNNI	Private Network to Network Interface
PPP	Point-to-Point Protocol
PVC	permanent Virtual Circuit
RARP	Reverse Address Resolution Protocol
RIP	Routing Information Protocol
SCSP	Server Cache Synchronization Protocol

Acronym	Description
SVC	switched virtual circuit
TCP	Transmission Control Protocol
TOS	type of service
TTL	time to live (threshold)
UDP	User Datagram Protocol
VCI	Virtual Channel Identifier
VIF	virtual interface
VNN	Virtual Network Navigator
VPI	virtual path identifier
VPN	Virtual Private Network

1

Overview

This chapter provides an overview of Lucent's IP switching technology for core switches, called IP Navigator™, and describes how Lucent uses IP Navigator to implement the TCP/IP protocol suite into its multiservice core switching platforms.

About IP Switching

Internet Protocol (IP) switching technology allows Lucent's multiservice core switching platforms to assume the characteristics and role of an IP router. The main difference between Lucent's IP switching and traditional IP routing is that in the core of the Lucent network, IP packets are switched instead of routed. In other words, instead of examining IP headers at each hop, Lucent switches examine the IP header only at the ingress and egress ports to the Lucent network. In the core of the network, the switches function as IP hardware forwarding engines. The advantages to implementing IP switching technology over traditional routing include lower-layer packet handling, improved traffic management and throughput, increased performance, and end-to-end Quality of Service (QoS).

In existing Internet Service Provider (ISP) networks, the addition of IP switching allows service providers to optimize data traffic flow by eliminating the need for all data packets to flow through the core router. IP switching also eases the management and control duties of the core routers by reducing the number of routing sessions, eliminating IP table lookups, and in some cases, removing the need for the core router completely.

Lucent's Implementation of IP Switching

Lucent adds its IP Navigator software to existing multiservice WAN platforms enabling service providers to offer standard or enhanced IP services based on end-to-end Quality of Service (QoS). IP Navigator is a software upgrade to the NMS for Lucent's multiservice switch platforms. (For specific hardware and software requirements, refer to the software release notes that accompany your IP Navigator software package.)

IP Navigator enables a B-STDX and/or CBX switch on the edge of a WAN to run standard IP routing protocols. CBX and/or B-STDX switches are on the edge of the Lucent cloud, forwarding packets based on the IP address of the frames. Frame relay, ATM, and Ethernet interfaces are supported. Inside the cloud, a CBX 500 can be used to provide a high-speed ATM backbone, whereby packets are *switched* over automatically established virtual paths.

IP Forwarding

IP forwarding decisions are based on routes that are learned via standard routing protocols running on the switch. Inside the autonomous system (AS), the Open Shortest Path First (OSPF) protocol is the Interior Gateway Protocol (IGP). The Border Gateway Protocol (BGP) is the Exterior Gateway Protocol (EGP). BGP learns the routes to networks in other autonomous systems. On links to CPE routers, Routing Information Protocol (RIP), OSPF, or static routing can be used to learn routes to networks that are reached through these routers.

Routing Protocols

IP Navigator supports a variety of IP routing protocols that are required to communicate with traditional routers. IP Navigator includes all the necessary protocols a service provider needs to offer Internet and Intranet services. These protocols include:

- IGPs
- BGP-4 (used as the EGP)
- Internet and transport protocols
- Multicast protocols

Interior Gateway Protocols

RIP, RIP-2, and OSPF are interior gateway protocols (IGP). An IGP is used to develop the routing tables within a network that is administered by one company or organization. RIP is still widely used in smaller IP networks.

OSPF is the routing protocol that is typically used in new or large IP networks. An expanded version of OSPF is part of the Virtual Network Navigator (VNN), the connection-oriented routing technology used in Lucent switches.

▶ There are two instances of OSPF: one for IP Navigator and one for Virtual Network Navigator (VNN). See Chapter 9, "Configuring IP OSPF and VNN OSPF," for details.

Exterior Gateway Protocols

BGP is an EGP that exchanges routing information between autonomous systems. An AS is a set of routers that has a single routing policy running under a single technical administration. BGP advertises routes between external BGP neighbors or peers, unlike IGPs such as OSPF and RIP, which advertise routes within the same AS. When you configure a list of BGP neighbors and networks, you enable these peers and networks to exchange routing information with the BGP-configured switch. See Figure 8-1 for an example of AS relationships.

The Internet is a collection of autonomous systems. Interconnections among autonomous systems typically do not use IGPs; instead, they use protocols that are classified as EGPs, such as BGP.

▶ Lucent supports BGP-4.

Internet and Transport Protocols

IP Navigator supports the following Internet and transport protocols:

- Address Resolution Protocol (ARP)
- File Transfer Protocol (FTP)
- Internet Control Message Protocol (ICMP)
- Internet Protocol (IP)
- Inverse Address Resolution Protocol (InARP)
- Next Hop Resolution Protocol (NHRP)
- Simple Network Management Protocol (SNMP)
- Telnet Protocol (Telnet)
- Transmission Control Protocol (TCP)
- User Datagram Protocol (UDP)

Multicast Protocols

Multicasting allows a single packet of information to be sent to multiple destinations. Audio and videoconferencing are natural applications for this type of connection. IP Navigator supports Internet Group Management Protocol (IGMP), Multicast OSPF (MOSPF), and Distance Vector Multicast Routing Protocol (DVMRP).

Exchanging Routing Table Information

In any routed network, routers learn about the topology of the network by exchanging routing information. In an IP Navigator network, the same information must be exchanged so that every router-enabled switch shares the same network view.

Each switch running IP Navigator in the Lucent network communicates with every other switch running IP Navigator. Routing tables are maintained in each switch and a master table is maintained in the switch's control processor (CP) for B-STDX 8000/9000 switches and in the switch processor (SP) for CBX 500 switches. Each I/O processor (IOP) module stores routing table information, eliminating a single point of failure in the switch.

Mapping Routes to Virtual Circuits

The process of establishing and mapping routes to virtual circuits is the role of the Virtual Network Navigator (VNN). VNN establishes virtual circuits between entry and exit points in the network. The virtual circuits are then mapped to routes based on the egress node. The process of establishing the virtual circuits is automatic. Topology changes, such as the addition of a new switch, trigger a recalculation of all virtual circuits. More importantly, VNN continually monitors the performance of each virtual circuit, recalculating a new path if a better one exists. The monitoring process is a standard function of VNN, and its functions are expanded to include IP Navigator.

Label Switched Paths

► See "Terminology" on page 1-11 for information on terminology changes related to this section.

Label Switched Paths (LSPs) are a unique feature provided by Lucent. IP Navigator uses LSPs as a means of forwarding IP traffic over switched paths (no intermediate IP lookups) through the Lucent network. LSPs provide an efficient, fault-tolerant, high-performance protocol switching layer that is scalable to a large number of switches in a network. LSPs run across direct or OPTimum trunks which connect CBX 500s and B-STDX 9000s.

Depending on the type of LSP, an LSP can allow:

- Multiple nodes to share the same circuit for transmission to a single destination. This type of LSP is called a *multipoint-to-point tunnel (MPT) LSP.* You can think of this type of LSP as an inverse of the ATM point-to-multipoint (PMP) virtual circuit used to send packets from a source to multiple destinations. On adding a new switch to an IP Navigator network, the switch establishes MPT LSPs to all other switches running IP Navigator in the network. This provides every switch with a path to the new switch.

  MPT LSPs need less circuits than point-to-point tunnels. MPT LSPs require the number of circuits to be equal to the number of nodes, whereas point-to-point tunnels require the number of circuits to be equal to the number of nodes squared.

- A pair of nodes to share a point-to-point connection. This connection overrides a single root-to-leaf connection that would otherwise be part of an MPT LSP. This type of LSP is called a *point-to-point LSP.*

- A single node to use one circuit for transmission to multiple destinations. This type of LSP is called a *multicast LSP.* This type of LSP is similar to an ATM point-to-multipoint virtual circuit and is used to transmit IP multicast traffic.

You can configure end-to-end QoS guarantees for point-to-point LSPs. However, MPT LSPs and multicast LSPs are established using the best route available through the network. MPT LSP and multicast LSP connections do not have a QoS guarantee. Any information transferred over the link is sent with the lowest level priority on the link. For an ATM link, the information is classified as unspecified bit rate/available bit rate (UBR/ABR). The term *best effort* is used to describe this type of service.

Switches in an IP Navigator network automatically establish LSP circuits upon startup. However, you must enable LSPs on switches. For more information on this, see "Enabling MPT LSPs and Multicast LSPs" on page 12-11. VNN calculates the best path from each node to the new switch using its own extended version of OSPF. These paths are used to form the LSP. Major network changes cause VNN to recalculate the LSPs. If configured to do so, VNN will continually monitor the LSPs and recalculate them based on performance. VNN treats the LSPs the same way it treats any other circuit connection. No special configuration is required for LSP monitoring.

See Chapter 12, "Configuring Label Switched Paths," for more information on LSPs.

Absolute QoS

The absolute QoS feature allows IP Navigator to provide an on demand ATM SVC to a destination. IP Navigator uses fields within the IP packet such as source or destination address, IP protocol number, source and destination protocol ports, and the type of service field to map critical IP traffic on to a dynamically established ATM SVC. This feature is especially useful for applications that require a high degree of QoS, such as voice over IP.

The absolute QoS feature is implemented through NHRP, which is the protocol used to establish on demand ATM SVCs. See Chapter 13, "About Next Hop Resolution Protocol" and Chapter 14, "Configuring Next Hop Resolution Protocol" for more information on configuring NHRP.

Policy-Based Forwarding

► See "Terminology" on page 1-11 for information on terminology changes related to this section.

Policy-based forwarding is a technique for forwarding IP packets based on criteria defined in *forwarding policies*. Policy-based forwarding allows switches to forward packets based on policies rather than on destination IP addresses.

You associate forwarding policies with one or more ingress IP logical ports on a switch. If an IP packet received at IP logical port XYZ matches the criteria of a forwarding policy associated with that port, IP Navigator forwards the packet over a specified *policy PVC*. Otherwise, IP Navigator routes the packet on a best-effort basis according to the packet's destination IP address.

You associate IP forwarding policies with existing policy PVCs, which you set up with Quality of Service attributes (QoS class, traffic descriptors such as PCR, SCR).

See Chapter 5, "Configuring Policy-Based Forwarding," for more information on policy-based forwarding.

IP Virtual Private Networks

An IP Virtual Private Network (VPN) is a collection of IP network resources that a public carrier or service provider reserves for private use.

In a traditional IP enterprise network, all resources are owned and controlled by a single organization. To users within the organization, the network appears as a separate routing domain. However, when a public carrier or service provider reserves resources for IP VPNs, each IP VPN has its own view — that is, users of the IP VPN see only the resources reserved for them. Although multiple VPNs may share the same physical topology, each VPN appears to its users as if it were a separate routing domain.

To a network manager, a IP VPN also appears as a separate network. NavisCore allows the network manager to select a specific VPN to be managed. This allows the network manager to see only certain resources configured for the VPN (routes, route maps, and so on).

See Chapter 16, "Configuring IP Virtual Private Networks," for more information on IP VPNs.

Configuration and Management

With the addition of IP Navigator, NavisCore and associated network management server products provide the required support for all IP switching features.The protocols required to configure IP switching include: IP, OSPF, RIP, and BGP. In addition, IP Navigator adds new monitoring functions to enable network administrators to monitor their IP traffic parameters and routing-table contents.

IP Navigator supports the following standard MIBs:

- MIB II
- OSPF v2 MIB
- BGP-4 MIB
- Routing Table MIB
- RIP v2 MIB

In addition, IP Navigator supports the following Internet Engineering Task Force (IETF) draft MIBs:

- RFC Draft IGMP MIB
- RFC Draft DVMRP MIB
- RFC Draft Multicast MIB
- RFC Draft NHRP MIB

Logical Port Configuration

You can configure the following types of logical ports for IP routing:

- IP logical ports on B-STDX 8000/9000 and CBX 500 switches. IP logical ports are ports that support IP routing.
- IP Server logical ports on the CBX 500. IP Server logical ports provide a method of accepting or transmitting IP traffic to or from a cell-based logical port.

▶ NHRP logical ports are supported on B-STDX 8000/9000 and CBX 500 switches. NHRP logical ports support NHRP communications over IP logical ports. You add NHRP logical ports to IP logical ports.

Table 1-1 lists the logical ports and card types that support IP routing on the B-STDX 8000/9000. Table 1-2 lists the logical ports and examples of card types that support IP routing on the CBX 500. Contact your Lucent representative or see your software release notice for an exact list of supported card types.

See Chapter 3, "Configuring IP Logical Ports and IP Servers," and Chapter 16, "Configuring IP Virtual Private Networks," for further details about logical port configuration.

Table 1-1. Logical Ports That Support IP Routing on the B-STDX 8000/9000

Logical Port	Card Types	Encapsulation	Address Resolution
FR UNI-DCE FR UNI-DTE FR NNI	Frame cards[a]	RFC1490	InARP (RFC1293) ARP (RFC1490)
PPP	Frame cards[a]	PPP	N/A
ATM UNI DTE ATM UNI DCE	Frame cards[a]	RFC 1483	InATMARP
ATM UNI DTE ATM UNI DCE	ATM cards[b]	RFC 1483	InATMARP
Ethernet	2-port Ethernet card	IEEE SNAP Ethernet II	ARP
IP VPN Cloud	N/A	N/A	ARP

[a] Frame Card Examples = UIO, 4-T1, 4-E1, DSX-10, HSSI, Ch DS3, Ch DS 3/1/0, 12-E1

[b] ATM Card Examples = ATM CS, ATM DS3, ATM E3, ATM OC3

Table 1-2. Logical Ports Supporting IP Routing on the CBX 500

Logical Port	Card Types	Encapsulation	Address Resolution
FR UNI-DCE FR UNI-DTE FR NNI	6-Port DS3 FR/IP card 4-Port Channelized DS3/1 FR/IP card	RFC 1490	InARP ARP
PPP	4-Port Channelized DS3/1 FR/IP card	PPP	N/A
ATM UNI DTE ATM UNI DCE	ATM cards with an IP Server PVC connection	RFC 1483	InATMARP
Ethernet	4-port Ethernet card	IEEE SNAP Ethernet II	ARP
IP VPN Cloud	N/A	N/A	ARP

Terminology

There have been terminology changes since the last revision of this guide. These changes do not reflect any functional changes.

Table 1-3 describes terminology changes:

Table 1-3. Terminology Changes

Old Term	New Term
IP QoS PVC	Policy PVC
IP Flow Profile	Forwarding Policy
Multipoint-to-Point Tunneling (MPT)	Label Switched Path (LSP)
Multipoint-to-Point (MPT) Point-to-Point Connections	Point-to-Point Label Switched Path (LSP) Connections
Reverse Multipoint-to-Point Tunneling (RMPT)	Multipoint-to-Point Tunnel Label Switched Path (MPT LSP)
Forward Multipoint-to-Point Tunneling (FMPT)	Multicast Label Switched Path (MPT LSP)

2

Configuring Ethernet Logical Ports

This chapter describes how to configure logical ports for the following cards:

- 4-port Ethernet for use on the CBX
- 2-port Ethernet for use on the B-STDX

Each of these Ethernet cards support speeds of up to 100 Mbps full duplex.

Prerequisites

Before you configure an Ethernet logical port, check to make sure that you have:

- Set the Ethernet card’s attributes.
- Defined the physical ports on which the logical port(s) will reside.

For more information about these two tasks, refer to the *NavisCore Physical Interface Configuration Guide*.

Accessing the Logical Port Functions

To access the Logical Port functions in NavisCore:

1. Select the switch to which you want to add a logical port.
2. Log in to NavisCore using either a provisioning or operator password.
3. From the Administer menu, select Lucent Parameters ⇒ Set Parameters.

   The Switch Back Panel dialog box appears. Figure 2-1 illustrates the Switch Back Panel dialog box for a CBX 500.

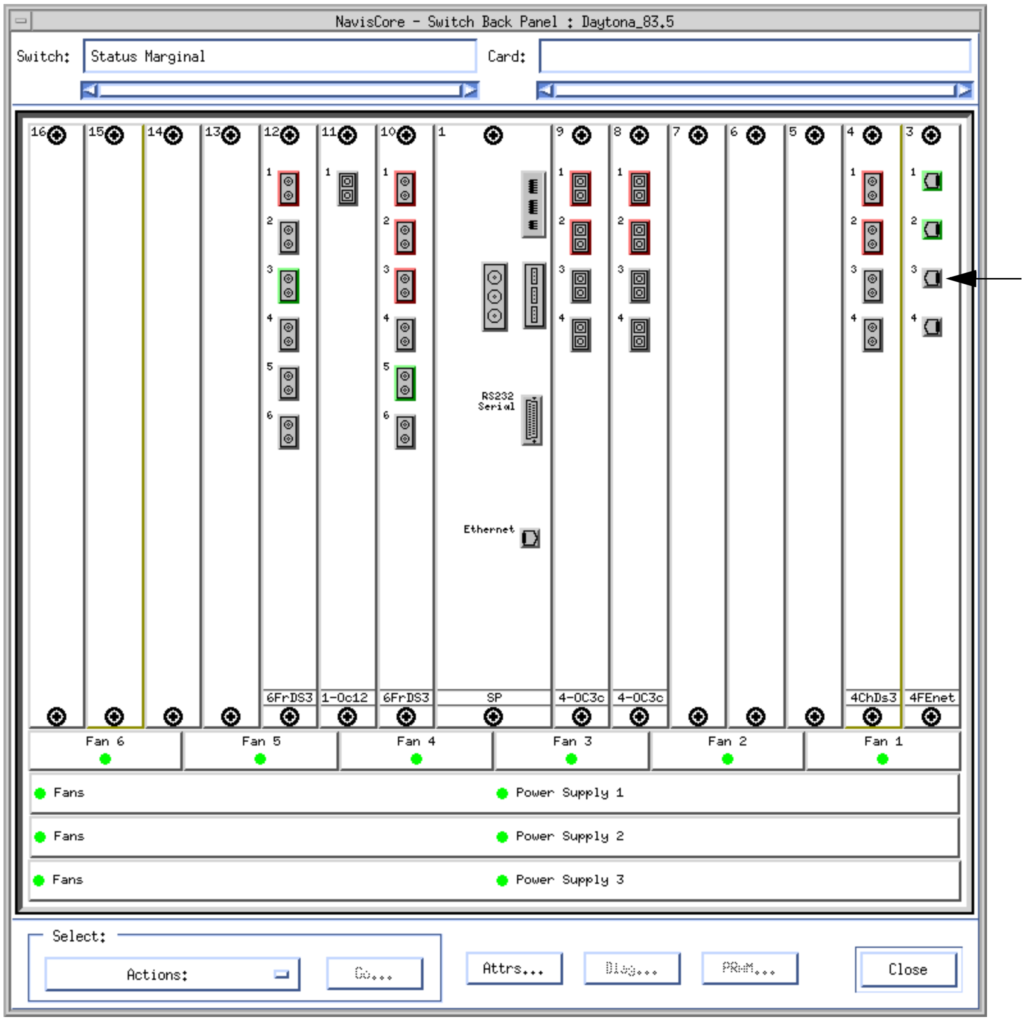

NavisCore - Switch Back Panel : Daytona_83.5

Switch: Status Marginal

Card:

Select the physical port and choose Attrs.

Figure 2-1. Switch Back Panel (CBX) Dialog Box

4. Select the physical port you want to configure. Choose the Attrs button. The Set Physical Port Attributes dialog box appears (see Figure 2-2).

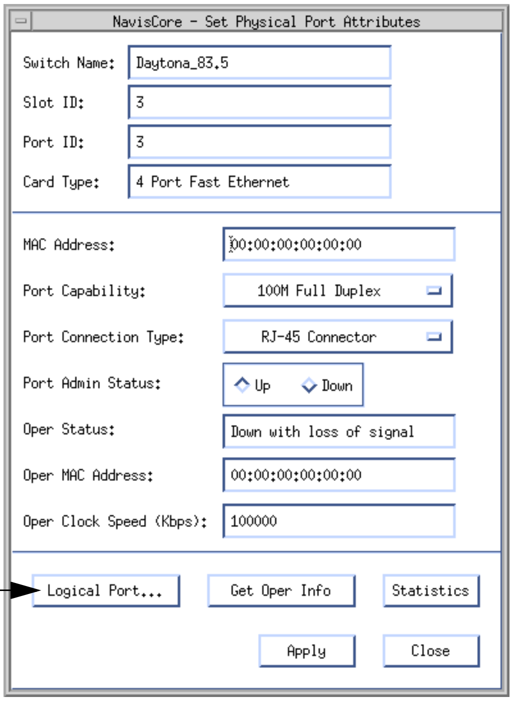

NavisCore - Set Physical Port Attributes

Switch Name: Daytona_83.5

Slot ID: 3

Port ID: 3

Card Type: 4 Port Fast Ethernet

MAC Address: 00:00:00:00:00:00

Port Capability: 100M Full Duplex

Port Connection Type: RJ-45 Connector

Port Admin Status: ◇ Up ◇ Down

Oper Status: Down with loss of signal

Oper MAC Address: 00:00:00:00:00:00

Oper Clock Speed (Kbps): 100000

Choose Logical Port.

Figure 2-2. Set Physical Port Attributes Dialog Box

5. Choose Logical Port. The Set All Logical Ports in PPort dialog box appears (see Figure 2-3).

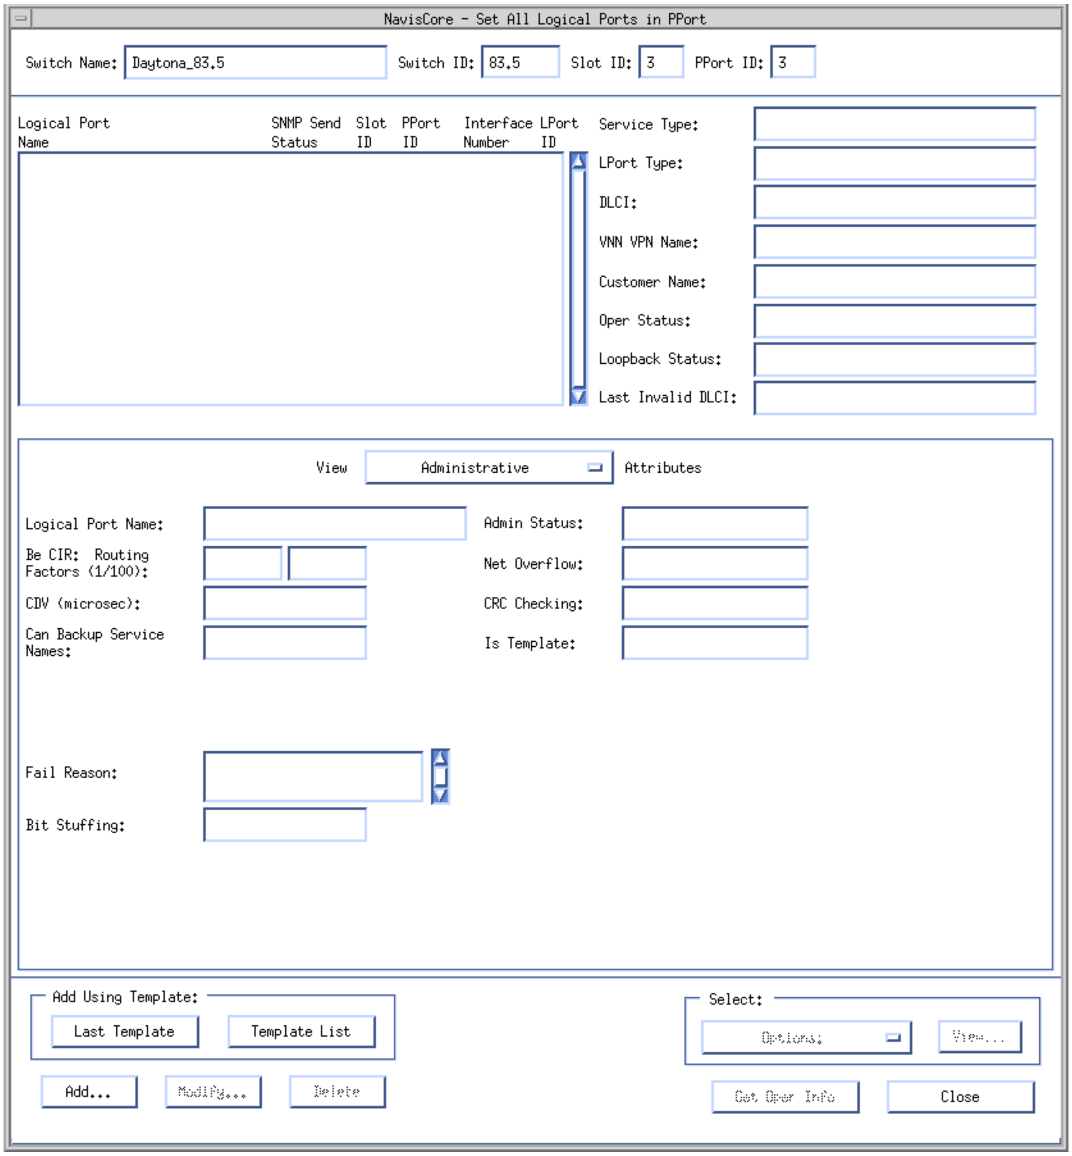

NavisCore - Set All Logical Ports in PPort

Switch Name: Daytona_83.5

Switch ID: 83.5

Slot ID: 3

PPort ID: 3

View Administrative Attributes

Figure 2-3. Set All Logical Ports in PPort Dialog Box

About the Set All Logical Ports in PPort Dialog Box

The Set All Logical Ports In PPort dialog box displays information about an existing logical port or enables you to add a new logical port. It also provides several command buttons that you can use to access additional logical port functions, such as add, modify, and delete logical ports.

Table 2-1 describes the Set All Logical Ports in PPort command buttons.

Table 2-1. Set All Logical Ports in PPort Command Buttons

Field/Command	Action/Description
Add	Adds a new logical port.
Modify	Modifies the selected logical port. The Modify command displays dialog boxes which are similar to those displayed when you Add a logical port; however, you cannot modify the logical port name and the logical port type.
Delete	Deletes the selected logical port.
Get Oper Info	Displays a brief status message of the logical port state.
Add Using Template	If you have already defined a logical port configuration and saved it as a template, use this option to define a new logical port using similar parameters. • Choose Last Template to use the last template you defined for this switch. • Choose Template List to display a list of templates previously defined for this map.
Options	Use the Select: Options button to select the following logical port options for Ethernet logical ports. **IP Parameters** — Displays the Set IP Parameters dialog box (see Figure 3-5 on page 3-10). **Statistics** — Displays the summary statistics for the selected logical port. For more information about summary statistics, see the *NavisCore Diagnostics Guide.* **Diagnostics** — Accesses diagnostic tests for the selected logical port. For more information about diagnostics, see the *NavisCore Diagnostics Guide.* Once you select an option from this list, choose View to access the information.

6. Choose Add to define a new logical port. The Add Logical Port Type dialog box appears (see Figure 2-4).

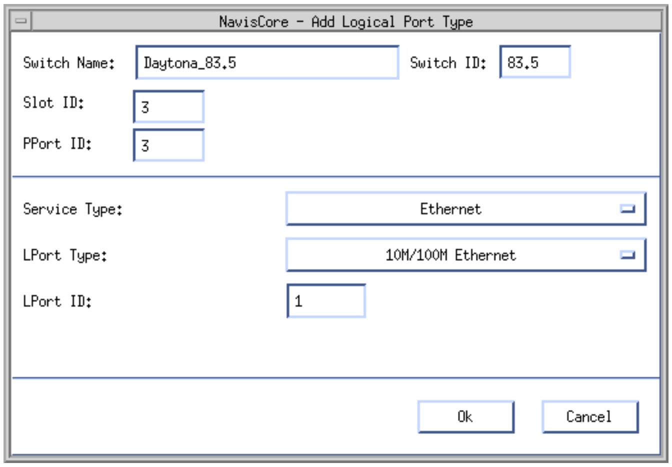
NavisCore - Add Logical Port Type

Switch Name: Daytona_83.5
Switch ID: 83.5
Slot ID: 3
PPort ID: 3

Service Type: Ethernet
LPort Type: 10M/100M Ethernet
LPort ID: 1

Figure 2-4. Add Logical Port Type Dialog Box

7. Accept the displayed values and choose OK. The Add Logical Port dialog box appears.

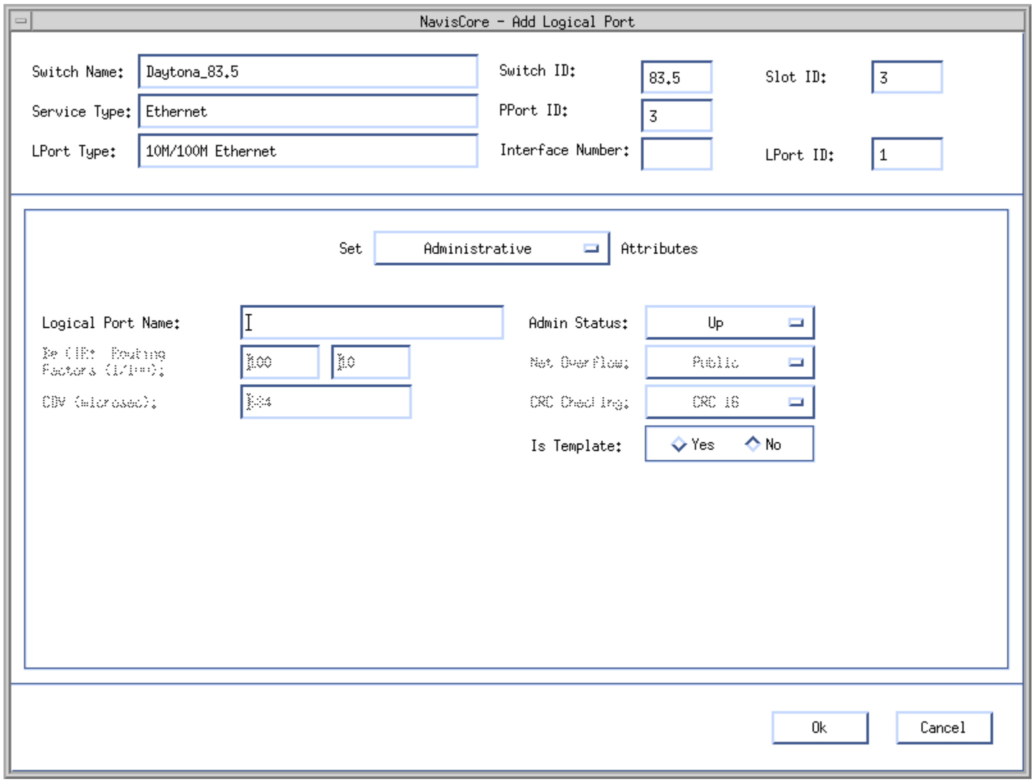
NavisCore - Add Logical Port

Switch Name: Daytona_83.5
Switch ID: 83.5
Slot ID: 3
Service Type: Ethernet
PPort ID: 3
LPort Type: 10M/100M Ethernet
Interface Number:
LPort ID: 1

Set Administrative Attributes

Logical Port Name:
Admin Status: Up
Bw CIR: Routing Factors (1/1==): 100 10
Net Overflow: Public
CDV (microsec): 504
CRC Checking: CRC 16
Is Template: ◇ Yes ◇ No

Figure 2-5. Add Logical Port Dialog Box

The Set Attributes Options Menu

When you define a new logical port, the Add Logical Port dialog box displays a Set Attributes option menu that enables you to set different attributes for each type of logical port. Attributes that you can set include:

Administrative — Sets the logical port name and admin status.

Trap Control — Sets the congestion threshold percentage in which traps are generated and the number of frame errors per minute for each logical port. The supported logical port types are different for each I/O module.

Ethernet Frame — Sets the encapsulation type for transmitted IP frames on an Ethernet port.

Administrative Attributes

Use the Set Administrative Attributes option to complete the fields described in Table 2-2.

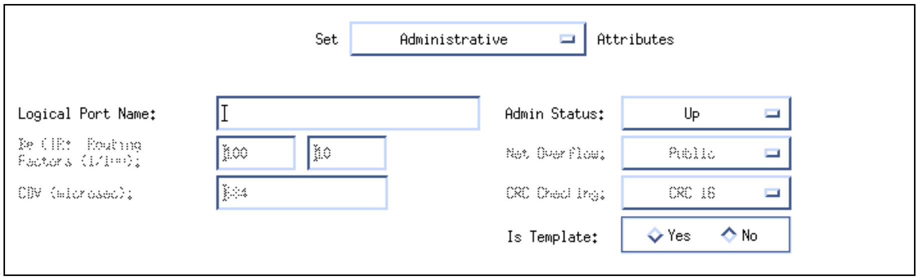

Set [Administrative] Attributes

Logical Port Name:

Be CIR: Routing Factors (1/1==): 100 10

CDV (microsec): 64

Admin Status: Up

Net Overflow: Public

CRC Checking: CRC 16

Is Template: ◇ Yes ◇ No

Figure 2-6. Administrative Attributes for Ethernet Logical Ports

Table 2-2. Administrative Attributes (Ethernet Ports) Fields

Field	Action/Description
Logical Port Name	Enter a unique alphanumeric name for this port. NavisCore uses this name to reference the logical port.
Admin Status	Set the Admin Status to Up (the default) to make the port active. Set the Admin Status to Down to make the port inactive.
Is Template	*(Optional)* Saves these settings as a template to configure another logical port with similar options. To create a template, choose Yes.

Trap Control Attributes

Use the Set Trap Control Attributes option to complete the fields described in Table 2-3.

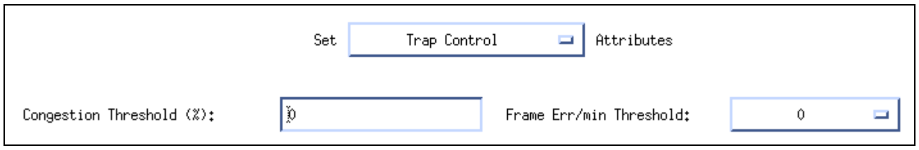

Congestion Threshold (%): 0

Frame Err/min Threshold: 0

Figure 2-7. Trap Control Attributes for Ethernet Logical Ports

Table 2-3. Set Trap Control Attributes (Ethernet Ports) Fields

Field	Action/Description
Congestion Threshold (%)	Enter a value between 0 and 100 to indicate the threshold percentage for generating and sending traps to the NMS for this logical port. A congestion trap is generated and sent to the NMS if the rate of congestion over a one-minute period exceeds the percentage value you enter. Adjust the entered value according to how sensitive this port needs to be to network congestion. Options include: *Low* – Generates a trap at the first sign of congestion. *High* – Generates traps for serious network congestion. *Zero* – (default) Disables the congestion threshold. If you enter zero, no traps are generated for this logical port.
Frame Err/min Threshold	Enter a value from 0 to 16384 to configure the threshold of frame errors on this logical port. If the number of frame errors received in one minute exceeds the specified number, a trap is sent to the NMS. Adjust the entered value according to how sensitive this port needs to be to frame errors. Options include: *Low* – Port is sensitive to frame errors. *High* – Generates traps when a significant number of frame errors occur within a one-minute period. *Zero* – (default) Disables this feature, which prevents traps from being generated for this logical port.

Ethernet Attributes

Use the Set Ethernet Frame Attributes option to complete the fields described in Table 2-4.

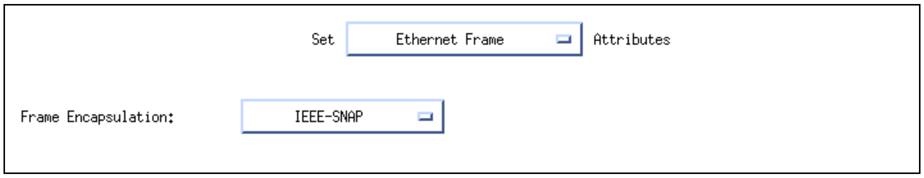

Set Ethernet Frame Attributes

Frame Encapsulation: IEEE-SNAP

Figure 2-8. Ethernet Frame Attributes

Table 2-4. Set Ethernet Frame Attributes (Ethernet Ports) Fields

Field	Action/Description
Frame Encapsulation	Select the Ethernet frame encapsulation type: *Ethernet-II* – (default) The original Ethernet standard frame type. Figure 2-9 illustrates an Ethernet II frame. The value of the Ethernet Type (specified in bytes 13 and 14) indicates the frame type for the packet. If the value of Ethernet type is greater than or equal to hexadecimal value 0x0600 (decimal value 1800), the packet uses an Ethernet II frame type. *IEEE-SNAP* – An IEEE subnetwork access protocol (SNAP) standard frame type that is 8 bytes larger than the Ethernet II type. Figure 2-10 illustrates an IEEE standard frame. The value of the Length (specified in bytes 13 and 14) indicates the frame type for the packet. If the value of the Length is less than or equal to hexadecimal value 0x05DC (decimal value 1500), the packet uses an IEEE-SNAP frame type.

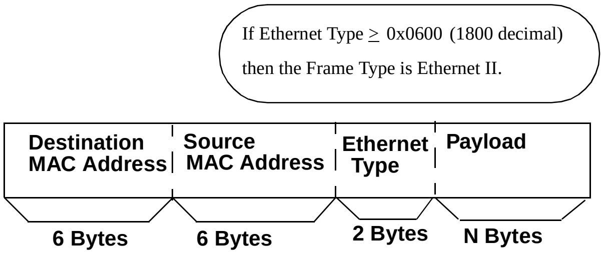

Figure 2-9. Ethernet II Frame Type

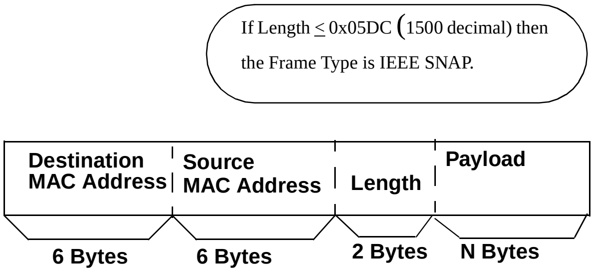

Figure 2-10. IEEE SNAP Frame Type

3

Configuring IP Logical Ports and IP Servers

This chapter describes how to configure:

- **IP logical ports** on B-STDX 8000/9000 and CBX 500 switches. IP logical ports are ports that support IP routing. See Figure 3-1 on page 3-5 for a summary of the IP logical port configuration process.
- **IP server logical ports** on the CBX 500 switch. IP server logical ports provide a method of accepting or transmitting IP traffic on a cell-based card. See Figure 3-14 on page 3-28 for a summary of the IP server logical port configuration process.

► This chapter does not discuss managing IP logical ports and IP server logical ports for private IP virtual private networks (IP VPNs). For more information on IP VPN management, see Chapter 16, "Configuring IP Virtual Private Networks."

Table 3-1 lists the logical ports that support IP routing. Contact your Lucent representative or see your software release notice for an exact list of supported card types.

Table 3-1. Logical Ports that Support IP Routing

Switch	Logical Port	Card Types
B-STDX 8000/9000	FR UNI-DCE FR UNI-DTE FR NNI	Frame cards[a]
	PPP	Frame cards[a]
	ATM UNI DTE ATM UNI DCE	Frame cards[a]
	ATM UNI DTE ATM UNI DCE	ATM cards[b]
	Ethernet	2-port Ethernet cards
CBX 500	Frame Relay	4-port channelized DS3/1 FR/IP card
	Frame Relay	6-port DS3 FR/IP card
	PPP	4-port channelized DS3/1 FR/IP card
	Ethernet	4-port Ethernet card
	ATM UNI DTE ATM UNI DCE	ATM cards with an IP server PVC connection

[a] Frame Card Examples = UIO, 4-T1, 4-E1, DSX-10, HSSI, Ch DS3, Ch DS 3/1/0, 12-E1

[b] ATM Card Examples = ATM CS, ATM DS3, ATM E3, ATM OC3

Prerequisites

Prior to configuring IP services, verify that the following tasks are complete.

- Create a network map.
- Configure the switch parameters. If you are configuring an IP logical port on a CBX 500, you must install and configure a 6-port DS3 FR/IP card, a 4-port Channelized DS3/1 FR/IP card, or a 4-port Ethernet card.
- Configure the physical port parameters.
- If you are configuring an IP logical port on a B-STDX 8000/9000, configure the logical port for Frame Relay, ATM, or Ethernet service.

For more details about these tasks, see Chapter 2, "Configuring Ethernet Logical Ports" and the following guides:

- *NavisCore Frame Relay Configuration Guide* for information about configuring logical ports for frame relay service.
- *NavisCore ATM Configuration Guide* for information about configuring logical ports for ATM service.
- *NavisCore Physical Interface Configuration Guide* for information about configuring cards and physical ports.

About IP Addresses

When you specify the IP address, you must specify the type of IP forwarding the logical port will use. The following two types of IP forwarding are automatically enabled by default:

Unicast — Enables IP forwarding from this logical port to a unicast address.

Broadcast — Enables IP forwarding from this logical port to a broadcast address.

Address Resolution Protocol

A node requires the following information to communicate with another node:

- IP address of the destination node
- Hardware address of the destination node (DLCI for Frame Relay and VPI/VCI for ATM)

IP Navigator uses one of the following protocols to resolve an unknown hardware or IP address:

Address Resolution Protocol (ARP) — Used for Frame Relay, Ethernet, and IP VPN cloud interface configurations when an IP address of a given destination is known, but the destination hardware address (e.g., MAC address or DLCI) is not.

Inverse Address Resolution Protocol (InARP) — Used for Frame Relay and ATM configurations when the destination hardware address (DLCI or VPI/VCI) is known, but the destination IP address is not.

The ARP table resides in the CP/SP memory. An ARP entry is stored for 25 minutes (the same amount of time as a BSD IP stack). All statically configured ARP entries are stored in PRAM. If there is a change in the ARP table, it is sent to the IOP.

Configuring IP Logical Ports

Figure 3-1 illustrates the steps for configuring an IP logical port on a B-STDX 8000/9000 switch or on a CBX 500 switch.

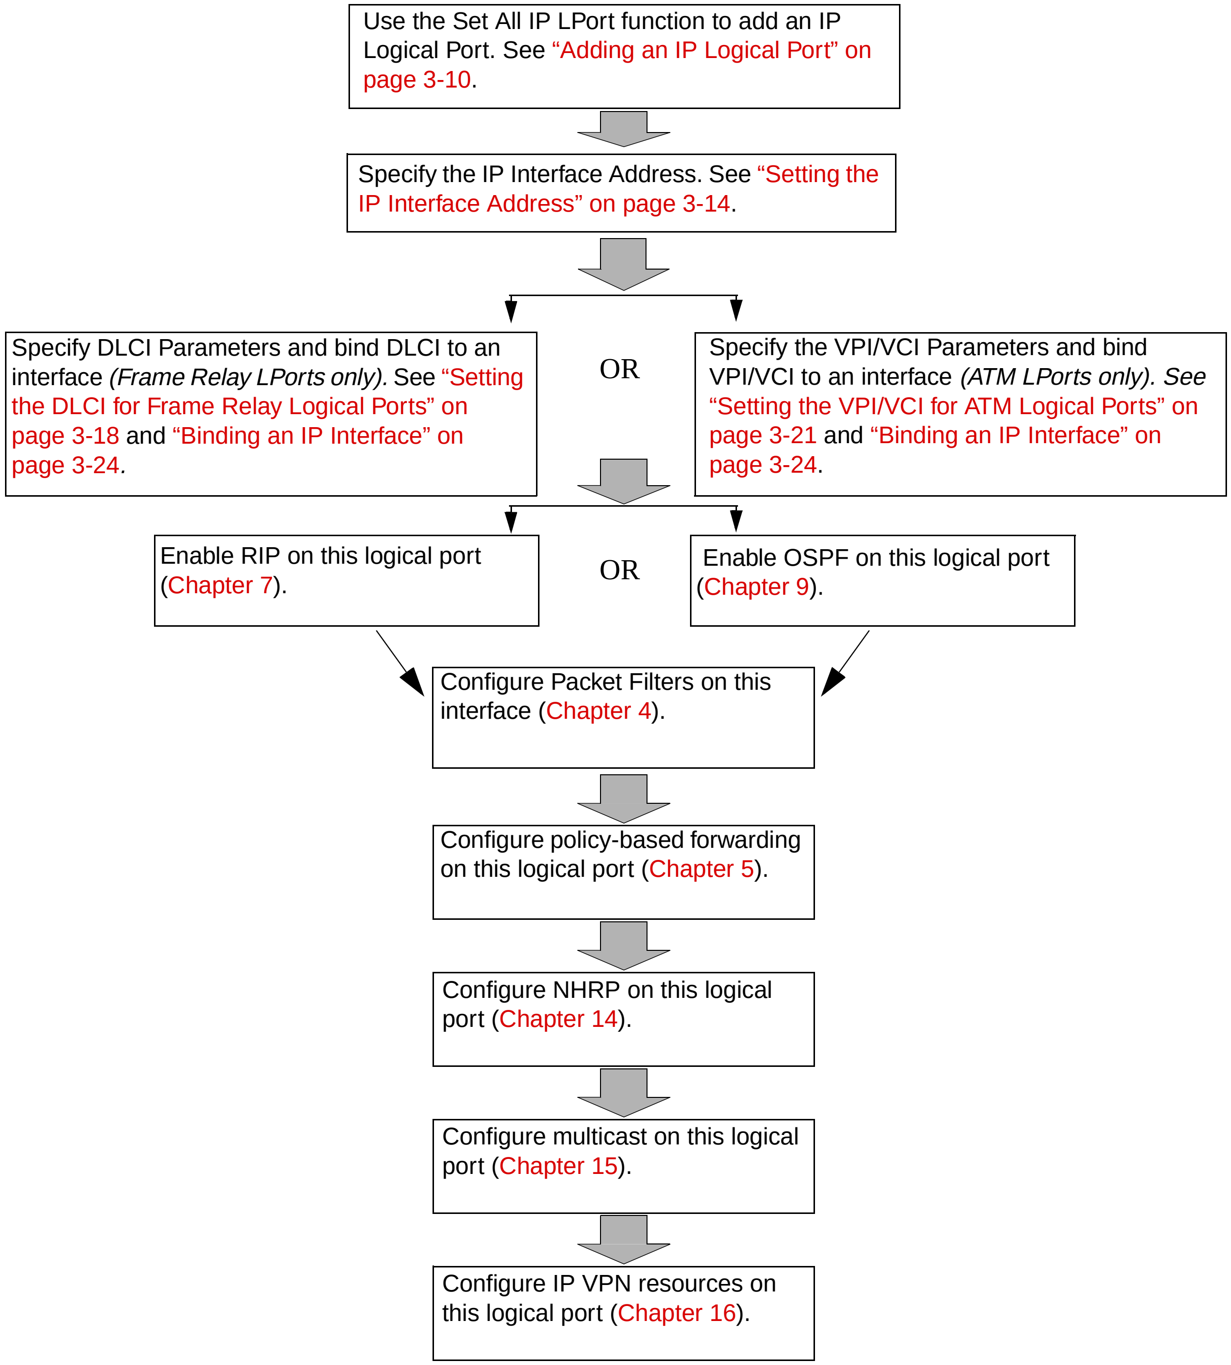
Use the Set All IP LPort function to add an IP Logical Port. See "Adding an IP Logical Port" on page 3-10.
Specify the IP Interface Address. See "Setting the IP Interface Address" on page 3-14.
Specify DLCI Parameters and bind DLCI to an interface (Frame Relay LPorts only). See "Setting the DLCI for Frame Relay Logical Ports" on page 3-18 and "Binding an IP Interface" on page 3-24.
OR
Specify the VPI/VCI Parameters and bind VPI/VCI to an interface (ATM LPorts only). See "Setting the VPI/VCI for ATM Logical Ports" on page 3-21 and "Binding an IP Interface" on page 3-24.
Enable RIP on this logical port (Chapter 7).
OR
Enable OSPF on this logical port (Chapter 9).
Configure Packet Filters on this interface (Chapter 4).
Configure policy-based forwarding on this logical port (Chapter 5).
Configure NHRP on this logical port (Chapter 14).
Configure multicast on this logical port (Chapter 15).
Configure IP VPN resources on this logical port (Chapter 16).

Figure 3-1. IP Logical Port Configuration Process

Accessing Logical Port Parameters

The following section describes how to access the screens that you will use to configure the IP Logical Port Parameters. You can use either of the following methods to access the Set IP Parameters dialog box:

- From the NavisCore Menu. See the following section, "Accessing the Set IP Parameters Dialog Box from the NavisCore Menu" for details about this method of access.
- From the Set All Logical Ports in PPort dialog box. See "Accessing the Set IP Parameters Dialog Box from the Set All Logical Ports Dialog Box."

Accessing the Set IP Parameters Dialog Box from the NavisCore Menu

To access the Set IP Parameters dialog box from the NavisCore menu:

1. Select the appropriate switch icon from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP Lports. The Set All IP LPorts dialog box appears (see Figure 3-2).

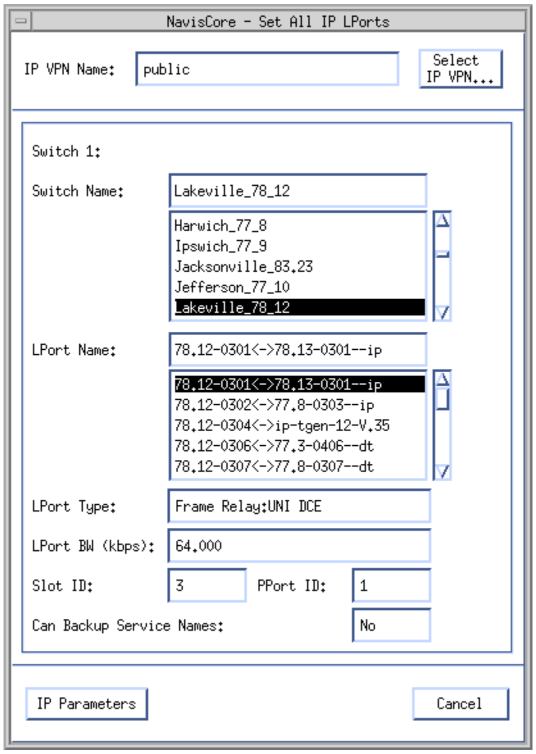

NavisCore - Set All IP LPorts

IP VPN Name: public

Switch 1:

Switch Name: Lakeville_78_12

Harwich_77_8
Ipswich_77_9
Jacksonville_83.23
Jefferson_77_10
Lakeville_78_12

LPort Name: 78.12-0301<->78.13-0301--ip

78.12-0301<->78.13-0301--ip
78.12-0302<->77.8-0303--ip
78.12-0304<->ip-tgen-12-V.35
78.12-0306<->77.3-0406--dt
78.12-0307<->77.8-0307--dt

LPort Type: Frame Relay:UNI DCE

LPort BW (kbps): 64.000

Slot ID: 3 PPort ID: 1

Can Backup Service Names: No

Figure 3-2. Set All IP LPorts Dialog Box

3. Select the LPort name from the list of LPorts.
4. Choose IP Parameters.

If an IP logical port has not defined for the logical port, the Set IP Parameters dialog box appears with the Add IP LPort button displayed (see Figure 3-3).

If an IP logical port has been previously defined for the logical port, the Set IP Parameters dialog box appears without the Add IP LPort button displayed (see Figure 3-5 on page 3-10).

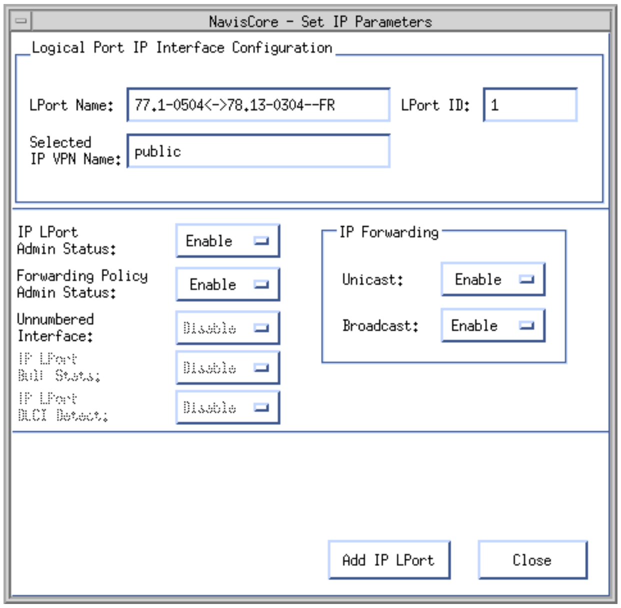

NavisCore - Set IP Parameters

Logical Port IP Interface Configuration

LPort Name: 77.1-0504<->78.13-0304--FR

LPort ID: 1

Selected IP VPN Name: public

IP LPort Admin Status: Enable

Forwarding Policy Admin Status: Enable

Unnumbered Interface: Disable

IP LPort Bulk Stats: Disable

IP LPort DLCI Detect: Disable

IP Forwarding

Unicast: Enable

Broadcast: Enable

Figure 3-3. Set IP Parameters Dialog Box (No IP LPort)

See "Adding an IP Logical Port" on page 3-10 for instructions on adding an IP LPort from the Set IP Parameters dialog box.

Accessing the Set IP Parameters Dialog Box from the Set All Logical Ports Dialog Box

To access the Set IP Parameters dialog box from the Set All Logical Ports in PPort dialog box:

1. From the network map select the appropriate switch icon.
2. From the Administer menu select Lucent Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
3. Select the physical port and choose Attrs. The Set Physical Port Attributes dialog box appears.
4. Perform one of the following actions:
   - **a.** For unchannelized physical ports, choose Logical Port. The Set All Logical Ports in PPort dialog box appears (see Figure 3-4).
   - **b.** For physical ports on channelized DS 3/1/0 modules on B-STDX switches, double click on the channel. The Set Channel Attributes dialog box appears. Choose Logical Port. The Set All Logical Ports in PPort dialog box appears (see Figure 3-4).
   - **c.** For channelized physical ports other than channelized DS 3/1/0 ports on B-STDX switches, select a channel and choose Logical Port. The Set All Logical Ports in PPort dialog box appears (see Figure 3-4).

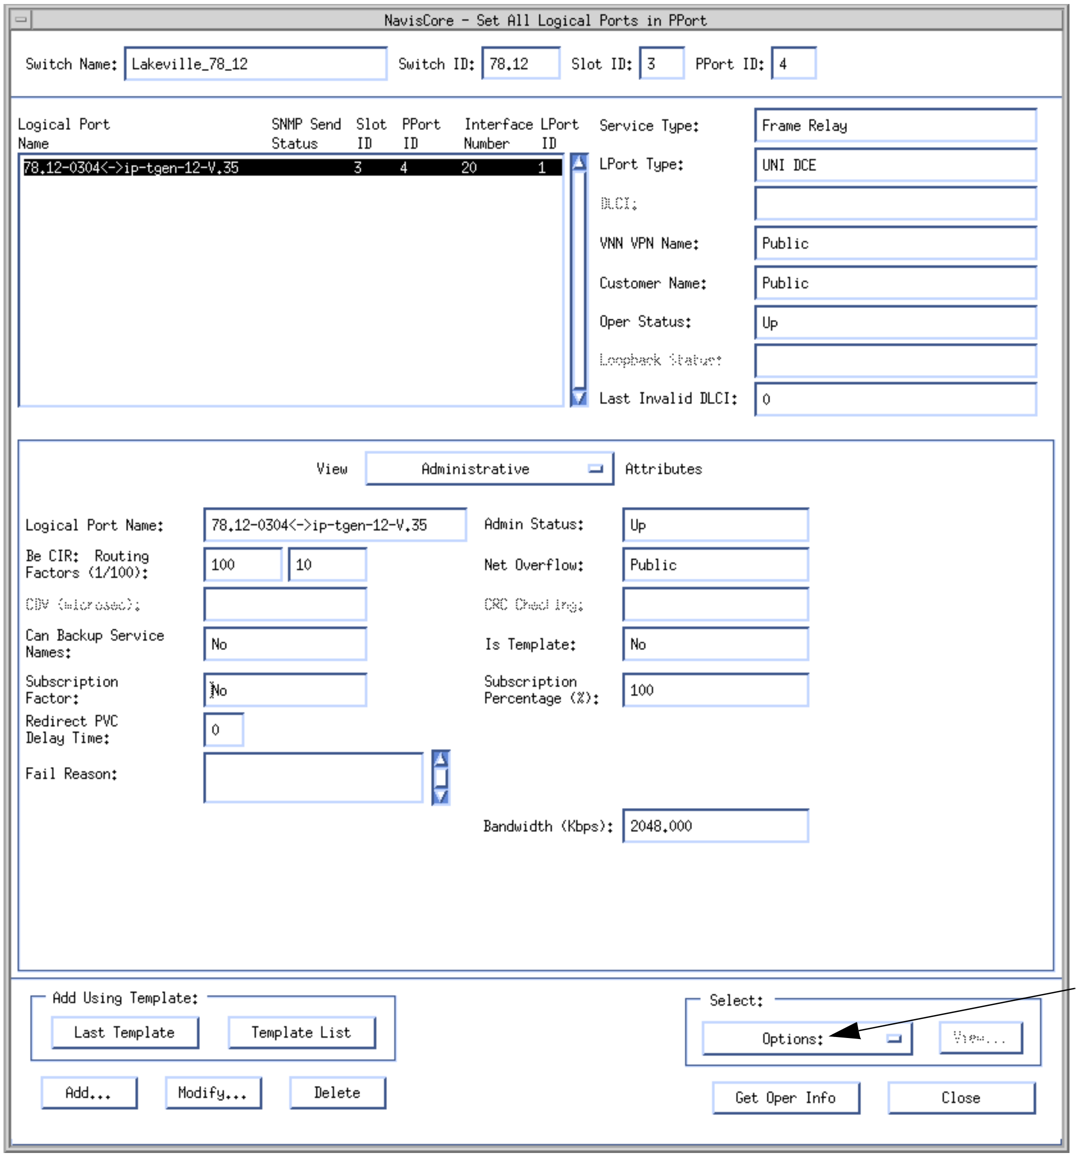

NavisCore - Set All Logical Ports in PPort

Switch Name: Lakeville_78_12
Switch ID: 78.12
Slot ID: 3
PPort ID: 4

Logical Port Name	SNMP Send Status	Slot ID	PPort ID	Interface Number	LPort ID
78.12-0304<->ip-tgen-12-V.35		3	4	20	1

Service Type: Frame Relay
LPort Type: UNI DCE
DLCI:
VNN VPN Name: Public
Customer Name: Public
Oper Status: Up
Loopback Status:
Last Invalid DLCI: 0

View: Administrative Attributes

Logical Port Name: 78.12-0304<->ip-tgen-12-V.35
Be CIR: Routing Factors (1/100): 100 10
CDV (microsec):
Can Backup Service Names: No
Subscription Factor: No
Redirect PVC Delay Time: 0
Fail Reason:
Admin Status: Up
Net Overflow: Public
CRC Checking:
Is Template: No
Subscription Percentage (%): 100
Bandwidth (Kbps): 2048.000

Select IP Parameters from the Options pull-down menu.

Figure 3-4. Set All Logical Ports in PPort

5. Select the logical port name from the list.
6. Select IP Parameters from the Options pull-down menu.
7. Choose Set. The Set IP Parameters dialog box appears (see Figure 3-3 on page 3-7).

Adding an IP Logical Port

To add an IP logical port:

1. Choose Add IP LPort from the Set IP Parameters dialog box (see Figure 3-3 on page 3-7). The second Set IP Parameters dialog box appears (see Figure 3-5).

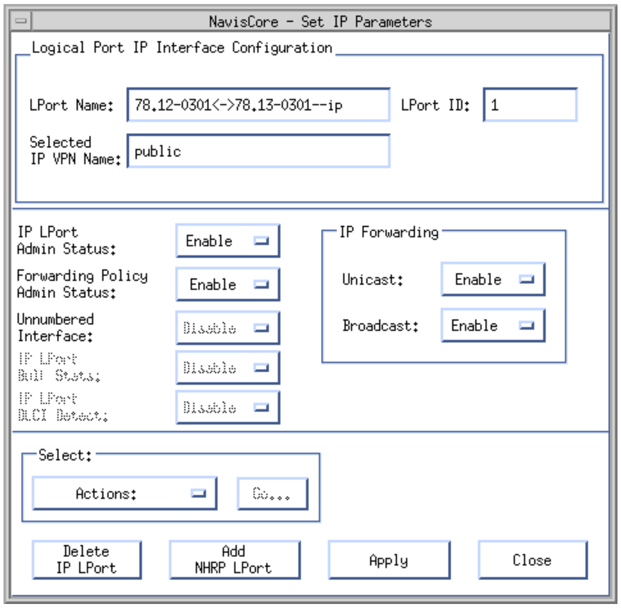

NavisCore - Set IP Parameters

Logical Port IP Interface Configuration

LPort Name: 78.12-0301<->78.13-0301--ip

LPort ID: 1

Selected IP VPN Name: public

IP LPort Admin Status: Enable

Forwarding Policy Admin Status: Enable

Unnumbered Interface: Disable

IP LPort Bulk Status: Disable

IP LPort DLCI Detect: Disable

IP Forwarding

Unicast: Enable

Broadcast: Enable

Figure 3-5. Set IP Parameters Dialog Box (IP LPort Already Added)

See Table 3-2 for a description of each of the buttons on the Set IP Parameters dialog box.

Table 3-2. Set IP Parameters Buttons

Button	Function
Options	Use the Select: Actions button to select the following IP parameter options. **IP Interface** – Displays the Set IP Interface Addresses dialog box, which enables you to configure the IP interface address. See page 3-14 for details. **Packet Filter** – Displays the Assign Logical Port IP Filter dialog box, which enables you to specify inbound and outbound packet filters. See the Chapter 4, "Configuring IP Packet Filters" for more details on this function. **Forwarding Policy** – Displays the Associate LPort Forwarding Policy dialog box, which enables you to add and associate IP forwarding policies. See Chapter 5, "Configuring Policy-Based Forwarding" for more details on this function. **DLCI** – *(For Frame Relay modules only)* Displays the Set All IP Interface Data Link IDs dialog box, which enables you to specify the Data Link Connection Identifier (DLCI) for the IP logical port. See "Setting the DLCI for Frame Relay Logical Ports" on page 3-18 for more details on this function. **VPI/VCI** – *(For ATM modules only)* Displays the Set All IP Interface Data Link IDs dialog box, which enables you to specify the Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) for the IP logical port. See"Setting the VPI/VCI for ATM Logical Ports" on page 3-21 for more details on this function. **Statistics** – Displays the IP Lport Statistics dialog box. See the *NavisCore Diagnostics Guide* for more information about IP logical port statistics. **DVMRP** (*For PPP and Fast Ethernet modules only*) – Displays the Set DVMRP Interface dialog box, which enables you to configure a DVMRP virtual interface for the IP logical port. See Chapter 15, "Configuring IP Multicast Routing" for more information. **IGMP** (*For Fast Ethernet modules only*) – Displays the Set IGMP dialog box, which enables you to configure IGMP for the IP logical port. See Chapter 15, "Configuring IP Multicast Routing" for more information. **Bind IP VPN** (*For PPP and Fast Ethernet modules only*) – Displays the Select IP VPN dialog box, which enables you to assign the IP logical port to an IP VPN. See Chapter 16, "Configuring IP Virtual Private Networks" for more information. Once you select an option from this list, choose Go to access the information.
Select IP VPN	Displays the Select IP VPN dialog box, which enables you to select an IP VPN. Once you select an IP VPN, you manage resources that apply to the selected IP VPN only. See Chapter 16, "Configuring IP Virtual Private Networks" for more information.
Delete IP Lport	Choose this option to delete the IP configuration values for this logical port so that the port is no longer an IP logical port.
Add NHRP LPort	Adds an NHRP logical port. See "Adding and Deleting NHRP Logical Ports" on page 14-10 for more information.

Table 3-2. Set IP Parameters Buttons (Continued)

Button	Function
Apply	Applies any modifications made to the IP logical port parameters. If you make changes to the IP logical port parameters on the Set IP Parameters dialog box, the changes are not actually made until you choose Apply.

2. Specify the necessary IP Parameter values listed in Table 3-3.

Table 3-3. IP Parameter Fields

Field	Action/Description
Lport Name	Displays the name assigned to the LPort at configuration. If you plan to use this logical port as an IP policy PVC, it is suggested that the Lport Name identify the port as an IP policy logical port. When you later associate this logical port with the IP policy PVC, you will have to select the logical port from a list of Lport Names. Chapter 5, “Configuring Policy-Based Forwarding” provides details about IP forwarding policies and policy PVCs.
Lport ID	Displays the ID number that uniquely identifies each logical port.
IP LPort Admin Status	Select one of the following options: *Enable* – Indicates that the port is activated for IP services. *Disable* – Indicates that the port has never been activated for IP services or that the port is offline for diagnostics. A logical port card with an IP LPort Admin Status of *Disable* is not operational for IP routing.
Forwarding Policy Admin Status	Select one of the following options: *Enable* – Enables the use of forwarding policies for the logical port. *Disable* – Disables the use of forwarding policies for the logical port.
Unnumbered Interface	Select one of the following options: *Enable* – Indicates that this IP logical port is not part of a subnet. It does not have a specific address and instead uses the router ID as its source address in IP packets that originate from the interface and are forwarded out of the interface. (The router ID is always the internal address, regardless of whether or not loopbacks are configured.) *Disable* – Indicates that this IP logical port is part of a subnet.
IP LPort Bulk Stats	Select one of the following options: *Enable* – Enables IP bulk statistics on this logical port. *Disable* – Disables IP bulk statistics on this logical port.

Table 3-3. IP Parameter Fields (Continued)

Field	Action/Description
IP LPort DLCI Detect (*Frame Relay UNI DTE Ports Only*)	Select one of the following options: *Enable* – Enables the ability to automatically recognize DLCIs received from the DCE in link management interface (LMI) status messages. The received DLCIs become the equivalent of pre-configured DLCIs. This eliminates the need to pre-configure DLCIs. **Do not enable the ability to automatically recognize DLCIs if the IP logical port is used by private IP VPNs. This feature should be enabled on IP logical ports that handle public traffic only.** See Chapter 16, “Configuring IP Virtual Private Networks” for more information on IP VPNs. *Disable* – (default) Disables the ability to automatically recognize DLCIs received from the DCE in link management interface (LMI) status messages. You must manually configure the DLCIs. **Disable the ability to automatically recognize DLCIs if the IP logical port is used by private IP VPNs.**
Unicast	Select one of the following options: *Enable* – Specifies that IP forwarding will be allowed from this logical port to a unicast address. *Disable* – Indicates that IP forwarding will not be allowed from this logical port to a unicast address. The specific unicast addresses are specified for each IP interface.
Broadcast	Select one of the following options: *Enable* – Specifies that IP forwarding will be allowed from this logical port to a broadcast address. *Disable* – Specifies that IP forwarding is not allowed from this logical port to a broadcast address. The specific broadcast addresses are specified for each IP interface.

The next step is to specify the IP interface address for the IP logical port. See the following section, “Setting the IP Interface Address,” for details.

Setting the IP Interface Address

To specify the IP Interface Address:

1. From the Set IP Parameters dialog box (see Figure 3-3 on page 3-7) select IP Interface and choose Go. The Set IP Interface Addresses dialog box appears (see Figure 3-6):

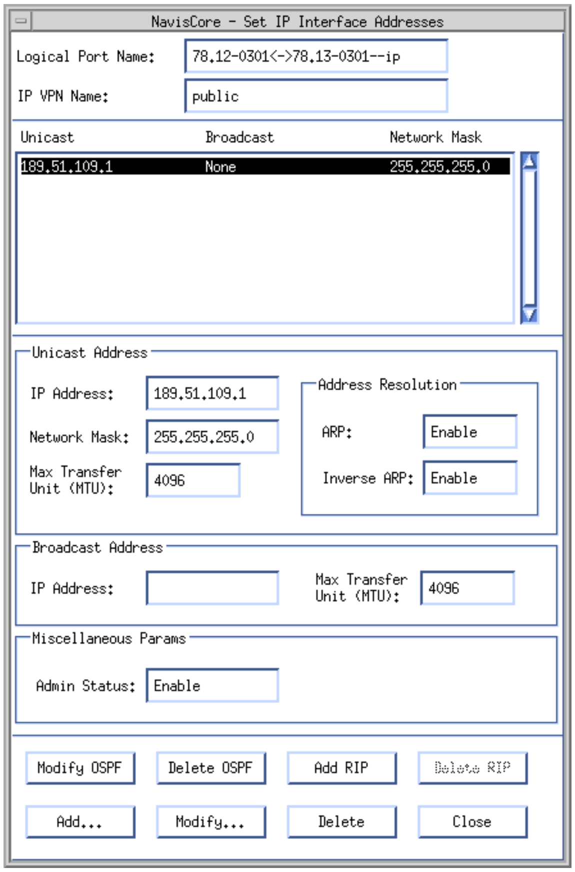

NavisCore - Set IP Interface Addresses

Logical Port Name: 78.12-0301<->78.13-0301--ip

IP VPN Name: public

Unicast	Broadcast	Network Mask
189.51.109.1	None	255.255.255.0

Unicast Address

IP Address: 189.51.109.1

Network Mask: 255.255.255.0

Max Transfer Unit (MTU): 4096

Address Resolution

ARP: Enable

Inverse ARP: Enable

Broadcast Address

IP Address:

Max Transfer Unit (MTU): 4096

Miscellaneous Params

Admin Status: Enable

Figure 3-6. Set IP Interface Addresses Dialog Box

Table 3-4 describes the buttons on the Set IP Interface Addresses dialog box.

Table 3-4. Set IP Interface Addresses Buttons

Button	Function
Add OSPF	Displays the Add OSPF Interface dialog box, which enables you to specify the OSPF parameters for the logical port. This button appears only if you have not yet specified any OSPF parameters for the logical port.
Add RIP	Displays the Add RIP Interface dialog box, which enables you to specify the RIP parameters for the logical port. This button does not appear if you have already configured RIP parameters for the logical port.
Modify OSPF	Displays the Modify OSPF Interface dialog box, which enables you to modify the OSPF parameters for the logical port. This button appears only if you have already specified the OSPF parameters for the logical port.
Modify RIP	Displays the Modify RIP Interface dialog box, which enables you to modify the RIP parameters for the logical port. This button appears only if you have already specified the RIP parameters for the logical port.
Delete OSPF	Displays the Delete OSPF Interface dialog box, which enables you to delete the OSPF parameters for the logical port. This button appears only if you have already specified the OSPF parameters for the logical port.
Delete RIP	Displays the Delete RIP Interface dialog box, which enables you to delete the RIP parameters for the logical port. This button appears only if you have already specified the RIP parameters for the logical port.
Add	Displays the Add Interface Address dialog box.
Modify	Displays the Modify Interface Address dialog box.
Delete	Displays the Delete Interface Address dialog box.

2. Choose Add to add an IP interface address. The Set IP Interface Address dialog box appears (see Figure 3-7).

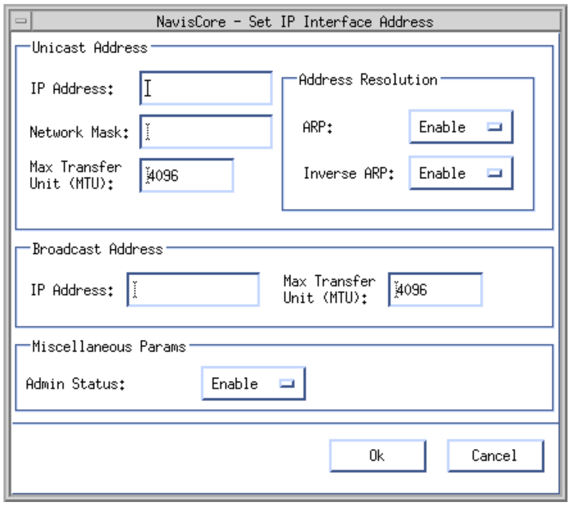

NavisCore - Set IP Interface Address

Unicast Address

IP Address:

Network Mask:

Max Transfer Unit (MTU): 4096

Address Resolution

ARP: Enable

Inverse ARP: Enable

Broadcast Address

IP Address:

Max Transfer Unit (MTU): 4096

Miscellaneous Params

Admin Status: Enable

Figure 3-7. Set IP Interface Address Dialog Box

3. Specify the IP Interface Address values described in Table 3-5.

Table 3-5. IP Interface Address Fields

Field	Action/Description
Unicast Address	
IP Address	The IP address for this interface. Interface addresses can be distributed across IP logical ports as required.
Network Mask	The mask used to determine the subnet of this IP interface. Once this value is set, you cannot use the Modify Interface Address function to modify the network mask value. In order to change the network mask, you must delete the IP interface and then add a new one using the correct network mask.
Max Transfer Unit (MTU)	The maximum size of a packet that can be sent through the physical port. The default value for this field varies depending on the logical port type as follows: **LPort Type** **Default** ATM 9180 Frame Relay 4096 Ethernet 1500
Address Resolution	
ARP (*Frame Relay, Ethernet, VPN Cloud only*)	Select one of the following options: *Enable* – (default) Enables the Address Resolution Protocol (ARP). *Disable* – Disables the ARP. See "Address Resolution Protocol" on page 3-4 for details.

Table 3-5. IP Interface Address Fields (Continued)

Field	Action/Description
Inverse ARP (*Frame Relay and ATM Only*)	Select one of the following options: *Enable* – (default) Enables the Inverse Address Resolution Protocol (InARP). *Disable* – Disables the InARP. See "Address Resolution Protocol" on page 3-4 for details.
Broadcast Address	
IP Address	The address used by this interface for subnet broadcasting.
Max Transfer Unit (MTU)	The maximum size of a packet that can be sent through the physical port. The default value for this field varies depending on the logical port type as follows: **LPort Type** **Default** ATM 9180 Frame Relay 4096 Ethernet 1500
Miscellaneous Params	
Admin Status	*Enable* – (default) Enables IP interface address status. *Disable* – Disables IP interface address status.

4. Choose OK.

After you assign the IP interface address you can then specify the DLCI (for Frame Relay logical ports) or the VPI/VCI (for ATM logical ports). See the following sections for more information on these tasks:

- "Setting the DLCI for Frame Relay Logical Ports" on page 3-18
- "Setting the VPI/VCI for ATM Logical Ports" on page 3-21

Setting the DLCI for Frame Relay Logical Ports

A data link connection identifier (DLCI) number is a 10-bit address that identifies PVCs. The range for an IP DLCI number is a value from 16 to 991.

To specify the DLCI for Frame Relay Logical Ports:

1. From the Administer menu choose Lucent IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears.
2. Select the LPort and choose IP Parameters. The Set IP Parameters dialog box appears.
3. Choose DLCI. The Set All IP Interface Data Link IDs dialog box appears (see Figure 3-8).

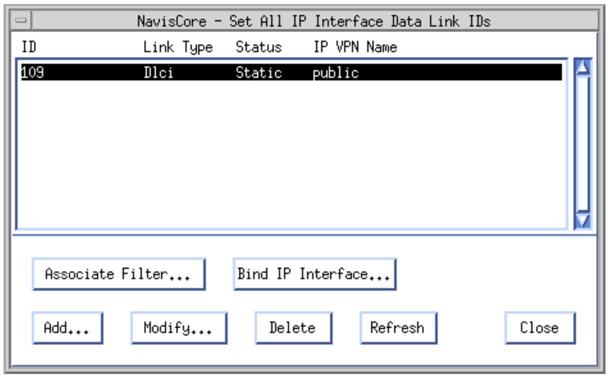

NavisCore - Set All IP Interface Data Link IDs

ID	Link Type	Status	IP VPN Name
109	Dlci	Static	public

Figure 3-8. Set All IP Interface Data Link IDs Dialog Box (FR LPorts)

> If the Status field displays “Static,” the DLCI was statically configured. If the Status field displays “Dynamic,” the DLCI was automatically recognized. DLCIs can be automatically recognized on Frame Relay UNI DTE ports if you enable the IP LPort DLCI Detect field on the Set IP Parameters dialog box. See Table 3-3 on page 3-12 for more information on this field.

The Set All IP Interface Data Link IDs dialog box provides the following buttons:

Table 3-6. Set All IP Interface Data Link ID Buttons

Button	Function
Associate Filter	Displays the Associate IP Circuit Filter List dialog box to enable you to associate a filter with a specific DLCI address. See "Assigning IP Packet Filters to Circuits" on page 4-18 for more information.
Bind IP Interface	See "Binding an IP Interface" on page 3-24 for more information.
Add	Displays the Set IP Protocol Connection ID dialog box to enable you to add a DLCI number.
Modify	Displays the Set IP Protocol Connection ID dialog box to enable you to modify a DLCI number.
Delete	Deletes a selected DLCI number.
Refresh	Refreshes the Status field.

4. Choose Add. The Add Protocol Connection ID dialog box appears (see Figure 3-9).

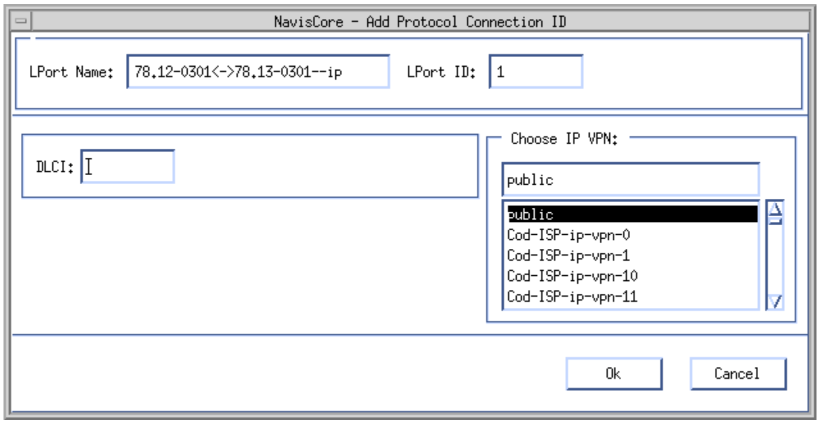

NavisCore - Add Protocol Connection ID

LPort Name: 78.12-0301<->78.13-0301--ip LPort ID: 1

DLCI:

Choose IP VPN: public

public
Cod-ISP-ip-vpn-0
Cod-ISP-ip-vpn-1
Cod-ISP-ip-vpn-10
Cod-ISP-ip-vpn-11

Figure 3-9. Add Protocol Connection ID Dialog Box (FR LPorts)

Specify the field values as described in Table 3-7.

Table 3-7. Add Protocol Connection ID Fields

Field	Action/Description
Lport Name	Displays the name assigned to the LPort at the time of configuration. The LPort ID uniquely identifies each logical port within the physical port.
Lport ID	Displays the ID number that uniquely identifies each logical port.
DLCI	The DLCI value for this IP interface. The range for an IP DLCI number is a value from 16-991. This range of DLCI numbers is available for most link management types. If link management is LMI-1, the maximum value is 1007. Note that DLCI numbers that range from 0 to 15 are reserved.
Choose IP VPN	Select the IP VPN to which you are assigning this DLCI. By default, the DLCI is a public network resource (that is, it is assigned to the public IP VPN). See "Configuring Ingress IP Interfaces for IP VPNs" on page 16-40 for more information.

5. Bind an IP interface to the DLCI. See "Binding an IP Interface" on page 3-24 for more information.

Setting the VPI/VCI for ATM Logical Ports

Virtual path identifiers (VPIs) and virtual channel identifiers (VCIs) are addressing identifiers (similar to Frame Relay's DLCI) that route cell traffic. Every ATM cell header contains both a VCI and a VPI. See the *NavisCore ATM Configuration Guide* for more information on VPIs and VCIs.

> ▶ The VPI and VCI are used only for establishing connections between two ATM entities, not the end-to-end connection.

To specify the VPI and VCI for ATM logical ports:

1. From the Administer menu choose Lucent IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears.

2. Select the LPort and choose IP Parameters. The Set IP Parameters dialog box appears.

3. Choose VPI/VCI. The Set All IP Interface Data Link IDs dialog box appears (see Figure 3-10).

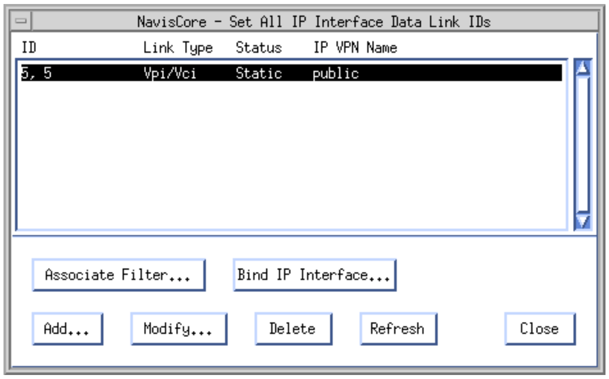

ID	Link Type	Status	IP VPN Name
5, 5	Vpi/Vci	Static	public

Figure 3-10. Set All IP Interface Data Link IDs Dialog Box (ATM LPorts)

The IP Protocol Connection ID dialog box provides the following buttons:

Table 3-8. IP Protocol Connection ID Buttons

Button	Function
Associate Filter	Displays the Associate IP Circuit Filter List dialog box to enable you to associate a filter with a specific VPI/VCI address. See "Assigning IP Packet Filters to Circuits" on page 4-18 for more information.
Bind IP Interface	See "Binding an IP Interface" on page 3-24 for more information.
Add	Displays the Set IP Protocol Connection ID dialog box to enable you to add a VPI/VCI number.
Modify	Displays the Set IP Protocol Connection ID dialog box to enable you to modify a VPI/VCI number.
Delete	Deletes a selected VPI/VCI number.
Refresh	Refreshes the Status field.

4. Choose Add. The Add Protocol Connection ID dialog box appears (see Figure 3-11).

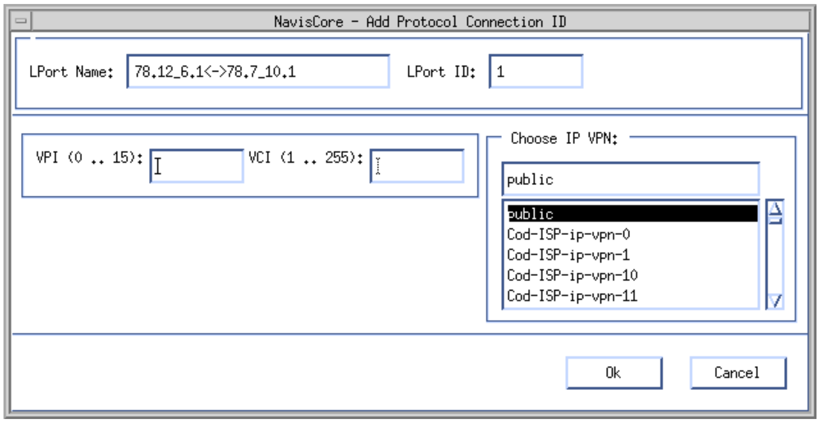
NavisCore - Add Protocol Connection ID

LPort Name: 78.12_6.1<->78.7_10.1
LPort ID: 1
VPI (0 .. 15):
VCI (1 .. 255):
Choose IP VPN: public

Figure 3-11. Add Protocol Connection ID Dialog Box (ATM LPorts)

5. Specify the field values as described in Table 3-9.

Table 3-9. IP Protocol Connection ID For ATM LPorts Fields

Field	Action/Description
LPort Name	Displays the name assigned to the LPort at configuration. The LPort ID identifies the selected logical port.
LPort ID	Displays the unique ID number assigned to the selected logical port.
VPI	The virtual path identifier (VPI). A virtual path (VP) is a group of virtual channels (VCs) carried between two points. VPs provide a way to bundle traffic headed in the same direction. The VPI is an addressing identifier that routes cell traffic. Switching equipment checks the VPI portion of the header to route traffic over certain trunks. This field displays a range of valid values based on the number of valid bits that are configured for this logical port. If the number of valid bits is set to 4, the valid range for the VPI value can be from 0 to 15. If VPI/VCIs already exist on the selected logical port, you can change the number of valid bits; however, the new values must be large enough to support the largest PVC configured for this ATM port. See the *NavisCore ATM Configuration Guide* for a complete description of the valid values for VPI. See "Forwarding Engines on IP Server Cards" on page 3-25 for specific information about VPI values for IP server logical ports.
VCI	The virtual channel identifier (VCI). A virtual channel (VC) is a connection between two communicating ATM entities. This field displays a range of valid values based on the number of valid bits that are configured for this logical port. If the number of valid bits is set to 8, the VCI value can be from 32 to 255 (VCI 0 - 31 are reserved and cannot be used per ATM Forum standards). If VPI/VCIs already exist on the selected logical port, you can change the number of valid bits; however, the new values must be large enough to support the largest PVC configured for this ATM port. See the *NavisCore ATM Configuration Guide* for a complete description of the valid values for VPI. See the "Forwarding Engines on IP Server Cards" on page 3-25 for specific information about VCI values for IP server logical ports.
Choose IP VPN	Select the IP VPN to which you are assigning this VPI/VCI. By default, the VPI/VCI is a public network resource (that is, it is assigned to the public IP VPN). See "Configuring Ingress IP Interfaces for IP VPNs" on page 16-40 for more information.

6. Bind an IP interface to the VPI/VCI. See "Binding an IP Interface" on page 3-24 for more information.

Binding an IP Interface

You must bind an IP interface to each DLCI or VPI/VCI. You can bind and create IP interfaces during the same process. The task of binding IP interfaces with DLCIs and VPI/VCIs is required due to Lucent's support of IP VPNs. See Chapter 16, "Configuring IP Virtual Private Networks" for more information.

To bind an IP interface to a DLCI or VPI/VCI:

1. At the Set All IP Interface Data Link IDs dialog box (see Figure 3-8 on page 3-18 or Figure 3-10 on page 3-21), select the DLCI or VPI/VCI to which you want to bind an IP interface.
2. Choose Bind IP Interface. The Bind IP Interface Address to Protocol ID dialog box appears (see Figure 3-12). Note that, although the dialog box in Figure 3-12 is for a DLCI, the dialog box for a VPI/VCI is almost identical. The only difference is that the dialog box for a VPI/VCI displays the VPI and VCI instead of the DLCI.

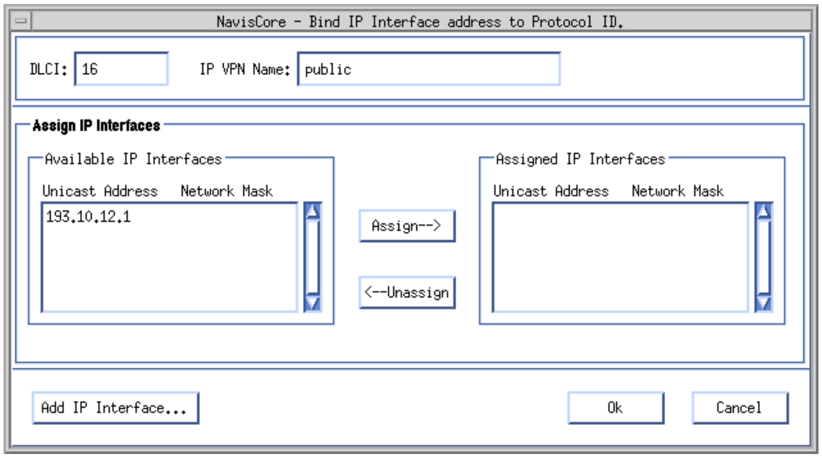
NavisCore - Bind IP Interface address to Protocol ID.

DLCI: 16

IP VPN Name: public

Available IP Interfaces

Unicast Address	Network Mask
193.10.12.1	

Figure 3-12. Bind IP Interface Address to Protocol ID Dialog Box (DLCI)

3. Select the IP interface you want to bind to the DLCI or VPI/VCI in the Available IP Interfaces column. Keep in mind that you can also create an IP interface at this time by choosing the Add IP Interface button. See "Setting the IP Interface Address" on page 3-14 for more information.
4. Choose Assign. The address moves to the Assigned IP Interfaces column.
5. Choose OK.

▶ At any time, you can delete the IP interface binding by selecting the IP interface from the list of assigned IP interfaces and choosing Unassign. Choose OK after you have deleted the binding.

About IP Server Logical Ports on the CBX 500

You can configure logical ports as IP server logical ports on either of the following CBX 500 cards:

- 6-port DS3 FR/IP Card
- 4-port Channelized DS3/1 FR/IP Card
- 4-port Ethernet Card

The purpose of an IP server logical port is to provide a method of accepting or transmitting IP traffic from or to a cell-based port. All IP traffic entering and exiting a CBX 500 ATM cell card must be transmitted through an IP server logical port that is configured on either a 6-port DS3 FR/IP card, 4-port DS3/1 FR/IP card, or on a 4-port Ethernet card.

The number of IP server cards is only limited by the number of slots in the switch.

Forwarding Engines on IP Server Cards

There are two forwarding engines (FEs) on a 6-port DS3 FR/IP card and on a 4-port Ethernet card. There is one FE on a 4-port Channelized DS3/1 FR/IP card. FEs reassemble cells and perform IP lookups. The NMS identifies the two FEs on an IP server card as Server 1 and Server 2 (or as just Server 1 for a 4-port Channelized DS3/1 FR/IP card).

Each FE has hard-coded values for the VPI and VCI bit parameters. The value for the VPI bits parameter is permanently set to 6. For this reason, **the maximum number of IP logical ports that you can create for each IP server card is 64 (0-63).** Therefore, on 6-port DS3 FR/IP cards and 4-port Ethernet cards, if you define 64 IP logical ports on the first FE of the card (which is identified in the NMS as Server 1), you cannot define any IP logical ports on the second FE (Server 2). You can create any number of IP logical ports (up to a maximum of 64) between the two FEs.

On 4-port Channelized DS3/1 FR/IP cards, you have only one FE; therefore, you create all 64 IP logical ports on that FE.

The value for the VCI bits parameter is permanently set to 8. For this reason, you can create up to 225 (2^8 - 31) PVCs for each IP logical port.

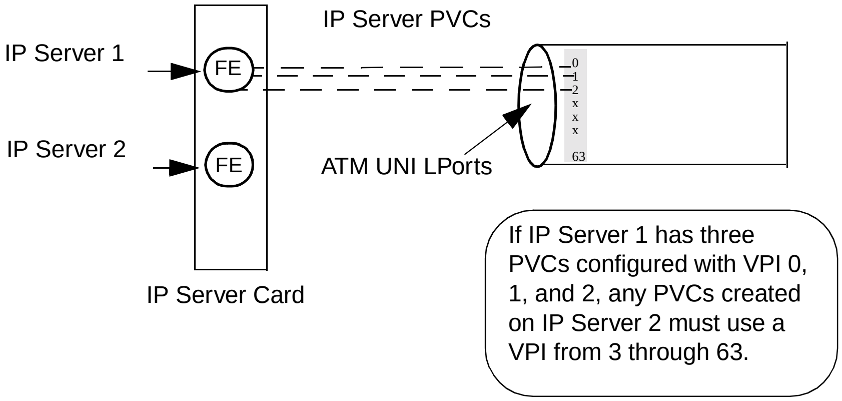

If IP Server 1 has three PVCs configured with VPI 0, 1, and 2, any PVCs created on IP Server 2 must use a VPI from 3 through 63.

Figure 3-13. VPI Parameters for IP Server Cards with Two FEs

IP Server Logical Ports

IP server logical ports on IP server cards are virtual ports. For this reason, physical port numbers on 6-port DS3 FR/IP cards start at 7 and physical port numbers on 4-port Ethernet cards and 4-port Channelized DS3/1 FR/IP cards start at 5.

You use the Set IP Server function on the Administer menu to configure IP server logical ports on a CBX 500. You can create multiple IP interfaces for each IP server logical port. See "Creating an IP Server Logical Port" on page 3-29 for more information.

Bandwidth Allocation

Multiple PVCs can be defined for an IP server logical port.

Logical port bandwidth on an IP server logical port must be sufficient to support all of the PVCs traversing the port. For this reason, before you configure IP server logical ports, **you must plan for the total amount of bandwidth that will be required for all of the PVCs that are associated with the IP server logical port.** If you assign all of the bandwidth to the first IP server logical port, there will be no bandwidth available for PVCs that are configured for subsequent IP server logical ports.

► If you want to create multiple IP server logical ports on an IP server card, you must be aware of the following requirements:

- All of the logical port bandwidth cannot be allocated to the first IP server logical port that you define (a Direct UNI-DCE LPort).
- For subsequent IP logical ports, a VPI start/stop range must be defined. There cannot be any overlap between logical ports. For example, if the first virtual IP server logical port is configured with a VPI start/stop range of 1 and 4, the second virtual IP server logical port cannot use a VPI start/stop range that includes any VPIs already defined.

Configuring IP Server Logical Ports

Figure 3-14 illustrates the steps for configuring an IP server logical port on a CBX 500 switch.

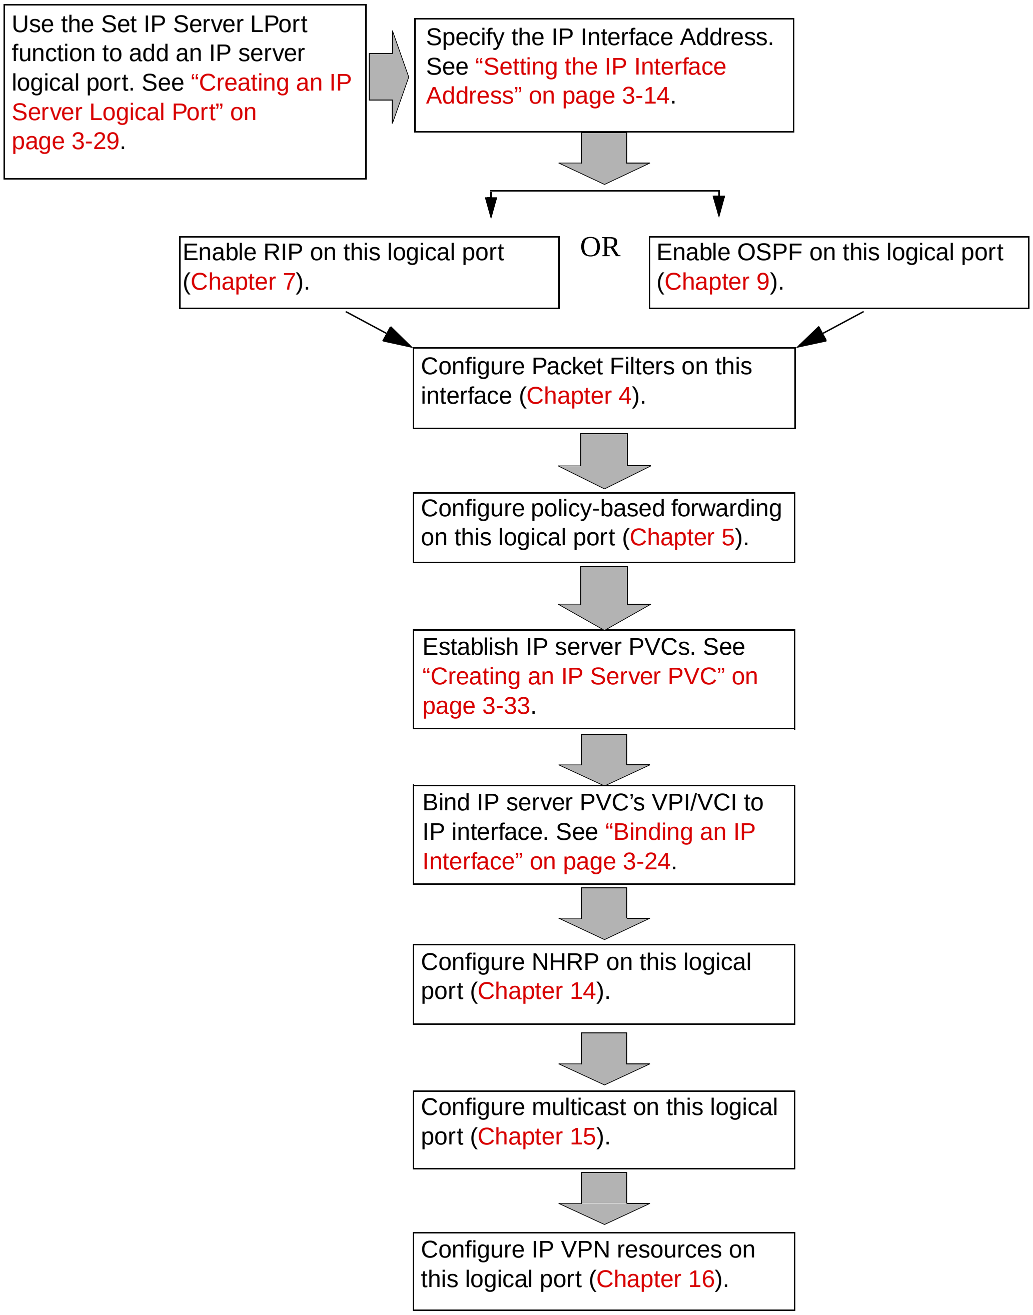
Use the Set IP Server LPort function to add an IP server logical port. See "Creating an IP Server Logical Port" on page 3-29.
Specify the IP Interface Address. See "Setting the IP Interface Address" on page 3-14.
Enable RIP on this logical port (Chapter 7).
OR
Enable OSPF on this logical port (Chapter 9).
Configure Packet Filters on this interface (Chapter 4).
Configure policy-based forwarding on this logical port (Chapter 5).
Establish IP server PVCs. See "Creating an IP Server PVC" on page 3-33.
Bind IP server PVC's VPI/VCI to IP interface. See "Binding an IP Interface" on page 3-24.
Configure NHRP on this logical port (Chapter 14).
Configure multicast on this logical port (Chapter 15).
Configure IP VPN resources on this logical port (Chapter 16).

Figure 3-14. Configuring IP Server Logical Ports on the CBX 500

Creating an IP Server Logical Port

To set an IP server logical port from the NavisCore menu:

1. Select the appropriate CBX 500 switch icon from the network map.
2. Select Lucent IP Parameters ⇒ Set IP Servers ⇒ Set IP Server LPorts from the Administer menu. The Show IP Servers dialog box appears (Figure 3-15).

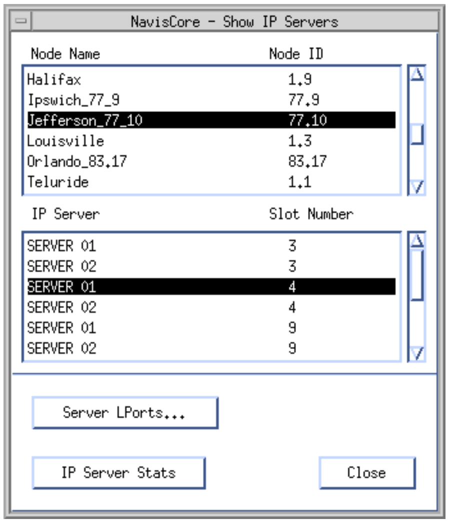

Node Name	Node ID
Halifax	1.9
Ipswich_77_9	77.9
Jefferson_77_10	77.10
Louisville	1.3
Orlando_83.17	83.17
Teluride	1.1

IP Server	Slot Number
SERVER 01	3
SERVER 02	3
SERVER 01	4
SERVER 02	4
SERVER 01	9
SERVER 02	9

Figure 3-15. Show IP Servers Dialog Box

The Show IP Servers dialog box lists all of the CBX 500 switches in your network. If a switch has one or more IP server cards installed, the cards are listed in the IP Server group box. The switch must have an IP server card installed before you can set an IP logical port. For detailed instructions about how to configure an IP server card and physical port, see the *NavisCore Physical Interface Configuration Guide.*

The buttons on the Show IP Servers dialog box are described in Table 3-10.

Table 3-10. Show IP Servers Buttons

Button	Function
Server LPorts	Displays the Set All Logical Ports in IP Server PPort dialog box (see Figure 3-16) to enable you to add an IP server logical port.
IP Server Stats	Displays the Physical Port Summary Statistics dialog box to enable you to display the statistics for a selected IP server port. See the *NavisCore Diagnostics Guide* for more details about physical port statistics.

3. Select the IP server name from the list. There are two IP server FEs available for each 6-port DS3 FR/IP card and each 4-port Ethernet card on the switch. There is one IP server FE available for each 4-port Channelized DS3/1 FR/IP card on the switch.

4. Choose Server LPorts.The Set All Logical Ports in IP Server PPort dialog box appears (see Figure 3-16).

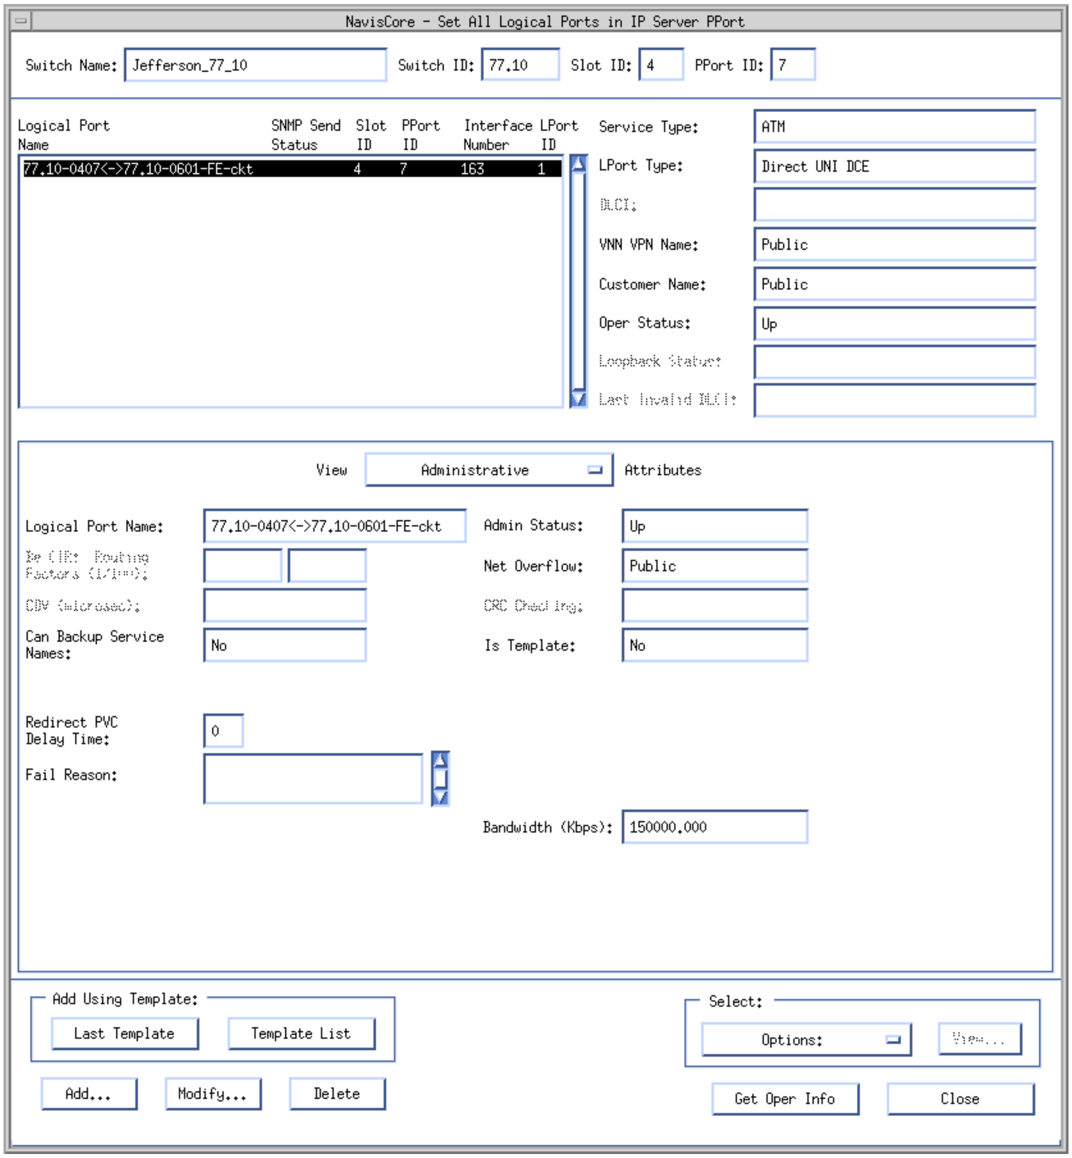

NavisCore - Set All Logical Ports in IP Server PPort

Switch Name: Jefferson_77_10 Switch ID: 77.10 Slot ID: 4 PPort ID: 7

Logical Port Name	SNMP Send Status	Slot ID	PPort ID	Interface Number	LPort ID
77.10-0407<->77.10-0601-FE-ckt		4	7	163	1

Service Type: ATM
LPort Type: Direct UNI DCE
DLCI:
VNN VPN Name: Public
Customer Name: Public
Oper Status: Up
Loopback Status:
Last Invalid DLCI:

View: Administrative Attributes

Logical Port Name: 77.10-0407<->77.10-0601-FE-ckt
Be CIR: Routing Factors (1/1**):
CDV (microsec):
Can Backup Service Names: No
Redirect PVC Delay Time: 0
Fail Reason:

Admin Status: Up
Net Overflow: Public
CRC Checking:
Is Template: No
Bandwidth (Kbps): 150000.000

Figure 3-16. Set All Logical Ports in IP Server PPort

5. Choose Add. The Add Logical Port Type dialog box appears (see Figure 3-17).

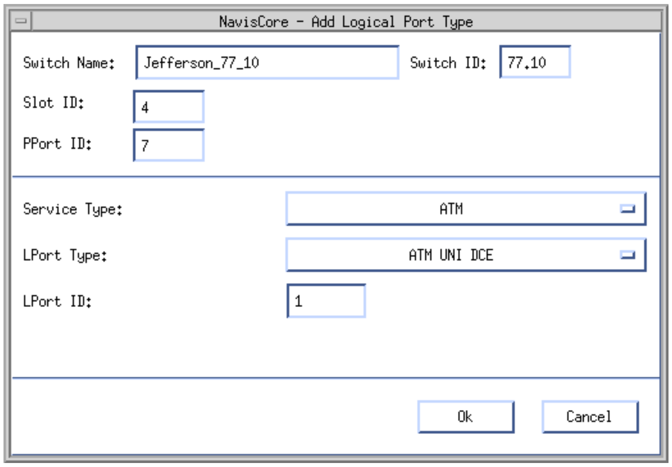

Switch Name: Jefferson_77_10

Switch ID: 77.10

Slot ID: 4

PPort ID: 7

Service Type: ATM

LPort Type: ATM UNI DCE

LPort ID: 1

Figure 3-17. Add Logical Port Type

6. Complete the Add Logical Port Type dialog box as follows:

LPort Type — Select ATM UNI DCE.

LPort ID — Defaults to 1 for this type of configuration and cannot be changed.

7. Choose OK. The Add Logical Port dialog box appears (see Figure 3-18).

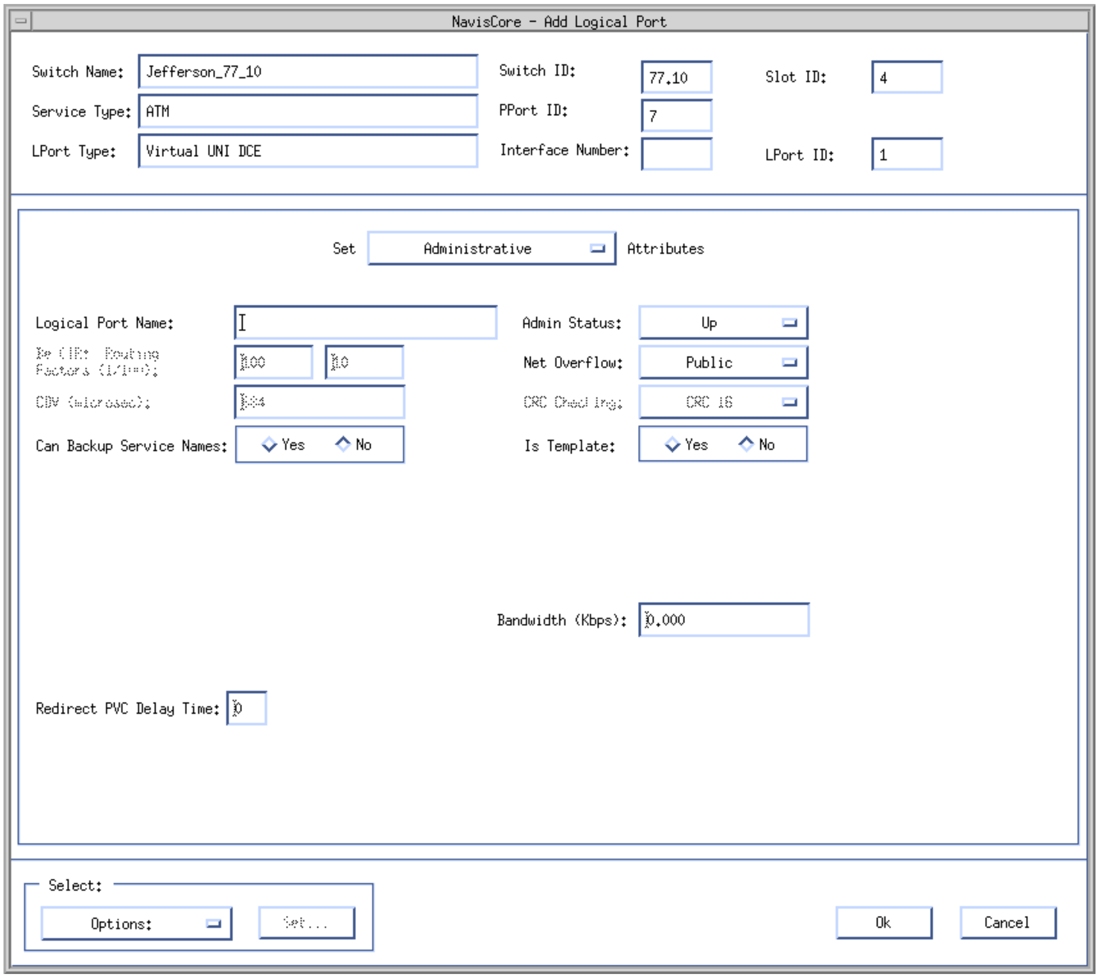

NavisCore - Add Logical Port

Switch Name: Jefferson_77_10
Service Type: ATM
LPort Type: Virtual UNI DCE
Switch ID: 77.10
PPort ID: 7
Interface Number:
Slot ID: 4
LPort ID: 1

Set Administrative Attributes

Logical Port Name:
Be CIR: Routing Factors (1/1==): 100 10
CDV (microsec): 104
Can Backup Service Names: ◇ Yes ◇ No
Admin Status: Up
Net Overflow: Public
CRC Checking: CRC 16
Is Template: ◇ Yes ◇ No

Bandwidth (Kbps): 0.000

Redirect PVC Delay Time: 0

Figure 3-18. Add Logical Port Administrative Attributes Dialog Box

8. Configure the logical port as an ATM UNI DCE port using the configuration instructions in the *NavisCore ATM Configuration Guide.*

9. The next step is to configure IP interfaces for the IP logical port. See "Setting the IP Interface Address" on page 3-14.

IP Server PVCs on the CBX 500

For each IP-enabled ATM logical port, you need a working PVC from an ATM logical port to an IP server logical port.

You can also create an IP server PVC between IP server endpoints on the same CBX 500 switch or two different CBX 500 switches. The CBX 500 switches may be separated by an intermediate ATM network. You can use this feature as an alternative to ATM OPTimum trunks.

Creating an IP Server PVC

To create an IP server PVC from the NavisCore menu:

1. Select the appropriate CBX 500 switch icon from the network map.
2. Select Lucent IP Parameters ⇒ Set IP Servers ⇒ Set IP Server PVCs from the Administer menu. The Set All IP Server PVCs on Map dialog box appears (see Figure 3-19).

   The Set All IP Server PVCs on Map dialog box initially displays no defined circuit names. To display all of the defined circuit names, position the cursor in the *Search by Name* field and press Enter. This search may take several minutes depending on your configuration.

   For a partial search, enter the selected search criteria in the *Search by Name* field. To use a wildcard search to find a specific circuit name, you can:

   - Use an * to match any number of characters
   - Use a ? to match a single character
   - Use a * to match the * character
   - Use a \? to match the ? character
   - Use a \\ to match the \ character

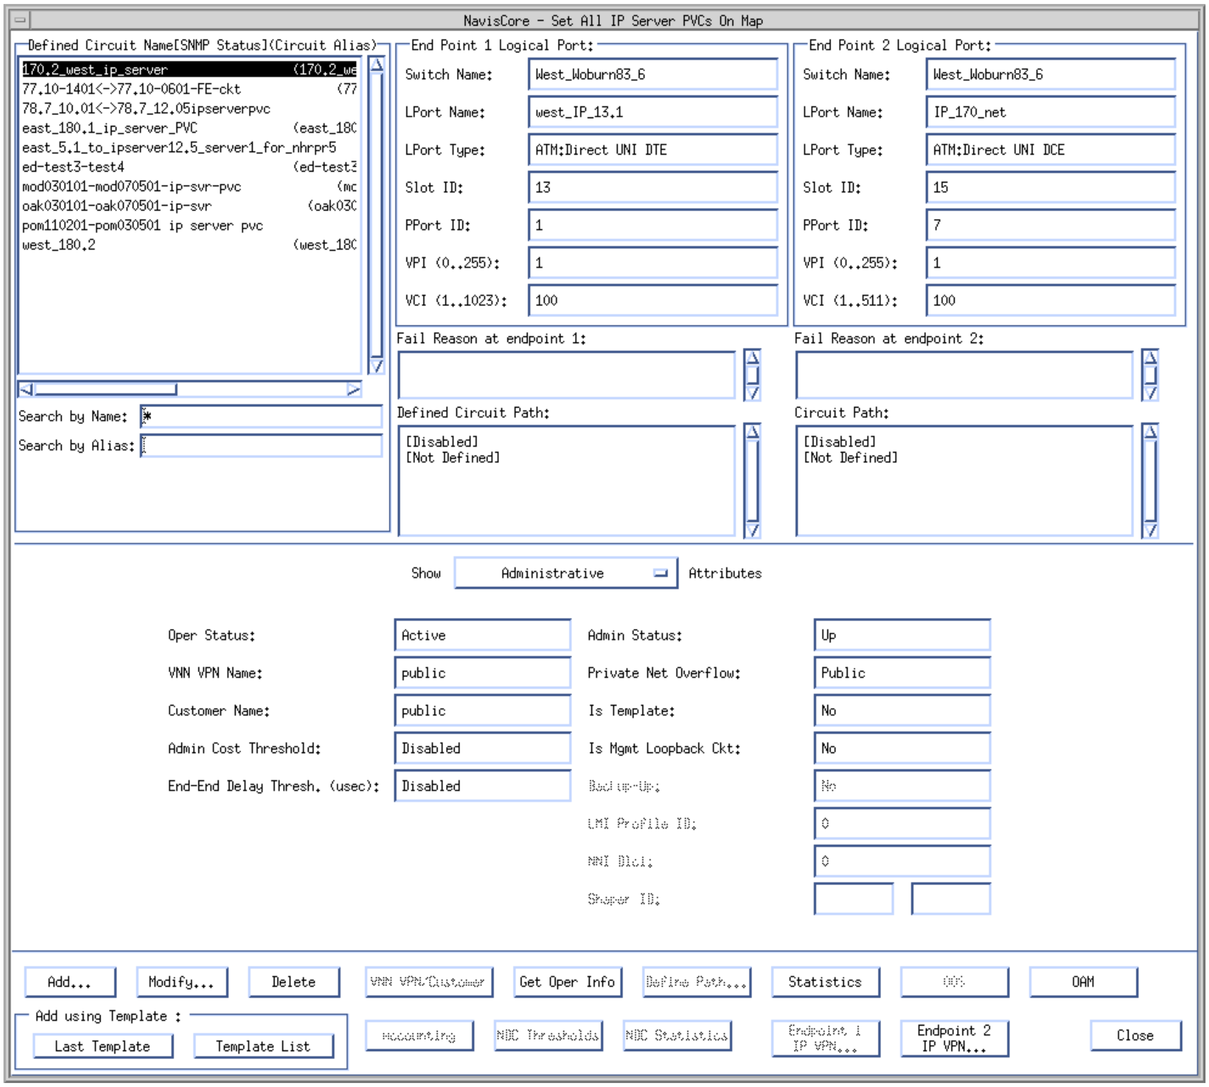

Defined Circuit Name[SNMP Status](Circuit Alias)

```
170.2_west_ip_server                   (170.2_we
77.10-1401<->77.10-0601-FE-ckt              (77
78.7_10.01<->78.7_12.05ipserverpvc
east_180.1_ip_server_PVC               (east_18C
east_5.1_to_ipserver12.5_server1_for_nhrpr5
ed-test3-test4                         (ed-test3
mod030101-mod070501-ip-svr-pvc              (mc
oak030101-oak070501-ip-svr               (oak03C
pom110201-pom030501 ip server pvc
west_180.2                             (west_18C
```

Search by Name: *

Search by Alias:

	End Point 1 Logical Port:		End Point 2 Logical Port:
Switch Name:	West_Woburn83_6	Switch Name:	West_Woburn83_6
LPort Name:	west_IP_13.1	LPort Name:	IP_170_net
LPort Type:	ATM:Direct UNI DTE	LPort Type:	ATM:Direct UNI DCE
Slot ID:	13	Slot ID:	15
PPort ID:	1	PPort ID:	7
VPI (0..255):	1	VPI (0..255):	1
VCI (1..1023):	100	VCI (1..511):	100

Fail Reason at endpoint 1:

Fail Reason at endpoint 2:

Defined Circuit Path:
[Disabled]
[Not Defined]

Circuit Path:
[Disabled]
[Not Defined]

Show Administrative Attributes

Oper Status:	Active	Admin Status:	Up
VNN VPN Name:	public	Private Net Overflow:	Public
Customer Name:	public	Is Template:	No
Admin Cost Threshold:	Disabled	Is Mgmt Loopback Ckt:	No
End-End Delay Thresh. (usec):	Disabled	Backup-Up:	No
		LMI Profile ID:	0
		NNI Dlci:	0
		Shaper ID:	

Figure 3-19. Set All IP Server PVCs on Map Dialog Box

3. Choose Add. The Select End Logical Ports dialog box appears.

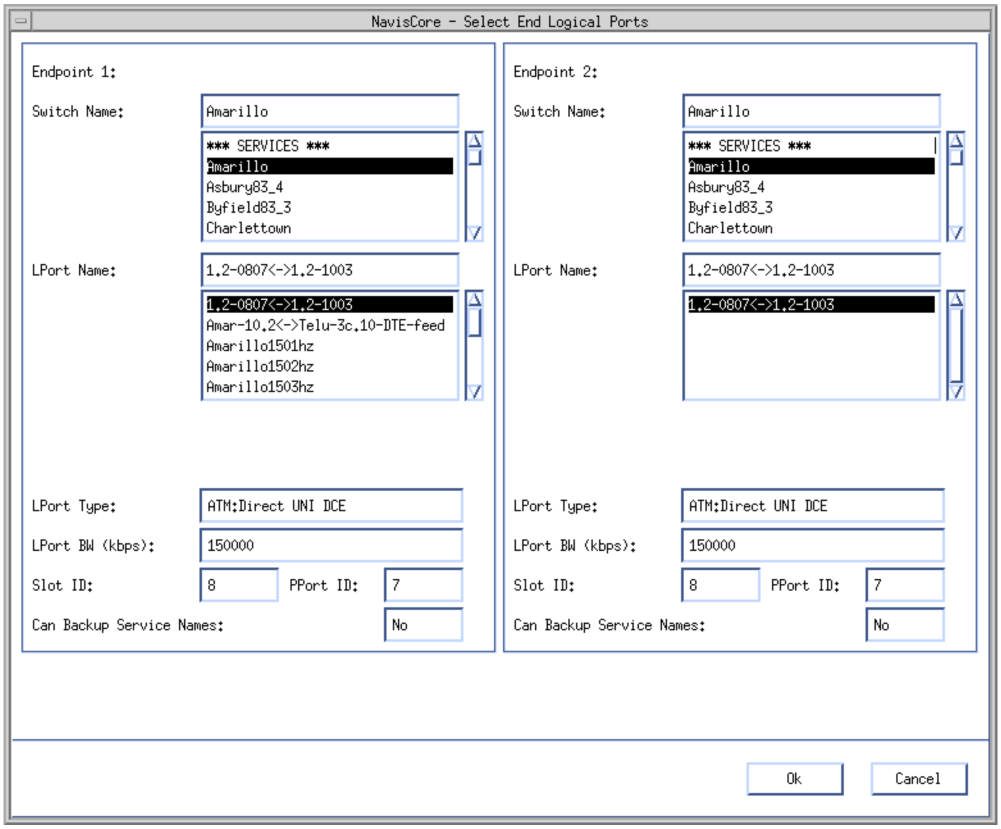

NavisCore - Select End Logical Ports

Endpoint 1:

Switch Name: Amarillo

*** SERVICES ***
Amarillo
Asbury83_4
Byfield83_3
Charlettown

LPort Name: 1.2-0807<->1.2-1003

1.2-0807<->1.2-1003
Amar-10.2<->Telu-3c.10-DTE-feed
Amarillo1501hz
Amarillo1502hz
Amarillo1503hz

LPort Type: ATM:Direct UNI DCE

LPort BW (kbps): 150000

Slot ID: 8 PPort ID: 7

Can Backup Service Names: No

Endpoint 2:

Switch Name: Amarillo

*** SERVICES ***
Amarillo
Asbury83_4
Byfield83_3
Charlettown

LPort Name: 1.2-0807<->1.2-1003

1.2-0807<->1.2-1003

LPort Type: ATM:Direct UNI DCE

LPort BW (kbps): 150000

Slot ID: 8 PPort ID: 7

Can Backup Service Names: No

Figure 3-20. Select End Logical Ports

4. Select the name of the switch where Endpoint 1 resides, then select the name of the switch where Endpoint 2 resides.

   In order to accept or transmit IP traffic on a cell-based card, you must configure a PVC connection from the cell-based card to an IP server card for all traffic entering the IP port. In this case, select the endpoints as follows:

   - Endpoint 1 is the ATM cell endpoint.
   - Endpoint 2 is the IP server endpoint.

   Keep in mind that you can also create an IP server PVC between two IP server endpoints on the same CBX 500 switch or different CBX 500 switches. The CBX 500 switches may be separated by an intermediate ATM network. You can use this feature as an alternative to ATM OPTimum trunks.

5. Select the name of the logical port for Endpoint 1, then select the name of the logical port for Endpoint 2. The Select End Logical Ports dialog box displays information for both Endpoint 1 and Endpoint 2. Table 3-11 describes each of these displayed fields.

Table 3-11. Information Displayed for Endpoints 1 and 2

Field	Action/Description
Primary Switch Name	Displays the name of the switch on which the primary (active) logical port resides.
Primary LPort Name	Displays the name of the primary (active) logical port endpoint.
LPort Type	Displays the logical port type for the selected logical ports.
LPort Bandwidth	Displays the logical port bandwidth for the selected logical ports. At each endpoint, logical ports may have a different bandwidth.
Slot ID	Displays the I/O slot (number) where the IOMs for the selected logical ports reside.
PPort ID	Displays the port ID numbers for the selected logical ports.
Can Backup Service Names	Displays either Yes or No to specify whether or not this logical port can be backed up to a service name binding.

6. Choose OK from the Select End Logical Ports dialog box. The Add PVC dialog box appears.

7. See the *NavisCore ATM Configuration Guide* to define the following attributes for each PVC:

   - Administrative
   - Traffic Type
   - User Preference
   - Frame Discard

▶ You can now configure traffic descriptors that are better than UBR (for example, VBRnrt) for IP server PVCs. Traffic shaping and policing is also supported for IP Server PVCs.

4

Configuring IP Packet Filters

This chapter describes the following tasks:

- Configuring IP packet filters
- Assigning IP packet filters to a logical port
- Assigning IP packet filters to the host
- Assigning IP packet filters to a circuit
- Viewing an IP packet filter configuration

About Packet Filters

Packet filtering enables a switch to accept or reject inbound or outbound packets by comparing a packet's IP upper-layer header information (see below for the IP header fields) to configured parameters called *filters*, which you define in NavisCore.

You define packet filters based on the following fields in the IP packet header:

- IP Header
- UDP/TCP Header

The following sections describe each of these fields.

IP Header

Source Address — The source address field contains the IP address that sends the packet.

Destination Address — The destination address field contains the IP address that receives the packet.

Type of Service (TOS) — The TOS field indicates the packet's priority.

Protocol — The transport field specifies the protocol (TCP or UDP) that enables the packet to be delivered to the correct destination protocol.

UDP/TCP Header

Source Port — This field contains the 16-bit protocol port number used to demultiplex datagrams among processes waiting to receive them. The source port is optional. When used, it specifies the port to which replies should be sent. If not used, the field should be left blank.

Destination Port — This field contains the 16-bit protocol port number used to demultiplex datagrams among processes waiting to receive them.

For inbound filters, when a packet is received, the forwarding code checks the packet against the interface's list of filters. If the packet matches a filter in the filter list, the packet is accepted or rejected and further filtering is terminated. The packet goes through a similar process for outbound filters, however, the process occurs only after the packet is received and routed to an interface.

Configuring IP Packet Filters

When you define an IP packet filter, you specify parameters that control the processing of inbound and/or outbound packets. After you define the filter, you can assign it to IP logical ports, the switch itself (host), or PVCs.

This section describes how to:

- Define an IP packet filter
- Assign an IP packet filter to a logical port
- Assign an IP packet filter to a host (switch)
- Assign an IP packet filter to a circuit
- View an IP packet filter's configuration and its associated logical port and/or circuit

► You can create a maximum of 1024 packet filters per switch.

You can define 128 logical port/circuit filter bindings per IOP.

You can assign a maximum of 32 inbound and 32 outbound filters per logical port.

You can assign a maximum of 32 inbound and 32 outbound filters per circuit.

Defining an IP Packet Filter

To define an IP packet filter:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Packet Filters ⇒ Set All Packet Filters. The Set All Packet Filters dialog box appears (see Figure 4-1).

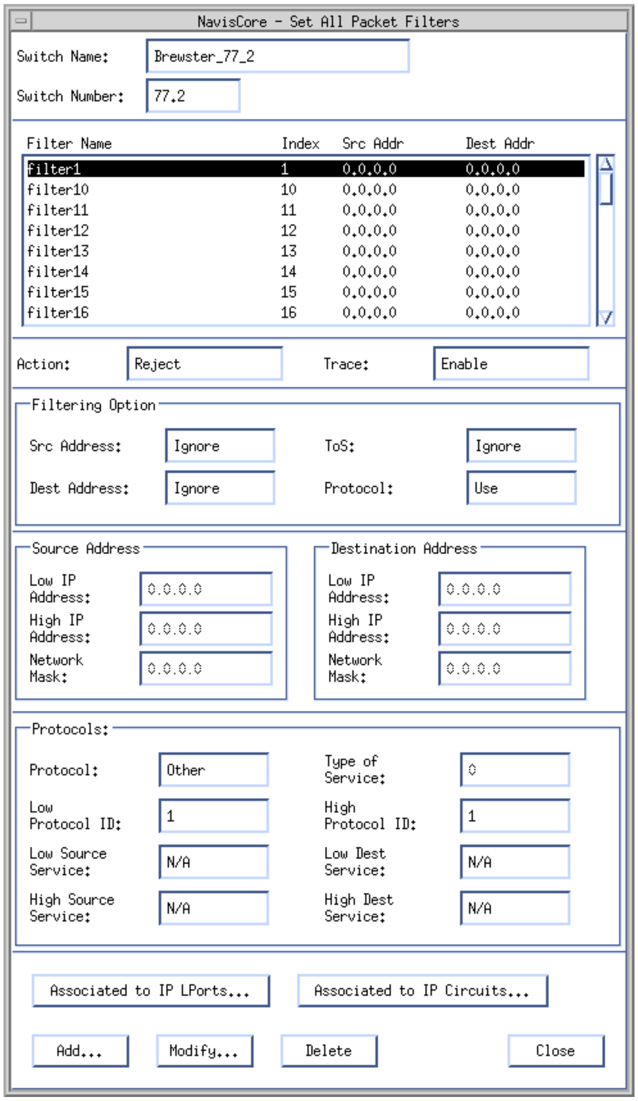

NavisCore - Set All Packet Filters

Switch Name: Brewster_77_2

Switch Number: 77.2

Filter Name	Index	Src Addr	Dest Addr
filter1	1	0.0.0.0	0.0.0.0
filter10	10	0.0.0.0	0.0.0.0
filter11	11	0.0.0.0	0.0.0.0
filter12	12	0.0.0.0	0.0.0.0
filter13	13	0.0.0.0	0.0.0.0
filter14	14	0.0.0.0	0.0.0.0
filter15	15	0.0.0.0	0.0.0.0
filter16	16	0.0.0.0	0.0.0.0

Action: Reject

Trace: Enable

Filtering Option

Src Address: Ignore

ToS: Ignore

Dest Address: Ignore

Protocol: Use

Source Address

Low IP Address: 0.0.0.0

High IP Address: 0.0.0.0

Network Mask: 0.0.0.0

Destination Address

Low IP Address: 0.0.0.0

High IP Address: 0.0.0.0

Network Mask: 0.0.0.0

Protocols:

Protocol: Other

Type of Service: 0

Low Protocol ID: 1

High Protocol ID: 1

Low Source Service: N/A

Low Dest Service: N/A

High Source Service: N/A

High Dest Service: N/A

Figure 4-1. Set All Packet Filters Dialog Box

The Set All Packet Filters dialog box displays the following buttons.

Table 4-1. Set All Packet Filters Buttons

Button	Function
Associated to IP LPorts	Displays the logical ports that are using a selected packet filter. For more information, see "Viewing an IP Packet Filter's Configuration" on page 4-21.
Associated to IP Circuits	Displays the circuits that are using a selected packet filter. For more information, see "Viewing an IP Packet Filter's Configuration" on page 4-21.
Add	Enables you to add a filter.
Modify	Enables you to modify an existing filter.
Delete	Enables you to delete an existing filter.

3. Choose Add. The Set Filter dialog box appears (see Figure 4-2).

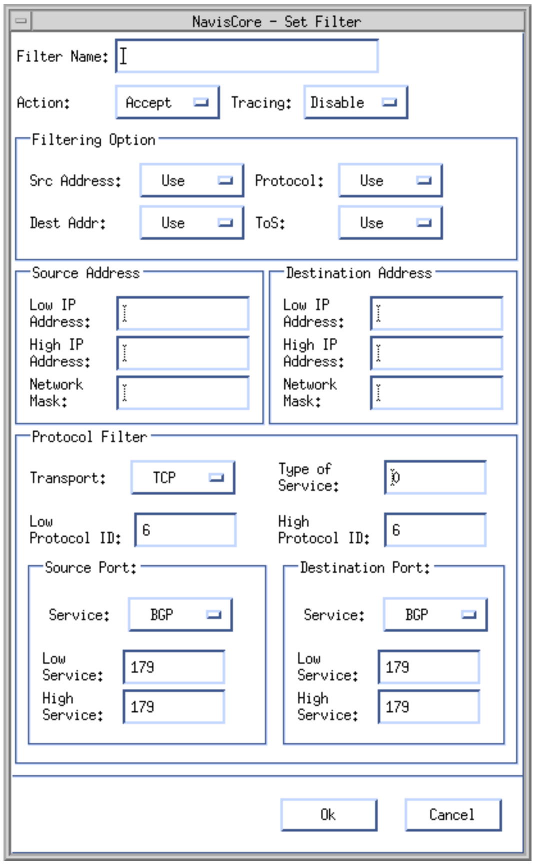

Filter Name:

Action: Accept
Tracing: Disable

Filtering Option
Src Address: Use
Protocol: Use
Dest Addr: Use
ToS: Use

Source Address
Low IP Address:
High IP Address:
Network Mask:

Destination Address
Low IP Address:
High IP Address:
Network Mask:

Protocol Filter
Transport: TCP
Type of Service: 0
Low Protocol ID: 6
High Protocol ID: 6

Source Port:
Service: BGP
Low Service: 179
High Service: 179

Destination Port:
Service: BGP
Low Service: 179
High Service: 179

Figure 4-2. Set Filter Dialog Box

4. Complete the fields as described in Table 4-2.

Table 4-2. Set Filter Fields

Field	Action/Description
Filter Name	Enter a filter name to identify the filter.
Action	Select one of the following options: *Accept* – This parameter instructs the switch to accept packets that match the filtering criteria. *Reject* – This parameter instructs the switch to reject packets that match the filtering criteria.
Tracing	Select one of the following options: *Enable* – This parameter instructs the switch to pass matched packets to the trace manager. *Disable* – This parameter instructs the switch not to pass matched packets to the trace manager.
Filtering Option	
Src Address	Select one of the following options: *Use* – To filter packets based on the source address field in the IP packet header. *Ignore* – To ignore filtering based on the source address field in the IP packet header. If you choose Ignore, the source address fields are grayed out and cannot be defined.
Protocol	Select one of the following options: *Use* – To filter packets based on the source address field in the IP packet header. *Ignore* – To ignore filtering based on the source address field in the IP packet header. If you choose Ignore, the protocol fields are grayed out and cannot be defined.
Dest Addr	Select one of the following options: *Use* – To filter packets based on the destination address field in the IP packet header. *Ignore* – To ignore filtering based on the destination address field in the IP packet header. If you choose Ignore, the Destination Address fields are grayed out and cannot be defined.

Table 4-2. Set Filter Fields (Continued)

Field	Action/Description
ToS	Select one of the following options: *Use* – To filter packets based on the destination address field in the IP packet header. *Ignore* – To ignore filtering based on the Type of Service field in the IP packet header. If you choose Ignore, the Type of Service field is grayed out and cannot be defined.
Source Address	
Low IP Address	Enter the low IP address of the node that sends the packet. When you specify the source address you specify one IP address (in the Low IP Address field) or you can specify a range between the lowest and highest IP address. If a packet's source address is within the range, there is a match. ***Note:*** *If you want to filter packets coming from one IP address, specify the IP address in the low IP address field. You do not have to specify a value in the high IP address field.*
High IP Address	Enter the high IP address of the node that sends the packet (the default is *high IP address=low IP address).*
Network Mask	Enter the Network Mask that applies to the source address.
Destination Address	
Low IP Address	Enter the low IP address of the node that receives the packet. When you specify the destination address you specify one IP address (in the Low IP Address field) or you can specify a range between the lowest and highest IP address. If a packet's source address is within the range, there is a match.
High IP Address	Enter the high IP address of the node that receives the packet (the default is *high IP address=low IP address).*
Network Mask	Enter the Network Mask that applies to the destination address.

Table 4-2. Set Filter Fields (Continued)

Field	Action/Description
Protocol Filter	
Transport	Select the packet's transport protocol type: *TCP* – Transmission Control Protocol. *UDP* – User Datagram Protocol *Others* – You must specify protocol IDs in the low and high protocol ID fields. Transport refers to the protocol (TCP, UDP, or Others) that enables the packet to be delivered to the correct destination protocol. ***Note:*** *When you select TCP or UDP, the low and high protocol fields are automatically filled in with the protocol's corresponding protocol ID.* *In addition, if you select TCP or UDP, you must specify the source and destination port fields. However, if you select Others, the source and destination port sections are grayed out and cannot be defined.*
Type of Service	Enter a value between 0 and 254. Protocols use the type of service value to specify the packet's priority.
Low Protocol ID	If you selected Others in the Transport field, enter the low protocol ID. See *RFC 1700* for protocol ID numbers. You can either specify *one value* (in this field) or you can enter a range between this value and the high protocol ID. If the packet's protocol ID is between the low and high protocol ID, there is a match.
High Protocol ID	If you selected Others in the Transport field, enter the high protocol ID. See *RFC 1700* for protocol ID numbers. When you enter this value, you enter a range between the low protocol ID and this value. If the packet's protocol ID is between the low and high protocol ID, there is a match.

Table 4-2. Set Filter Fields (Continued)

Field	Action/Description
Source Port	
Service	*If you selected TCP* in the Transport field of the Protocol Filter section, select one of the following protocols: *BGP* – Border Gateway Protocol. *FTP* – File Transfer Protocol. *Gopher* – Protocol that facilitates internet access. *IRC* – Internet Relay Chat Protocol. *Talk – Unit talk application.* *Telnet* – Standard terminal emulation protocol in the TCP/IP protocol stack. Telnet is used for remote terminal connection. *WWW* – World Wide Web. *Ignore* – Enables you to filter on all UDP packets. *Other* – You must specify port numbers in the low and high service fields. *If you selected UDP* in the Transport field of the Protocol Filter section, select one of the following protocols: *RIP* – Routing Information Protocol. *SNMP* – Simple Network Management Protocol. *Traps (SNMP)* – Message sent by an SNMP agent to an NMS station to indicate that an event occurred. *TFTP* – Trivial File Transfer Protocol. *Ignore* – Enables you to filter on any service. *Other* – You must specify port numbers in the low and high service fields. ***Note:*** *When you select a service, the low and high service fields in the source port field are automatically filled in with the service's corresponding port number.*
Low Service	If you selected Other in the Service field, enter the low service port number. See *RFC 1700* for the port numbers. ***Note:*** *To filter packets that have the same service port number, specify the port number in the low service field. You do not have to specify a value in the high service field.*

Table 4-2. Set Filter Fields (Continued)

Field	Action/Description
High Service	If you selected Other in the Service field, enter the high service port number. See *RFC 1700* for the port numbers. When you enter this value, you enter a range between the low service port number and this value. If the packet's service port number is between the low and high service port numbers, there is a match.
Destination Port	
Service	See the "Service" field description in the Source Port field section.
Low Service	See the "Low Service" field description in the Source Port field section.
High Service	See the "High Service" field description in the Source Port field section.

5. Choose OK.

6. At the Set All Packet Filters dialog box, choose Close.

Packet Filter Configuration Example

The following configuration is an example of a filter that restricts packets coming from a specified source IP address.

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Packet Filters ⇒ Set All Packet Filters. The Set All Packet Filters dialog box appears (see Figure 4-1).
3. Choose Add. The Set Filter dialog box appears (see Figure 4-2).
4. In the Filter Name field, enter:

   `reject152.148.51.118`

5. In the Action field, select Reject.
6. In the Tracing field, select Disable to disable the trace manager.
7. In the Filtering Option fields;
   - Select Use in the Src Address field.
   - Select Ignore in the Dest Addr, Protocol, and ToS fields.

> ▶ When you select Ignore in these fields, the fields in the Protocol Filter, Source Port, and Destination Port sections are grayed out. These fields are disabled and are not used to filter packets.

8. In the Low IP Address field in the Source Address section, enter:

   `152.148.51.118`

> ▶ You do not have to specify the high IP address for the source address because you are restricting packets coming from one IP address. However, if you want to restrict packets coming from a range of IP addresses, specify both the low and high IP addresses.

9. In the Network Mask field, enter:

   `255.255.255.255`

Figure 4-3 displays the specified fields.

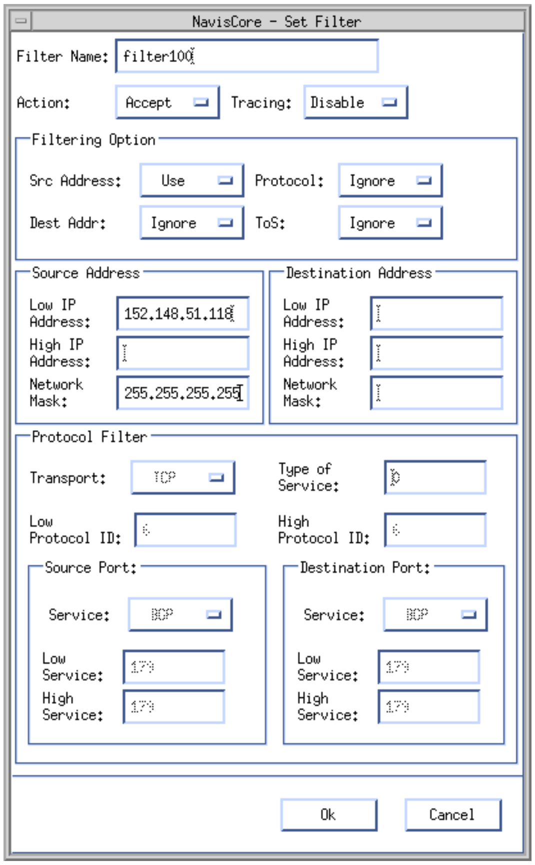

Filter Name: filter100
Action: Accept
Tracing: Disable
Src Address: Use
Protocol: Ignore
Dest Addr: Ignore
ToS: Ignore

Source Address
Low IP Address: 152.148.51.118
High IP Address:
Network Mask: 255.255.255.255

Destination Address
Low IP Address:
High IP Address:
Network Mask:

Protocol Filter
Transport: TCP
Type of Service: 0
Low Protocol ID: 6
High Protocol ID: 6

Source Port:
Service: BGP
Low Service: 179
High Service: 179

Destination Port:
Service: BGP
Low Service: 179
High Service: 179

Specify 152.148.51.118 in the Low IP Address field and 255.255.255.255 in the Network Mask field.

The fields in the Protocol Filter, Source Port, and Destination Port sections are grayed out because you selected Ignore in the Filtering Option fields.

Figure 4-3. Set Filter Dialog Box (Sample Packet Filter Settings)

10. Choose OK.

When you assign this filter to a specific logical port or host, all packets coming from *152.148.51.118* are not allowed to pass through.

Assigning IP Packet Filters to Logical Ports

To assign an IP packet filter to a logical port:

1. Select the switch icon from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Packet Filters ⇒ Set All Logical Port Filters. The Set All Logical Port Filters dialog box appears (see Figure 4-4).

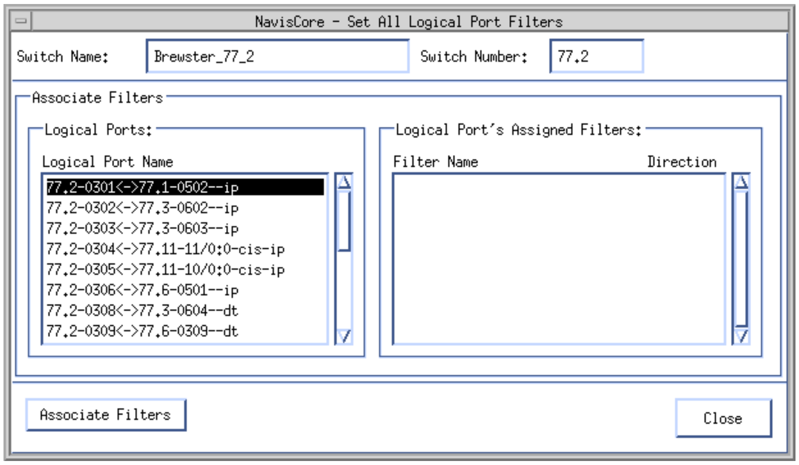

NavisCore - Set All Logical Port Filters

Switch Name: Brewster_77_2

Switch Number: 77.2

Logical Port Name

77.2-0301<->77.1-0502--ip
77.2-0302<->77.3-0602--ip
77.2-0303<->77.3-0603--ip
77.2-0304<->77.11-11/0:0-cis-ip
77.2-0305<->77.11-10/0:0-cis-ip
77.2-0306<->77.6-0501--ip
77.2-0308<->77.3-0604--dt
77.2-0309<->77.6-0309--dt

Figure 4-4. Set All Logical Port Filters Dialog Box

3. In the Logical Ports list box, select the logical port with which you want to associate a filter.
4. Choose the Associate Filters button. The Assign Logical Port IP Filter dialog box appears (see Figure 4-5).

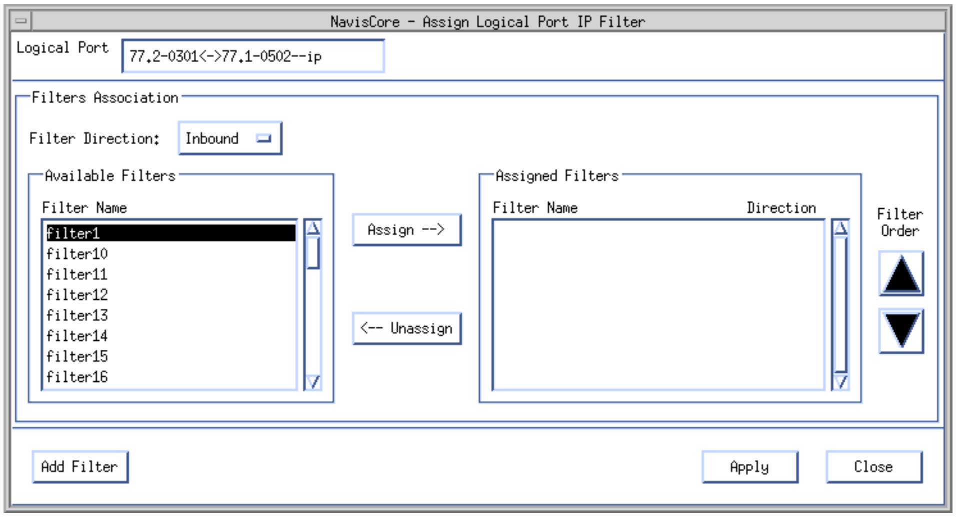

NavisCore - Assign Logical Port IP Filter

Logical Port: 77.2-0301<->77.1-0502--ip

Filter Direction: Inbound

Available Filters

Filter Name

filter1
filter10
filter11
filter12
filter13
filter14
filter15
filter16

You can view a packet filter's configuration by double-clicking the desired filter in either the Available or Assigned Filters fields.

Figure 4-5. Assign Logical Port IP Filter Dialog Box

▶ You can view a packet filter's configuration by double-clicking the desired filter in either the Available or Assigned Filters fields. The Show IP Filter Configuration dialog box appears with the filter's configuration.

Table 4-3 describes the fields on the Assign Logical Port IP Filter dialog box.

Table 4-3. Assign Logical Port IP Filter Fields

Field	Description
Logical Port	Displays the name of the logical port.
Filter Direction	Enables you to indicate the direction (inbound or outbound) in which you want the packets filtered through this logical port.
Available Filters	
Filter Name	The name that identifies the filter.
Assigned Filters	
Filter Name	The name that identifies the filter.
Direction	The direction (inbound or outbound) in which you want the packets filtered.

The Assign Logical Port IP Filter buttons are described in Table 4-4.

Table 4-4. Assign Logical Port IP Filter Buttons

Button	Function
Assign	Enables you to assign an IP packet filter to this logical port.
Unassign	Enables you to remove an IP packet filter from this logical port.
Filter Order	Enables you to specify the order in which the defined packet filters are applied. When a match occurs, the filtering process ends.
Add Filter	Enables you to configure an additional IP packet filter.
Apply	Applies any of the changes that you have made on this dialog box.

5. In the Filter Direction field, select either Inbound or Outbound to indicate the direction in which you want the packets filtered through this logical port.
6. From the Available Filters list box, select the filter and choose Assign to assign the IP packet filter to this logical port.
7. Repeat step 6 until you have assigned all the necessary IP packet filters to this logical port.
8. When you are done, choose Apply.
9. To configure an additional IP packet filter, choose Add Filter. See "Defining an IP Packet Filter" on page 4-4 for more information.

Assigning IP Filters to the Host (Switch)

▶ You can assign a maximum of 32 IP packet filters per host.

To assign a filter to the host (switch):

1. Select the switch icon from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Packet Filters ⇒ Set All Host Filters. The Set All Host filters dialog box appears (see Figure 4-6).

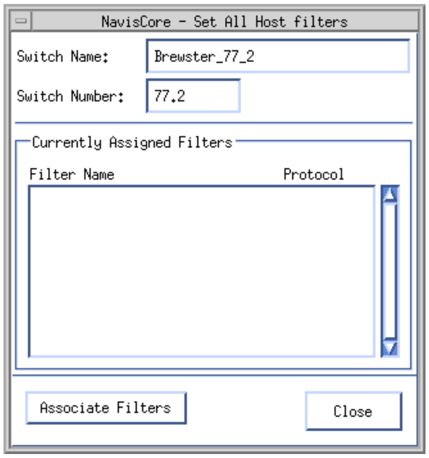

NavisCore - Set All Host filters

Switch Name: Brewster_77_2

Switch Number: 77.2

Figure 4-6. Set All Host filters Dialog Box

3. Choose Associate Filters. The Associate Host Filters dialog box appears (see Figure 4-7).

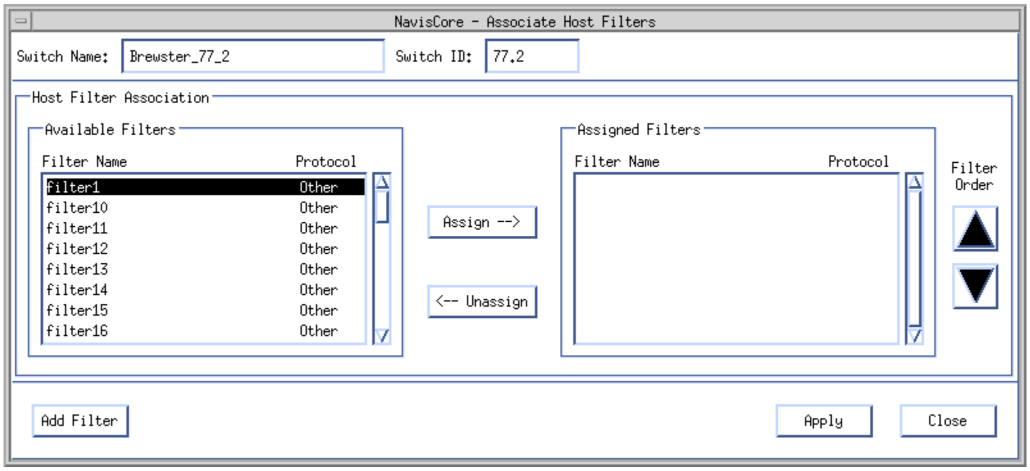

Switch Name: Brewster_77_2
Switch ID: 77.2

Filter Name	Protocol
filter1	Other
filter10	Other
filter11	Other
filter12	Other
filter13	Other
filter14	Other
filter15	Other
filter16	Other

The Filter Order buttons enable you to change the order in which filters are applied.

Figure 4-7. Associate Host Filters Dialog Box

> ▶ You can view a packet filter's configuration by double-clicking on the desired filter in either the Available or Assigned Filters fields. The Show IP Filter Configuration dialog box appears with the filter's configuration.

Table 4-5 describes the fields on the Associate Host Filters dialog box.

Table 4-5. Associate Host Filters Fields

Field	Description
Switch Name	Displays the name of the switch.
Switch ID	Displays the switch ID.
Available Filters	
Filter Name	The name that identifies each of the defined filters.
Protocol	Displays TCP, UDP, or Others to indicate the filter's transport protocol. See Table 4-2 for a description of each of these protocol types.
Assigned Filters	
Filter Name	The name that identifies each of the defined filters that are assigned to the switch.
Protocol	Displays TCP, UDP, or Others to indicate the filter's transport protocol. See Table 4-2 for a description of each of these protocol types.

The Associate Host Filters dialog box provides the following buttons:

Table 4-6. Associate Host Filters Buttons

Button	Function
Assign	Enables you to assign an IP packet filter to this host.
Unassign	Enables you to delete an IP packet filter from this host.
Filter Order	Enables you to specify the order in which the defined packet filters are applied. When a match occurs, the filtering process terminates.
Add Filter	Enables you to configure an additional IP packet filter.
Apply	Applies any of the changes that you have made on this dialog box.

4. From the Available Filters list box, select the filter and choose Assign to assign the IP packet filter to this switch.
5. Repeat step 4 until you have added the necessary IP packet filters to this switch.
6. When you are done, choose Apply.
7. To configure an additional IP packet filter, choose Add Filter. See "Defining an IP Packet Filter" on page 4-4 for more information.

Assigning IP Packet Filters to Circuits

Circuit filters are similar to logical port filters but differ in that you apply circuit filters to individual DLCIs (for Frame Relay circuits) or individual VPIs/VCIs (for ATM circuits). Before you assign circuit filters to PVCs, you must define these PVCs. For more information, refer to Chapter 3, “Configuring IP Logical Ports and IP Servers.”

> If you assign packet filters to both logical ports and circuits, the order in which packets are filtered is as follows:
> **Inbound —** Circuit Filters ⇒ Logical Port Filters
> **Outbound —** Logical Port Filter ⇒ Circuit Filters

To assign a filter to the circuit:

1. Select the switch icon from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Packet Filters ⇒ Set All Circuit Filters. The Set All IP Circuit Filters dialog box appears.

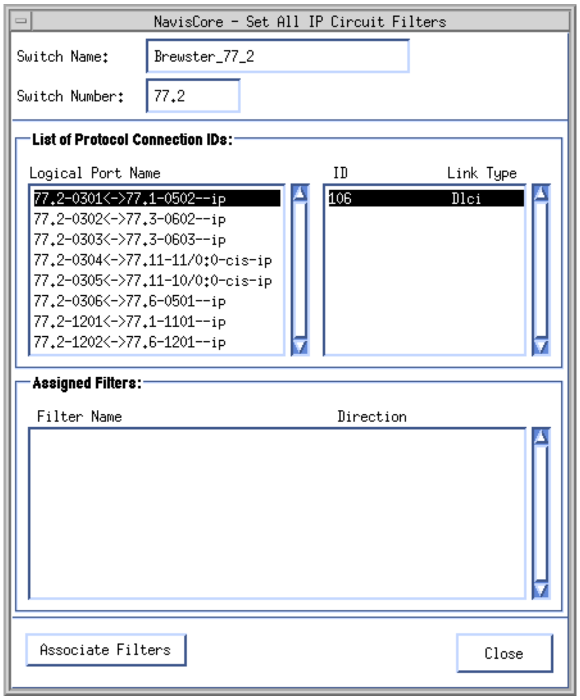

Switch Name: Brewster_77_2
Switch Number: 77.2

List of Protocol Connection IDs:

Logical Port Name	ID	Link Type
77.2-0301<->77.1-0502--ip	106	Dlci
77.2-0302<->77.3-0602--ip		
77.2-0303<->77.3-0603--ip		
77.2-0304<->77.11-11/0:0-cis-ip		
77.2-0305<->77.11-10/0:0-cis-ip		
77.2-0306<->77.6-0501--ip		
77.2-1201<->77.1-1101--ip		
77.2-1202<->77.6-1201--ip		

Assigned Filters:

Filter Name	Direction

Figure 4-8. Set All IP Circuit Filters Dialog Box

3. Choose Associate Filters. The Associate IP Circuit Filter List dialog box appears (see Figure 4-9).

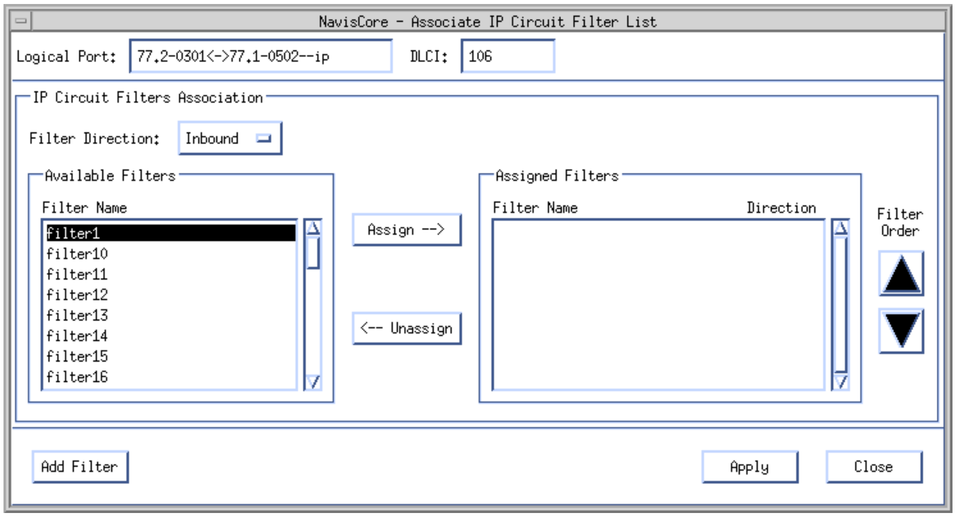

NavisCore - Associate IP Circuit Filter List

Logical Port: 77.2-0301<->77.1-0502--ip
DLCI: 106

IP Circuit Filters Association

Filter Direction: Inbound

Available Filters

Filter Name
filter1
filter10
filter11
filter12
filter13
filter14
filter15
filter16

Assigned Filters

Filter Name | Direction

Figure 4-9. Associate IP Circuit Filter List Dialog Box

You can view a packet filter’s configuration by double-clicking the desired filter in either the Available or Assigned Filters fields. The Show IP Filter Configuration dialog box appears with the filter’s configuration.

Table 4-7 describes the Associate IP Circuit Filter List dialog box fields.

Table 4-7. Associate IP Circuit Filter List Fields

Field	Description
Logical Port	Displays the circuit name.
VPI/VCI (for ATM logical ports)	Displays the circuit’s VPI/VCI.
DLCI (for Frame Relay logical ports)	Displays the circuit’s DLCI.
Filter Direction	Enables you to indicate the direction (inbound or outbound) in which you want the packets filtered through this circuit.
Available Filters	
Filter Name	The name that identifies the filters available to this circuit.
Assigned Filters	
Filter Name	The name that identifies the filter(s) assigned to this circuit.
Direction	Indicates the direction (inbound or outbound) in which you want the packets filtered.

The Associate IP Circuit Filter List dialog box provides the following buttons:

Table 4-8. Associate IP Circuit Filter List Buttons

Button	Function
Assign	Enables you to assign an IP packet filter to this host.
Unassign	Enables you to delete an IP packet filter from this circuit.
Filter Order	Enables you to specify the order in which the defined packet filters are applied. When a match occurs, the filtering process terminates.
Add Filter	Enables you to configure an additional IP packet filter.
Apply	Applies any of the changes that you have made on this dialog box.

4. From the Available Filters List box, select the filter and choose Assign to assign the IP packet filter to this switch.
5. Repeat step 4 until you have added the necessary IP packet filters to this switch.
6. When you are done, choose Apply.
7. To configure an additional IP packet filter, choose Add Filter. See "Defining an IP Packet Filter" on page 4-4 for more information.

Viewing an IP Packet Filter’s Configuration

Once you define an IP packet filter and associate it, you can view its configuration and associated logical port or circuit. Use the following steps to view an IP packet filter configuration for a logical port:

1. From the network map, select the appropriate switch icon.

2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Packet Filters ⇒ Set All Packet Filters. The Set All Packet Filters dialog box appears (see Figure 4-10).

First select a filter in the Filter Name field.

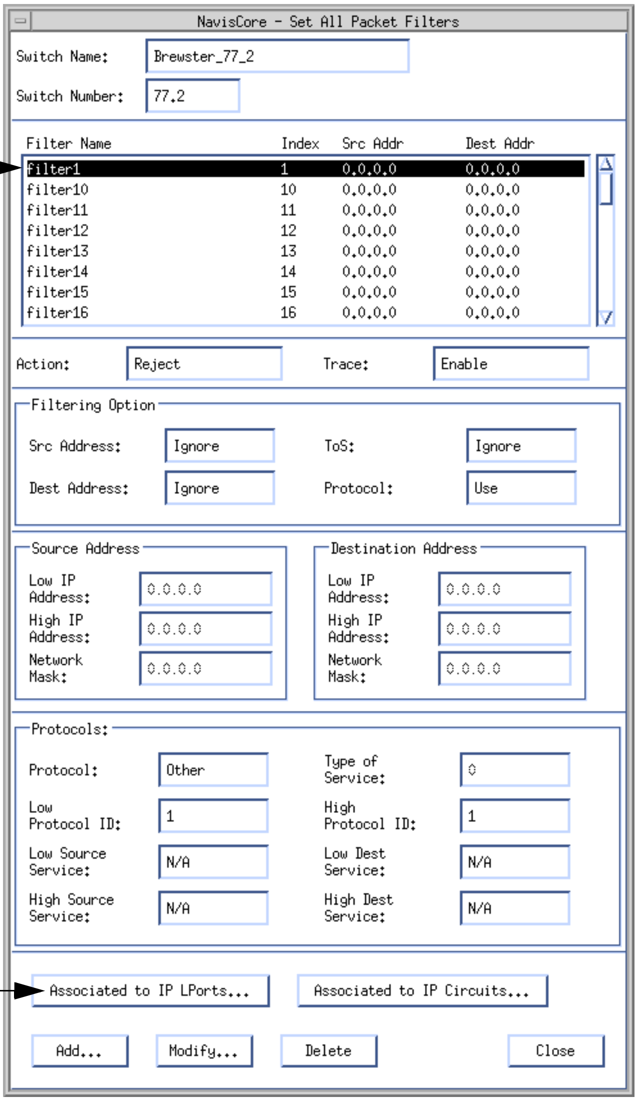

NavisCore - Set All Packet Filters

Switch Name: Brewster_77_2

Switch Number: 77.2

Filter Name	Index	Src Addr	Dest Addr
filter1	1	0.0.0.0	0.0.0.0
filter10	10	0.0.0.0	0.0.0.0
filter11	11	0.0.0.0	0.0.0.0
filter12	12	0.0.0.0	0.0.0.0
filter13	13	0.0.0.0	0.0.0.0
filter14	14	0.0.0.0	0.0.0.0
filter15	15	0.0.0.0	0.0.0.0
filter16	16	0.0.0.0	0.0.0.0

Action: Reject
Trace: Enable

Filtering Option
Src Address: Ignore
ToS: Ignore
Dest Address: Ignore
Protocol: Use

Source Address
Low IP Address: 0.0.0.0
High IP Address: 0.0.0.0
Network Mask: 0.0.0.0

Destination Address
Low IP Address: 0.0.0.0
High IP Address: 0.0.0.0
Network Mask: 0.0.0.0

Protocols:
Protocol: Other
Type of Service: 0
Low Protocol ID: 1
High Protocol ID: 1
Low Source Service: N/A
Low Dest Service: N/A
High Source Service: N/A
High Dest Service: N/A

Then choose the Associated to IP LPorts button to view the IP logical port(s) associated with the IP packet filter.

Figure 4-10. Set All Packet Filters Dialog Box

3. Do the following:

   a. Select a filter in the filter field.

   b. Choose the Associated to IP LPorts button. The Logical ports using the Packet Filter dialog box appears (see Figure 4-11).

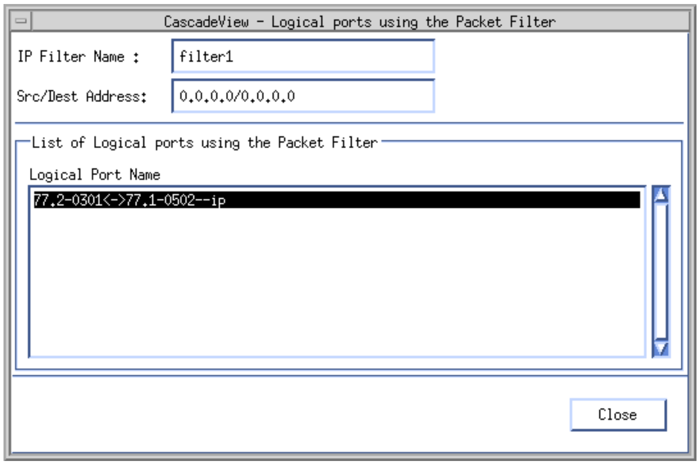

CascadeView - Logical ports using the Packet Filter

IP Filter Name : filter1

Src/Dest Address: 0.0.0.0/0.0.0.0

List of Logical ports using the Packet Filter

Logical Port Name

77.2-0301<->77.1-0502--ip

Figure 4-11. Logical ports using the Packet Filter Dialog Box

If you selected a packet filter that is not assigned to an IP logical port, the following message appears:

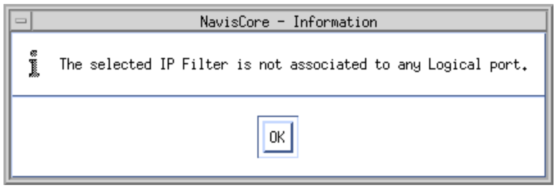

NavisCore - Information

The selected IP Filter is not associated to any Logical port.

Figure 4-12. Packet Filter Error Message Dialog Box

4. Choose Close.

5. To view a packet filter that is assigned to a circuit, perform step 1 through step 3, except choose the Associated to IP Circuits button instead of the Associated to IP LPorts button.

5

Configuring Policy-Based Forwarding

This chapter provides an overview of policy-based forwarding and describes the following tasks:

- Configuring a policy PVC
- Defining a forwarding policy
- Assigning a forwarding policy to a logical port
- Enabling policy-based forwarding

About Forwarding Policies

Policy-based forwarding is a technique for forwarding IP packets based on criteria defined in *forwarding policies*. Policy-based forwarding allows switches to forward packets based on criteria other than destination IP addresses. For example, you can use policy-based forwarding to route packets based on source address, source or destination port number, Type of Service (ToS), protocol type, or other user-configurable criteria.

You associate forwarding policies with one or more ingress IP logical ports on a switch. If an IP packet received at an ingress IP logical port matches the criteria of a forwarding policy associated with that port, IP Navigator forwards the packet over a specified *policy PVC*. Otherwise, IP Navigator routes the packet on a best-effort basis according to the packet’s destination IP address.

When you define a forwarding policy, you specify:

- Forwarding criteria, including source and destination IP addresses and masks, an IP protocol, a Type of Service (ToS), source and destination TCP/UDP ports, and a Virtual Private Network (VPN) ID.
- Policy PVCs and associated Quality of Service (QoS) attributes, including a QoS class and traffic descriptors such as peak cell rate (PCR) and sustainable cell rate (SCR).

You can associate more than one forwarding policy with an ingress IP logical port. IP Navigator enforces the policies on any IP packets received at that port. If multiple policies are associated with one IP logical port, incoming IP packets are processed according to the first matching policy found. You can specify the sequence in which forwarding policies are applied to an IP logical port.

Remember the following important facts about IP policy PVCs and forwarding policies:

- Policy-based forwarding replaces IP QoS PVCs completely. However, because IP QoS is a subset of policy-based forwarding, you can convert any existing applications to policy-based forwarding. This means that you can convert IP QoS flow profiles to forwarding policies, and convert IP QoS PVCs to policy PVCs.
- One policy PVC may be shared by multiple forwarding policies.
- One forwarding policy may be shared by multiple logical ports on a switch (B-STDX 8000/9000) or on a Frame Engine (CBX 500). Forwarding PVCs cannot be shared by two Frame Engines (FEs) on one CBX 500 card or by FEs on different cards.
- You can assign as many as 100 forwarding policies to one ingress logical port.
- A switch (B-STDX 8000/9000) or Frame Engine (CBX 500) may have as many as 1500 forwarding policies.
- You can have as many as 3,736 logical ports per switch (B-STDX 8000/9000) or Frame Engine (CBX 500) that have associated forwarding policies.
- You can modify an IP address in a forwarding policy with a Classless InterDomain Routing (CIDR) mask. For example, if the address field is set to 1.2.3.4 with a mask of 255.255.255.0, an address of 1.2.3.*n* causes a match for any value of *n*.
- You can specify the order in which forwarding policies are applied to incoming packets.
- The switch processor (CP/SP) maintains the forwarding policy and policy PVC information, but then distributes the information to the appropriate IOMs/IOPs so that the switch processor is not required for proper forwarding policy operation. This step eliminates a single point of failure.
- IP Navigator checks packet filters before it checks forwarding policies. Because packet filtering occurs first, a packet could be dropped before policy-based forwarding ever checks it.
- IP Navigator checks NHRP flow profiles before it checks forwarding policies. Because NHRP flow detection occurs first, a packet could get forwarded over an NHRP SVC shortcut before policy-based forwarding ever checks it.
- You are responsible for assuring that the bandwidth allocated for a policy PVC supports the amount of traffic travelling over the policy PVC.

About Policy PVCs

A policy PVC is like an ordinary PVC, with some exceptions. This section describes the exceptions.

A policy PVC between two B-STDX 8000/9000 switches is either a switch-to-switch PVC or a switch-to-logical-port PVC (not logical port-to-logical port). IP packets must enter a policy PVC at the switch, and exit the PVC at either another switch (in which case they are routed based on a routing table lookup) or a specific egress logical port.

A policy PVC between two CBX 500 switches is either a FE-to-FE PVC or a FE-to-logical-port PVC (not logical port-to-logical port). IP packets must enter a policy PVC on a CBX 500 at the FE, and exit the PVC at either another FE (in which case they are routed based on a routing table lookup) or a specific egress logical port.

When a policy PVC exits the Lucent network at a switch or FE, the policy PVC is bi-directional. When a policy PVC exits the Lucent network at a logical port, the policy PVC is uni-directional. To handle traffic in the opposite direction, you can:

- Create a policy PVC in the opposite direction.
- Route traffic in the opposite direction to an IP logical port. In this case, traffic will be routed using routing table lookups.

The following types of user logical ports may act as egress ports for policy PVCs:

- IP logical port
- ATM UNI logical port
- Frame Relay UNI logical port
- PPP logical port
- Ethernet logical port

On the B-STDX 8000/9000 switch, some or all of the IP logical ports can share the same policy PVC. On the CBX 500 switch, some or all of the IP logical ports on one FE can share the same policy PVC. A policy PVC cannot be shared by logical ports on more than one FE.

Policy-Based Forwarding Tutorial

This section provides a tutorial on policy-based forwarding. As an example, this section presents a common use of policy-based forwarding called *source-based routing*. Source-based routing is a scheme for forwarding IP packets based on their *source* IP address instead of their *destination* IP address.

In this tutorial, traffic intended for host Boston must be forwarded through different Internet Service Providers (ISPs) according to the source address of the traffic:

- IP packets that have a source IP address of 152.148.10.*x* must be forwarded through ISP 1 to host Boston.
- IP packets that have a source IP address of 152.148.20.*x* must be forwarded through ISP 2 to host Boston.

Figure 5-1 illustrates the source-based routing example.

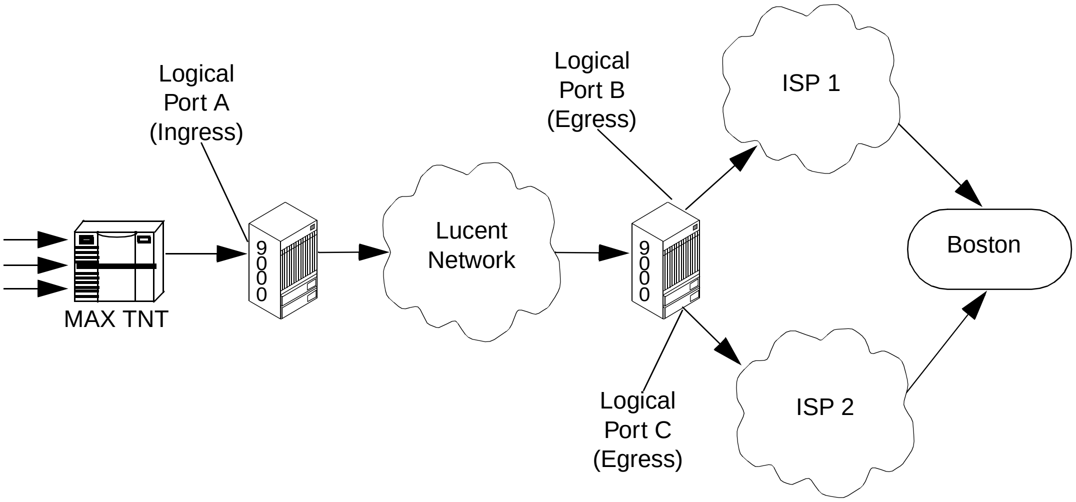

Figure 5-1. Source-Based Routing Example

How Do I Configure Policy-Based Forwarding?

This section describes the five general steps for configuring policy-based forwarding:

1. Plan the policy PVCs.
2. Create the policy PVCs.
3. Create the forwarding policies.
4. Assign the forwarding policies to IP logical ports.
5. If necessary, enable policy-based forwarding.

Step 1: Plan the Policy PVCs

When you plan a policy PVC, identify the problem that you want the policy PVC to solve. For example, in Figure 5-1 on page 5-4, the policy PVC must provide a source-based routing solution.

Step 2: Create the Policy PVCs

In Figure 5-1 on page 5-4, you want to create two policy PVCs:

POL-PVC10 — IP packets enter this policy PVC at the ingress switch and exit the policy PVC at Logical Port B in the egress switch. You use this policy PVC to forward IP packets with a source address of 152.148.10.*x* over Logical Port B to ISP 1. ISP 1 will then be responsible for delivering the packets to host Boston.

POL-PVC20 — IP packets enter this policy PVC at the ingress switch and exit the PVC at Logical Port C in the egress switch. You use this policy PVC to forward packets with a source address of 152.148.20.*x* over Logical Port C to ISP 2. ISP 2 will then be responsible for delivering the packets to host Boston.

See "Configuring a Policy PVC" on page 5-12 for instructions on creating a policy PVC. When you create a policy PVC, keep the following in mind:

- The ingress side of a policy PVC is always a switch (on a B-STDX 8000/9000) or Frame Engine (on a CBX 500).
- The egress side of a policy PVC may be a B-STDX switch or CBX 500 FE (that is, the "dummy" logical port that represents a switch or FE), or it may be a logical port on a switch.

Step 3: Create the Forwarding Policies

To set up the forwarding policies illustrated in Figure 5-1 on page 5-4, you would:

- Create a forwarding policy for POL-PVC10 with:
  - The Source IP Address field set to 152.148.10.0
  - The Network Mask field set to 255.255.255.0
  - The Policy PVC field set to POL-PVC10
- Create a forwarding policy for POL-PVC20 with:
  - The Source IP address field set to 152.148.20.0
  - The Network Mask field set to 255.255.255.0
  - The Policy PVC field set to POL-PVC20

See "Defining a Forwarding Policy" on page 5-28 for more information on creating forwarding policies.

Step 4: Assign the Forwarding Policies to IP Logical Ports

Assign the forwarding policies to the appropriate ingress logical port (for example, Logical Port A in Figure 5-1). See "Assigning a Forwarding Policy to a Logical Port" on page 5-33 for detailed instructions. Choose the appropriate logical port and assign the policies to that port.

Use the following switch console command to verify that the assignment was made properly:

```
show ip policy forwarding attributes <interface #>
```

This command shows all forwarding policies assigned to the specified port. The **interface #** is the logical port's interface ID (that is, the **LPort ID**).

Step 5: Enable Policy-Based Forwarding

Enable policy-based forwarding on the appropriate ingress port (logical port A in Figure 5-1). See "Enabling Policy-Based Forwarding" on page 5-34 for detailed instructions.

Use the following switch console command to verify that policy-based forwarding is enabled for the specified port:

```
show ip policy forwarding state <interface #>
```

This command shows whether policy-based forwarding is enabled for the specified logical port. The **interface #** is the logical port's interface ID (that is, the **LPort ID**).

Policy-Based Forwarding Configuration Flowchart

The flowchart in Figure 5-2 summarizes the policy-based forwarding configuration process.

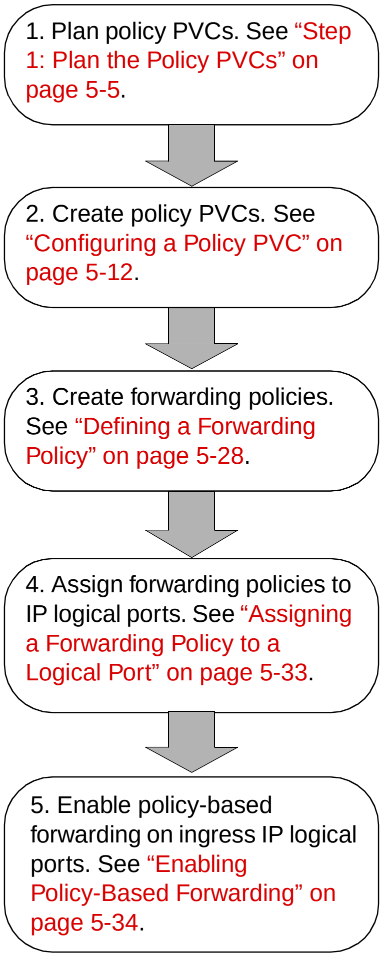

1. Plan policy PVCs. See “Step 1: Plan the Policy PVCs” on page 5-5.

2. Create policy PVCs. See “Configuring a Policy PVC” on page 5-12.

3. Create forwarding policies. See “Defining a Forwarding Policy” on page 5-28.

4. Assign forwarding policies to IP logical ports. See “Assigning a Forwarding Policy to a Logical Port” on page 5-33.

5. Enable policy-based forwarding on ingress IP logical ports. See “Enabling Policy-Based Forwarding” on page 5-34.

Figure 5-2. Policy-Based Forwarding Configuration Flowchart

How Do I Verify that Policy-Based Forwarding is Working?

After policy PVCs are active and forwarding policies are associated with ingress logical ports, you can verify that policy-based forwarding is working in three ways:

- Use the following switch console command to view forwarding statistics:

  ```
  show ip policy forwarding statistics <interface #>
  ```

  This command shows a count of received IP packets that matched a forwarding policy associated with the specified logical port, and, as a result, were forwarded over the associated policy PVC. Note that if a received packet matches a policy, but the associated PVC is not active, the counter is not incremented. The packet is only counted if it is forwarded over the PVC. A separate counter is kept for each forwarding policy associated with the specified logical port.

- Use the following Monitor menu selection to view forwarding statistics: Lucent IP Objects ⇒ Show All Forwarding Policies ⇒ Show All Logical Port Forwarding Policies.

- Use the following switch console command to verify that a policy PVC is active:

  ```
  show ip policy forwarding pvc
  ```

  This command identifies those policy PVCs that originate on the switch or FE where the command is issued (that is, the ingress side of the policy PVC).

What Happens When a Policy PVC Goes Down?

If a logical port receives an IP packet that matches an associated forwarding policy, but the policy PVC specified in the forwarding policy is not active, then the packet is routed on a best-effort basis according to its destination IP address.

For example, in Figure 5-1 on page 5-4, best-effort routing could forward the packet to either ISP 1 or ISP 2, depending upon routing table information. If you assume that best-effort routing is not acceptable, you can assign packet filters to egress logical ports B and C to eliminate the possibility of the wrong packet being sent to the wrong ISP. The filters would cause the unwanted packets to be dropped. Specifically, the filters would dictate that any packets with a source IP address of 152.148.20.*x* would not be forwarded out logical port B, and any packets with source IP address of 152.148.10.*x* would not be forwarded out logical port C.

Another way to solve the problem of inactive policy PVCs is to create a backup policy PVC. See "How Can I Set Up PVCs that are Redundant and Share Traffic Load?" on page 5-9 for details.

How Can I Set Up PVCs that are Redundant and Share Traffic Load?

Assuming the existence of two logical ports between the egress switch and one of the ISPs in Figure 5-1 on page 5-4, you could then create two policy PVCs that start in the ingress switch, and terminate at their respective egress switch logical ports. You could then create two forwarding policies, which would be identical except that they use different policy PVCs.

Figure 5-3 illustrates the new configuration. In this figure, two redundant policy PVCs are created for ISP 1. Each PVC is associated with a different egress IP logical port (B1 and B2). Two forwarding policies are created and associated with the ingress IP logical port.

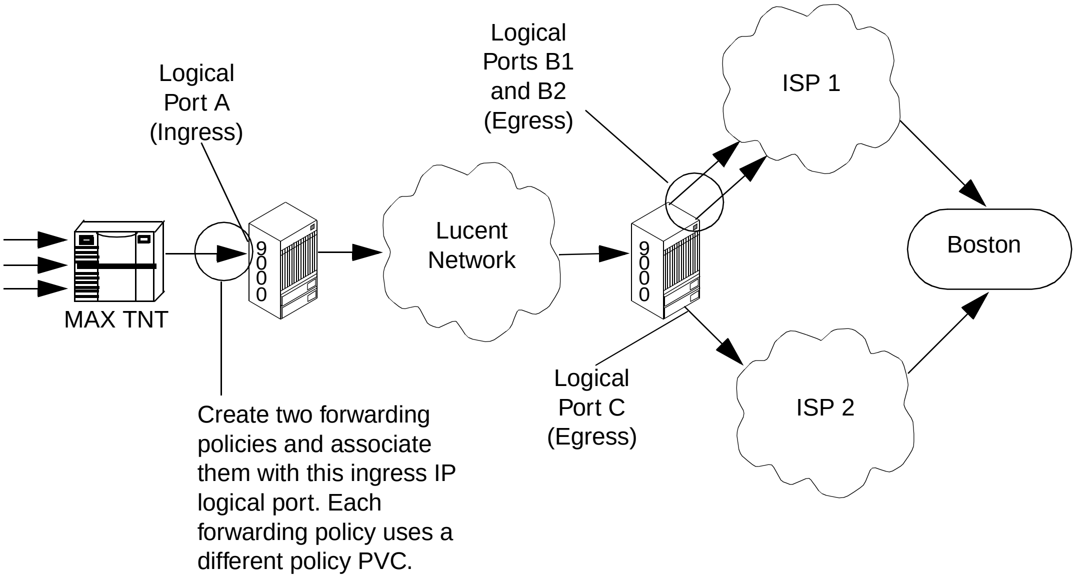
Logical
Port A
(Ingress)
Logical
Ports B1
and B2
(Egress)
ISP 1
MAX TNT
9000
Lucent
Network
9000
Boston
Logical
Port C
(Egress)
ISP 2
Create two forwarding policies and associate them with this ingress IP logical port. Each forwarding policy uses a different policy PVC.

Figure 5-3. Redundant PVCs

The following sequence of events will occur when a packet is received by the ingress switch logical port:

1. The packet matches the first of the two forwarding policies, but the PVC specified by that policy is down.
2. The second forwarding policy is checked for a match, and the packet matches that policy.
3. The packet is successfully forwarded over the policy PVC specified by the second forwarding policy, since the second policy PVC is up.

You can specify the sequence in which forwarding policies are applied to packets coming in over a logical port. Therefore, you can determine which PVC/egress port combination should be used as the primary path and which should be used as the secondary path.

Now, assume that you want to associate eight logical ports in the ingress switch or FE with the same forwarding policies. For four of the logical ports, you can specify:

- Forwarding policy A (which specifies policy PVC A) will be applied first
- Forwarding policy B (which specifies policy PVC B) will be applied if policy PVC A is not functioning

For the remaining four ports, you can reverse the policy order — forwarding policy B will be applied first, and forwarding policy A will be applied if policy PVC A is not functioning. In this way, both PVC/egress port combinations are used. If either policy PVC fails, the other policy PVC can support the traffic from all eight ingress ports.

The fault-tolerant UNI PVC feature allows the logical port side of a PVC to have two ports associated with it — a primary logical port and a secondary (backup) logical port. This may be considered as an alternative to the method previously described in this section. However, fault tolerant PVCs have a limitation: the PVC stays down temporarily while the switchover from primary logical port to backup logical port takes place. Packets are lost during that time. See the *NavisCore Frame Relay Configuration Guide* or the *NavisCore ATM Configuration Guide* for more information on fault-tolerant PVCs.

When Should I Not Use Switch/FE-to-Switch/FE Policy PVCs?

When a policy PVC terminates at a switch or FE instead of a logical port, packets arriving over the PVC at the egress switch are *routed* based on the routing tables in that egress switch. This type of policy PVC would not work in the example in Figure 5-1 on page 5-4. If the policy PVC terminated at the switch instead of at a logical port, packets would be sent to either ISP 1 or ISP 2 according to the current information in the routing table rather than according to the PVC policy.

When Should I Use Switch/FE-to-Switch/FE Policy PVCs?

Use of a switch-to-switch or FE-to-FE policy PVC could make sense in another configuration. Suppose that two egress switches are available, one connected to ISP 1 and the other to ISP 2. In each switch, multiple logical ports for reaching each ISP exist. Figure 5-4 illustrates this configuration.

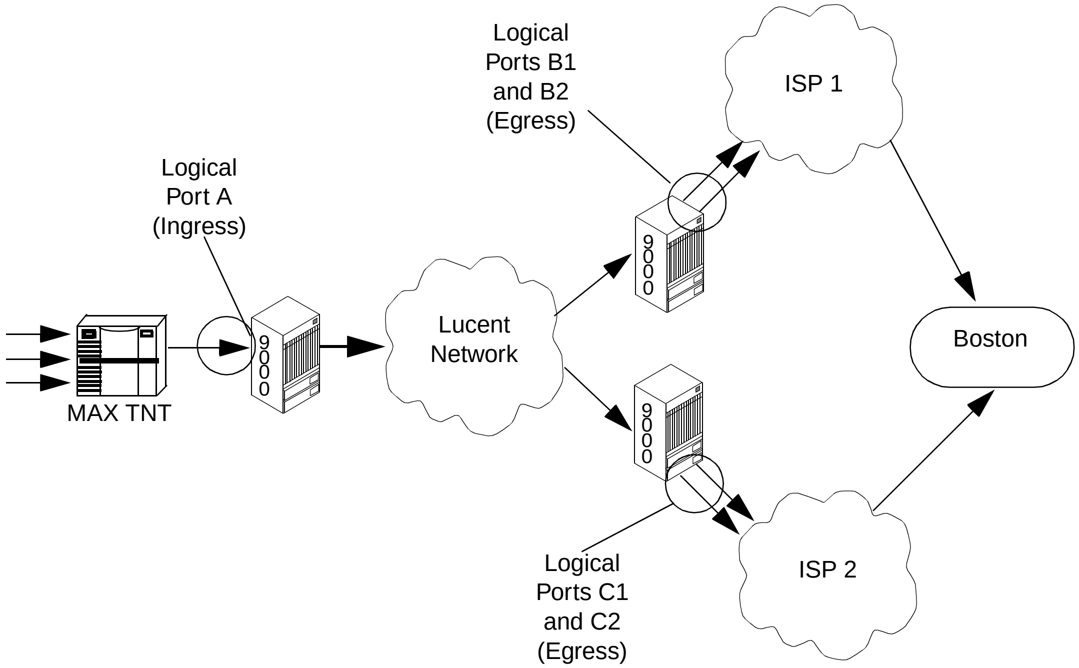

Figure 5-4. Switch-to-Switch Policy PVC Configuration

Since the switches share the same routing information, you still encounter the problem of packets being routed to either ISP 1 or ISP 2. However, if you configure static routes within the egress switches, you can guarantee that packets are sent to the ISP directly connected to the switch you specify, and not to another switch in the network. You must make sure that the cost of the static route is lower than other possible routes. If all ports between a switch and its corresponding ISP are down, however, packets will be routed to the other ISP via the other switch.

Can I Use Policy PVCs for QoS Bandwidth Reservation?

Although policy PVCs replace IP QoS PVCs, policy PVCs can meet the same application requirements that were met by IP QoS PVCs.

Configuring a Policy PVC

To configure a policy PVC, you first create the circuit endpoints, then define and name the circuit connection. You also assign administrative, user preference, and other attributes to the circuit.

▶ A number of seconds may pass between the time a policy PVC is created and the time it is active. A policy PVC is active when no "Fail Reason at endpoint *x*" messages appear on the Set All Policy PVCs On Map dialog box (see Figure 5-5). You can also use the **`show ip policy forwarding pvc command`** to verify that a PVC is active. See "How Do I Verify that Policy-Based Forwarding is Working?" on page 5-8 for more information.

To configure a policy PVC:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set Policy PVCs. The Set All Policy PVCs on Map dialog box appears (see Figure 5-5).

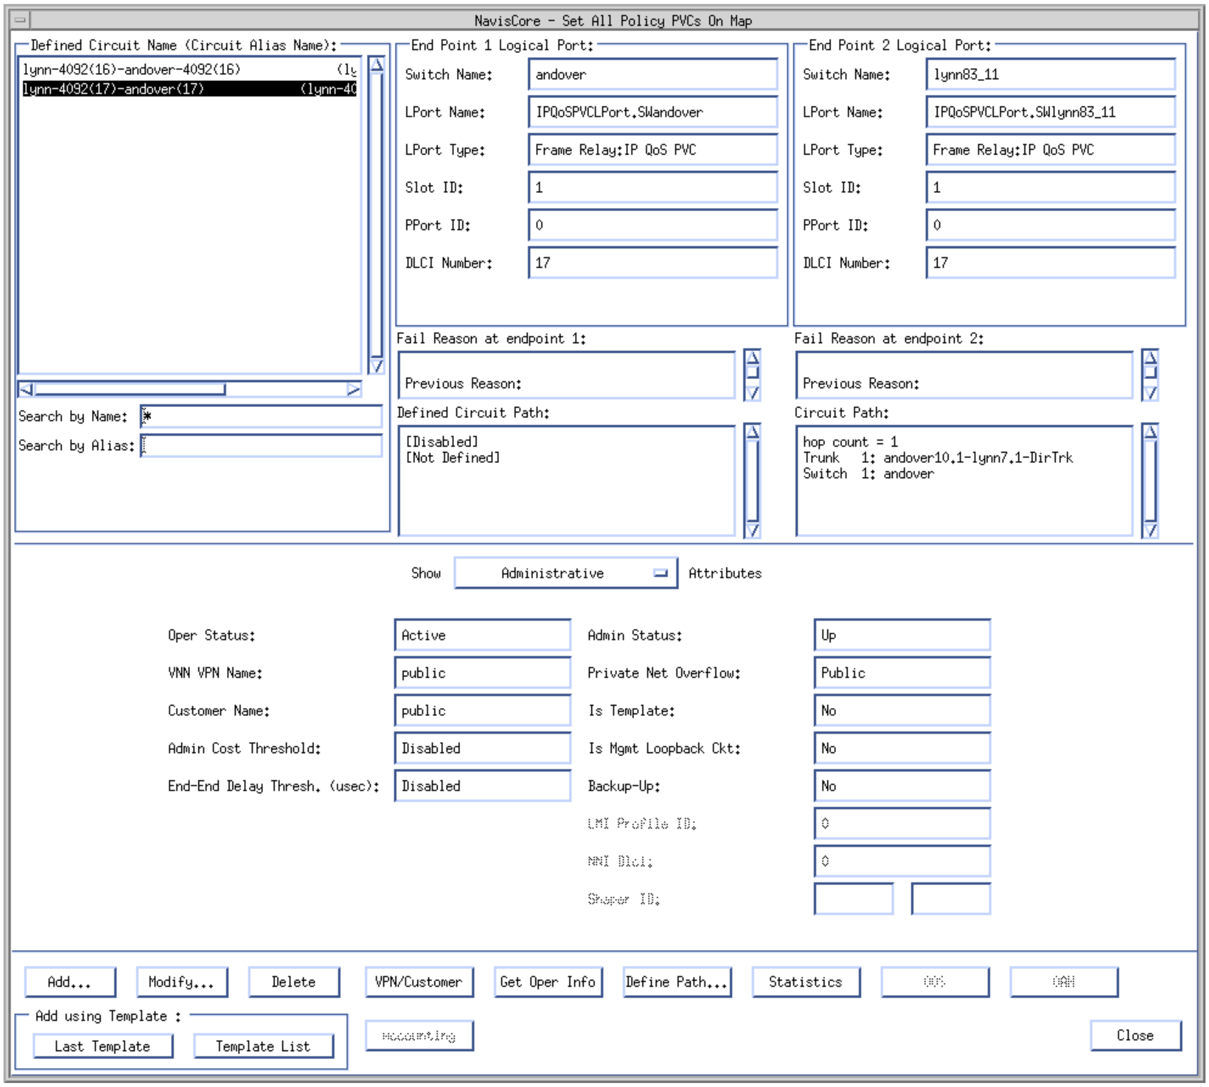

NavisCore - Set All Policy PVCs On Map

Defined Circuit Name (Circuit Alias Name):

lynn-4092(16)-andover-4092(16) (ly
lynn-4092(17)-andover(17) (lynn-40

Search by Name: *
Search by Alias:

End Point 1 Logical Port:

Switch Name: andover
LPort Name: IPQoSPVCLPort.SWandover
LPort Type: Frame Relay:IP QoS PVC
Slot ID: 1
PPort ID: 0
DLCI Number: 17

End Point 2 Logical Port:

Switch Name: lynn83_11
LPort Name: IPQoSPVCLPort.SWlynn83_11
LPort Type: Frame Relay:IP QoS PVC
Slot ID: 1
PPort ID: 0
DLCI Number: 17

Fail Reason at endpoint 1:
Previous Reason:

Fail Reason at endpoint 2:
Previous Reason:

Defined Circuit Path:
[Disabled]
[Not Defined]

Circuit Path:
hop count = 1
Trunk 1: andover10.1-lynn7.1-DirTrk
Switch 1: andover

Show Administrative Attributes

Oper Status: Active
VNN VPN Name: public
Customer Name: public
Admin Cost Threshold: Disabled
End-End Delay Thresh. (usec): Disabled
Admin Status: Up
Private Net Overflow: Public
Is Template: No
Is Mgmt Loopback Ckt: No
Backup-Up: No
LMI Profile ID: 0
NNI Dlci: 0
Shaper ID:

Figure 5-5. Set All Policy PVCs On Map Dialog Box

The Set All Policy PVCs on Map dialog box is blank initially. To display the complete list of defined circuit names, position the cursor in the *Search by Name* field and press Enter. This search may take several minutes depending on your configuration.

To display a filtered list of defined circuit names, enter the selected search criteria in the *Search by Name* field. To use a wildcard search to find a specific circuit name, you can:

- Use an * to match any number of characters.
- Use a ? to match a single character.
- Use a * to match the * character.
- Use a \? to match the ? character.
- Use a \\ to match the \ character.

Table 5-1 describes the buttons on the Set All Policy PVCs on Map dialog box.

Table 5-1. Set All Policy PVCs on Map Dialog Box: Buttons

Button	Function
Add/Modify/Delete	Enables you to add a new circuit or modify or delete an existing circuit. **Note**: If the PVC loopback status field does not display NONE, do not attempt to modify or delete the selected circuit.
VPN/Customer	Displays the virtual private network customer's name.
Get Oper Info	Displays a status message in the Oper Status field about the selected circuit. For more information, see the *NavisCore Diagnostics Guide*.
Define Path	Enables you to define a circuit path manually.
Statistics	Displays the summary statistics for the selected circuit. For information, see the *NavisCore Diagnostics Guide*.
QoS	Displays the Quality of Service (QoS) values for the selected circuit.
OAM Alarms (*ATM CS and IWU modules only*)	Displays the OAM alarms, which indicate whether the circuit is up or down. These alarms send a signal to the logical port whenever the circuit goes down or comes back up.
Add using Template	If you have already defined a circuit configuration and saved it as a template, use this option to define a new circuit. Choose *Last Template* to use the last template you defined for this switch. Choose *Template List* to display a list of templates previously defined for this map.
Accounting	Accesses the accounting functions for a PVC. For more information, see the *NavisXtend Accounting Server Administrator's Guide.*
NDC Thresholds *(CBX 500 only)*	Displays the configured network data collection (NDC) thresholds for the selected circuit.
NDC Statistic *(CBX 500 only)*	Displays the NDC statistics for the selected circuit.
Close	Exits the dialog box and returns you to the network map.

2. Choose Add. The Select End Logical Ports dialog box appears (see Figure 5-6).

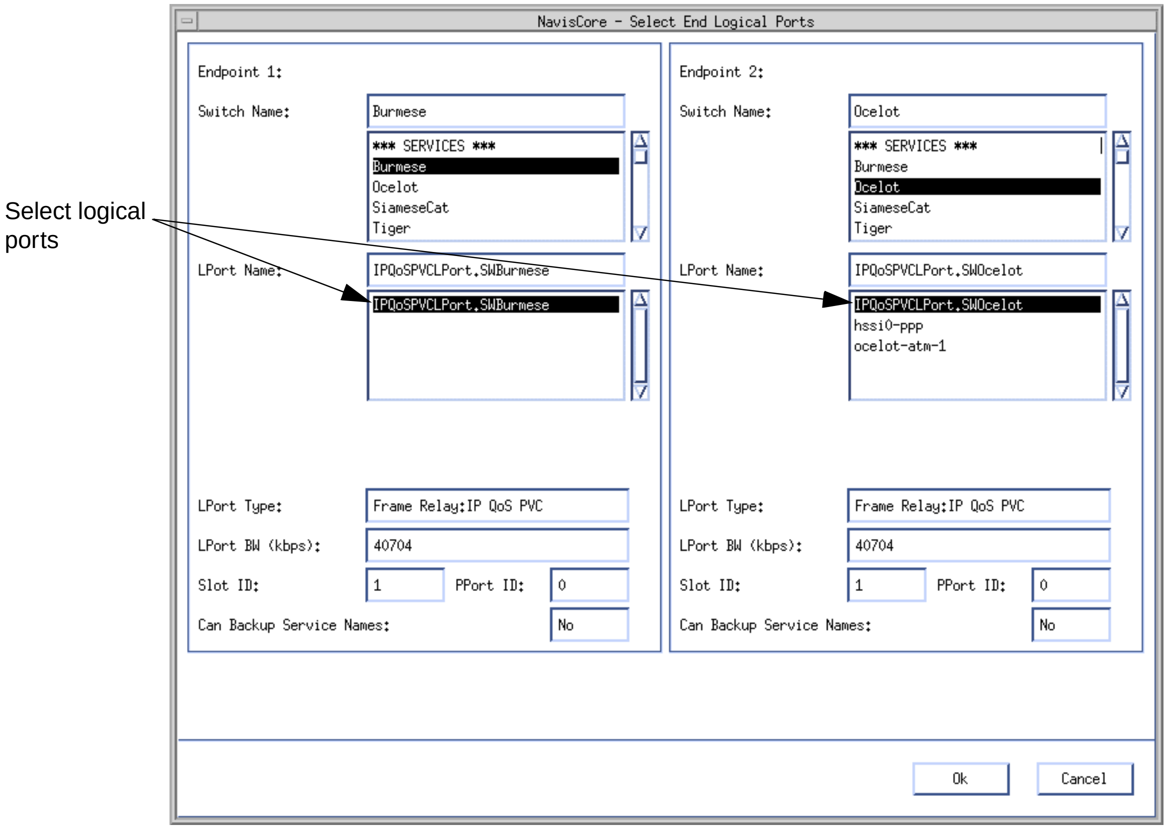

Select logical ports

NavisCore - Select End Logical Ports

Endpoint 1:

Switch Name: Burmese

*** SERVICES ***
Burmese
Ocelot
SiameseCat
Tiger

LPort Name: IPQoSPVCLPort.SWBurmese

IPQoSPVCLPort.SWBurmese

LPort Type: Frame Relay:IP QoS PVC

LPort BW (kbps): 40704

Slot ID: 1 PPort ID: 0

Can Backup Service Names: No

Endpoint 2:

Switch Name: Ocelot

*** SERVICES ***
Burmese
Ocelot
SiameseCat
Tiger

LPort Name: IPQoSPVCLPort.SWOcelot

IPQoSPVCLPort.SWOcelot
hssi0-ppp
ocelot-atm-1

LPort Type: Frame Relay:IP QoS PVC

LPort BW (kbps): 40704

Slot ID: 1 PPort ID: 0

Can Backup Service Names: No

Figure 5-6. Select End Logical Ports Dialog Box

The Select End Logical Ports dialog box displays information based on configuration selections you made. Table 5-2 describes each field.

Table 5-2. Select End Logical Ports Dialog Box: Fields

Field	Description
Primary Switch Name	Displays the name of the primary switch if the PVC is a fault-tolerant PVC. See the *NavisCore Frame Relay Configuration Guide* or the *NavisCore ATM Configuration Guide* for more information on fault-tolerant PVCs.
Primary LPort Name	Displays the name of the primary logical port if the PVC is a fault-tolerant PVC. See the *NavisCore Frame Relay Configuration Guide* or the *NavisCore ATM Configuration Guide* for more information on fault-tolerant PVCs.
LPort Name	Displays the name of the logical port associated with the endpoint. For Endpoint 1, a single ingress logical port endpoint appears and is used by the entire switch. This is a special, automatically-created logical port reserved for policy PVCs. The logical port is similar to Frame Relay logical ports in terms of the transmission service that they support. For Endpoint 2, multiple egress logical ports may appear. These ports may include the special Frame Relay logical port that is reserved for policy PVCs, the special logical port for one of the two FEs on a CBX 500 card, and any user-created logical ports.
LPort Type	Displays the logical port type for each port in the circuit configuration.
LPort BW (kbps)	Displays the bandwidth for each logical port in the trunk configuration.
Slot ID	Displays the I/O slot (number) in which the module resides.
PPort ID	Displays the port number for the physical port.
Can Backup Service Names	Displays either Yes or No to specify whether or not this logical port can be backed up to a service name binding.

3. Configure Endpoint 1 and Endpoint 2 as follows:

a. Select a switch name from the list.

b. Select the logical port.

4. Choose OK. The Add PVC dialog box (Figure 5-7) displays the current parameters.

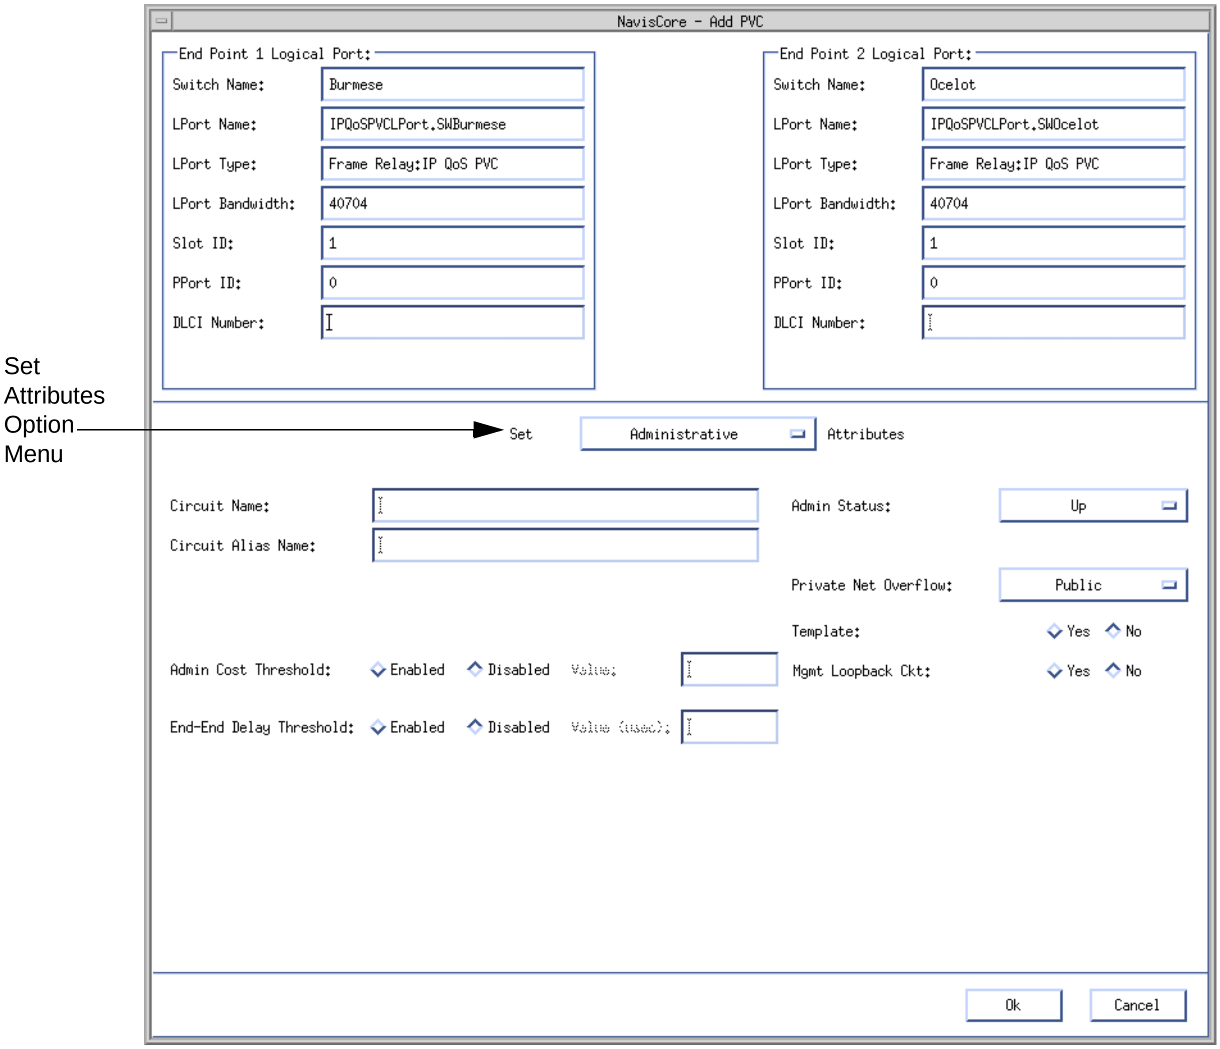

NavisCore - Add PVC

End Point 1 Logical Port:

Switch Name: Burmese
LPort Name: IPQoSPVCLPort.SWBurmese
LPort Type: Frame Relay:IP QoS PVC
LPort Bandwidth: 40704
Slot ID: 1
PPort ID: 0
DLCI Number:

End Point 2 Logical Port:

Switch Name: Ocelot
LPort Name: IPQoSPVCLPort.SWOcelot
LPort Type: Frame Relay:IP QoS PVC
LPort Bandwidth: 40704
Slot ID: 1
PPort ID: 0
DLCI Number:

Set Administrative Attributes

Circuit Name:
Circuit Alias Name:
Admin Status: Up
Private Net Overflow: Public
Template: ◇ Yes ◇ No
Admin Cost Threshold: ◇ Enabled ◇ Disabled Value:
Mgmt Loopback Ckt: ◇ Yes ◇ No
End-End Delay Threshold: ◇ Enabled ◇ Disabled Value (usec):

Ok Cancel

Figure 5-7. Add PVC Dialog Box (Administrative Attributes)

5. Access the Set Attributes option menu and complete the circuit attributes as described in the following sections.

6. After you complete the circuit attributes, choose OK to accept the circuit parameters and send the configuration file to the switch (provided the switch is communicating with the NMS). The Set All PVCs on Map dialog box reappears.

Administrative Attributes

Complete the administrative attributes fields described in Table 5-3.

Table 5-3. Set Administrative Attributes Fields

Field	Action/Description
DLCI Number *(Frame Relay and PPP endpoints)*	Enter a unique Data Link Connection Identifier (DLCI) for this PVC. A Frame Relay PVC is identified by its DLCI. For more information on DLCIs, see the *NavisCore Frame Relay Configuration Guide*.
VPI *(ATM endpoints)*	Enter a value from 0-*nnnn* to represent the Virtual Path Identifier (VPI) for the PVC. The maximum value you can enter is based on the valid bits in VPI that are configured for the logical port. Note that zero is not a valid value for a management PVC. See the *NavisCore ATM Configuration Guide* for information about setting this value.
VCI *(ATM endpoints)*	*(VCCs only)* Depending on the circuit configuration, enter a value to represent the Virtual Channel Identifier (VCI) for an ATM PVC. Note that VCIs apply to Virtual Channel Connections (VCCs) only. See the *NavisCore ATM Configuration Guide* for information about setting this value.
Next Hop IP Address *(egress Ethernet endpoints only)*	Enter the IP address of a node on the attached Ethernet LAN, such as the IP address of a router. The switch will forward IP packets from the egress port to the specified node. The IP address allows the switch to lookup the Ethernet Media Access Control (MAC) address of the node in its ARP cache. This field only appears when the endpoint is an Ethernet logical port.
Circuit Name	Enter any unique, continuous, alphanumeric name for the policy PVC. Do not use parentheses and asterisks.
Circuit Alias Name	(*Optional*) The circuit alias is used by service providers to identify the circuit in a way that is meaningful to their customers. This option is often used in conjunction with the NavisXtend Report Generator. Enter any unique alphanumeric name to identify the circuit. Do not use parentheses and asterisks. This name must be unique to the entire map.
Admin Status	Select Up or Down to activate or deactivate the circuit. *Up* – (default) Activates the circuit. *Down* – Takes the circuit off-line to run diagnostics such as PVC loopback.

Table 5-3. Set Administrative Attributes Fields (Continued)

Field	Action/Description
Private Net Overflow	(*For VNN Virtual Private Networks*) Set the Private Net Overflow parameters, which determine whether circuits originating from an LPort will be restricted to trunks of their own VNN VPN or use public (shared) trunks during overflow conditions. Options include: *Public* – (default) Enables the circuit to use public trunks during traffic overflow or trunk failure conditions. *Restrict* – Restricts trunks to their own virtual private network.
Template	(*Optional*) Save these settings as a template to use when configuring another circuit with the same options. To create a template, choose Yes in the *Is Template* field.
Mgmt Loopback Ckt	Choose Yes to include this PVC configuration in the NMS initialization script file. This file contains all the SNMP set requests necessary to replicate the entire switch configuration. Once you download this file to the switch, this PVC can be used to establish NMS-to-switch connectivity. This option is especially useful in some management DLCI configurations. The default value is No. (For more information about management DLCIs, see the *NavisCore Frame Relay Configuration Guide*.)
Admin Cost Threshold	If you enable this option, the PVC will not be routed over a path whose total administrative cost exceeds the entered value. This means that if you enable this field and enter a value of 1000, the PVC will not be routed over a path whose total administrative cost exceeds 1000. The total administrative cost for a path is calculated by summing the administrative cost for each trunk in the path. The valid range of values for this field is 1 - 65534. (Do not enable this option if you use End-End Delay routing.)
End-End Delay Threshold	If you enable this option, the PVC will not be routed over a path whose total end-to-end delay exceeds the entered value. This means that if you enable this field and enter a value of 500 µsec., the PVC will not be routed over a path whose total end-to-end delay exceeds 500 µsec. The total end-to-end delay for a path is calculated by summing end-to-end delay for each trunk in the path. The valid range for this field is 0 - 167777214 µsec. ***Note:*** *The value you enter should reflect your network topology. If a PVC will typically traverse high speed trunks, set the delay rate lower; increase the delay if the PVC must use low-speed trunks.*

Traffic Type Attributes

1. Select Traffic Type from the Set Attributes option menu. The traffic type fields appear (see Figure 5-8).

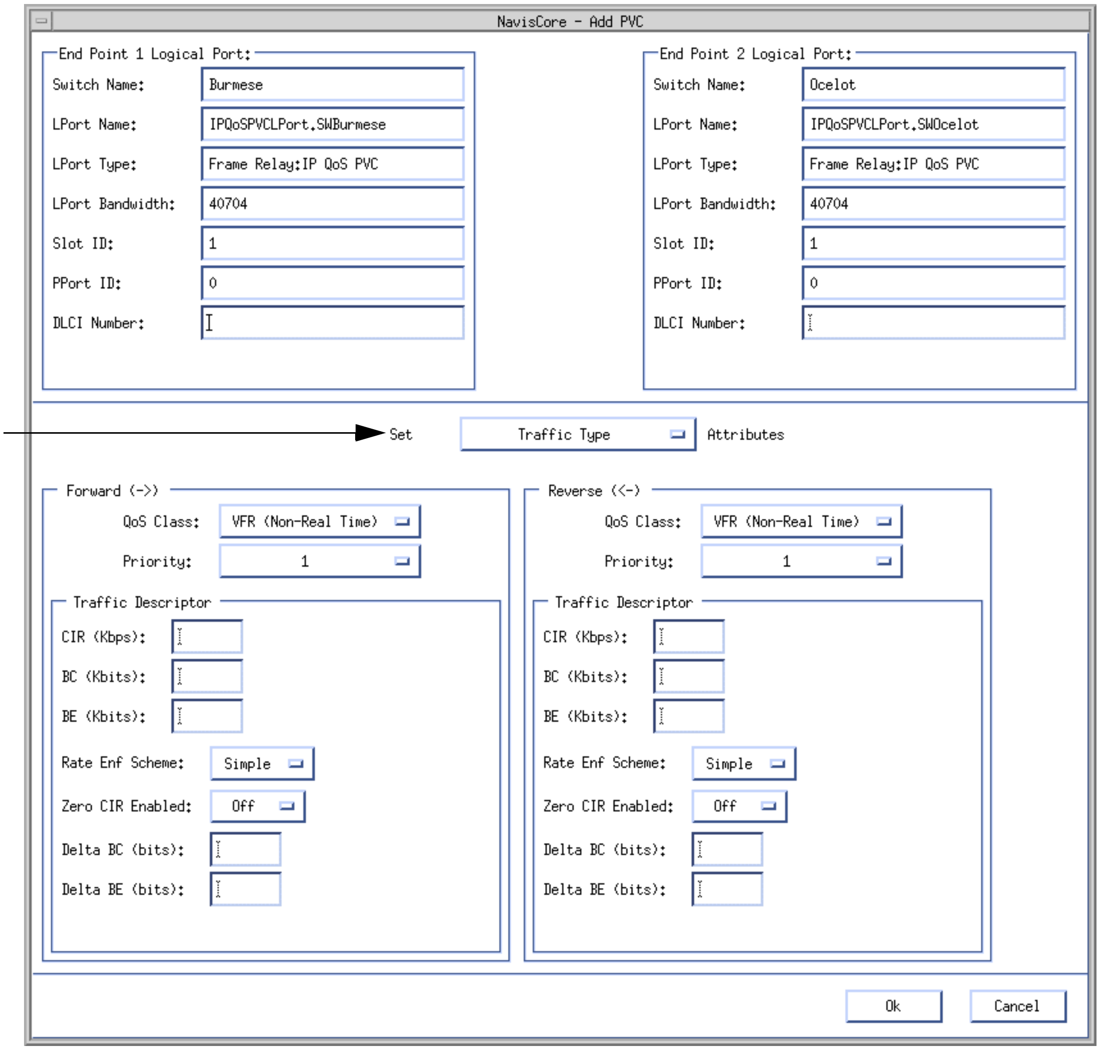

Choose Traffic Type

NavisCore - Add PVC

End Point 1 Logical Port:

Switch Name:	Burmese
LPort Name:	IPQoSPVCLPort.SWBurmese
LPort Type:	Frame Relay:IP QoS PVC
LPort Bandwidth:	40704
Slot ID:	1
PPort ID:	0
DLCI Number:	

End Point 2 Logical Port:

Switch Name:	Ocelot
LPort Name:	IPQoSPVCLPort.SWOcelot
LPort Type:	Frame Relay:IP QoS PVC
LPort Bandwidth:	40704
Slot ID:	1
PPort ID:	0
DLCI Number:	

Set Traffic Type Attributes

	Forward (->)	Reverse (<-)
QoS Class:	VFR (Non-Real Time)	VFR (Non-Real Time)
Priority:	1	1
Traffic Descriptor		
CIR (Kbps):		
BC (Kbits):		
BE (Kbits):		
Rate Enf Scheme:	Simple	Simple
Zero CIR Enabled:	Off	Off
Delta BC (bits):		
Delta BE (bits):		

Figure 5-8. Add PVC Dialog Box (Traffic Type Attributes)

2. Complete the fields described in Table 5-4 (for Frame Relay logical ports) and in Table 5-5 (for ATM logical ports). Keep in mind that the fields on the dialog box will vary, depending on whether the logical port endpoints support Frame Relay or ATM.

▶ The left column beneath the Forward (–>) arrow represents the logical port for the circuit that connects Endpoint 1 to Endpoint 2. The right column beneath the Reverse (<–) arrow represents the logical port for the circuit that connects Endpoint 2 to Endpoint 1. Enter values in both columns.

Table 5-4. Set Traffic Type Attributes Fields (Frame Relay Endpoints)

Field	Action/Description
QoS Class (Fwd/Rev)	Select one of the following Quality of Service (QoS) classes: *VFR (Real-Time)* – Variable Frame Rate (VFR). Used for special delay-sensitive applications that require low delay variation between endpoints. *VFR (Non-Real Time)* – Handles transfer of long, bursty data streams over a pre-established connection. This service provides low data loss but no delay guarantee. Also used for short, bursty data, such as LAN traffic. CPE protocols adjust for any delay or loss incurred through the use of VFR-NRT. *UFR* – Unspecified Frame Rate (UFR). Used for LAN traffic, primarily. The CPE should compensate for any delay or frame loss.
Priority (Fwd/Rev)	Select *1*, *2*, or *3* to configure the priority of data being transmitted on this circuit. Circuit priority determines the data's forwarding priority. The highest priority is 1 (do not discard data); the lowest is 3 (discards data). The forward and reverse circuit priority values do not have to match.
CIR (Kbps) (Committed Information Rate)	Enter the rate in Kbps at which the network transfers data under normal conditions. Normal conditions refer to a properly designed network with ample bandwidth and switch capacity. The rate is averaged over a minimum increment of the committed rate measurement interval (Tc). The value on each PVC is asymmetric (you can set a different CIR in each direction), which provides more efficient use of bandwidth.
Bc (Kbits) (Committed Burst Size)	Enter the maximum amount of data, in Kbits, that the network attempts to transfer data under normal conditions during a specified time interval, Tc. Tc is calculated as Bc/CIR. This value must be greater than zero and is typically set to the same value as CIR.
Be (Kbits) (Excess Burst Size)	Enter the maximum amount of uncommitted data, in Kbits, the network will attempt to deliver during a specified time interval, Tc. Tc is calculated as Bc/CIR. The network treats this data as *discard eligible* (DE) data.
Rate Enf Scheme	Select *Simple (default)* or *Jump*. The configurable rate enforcement scheme provides more flexibility, increased rate enforcement accuracy, and improved switch performance. ***Note****: If you select the Simple scheme, the "bad" PVC detection feature is disabled.*

Table 5-4. Set Traffic Type Attributes Fields (Frame Relay Endpoints) (Continued)

Field	Action/Description
Zero CIR Enabled (Fwd/Rev)	Set the CIR parameter to *On* or *Off*. *On* – Indicates that the PVC has an assigned CIR value of zero and is a best-effort delivery service. Customer data that is subscribed to zero CIR service can burst to the port speed if there is network bandwidth available to deliver frames. However, no frame delivery guarantees are made. All frames entering the network on zero CIR PVCs have DE set to one. *Off* – (default) Disables zero CIR. ***Note****: If you set Zero CIR Enabled to On, you cannot set the CIR, Bc, and Be values.*
Delta Bc (bits)	Set the number of Delta Bc bits for this circuit between 0 - 65528 (*default* 65528). This value represents the maximum number of bits the network agrees to transfer over the circuit (as committed bits) during the measurement interval, provided there is positive committed bit (Bc) credits before receiving the frame, but negative Bc credits after accepting the frame.
Delta Be (bits)	Set the number of Delta Be bits for this circuit between 0 - 65528 (default 65528). This value represents the maximum number of bits the network agrees to transfer over the circuit (as excess bits) during the measurement interval, provided there is positive excess bit (Be) credits before receiving the frame, but negative Be credits after accepting the frame.

Table 5-5. Set Traffic Type Attributes Fields (ATM Endpoints)

Field	Action/Description
QoS Class (Fwd/Rev)	Select one of the following Quality of Service (QoS) classes: *CBR* – Constant Bit Rate handles digital information, such as video and digitized voice that is represented by a continuous bit stream. CBR traffic requires guaranteed throughput rates and service levels. *VBR (Real Time)* – For packaging special delay-sensitive applications, such as packet video, that require low cell delay variation between endpoints. *VBR (Non-Real Time)* – Handles packaging for transfer of long, bursty data streams over a pre-established ATM connection. This service is also used for short bursty data, such as LAN traffic. CPE protocols adjust for any delay or loss incurred through the use of a VBR non-real time service class.
Priority (Fwd/Rev)	Select *1*, *2*, or *3* to configure the priority of data being transmitted on this circuit. Circuit priority determines the data’s forwarding priority. The highest priority is 1 (do not discard data); the lowest is 3 (discards data).
Traffic Descriptor	

Table 5-5. Set Traffic Type Attributes Fields (ATM Endpoints) (Continued)

Field	Action/Description
Type	The type of traffic descriptor you can specify depends on your choice of QoS class (see Table 5-6 on page 5-24).
PCR	Peak Cell Rate (PCR) is the maximum allowed cell transmission rate (expressed in cells per second). It defines the shortest time period between cells and provides the highest guarantee that network performance objectives (based on cell loss ratio) will be met.
SCR	Sustained Cell Rate (SCR) is the maximum average cell transmission rate that is allowed over a given period of time on a given circuit. It allows the network to allocate sufficient resources (but fewer resources than would be allocated based on PCR) for guaranteeing that network performance objectives are met. This parameter applies only to VBR traffic; it does not apply to CBR or UBR/ABR traffic.
MBS	Maximum Burst Size (MBS) is the maximum number of cells that can be received at the Peak Cell Rate. This allows a burst of cells to arrive at a rate higher than the SCR. If the burst is larger than anticipated, the additional cells are either tagged or dropped. This parameter applies only to VBR traffic; it does not apply to the CBR or UBR traffic.
MCR	Minimum Cell Rate (MCR) is the rate at which the source switch is always allowed to send data. This parameter only applies to ABR traffic. For more information about Flow Control Processor features, see the *NavisCore ATM Configuration Guide*.

Table 5-6. QoS Class Traffic Descriptors

QoS Class	Traffic Descriptor	Description
Constant Bit Rate (CBR) (specified/ unspecified)	PCR CLP=0, PCR CLP=0+1, tagging	Traffic conformance is based on the Peak Cell Rate (PCR) of both the CLP=0 and CLP=0+1 cell streams with tagging enabled.
	PCR CLP=0, PCR CLP=0+1, no tagging	Traffic conformance is based on the PCR of both the CLP=0 and CLP=0+1 cell streams with no tagging.
	PCR CLP=0+1	Traffic conformance is based only on the PCR of the CLP=0+1 aggregate cell stream with no best effort.
VBR-RT/ VBR-NRT (specified/ unspecified)	PCR CLP=0+1, SCR CLP=0, MBS CLP=0, tagging	Traffic conformance is based on the PCR of the CLP=0+1 aggregate cell stream, as well as the Sustained Cell Rate (SCR) and maximum burst size (MBS) of the CLP=0 cell stream with tagging enabled.
	PCR CLP=0+1, SCR CLP=0, MBS CLP=0, no tagging	Traffic conformance is based on the PCR of the CLP=0+1 aggregate cell stream, as well as the SCR and MBS of the CLP=0 cell stream with no tagging.
	PCR CLP=0+1, SCR CLP=0+1, MBS CLP=0+1	Traffic conformance is based on the PCR, SCR, and MBS of the CLP=0+1 cell stream with no tagging.
UBR	PCR CLP=0+1	Traffic conformance is based only on the PCR of the CLP=0+1 aggregate cell stream with no best effort.
	Best Effort	No traffic conformance is applied to this cell steam. A “best effort” attempt is made to deliver all traffic, but there is no guarantee the switch will not drop cells due to congestion.
	Best Effort, Tagging	Traffic conformance is only applied to tag all cells as CLP1. A “best effort” attempt is made to deliver all traffic, but there is no guarantee the switch will not drop cells due to congestion.
ABR	PCR CLP=0, MCR CLP=0	Traffic conformance is based on the PCR of the CLP=0 cell stream, as well as the MCR of the CLP=0 cell stream with no tagging.

User Preference Attributes

1. Select User Preference from the Set Attributes option menu. The user preference fields appear (see Figure 5-9).

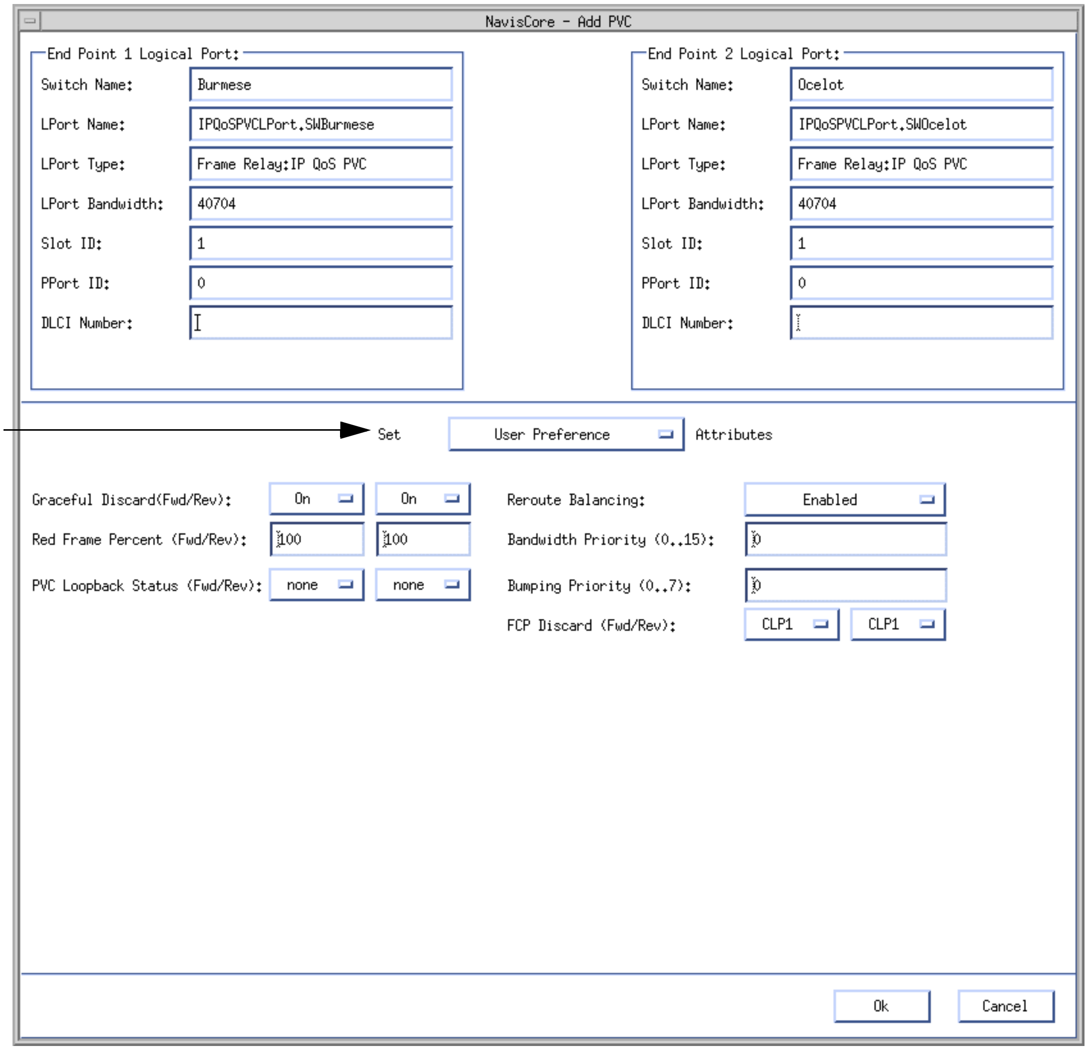

Choose User Preference

NavisCore - Add PVC

End Point 1 Logical Port:

Switch Name: Burmese
LPort Name: IPQoSPVCLPort.SWBurmese
LPort Type: Frame Relay:IP QoS PVC
LPort Bandwidth: 40704
Slot ID: 1
PPort ID: 0
DLCI Number:

End Point 2 Logical Port:

Switch Name: Ocelot
LPort Name: IPQoSPVCLPort.SWOcelot
LPort Type: Frame Relay:IP QoS PVC
LPort Bandwidth: 40704
Slot ID: 1
PPort ID: 0
DLCI Number:

Set User Preference Attributes

Graceful Discard(Fwd/Rev): On On
Red Frame Percent (Fwd/Rev): 100 100
PVC Loopback Status (Fwd/Rev): none none
Reroute Balancing: Enabled
Bandwidth Priority (0..15): 0
Bumping Priority (0..7): 0
FCP Discard (Fwd/Rev): CLP1 CLP1

Figure 5-9. Add PVC Dialog Box (User Preference Attributes)

2. Complete the fields described in Table 5-7. Keep in mind that the fields on the dialog box will vary, depending on whether the logical port endpoints support Frame Relay or ATM.

Table 5-7. User Preference Attributes

Field	Action/Description
Reroute Balancing	When enabled (default), the PVC conforms to the configured reroute tuning parameters. This means that when the PVC reroutes during trunk failure, it will migrate back to its original trunk at a rate and time determined by the configured reroute tuning parameters. When disabled, the PVC ignores the switch tuning parameters. For more information, see the *NavisCore NMS Getting Started Guide*.
Bandwidth Priority (0..15)	Set a value from 0 through 15 where 0 is the default and indicates the highest priority. See the *NavisCore ATM Configuration Guide* or the *NavisCore Frame Relay Configuration Guide* for more information.
Bumping Priority (0..7)	Set a number from 0 through 7 where 0 is the default and indicates the highest priority. See the *NavisCore ATM Configuration Guide* or the *NavisCore Frame Relay Configuration Guide* for more information.
FCP Discard (Fwd/Rev)	Displays only if you selected a QoS class that supports FCP Discard. Select one of the following options: *CLP1* – You can provision selective CLP1 discard for UBR, ABR, and VBR-NRT PVCs. If the current cell causes the queue for a PVC to exceed the discard thresholds, and the cell has CLP set to 1, the cell is discarded. Note that EPD is not performed in this case. *EPD* – Early Packet Discard. The ATM Flow-control processor can perform EPD for UBR, ABR, and VBR-NRT PVCs. If you select this option, when a cell causes the queue for a PVC to exceed the discard thresholds, the VC enters the EPD state. The cells in the current packet of the VC are admitted to the queue. However, when the end of the current packet is detected, all of the cells in the next packet are discarded for that PVC. See the *NavisCore ATM Configuration Guide* for more information. ***Note:*** *While the frame discard attribute is only applicable to a CBX 500 with an FCP, this attribute is offered as a selection on non-CBX endpoints. This is because even though one or both endpoints may not be on a CBX with FCP, the PVC might traverse a CBX 500 FCP trunk. In this case, the provisioned attribute is used.*
OAM Alarms (*ATM CS/IWU modules only*)	Set to *Enabled* (*default*) to use OAM alarms on this circuit. Set to Disabled to disable OAM alarms on this circuit. When enabled, the switch sends OAM F5 or F4 AIS (alarm indicator signal) cells out of each UNI logical port endpoint to indicate that the circuit is down. ***Note:*** *For a management PVC, this field is set to disabled and cannot be changed.*

Table 5-7. User Preference Attributes (Continued)

Field	Action/Description
UPC Function *(CBX 500 only)*	Enables (default) or disables the usage parameter control (UPC) function. When you enable UPC, the circuit tags or drops cells as they come into the port that do not conform to the configured traffic descriptors. When you disable UPC, the circuit allows all traffic, including non-conforming traffic, into the port. As a result, when you disable UPC, quality of service is no longer guaranteed for circuits in the network due to the potential for increasing the cell loss ratio because of port congestion. For this reason, *you should enable the UPC function on all circuits.* ***Note:*** *To use the UPC function for individual circuits, verify that the UPC function is enabled for both logical port endpoints on which you will define the circuit. Enabling UPC at the circuit level has no effect if you did not enable UPC at the logical port level. UPC is enabled by default for both logical ports and circuits.* ***Not****e: If both endpoints are configured as ATM CE endpoints, the UPC Function is not available.*
CDV Tolerance *(CBX 500 only)*	Configure the cell delay variation (CDV) tolerance. The UPC uses this value to police the requested traffic descriptor. Valid values are between 1–65535 microseconds. The default is 600 microseconds.
Graceful Discard (Fwd/Rev) (*ATM UNI endpoint on frame-based card*)	Select On (default) or Off to define how this circuit handles “red” packets. Red packets are designated as those bits received during the current time interval that exceed the committed burst size (Bc) and excess burst size (Be) thresholds, including the current frame. The Discard Eligible (DE) bit for a red packet is set to 1, meaning the network can discard this packet unless Graceful Discard is set to On. *On* – Forwards some red packets if there is no congestion. *Off* – Immediately discards red packets. ***Note:*** *For the ATM UNI DS3/E3, if you set this value for shaping purposes, the switch code ignores the PCR, SCR, and MBS values calculated from Set Traffic Descriptor Attributes; the switch instead picks the highest PCR queue available and sets the SCR to that PCR.*
Red Frame Percent (Fwd/Rev) (*ATM UNI endpoint on frame-based card*)	Set this value only if Graceful Discard is set to *On*. See the *NavisCore ATM Configuration Guide* for more information. The Red Frame Percent limits the number of red frames the network is responsible to deliver.
PVC Loopback Status (Fwd/Rev) (*ATM UNI endpoint on frame-based card*)	Displays the current loopback state. If **None** is not displayed in the PVC Loopback Status field, do not attempt to modify or delete the selected circuit. See the *NavisCore Diagnostics Guide* for more information about loopback testing.

Defining a Forwarding Policy

A forwarding policy contains a set of criteria. IP Navigator compares the values in the IP packet header fields (such as the IP protocol type) with the criteria defined in the IP forwarding policy. If the values in the IP packet header match all of the corresponding criteria in the forwarding policy, IP Navigator allows the IP packet to traverse the associated policy PVC. Otherwise, IP Navigator rejects the IP packet.

For example, suppose that you specify that only TCP traffic is allowed to traverse the policy PVC. Only IP packets that have a TCP protocol type value in the IP packet header are allowed to traverse the policy PVC. All other IP packets (e.g., UDP packets) are rejected.

To define a forwarding policy:

1. From the network map select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Forwarding Policies ⇒ Set All Forwarding Policies. The Set All Forwarding Policies dialog box appears (see Figure 5-10).

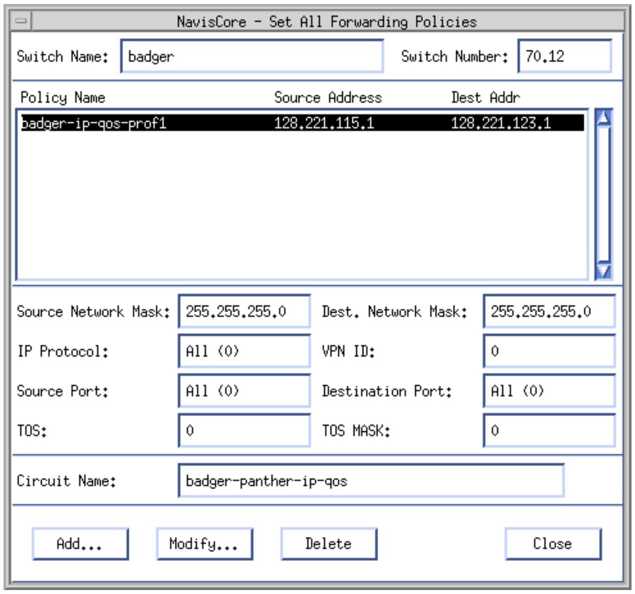

Switch Name: badger
Switch Number: 70.12

Policy Name	Source Address	Dest Addr
badger-ip-qos-prof1	128.221.115.1	128.221.123.1

Source Network Mask: 255.255.255.0
Dest. Network Mask: 255.255.255.0
IP Protocol: All (0)
VPN ID: 0
Source Port: All (0)
Destination Port: All (0)
TOS: 0
TOS MASK: 0
Circuit Name: badger-panther-ip-qos

Figure 5-10. Set All Forwarding Policies Dialog Box

The following table describes the buttons on the Set All Forwarding Policies dialog box.

Table 5-8. Set All Forwarding Policies Dialog Box: Buttons

Button	Function
Add	Enables you to add a forwarding policy.
Modify	Enables you to modify a forwarding policy.
Delete	Enables you to delete a forwarding policy.

3. Choose the Add button. The Add Forwarding Policy dialog box appears (see Figure 5-11).

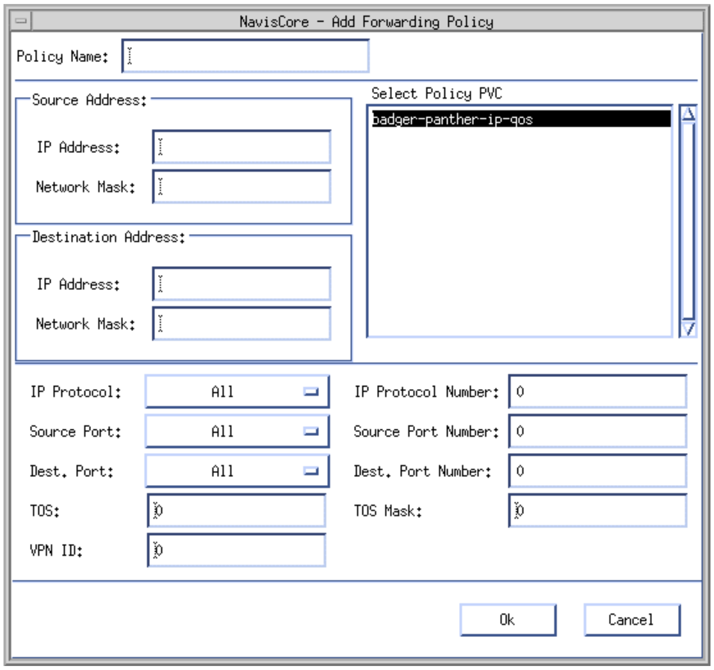

NavisCore - Add Forwarding Policy

Policy Name:

Source Address:
IP Address:
Network Mask:

Destination Address:
IP Address:
Network Mask:

Select Policy PVC
badger-panther-ip-qos

IP Protocol: All
Source Port: All
Dest. Port: All
TOS: 0
VPN ID: 0

IP Protocol Number: 0
Source Port Number: 0
Dest. Port Number: 0
TOS Mask: 0

Figure 5-11. Add Forwarding Policy Dialog Box

To allow **all** IP packets to traverse the associated policy PVC, accept the defaults when you define a forwarding policy. However, keep in mind that the default VPN ID (0) means that only public IP traffic (that is, IP packets that are not owned by a specific VPN customer) will traverse the associated policy PVC.

For example, if you enter a specific IP VPN ID (for example, 3) but retain the other default values, all IP packets belonging to the VPN customer who is assigned IP VPN ID 3 would traverse the policy PVC. Public IP packets and IP packets from other IP VPN customers would not traverse the policy PVC.

4. Complete the fields described in Table 5-9.

Table 5-9. Add Forwarding Policy Dialog Box: Fields

Field	Action/Description
Policy Name	Enter a policy name that will be used to associate the policy with a logical port.
Source Address	
IP Address	Enter the IP address of the network or host that sends the packet.
Network Mask	Enter the network mask that applies to the source address.
Destination Address	
IP Address	Enter the IP address of the network or host that receives the packet.
Network Mask	Enter the network mask that applies to the destination address.
Port/Protocol/ToS/VPN Information	
IP Protocol	Choose the IP protocol for the forwarding policy. The default (All) acts as a wildcard. You may choose from the following list of IP protocols: *All* – The forwarding policy accepts all types of IP protocols. If you leave the default unchanged, 0 (zero) appears in the IP Protocol Number field. *TCP* – Transmission Control Protocol. This protocol is connection-oriented, guaranteeing reliable transmission between applications that use it. Examples of applications that use TCP include FTP, Telnet, and HTTP (World-Wide Web protocol). If you choose TCP, 6 appears in the IP Protocol Number field. This number is the protocol’s official number as assigned by the Internet Assigned Numbers Authority (IANA). *UDP* – User Datagram Protocol. This protocol is connectionless — it makes a best effort to deliver datagrams between applications. Examples of applications that use UDP include RIP, TFTP, and SNMP. If you choose UDP, 17 appears in the IP Protocol Number field. This number is the protocol’s official number as assigned by the IANA. *ICMP* – Internet Control Message Protocol. This protocol provides dynamic routing support, such as routing redirects when routes are unavailable. If you choose ICMP, 1 appears in the IP Protocol Number field. This number is the protocol’s official number as assigned by the IANA. *User Specified* – Allows you to specify a protocol not included in this list, such as a proprietary protocol. If you choose *User Specified*, you must enter the IP protocol number assigned to the protocol in the IP Protocol Number field. This number is assigned by the IANA. For more information, refer to the IANA web site (http://www.iana.org) or RFC 1700, *Assigned Numbers*.
IP Protocol Number	Specify the number that identifies the IP protocol only if you chose *User Specified* for the IP Protocol. See the description of the IP Protocol field for more information. If you chose an IP Protocol type other than *User Specified*, NavisCore does not allow you to enter a value in this field and displays the IP protocol number associated with that choice.

Table 5-9. Add Forwarding Policy Dialog Box: Fields (Continued)

Field	Action/Description
Source Port (*Meaningful only if TCP or UDP is selected for IP Protocol*)	Choose the source application for the forwarding policy. Choose one of the following: *All* – (default) The forwarding policy applies to all application traffic from source to destination. *BGP* – BGP traffic. *FTP Data* – FTP data traffic. *FTP Control* – FTP control traffic. *Gopher* – Gopher traffic. *IRC* – IRC traffic. *Talk* – Talk traffic. *Telnet* – Telnet traffic. *WWW* – HTTP traffic. *RIP* – RIP traffic. *SNMP* – SNMP traffic. *SNMP Traps* – SNMP traps traffic. *TFTP* – TFTP traffic. *Ambiguous* – The traffic flows are tracked based on discrete source application values (1024 is used for the port value). *User Specified* – Allows you to specify an application that is not in the list (such as a third-party vendor application), and requires you to enter a port number that the application uses in the Source Port Number field. You can obtain this port number from the application developer. For more information on port numbers, refer to the IANA web site (http://www.iana.org) or RFC 1700, *Assigned Numbers*. You can find a listing of many port numbers at the following URL: http://www.isi.edu/in-notes/iana/assignments/port-numbers.
Source Port Number	Specify the port number used by the source application only if you chose *User Specified* for the Source Application. You can obtain this port number from the application developer. For more information on port numbers, refer to the IANA web site (http://www.iana.org) or RFC 1700, *Assigned Numbers*. You can find a listing of many port numbers at the following URL: http://www.isi.edu/in-notes/iana/assignments/port-numbers. If you chose a source port other than *User Specified*, NavisCore does not allow you to enter a value in this field and displays the port number associated with that choice. See the description of the Source Port field for more information.
Dest. Port (*Meaningful only if TCP or UDP is selected for IP Protocol*)	Choose the destination application for the forwarding policy. In most cases, your choices of source and destination application will be the same, since communication between different types of applications is rare. See the description of the Source Port field for information on the choices available to you.

Table 5-9. Add Forwarding Policy Dialog Box: Fields (Continued)

Field	Action/Description
Dest. Port Number	Specify the port number used by the destination application only if you chose *User Specified* for the Destination Port. You can obtain this port number from the application developer. For more information on port numbers, refer to the IANA web site (http://www.iana.org) or RFC 1700, *Assigned Numbers*. You can find a listing of many port numbers at the following URL: http://www.isi.edu/in-notes/iana/assignments/port-numbers. If you chose a Destination Port other than *User Specified*, NavisCore does not allow you to enter a value in this field and displays the port number associated with that choice. See the description of the Destination Port field for more information.
ToS	Specify the Type of Service (ToS) value. This value identifies a traffic flow associated with a particular ToS, such as voice. Values may range from 0 (the default) to 255. The default ToS mask (0) combined with the default ToS value (0) means that the forwarding policy will work with any ToS. If you specify a non-zero ToS value, you must specify a non-zero ToS mask. See the description of the ToS Mask field for more information. If you use a ToS mask and ToS value other than 0, make sure you specify values that are compatible with the network equipment (such as CPE) with which the switch exchanges service traffic (such as voice over IP).
ToS Mask	Specify the Type of Service (ToS) mask, in decimal, for the ToS value (see the description of the ToS field for more information). The mask determines the location of bits required for the ToS value. For example, if you specify a ToS value of 4 (100 in binary), you must specify a compatible ToS mask (such as decimal 4). Valid ToS masks range from 0 to 255. If you specify an invalid ToS mask/value combination, the NMS generates an error. The default ToS mask (0) combined with the default ToS value (0) means that the forwarding policy will work with any ToS. If you use a ToS mask and ToS value other than 0, make sure you specify values that are compatible with the network equipment (such as CPE) with which the switch exchanges service traffic (such as voice over IP).
VPN ID	Specify the ID of the IP VPN if you want to reserve the forwarding policy for use by a private IP VPN. The default is 0 (the forwarding policy is reserved for public use). To determine IP VPN IDs, display the Set All IP Virtual Private Networks dialog box. See "Adding an IP VPN" on page 16-28 for more information on displaying this dialog box. When you reserve a forwarding policy for use by a private IP VPN, no other IP VPN can use the forwarding policy.
Select Policy PVC	
Select Policy PVC	Select the PVC to which you want to assign the forwarding policy.

5. Choose OK.
6. At the Set All Forwarding Policies dialog box, choose Close.

Assigning a Forwarding Policy to a Logical Port

To assign a forwarding policy to a logical port:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, choose Lucent IP Parameters ⇒ Set All Forwarding Policies ⇒ Set All Logical Port Forwarding Policies. The Set All Logical Port Forwarding Policies dialog box appears (see Figure 5-12).

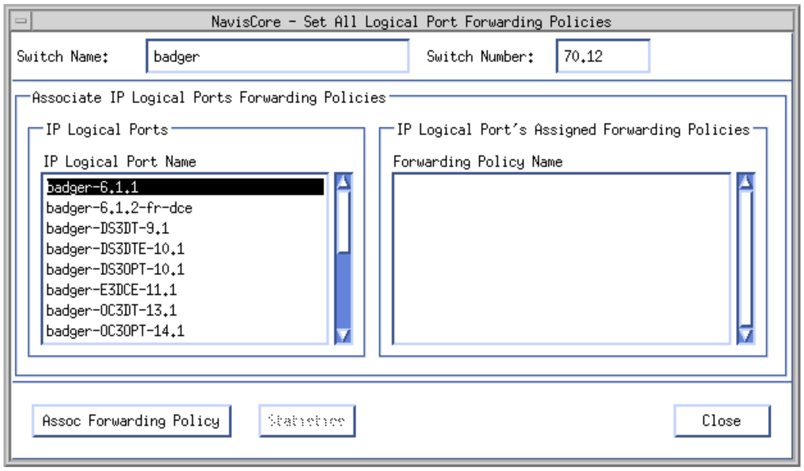

NavisCore - Set All Logical Port Forwarding Policies

Switch Name: badger

Switch Number: 70.12

IP Logical Port Name:

badger-6.1.1
badger-6.1.2-fr-dce
badger-DS3DT-9.1
badger-DS3DTE-10.1
badger-DS3OPT-10.1
badger-E3DCE-11.1
badger-OC3DT-13.1
badger-OC3OPT-14.1

Figure 5-12. Set All Logical Port Forwarding Policies Dialog Box

3. Select the IP logical port with which you want to associate a forwarding policy and choose the Assoc Forwarding Policy button. The Associate LPort Forwarding Policy dialog box appears (see Figure 5-13).

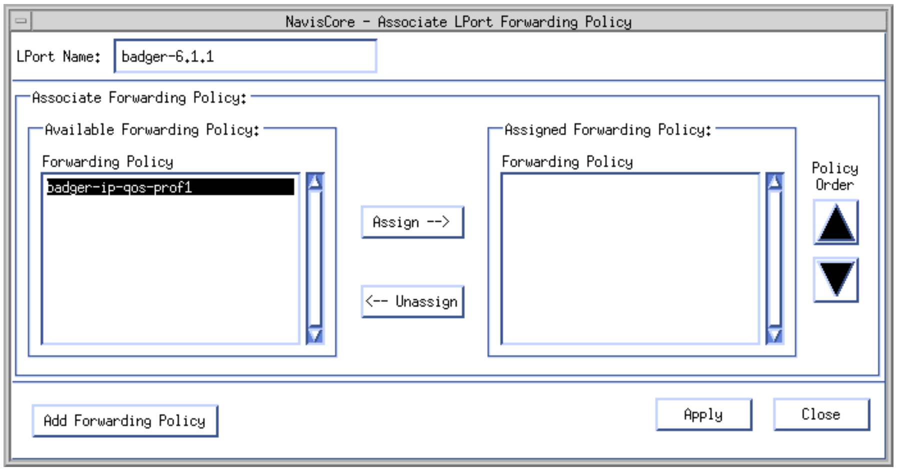

NavisCore - Associate LPort Forwarding Policy

LPort Name: badger-6.1.1

Available Forwarding Policy: badger-ip-qos-prof1

Figure 5-13. Associate LPort Forwarding Policy Dialog Box

4. Complete the fields on the Associate LPort Forwarding Policy dialog box as described in Table 5-10.

Table 5-10. Associate LPort Forwarding Policy Dialog Box: Fields

Field	Action Description
LPort Name	Displays the name of the logical port.
Available Forwarding Policy	Lists all current forwarding policies that are available.
Assigned Forwarding Policy	Lists all current forwarding policies assigned to this logical port.
Assign button	Enables you to assign a forwarding policy to the logical port.
Unassign button	Enables you to delete a forwarding policy from a logical port.
Policy Order buttons	Enables you to assign the forwarding policy's order of importance. In cases where the forwarding policy matches multiple policies, the first policy is always used.
Add Forwarding Policy	Enables you to define additional forwarding policies. See "Defining a Forwarding Policy" on page 5-28 for more information.

5. From the Available Forwarding Policy list box, select the policy to associate with the logical port and choose Assign.

   To delete a policy from the list, select a policy from the Assigned Forwarding Policy list box and choose Unassign.

6. Choose Apply.

   To create another forwarding policy, choose Add Forwarding Policy. For more information, see "Defining a Forwarding Policy" on page 5-28.

7. Choose Close.
8. At the Set All Logical Port Forwarding Policies dialog box, choose Close.

Enabling Policy-Based Forwarding

To enable policy-based forwarding on an IP logical port:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears.
2. Select the switch where the IP logical port resides from the Switch Name list.
3. Select the IP logical port from the LPort Name list.
4. Choose IP Parameters. The Set IP Parameters dialog box appears (see Figure 5-14).

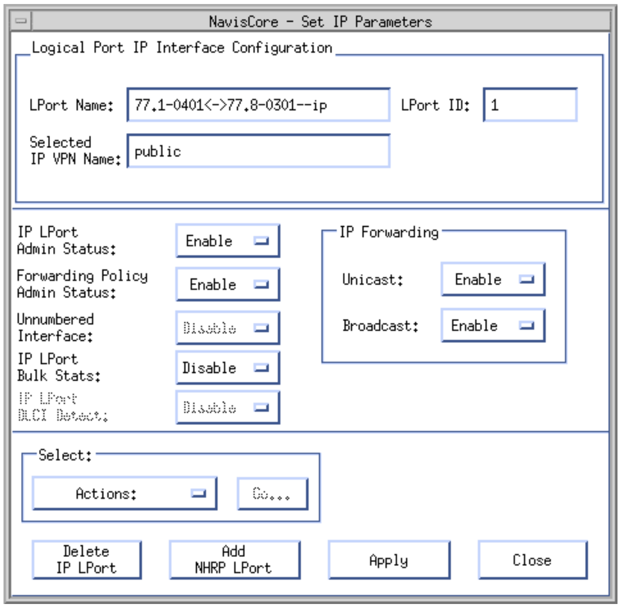

NavisCore - Set IP Parameters

Logical Port IP Interface Configuration

LPort Name: 77.1-0401<->77.8-0301--ip

LPort ID: 1

Selected IP VPN Name: public

IP LPort Admin Status: Enable

Forwarding Policy Admin Status: Enable

Unnumbered Interface: Disable

IP LPort Bulk Stats: Disable

IP LPort DLCI Detect: Disable

IP Forwarding

Unicast: Enable

Broadcast: Enable

Figure 5-14. Set IP Parameters Dialog Box (With Forwarding Policy Enabled)

5. Specify Enable in the Forwarding Policy Admin Status field.

6. Choose Apply.

Configuring the Recipient Switch or Router

When a policy PVC terminates on a Frame Relay or ATM logical port, the circuit ID of that logical port (DLCI or VPI/VCI) is used in the frames or cells sent over that policy PVC port instead of the circuit ID associated with the lport's IP address. You must configure the switch/router that receives datagrams from the policy PVC to accept frames or cells that use the circuit ID of the policy PVC or the recipient switch/router will ignore them.

For example, Figure 5-15 illustrates a policy PVC between two switches. The policy PVC is assigned a DLCI of 432, meaning that frames sent over the policy PVC have a DLCI of 432. However, the IP logical port on the egress switch has a DLCI of 123 associated with its IP address, meaning that frames sent out that port on a hop-by-hop basis have a DLCI of 123.

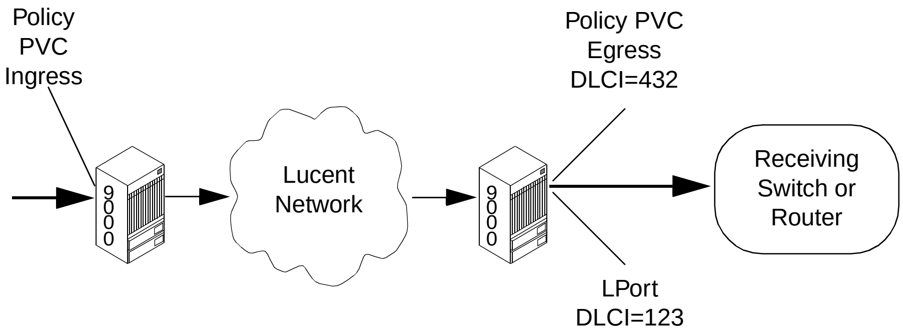

Figure 5-15. Configuring a Receiving Switch/Router

Because ARP/LMI does not distribute policy PVC circuit information automatically, you must configure the switch or router receiving the frames or cells to accept the policy PVC circuit ID or the receiving switch/router will discard frames/cells with unrecognized DLCI values.

The procedure for defining the policy PVC circuit ID in a receiving switch or router depends on the type of device. For example, if the device receiving the frames was a Cisco router, you would use the following port configuration parameters:

`frame-relay interface-dlci 432` (Frame Relay)

`atm pvc 2 4 32 aal5snap inarp 5` (ATM)

6

Configuring Static ARP Entries

This chapter describes how to define Static ARP Entries. When you define Static ARP entries, you create a table that matches IP addresses to specific hardware addresses (MAC, DLCI, VPI/VCI) or internal switch number addresses. The hardware address (or internal switch number) you define depends on the link type.

Address Resolution Protocol

A switch requires the following information in order to communicate with another node over Frame Relay, ATM, or Ethernet:

- IP address of the destination node
- Hardware address of the destination node (DLCI for Frame Relay, VPI/VCI for ATM, or MAC address for Ethernet)

To communicate with another switch over an IP VPN cloud interface, a switch requires:

- IP address of the destination switch's IP VPN cloud interface (see Chapter 16, "Configuring IP Virtual Private Networks" for more information on IP VPN cloud interfaces)
- Internal switch number of the destination switch

When an interface is configured for Ethernet, the IP addresses of the destination nodes are known (the hardware addresses are not known). When an interface is configured for Frame Relay, the hardware addresses of the destination nodes are known. IP services use ARP and InARP to resolve the lack of a hardware or IP address.

Defining a Static ARP Entry

Before you create a static ARP entry for Frame Relay or ATM interfaces, make sure that the hardware address that you plan to use as the Static ARP entry (either DLCI or VPI/VCI) has been already defined for the IP logical port on the Set IP Protocol Connection ID dialog box. See one of the following sections for details:

- "Setting the DLCI for Frame Relay Logical Ports"
- "Setting the VPI/VCI for ATM Logical Ports"

To define a static ARP entry:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Static ARP Entries. The Set All Static ARP Entries dialog box appears (see Figure 6-1).

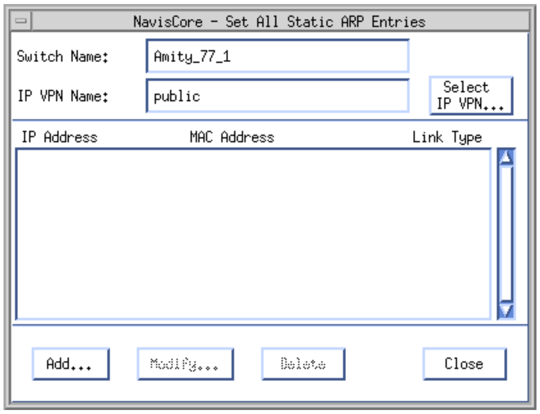

NavisCore - Set All Static ARP Entries

Switch Name: Amity_77_1

IP VPN Name: public

Figure 6-1. Set All Static ARP Entries List Dialog Box

The Set All Static ARP Entries dialog box displays the following buttons.

Table 6-1. Set All Static ARP Entries Buttons

Button	Function
Select IP VPN	Displays the Select IP VPN dialog box, allowing you to select an IP VPN. Once you select the IP VPN, all of the static ARP entries you define are for the selected IP VPN only. The name of the selected IP VPN appears in the IP VPN Name field. See "Selecting the IP VPN" on page 16-33 for more information.
Add	Enables you to add a static ARP entry.
Modify	Enables you to modify a static ARP entry.
Delete	Enables you to delete a static ARP entry.

▶ If you are creating an ARP entry for an IP VPN, make sure you are in the context of that IP VPN. For example, if you are adding an IP VPN cloud ARP entry for an IP VPN called "IPVPN1," make sure you are in the context of IPVPN1. To enter the context of an IP VPN, choose the Select IP VPN button and select the appropriate IP VPN.

3. Choose the Add button. The Set Static ARP dialog box appears (see Figure 6-2).

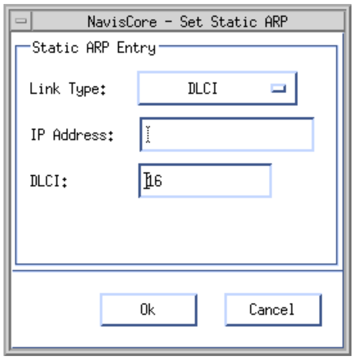
NavisCore - Set Static ARP

Static ARP Entry

Link Type: DLCI

IP Address:

DLCI: 16

Figure 6-2. Set Static ARP Dialog Box

4. Select the link type and complete the fields described in Table 6-2.

▶ For Frame Relay and ATM interfaces, the hardware address that you specify as the Static ARP entry (either DLCI or VPI/VCI) must have already been specified for the IP logical port on the Set IP Protocol Connection ID dialog box. See one of the following sections for details:

- "Setting the DLCI for Frame Relay Logical Ports"
- "Setting the VPI/VCI for ATM Logical Ports"

Table 6-2. Static ARP Fields

Link Type	Field	Action/Description
DLCI	IP Address	Enter the IP address of the neighbor.
	DLCI	Enter the DLCI used for the neighbor. Valid values range from 0 through 937. A DLCI is a 10-bit address that identifies PVCs.
VPI-VCI	IP Address	Enter the IP address of the neighbor.
	VPI	Enter the VPI used for the neighbor. Valid values range from 0 to 255. A VPI is an 8-bit field in the ATM cell header that is used as an addressing identifier to route cell traffic.
	VCI	Enter the VCI used for the neighbor. Valid values range from 0 to 255. A VCI is a 16-bit field in the ATM cell header that is used as an addressing identifer to route cell traffic.
Ethernet	IP Address	Enter the IP address of the neighbor.
	MAC Address	Enter the MAC Address used for the neighbor. A MAC address is a standardized data link layer address that is required for every port or device that connects to a LAN.
Cloud Interface	IP Address	Enter the IP address of the neighbor's IP VPN cloud interface. See Chapter 16, "Configuring IP Virtual Private Networks" for more information on IP VPN cloud interfaces.
	Cloud Address	Enter the IP address that the neighbor uses for its internal IP address. The internal IP address is configured when the switch is installed. To determine the internal IP address, access the switch console and issue the **`show system`** command. In the command output, the internal IP address appears in the Internal IP Addr field. For example: **Internal IP Addr: 150.202.77.2** In this example, the internal IP address is 150.202.77.2.

5. Choose OK. The ARP table entry is created for these addresses.
6. At the Set IP ARP List dialog box, choose Close.

7

Configuring RIP

This chapter describes how to configure Routing Information Protocol (RIP) parameters on an IP logical port and how to configure equal-cost multipath (ECMP) routing support for RIP. RIP is a distance vector protocol, which bases all routing decisions on the distance to the destination.

▶ This chapter does not discuss configuring RIP for IP virtual private networks (IP VPNs). For more information on IP VPN management, see Chapter 16, "Configuring IP Virtual Private Networks."

Configuring RIP at the Logical Port

To configure RIP at the logical port:

1. Add an IP logical port and interface. For more details on these procedures, see "Adding an IP Logical Port" on page 3-10.
2. Choose Add RIP from the Set IP Interface Addresses dialog box. The Add RIP Interface dialog box appears (see Figure 7-1).

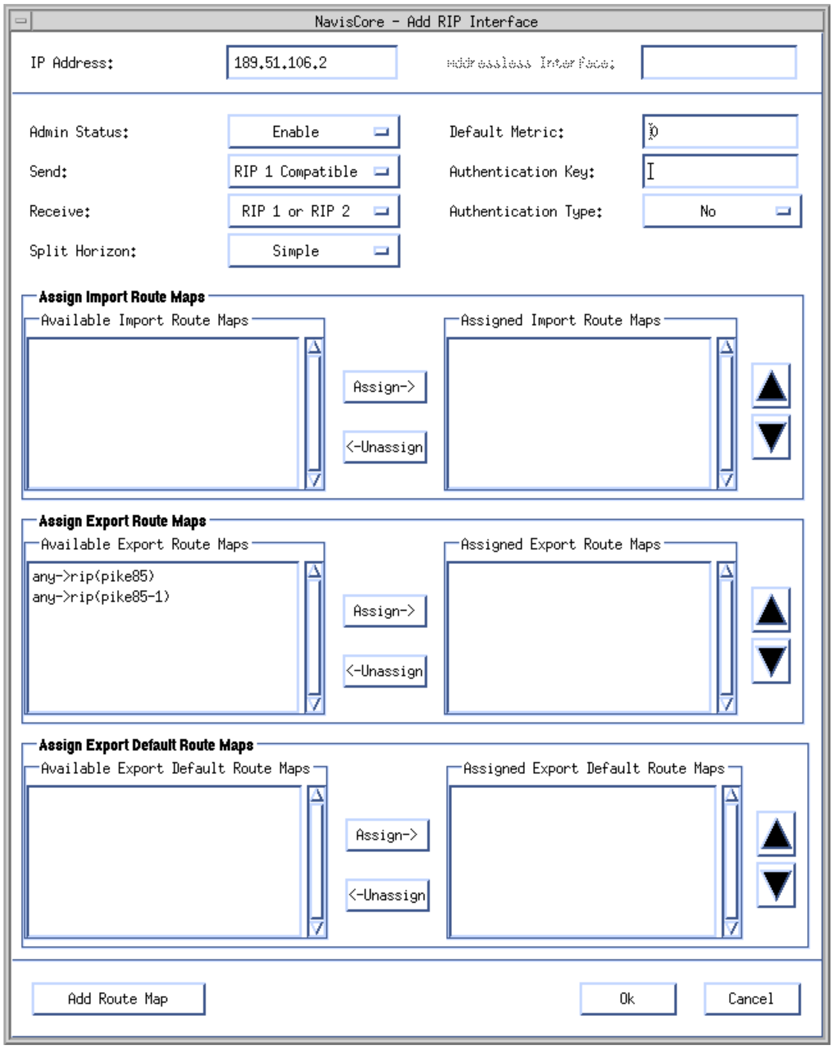

NavisCore - Add RIP Interface

IP Address: 189.51.106.2

Addressless Interface:

Admin Status: Enable

Send: RIP 1 Compatible

Receive: RIP 1 or RIP 2

Split Horizon: Simple

Default Metric: 0

Authentication Key:

Authentication Type: No

Available Export Route Maps: any->rip(pike85) any->rip(pike85-1)

Assign Import Route Maps
Available Import Route Maps
Assigned Import Route Maps
Assign->
<-Unassign
Assign Export Route Maps
Available Export Route Maps
Assigned Export Route Maps
Assign->
<-Unassign
Assign Export Default Route Maps
Available Export Default Route Maps
Assigned Export Default Route Maps
Assign->
<-Unassign
Add Route Map
Ok
Cancel

Figure 7-1. Add RIP Interface Dialog Box

3. To add a route map for this RIP interface, choose Add Route Map. See Chapter 11, "Configuring Route Policies" for more information on route maps.

4. To define the RIP interface, specify the field values as described in Table 7-1.

Table 7-1. RIP Interface Fields

Field	Action/Description
IP Address	The IP address for this interface.
Admin Status	Select one of the following options: *Enable* – Indicates that the port is activated for RIP and RIP packets can be exchanged over this logical port. *Disable* – Indicates that the port has not been activated for RIP or that the port is offline for diagnostics. An IP interface with an Admin Status of Disable cannot exchange RIP packets.
Send	Possible values are: *Disable, RIP 1, RIP 1 Compatible,* or *RIP 2.* RIP 1 Compatible is the default value.
Receive	Possible values are: *RIP 1, RIP2, RIP 1 or RIP 2,* or *Disable. RIP1 or RIP 2* is the default value.
Split Horizon	Split horizon is a method for avoiding common situations that require *counting to infinity.* Specify one of the following options: *Disable* – Indicates that split horizon will not be used. *Simple* – Indicates that split horizon will be used. The simple form of split horizon specifies that if a router learns of a route from an update received on the link, it does not advertise that route on updates that it transmits to the link. *Poison Reverse* – Is a stronger form of split horizon. In this form, routers do not omit destinations learned from an interface. Instead, they include these destinations, but advertise an infinite cost to reach them. This option increases the size of routing updates. In addition, it provides a positive indication that a specific location is not reachable through a router.
Default Metric	A variable that specifies the metric that is used for the default route entry in RIP updates that originate on this interface. A value of zero indicates that no default route should be originated.
Authentication Key	Do not specify this value if you specified a value of *No* as the authentication type. If you specified a value of *Simple* or *MD5* as the authentication type, you must specify the authentication password in this field.

Table 7-1. RIP Interface Fields (Continued)

Field	Action/Description
Authentication Type	This value specifies the type of authentication that RIP uses as a security measure to ensure that this logical port and router are exchanging information with proper neighbors. Possible values are *No, Simple,* or *MD5.* *No* – Specifies that no authentication will be performed. *Simple* – Specifies a simple password authentication method that enables you to designate a password that is part of all RIP messages on an interface-by-interface basis. When a router receives a message on an interface that is using simple password authentication, it checks the incoming RIP message to ensure that the proper password is included in the message. If the password is correct, the message is processed normally. If the password is not part of the incoming message or an incorrect password is used, the message is ignored and dropped. *MD5* – Specifies the Message Digest Algorithm Version 5 (MD5) authentication. This method is similar to the simple password method, however, the password is never transmitted. Instead, the router uses the MD5 algorithm to create a message digest of the password. The message digest is sent instead of the password. This method prevents the password from being read during transmission.
Available Import Route Maps	The import route maps that are available for assignment to this RIP interface. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, "Configuring Route Policies." To display the parameters for any listed route map, double-click on the map.
Assigned Import Route Maps	The import route maps that are assigned to this RIP interface. All incoming routes on this RIP interface are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.
Available Export Route Maps	The export route maps that are available for assignment to this RIP interface. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, "Configuring Route Policies." To display the parameters for any listed route map, double-click on the map.

Table 7-1. RIP Interface Fields (Continued)

Field	Action/Description
Assigned Export Route Maps	The export route maps that are assigned to this RIP interface. All outgoing routes on this RIP interface are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.

5. Choose OK. NavisCore adds the RIP interface and returns to the Set IP Interface Addresses dialog box.

Configuring ECMP Routing for RIP

This release supports equal-cost multipath (ECMP) routing, which load balances IP traffic over multiple routes of equal cost to the same destination. This feature applies to routes learned through BGP, RIP, OSPF, and manually configured static routes.

If multiple gateways of equal cost (that is, multiple equal-cost paths) can be used to reach the same destination, RIP will add up to four ECMP routes. If multiple equal-cost paths from a switch to the same destination are available, you can configure RIP on that switch to use them. Otherwise, RIP uses only one path. By taking advantage of multiple paths to a destination, you can load-balance network traffic.

By default, equal-cost multipath (ECMP) routing is disabled for RIP. Before you enable ECMP for RIP, note the following:

- If multiple neighbor gateways report the same cost for reaching the same destination, equal-cost routes are created.
- The maximum number of equal-cost paths to the same destination is four.
- Routing updates update the RIP equal-cost routes.
- Each route is aged independently.

To configure RIP to use multiple equal-cost paths:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set RIP Parameters. The Set RIP Parameters dialog box appears (see Figure 7-2).

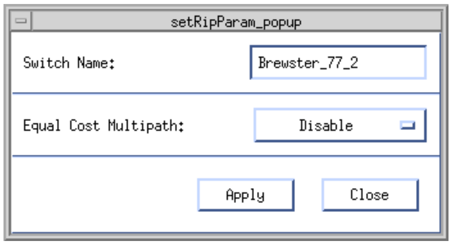

setRipParam_popup

Switch Name: Brewster_77_2

Equal Cost Multipath: Disable

Figure 7-2. Set RIP Parameters Dialog Box

3. Specify *Enable* (use multiple paths of equal cost) in the Equal Cost Multipath field. The default is *Disable* (do not use multiple paths of equal cost).
4. Choose Apply.

8

Configuring BGP Parameters

This chapter provides an overview of the Border Gateway Protocol (BGP) and describes how to perform the following tasks:

- Configuring BGP neighbors
- Configuring BGP aggregates
- Configuring BGP peer groups
- Configuring BGP route dampening
- Configuring IP loopback addresses

About BGP

BGP is a protocol that exchanges routing information between autonomous systems, which is a set of routers having a single routing policy running under a single technical administration. BGP advertises routes between external BGP neighbors or peers, unlike Interior Gateway Protocols (IGPs), which advertise routes between internal peers within the same autonomous system. Open Shortest Path First (OSPF) and Routing Information Protocol (RIP) are examples of IGPs.

See Figure 8-1 for an example of AS relationships.

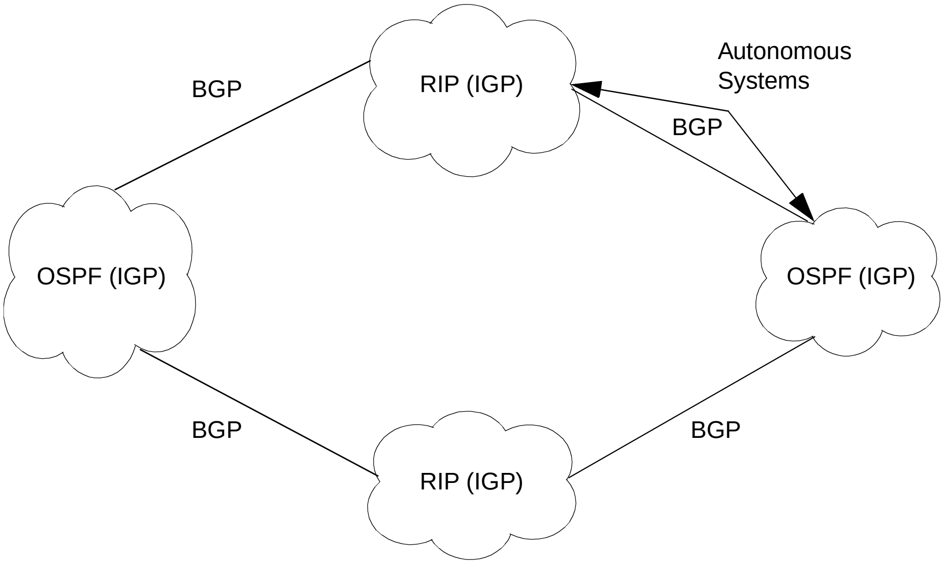

Figure 8-1. Autonomous System Examples

BGP Peers and Route Updates

BGP is considered a path-vector protocol because it carries a sequence of autonomous-system numbers that indicate the path a *route* has taken. When you define an autonomous system, you specify the networks to which a BGP peer sends route information. The network administrator uses a set of BGP parameters as tie breakers to indicate the routes that BGP should select as the best path.

BGP peers form a connection between each other, exchanging messages to open and confirm a TCP-based connection. Peers exchange route-update messages, which contain network reachability, path attributes, and preferred-route information. If there is disagreement between the peers, BGP sends an error to each peer and the connection is not established.

BGP updates route changes in a routing table. If routing information changes, BGP informs the peers by removing invalid routes and adding the new route information. If no changes occur, BGP peers exchange keep-alive messages to ensure the connection is alive.

BGP Peer Groups

A BGP *peer group* is a group of BGP neighbors that share the same route maps. Because you probably set the same routing policies for most BGP peers, the majority of your BGP neighbors are probably eligible for peer group membership.

Peer groups provide the following benefits:

Simplified Route Map Management — Peer groups simplify route map management. Instead of defining the same route maps for each neighbor, you can define one set of route maps and assign them to a peer group. Neighbors that are members of a peer group inherit all of the route map configuration options of the peer group. If a route map is applied to a BGP neighbor, and that neighbor is a member of a peer group, the individual route map is executed first, followed by the peer group policies.

Efficient Route Map Updates — Peer groups make the route map update process efficient. If you set up peer groups, IP Navigator no longer must parse route maps sequentially for each neighbor. Based on the route maps in the peer group, IP Navigator formulates the route map update once, and then floods the same route map update to all the group members.

Configuring IBGP

Typically, OSPF and RIP are used as the interior gateway protocol within the autonomous system. However, you can use BGP as the IGP. You can configure Interior Border Gateway Protocol (IBGP) the following ways:

- Full Mesh IBGP
- Route Reflection
- BGP Confederation

Full Mesh IBGP

In a full mesh IBGP, all IBGP neighbors within an autonomous system must be connected to exchange route update information. However, this is not the preferred configuration due to limited computing resources in a switch environment.

Figure 8-2 displays a full mesh IBGP.

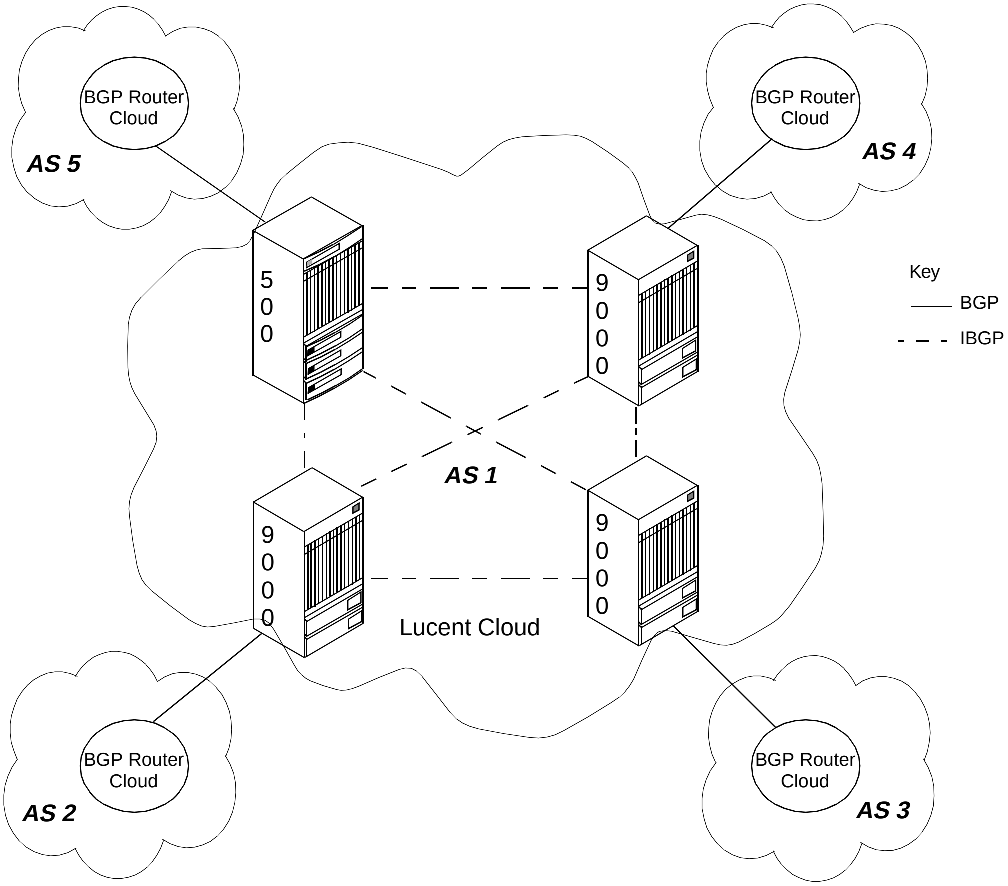

Figure 8-2. Full Mesh Interior Border Gateway Protocol Example

Route Reflection

Route reflection is a better alternative to full mesh IBGP. In route reflection, a BGP switch is designated as the route reflector, sending or *reflecting* received route information to all internal neighbors (or peers). There are two groups of route reflection peers:

- Client peers
- Non-client peers

When comparing the two groups, client peers do not have to be meshed, while non-client peers must be fully meshed together. Client peers are grouped into a *cluster* and communicate with each other. Client peers cannot communicate with non-client peers (peers outside of their cluster) but must communicate with the route reflector that belongs to the non-client peers' cluster.

Figure 8-3 illustrates an example of a route reflection configuration.

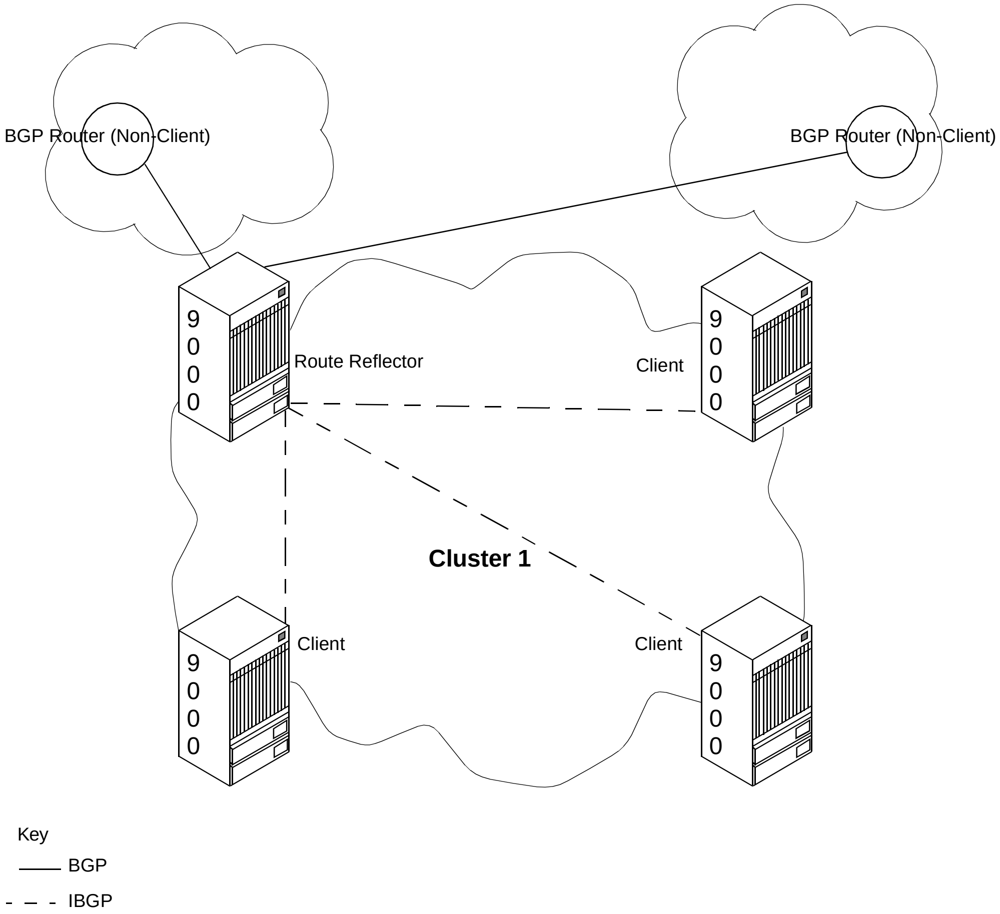

Figure 8-3. Route Reflection Example

For every route update received from an advertiser peer, the route reflector does one of the following (provided the best path selection is applied first):

- If the advertiser peer is a non-client, then the route reflector reflects the route to all non-clients.
- If the advertiser peer is a client peer, then the route reflector reflects the route to all non-client peers and all client peers other than the original advertiser.
- If the advertiser peer is an external BGP peer, then the route reflector reflects the route to all clients and non-clients (normal BGP operation).

Route reflection defines the following attributes for detection and avoidance of path loops:

ORIGINATOR_ID — This attribute is the router ID of the route originator in the local AS.

CLUSTER_LIST — This attribute is a sequence of cluster ID values that represent the reflection path the route passed.

Autonomous systems may have multiple route reflectors. Route reflectors communicating with each other are considered non-client peers and should be fully meshed.

BGP Confederation

Configuring IBGP to support BGP confederation provides an extension to BGP. This extension enables you to create a confederation of autonomous systems (AS), which are represented as a single autonomous system to BGP peers outside the confederation.

By subdividing an AS into smaller domains (called sub AS), you can control routing policy using information contained in the BGP AS_PATH attribute. The rules of IBGP apply inside each subAS. This means that within a subAS, all BGP routers must be fully meshed. Since each subAS uses a different AS number, the group of subASs that make up a confederation is using external BGP (EBGP). Even though EBGP runs between the subASs, routing within the confederation uses the same rules as IBGP routing within a single AS.

Subdividing a large AS provides a significant reduction in the total number of intra-domain BGP connections, since the connectivity requirements simplify to the model used for inter-domain connections. However, remember that dividing a large AS may unnecessarily lengthen the sequence portions of the AS_PATH attribute. Dividing an AS into separate systems may adversely affect the network's ability to route packets through the Internet.

You configure a BGP confederation using the following attributes:

AS Confederation — A collection of autonomous systems advertised as a single AS number to BGP speakers that are not members of the confederation.

AS Confederation Identifier — An externally visible autonomous system number that identifies the confederation as a whole.

Member-AS — An autonomous system that is contained in a given AS confederation.

The BGP confederation members use their confederation identifier in all transactions with peers outside the confederation. This confederation identifier is considered to be an "externally visible" AS number and this number is used in the AS_PATH attribute.

A member of a BGP confederation uses its routing domain identifier (the internally visible AS number) in all transactions with peers that belong to the same confederation as the given switch.

Figure 8-4 illustrates a BGP confederation configuration.

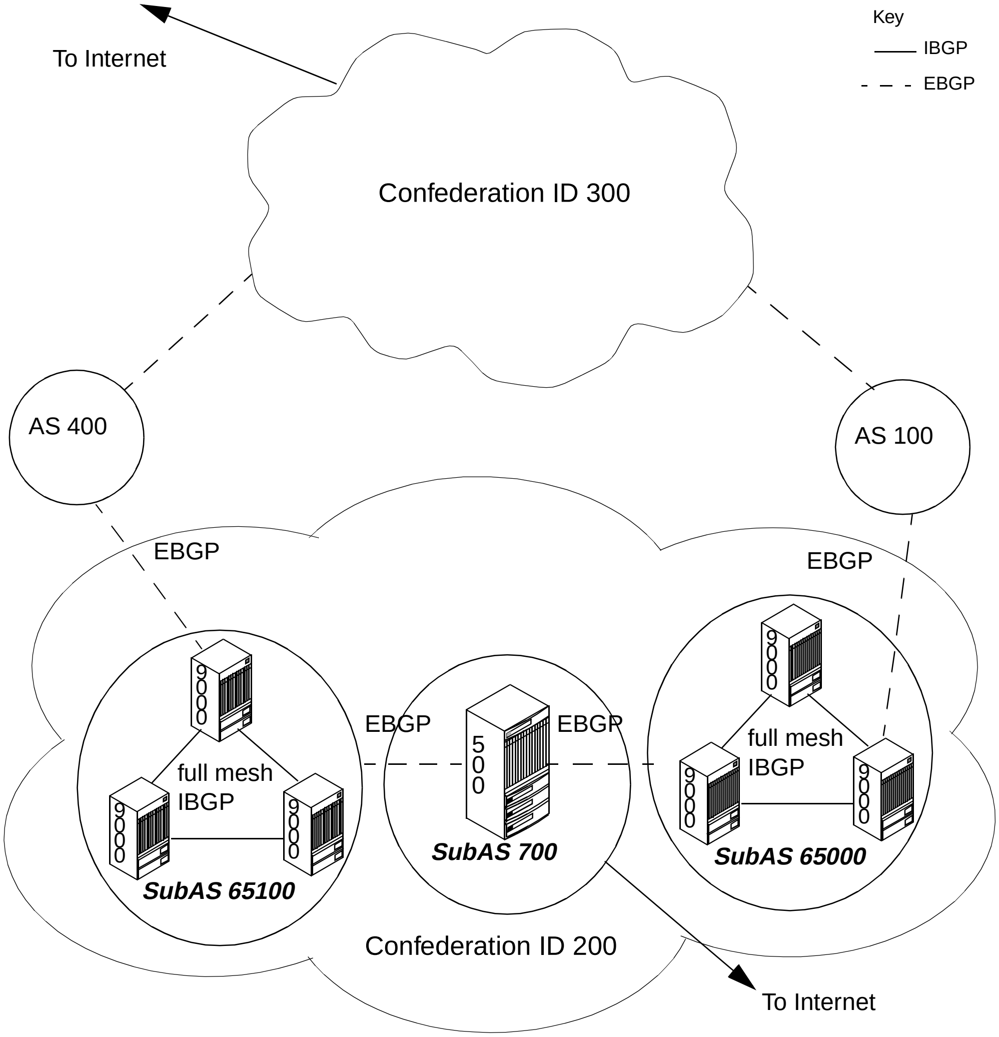

Figure 8-4. BGP Confederation Example

BGP Aggregates

An aggregate is formed by combining specific network addresses to less specific ones, thereby reducing the size of the routing table. Aggregates do the following:

- Reduce the size of the BGP routing table
- Provide better network control over network instability
- Provide a better mechanism to maintain route updates across areas

Aggregate networks use Classless Interdomain Routing (CIDR) addressing and are configured with a network prefix and mask. During the route update process, BGP scans the entire routing table for networks that are part of the configured aggregate network. If matches are found, BGP forms the aggregate networks and advertises aggregate routes to peers.

BGP Route Dampening

Because of hardware failures, software failures, system upgrades, insufficient network and system resources, and countless other problems, unstable routes are often introduced into routing tables across the network.

The intermittent appearance and disappearance of a route is the primary symptom of routing instability — a condition known as *flapping*. This condition causes the repeated propagation of BGP UPDATE messages (notifying the network of the appearance of a route) and WITHDRAWN messages (notifying the network of the route's disappearance) on the network. The large amount of routing traffic can consume significant network bandwidth and increase the CPU utilization of routers.

The establishment and maintenance of stable routes — both within networks and among networks — is a crucial goal for insuring reliable network connectivity. To achieve this goal, IP Navigator implements a mechanism called *route dampening*. The purpose of route dampening is to suppress routes that have become unstable. Route dampening reduces bandwidth utilization by reducing the amount of BGP UPDATE and WITHDRAWN messages that are transmitted as a result of route flapping. Routes that flap frequently are suppressed (that is, are not advertised) until a user-defined parameter expires.

To configure route dampening, network managers specify criteria to identify poorly behaved routes (routes that are announced and then withdrawn at a rapid rate). Depending on their degree of instability, flapping routes are penalized and prevented from being advertised to the network. If a route flaps at a low rate, the route is suppressed for a brief period of time or not at all.

Route suppression is dynamic, adapting to the frequency and duration with which particular routes appear to be flapping. The more a route flaps during a period of time, the longer it is suppressed based on several configurable parameters. See "Configuring BGP Route Dampening" on page 8-29 for more information on configuring these parameters.

Configuring BGP Switch Parameters

To configure BGP switch parameters:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All BGP ⇒ Set All BGP Parameters. The Set BGP dialog box appears (see Figure 8-5).

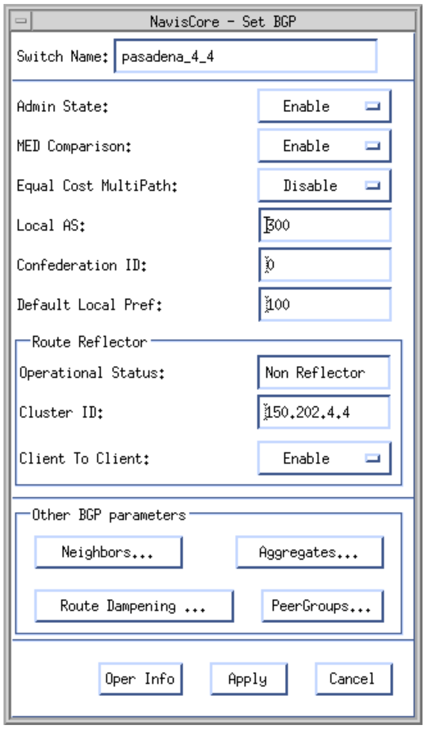

Switch Name: pasadena_4_4

Admin State: Enable

MED Comparison: Enable

Equal Cost MultiPath: Disable

Local AS: 300

Confederation ID: 0

Default Local Pref: 100

Operational Status: Non Reflector

Cluster ID: 150.202.4.4

Client To Client: Enable

Figure 8-5. Set BGP Dialog Box

The Set BGP dialog box provides the following buttons:

Table 8-1. Set BGP Buttons

Button	Function
Neighbors...	Choose this button to define BGP neighbor parameters. For more information, see "Configuring BGP Neighbors and Assigning Route Maps" on page 8-14.
Aggregates...	Choose this button to define BGP aggregate parameters. For more information, see "Configuring BGP Aggregates" on page 8-21.
Route Dampening	Choose this button to configure BGP route dampening. For more information, see "Configuring BGP Route Dampening" on page 8-29.
Peer Groups	Choose this button to configure BGP peer groups. For more information, see "Configuring BGP Peer Groups" on page 8-23.
Oper Info	Choose this button to update the Operational Status field.
Apply	Choose this button to put your changes into effect.

3. Complete the fields described in Table 8-2.

Table 8-2. BGP Parameter Fields

Field	Action/Description
Admin State	Select one of the following options: *Enable* – Allows the selected switch to exchange route updates using BGP. *Disable* – (default) Prevents the selected switch from exchanging route updates using BGP.
MED Comparison	Select one of the following options: *Enable* – (default) Allows you to use a multi-exit discriminator (MED) in the route selection process. MED allows BGP to communicate preferred path information to external neighbors when the autonomous system has multiple exits to another autonomous system. *Disable* – Prevents the use of MED in the route selection process.

Table 8-2. BGP Parameter Fields (Continued)

Field	Action/Description
Equal Cost Multipath	Select one of the following options: *Enable* – Enables equal-cost multipath (ECMP) routing for BGP. *Disable* – (default) Disables ECMP routing for BGP. If multiple, equal-cost AS paths from a switch to the same destination are available, you can configure BGP on that switch to use them. Otherwise, BGP uses only one path. By taking advantage of multiple paths to a destination, you can load-balance network traffic. By default, ECMP routing is disabled for BGP. Before you enable ECMP for BGP, note the following: • BGP can add up to four equal-cost paths for each item of Network Layer Reachability Information (NLRI). An NLRI is part of a route definition, consisting of an IP address and a prefix. Thus, BGP can add up to four equal-cost paths to the destination specified by the IP address/prefix combination in the NLRI. • When using multiple equal-cost paths to the same destination, BGP does not use the router ID to break the tie. • For NLRIs that have an ECMP route, the BGP update message for the NLRI specifies an aggregated AS_PATH of all the constituents of the ECMP route. • Both EBGP and IBGP routes are ECMP candidates. • Routes learned through both direct BGP peers and indirect BGP peers are ECMP candidates. • BGP peering to a loopback address is not an ECMP requirement.
Local AS	Enter a value between 1 and 65535. This parameter is the switch's Autonomous System (AS) number. The default is 0 (no AS number assigned).
Confederation ID	An externally visible Autonomous System number that identifies the confederation as a whole. If you are using BGP confederation, enter a value to specify a unique confederation identifier. The ID may be any unique number within the range of 1 to 65535. A value of 0 disables BGP confederation features.
Default Local Pref	Enter a value between 1 and 4294967295 (the default is 100). This value is exchanged between IBGP peers in the same AS. A local preference allows you to rank a route according to its importance. The local preference is compared to other routes that have the same destination. A higher local preference indicates the route is preferred. For example, a route with a local preference of 200 is preferred over a route with a local preference of 100.

Table 8-2. BGP Parameter Fields (Continued)

Field	Action/Description
Route Reflector	
Operational Status (READ ONLY)	This parameter identifies whether or not this peer is a route reflector. If it is, the peer forwards route information to all clients. The route reflector is implicitly defined when you define any of its peers to be a route reflector client.
Cluster ID	Enter the internal IP address of the selected switch if the switch is a route reflector in a cluster, which contains more than one route reflector. The default is the internal IP address of the switch. A cluster is a group of client peers that communicate with a BGP route reflector. A cluster ID specifies the cluster.
Client To Client (For Route Reflector peers only)	Select one of the following options: *Enable* – (default) If you enable this parameter, any routes that are received by the selected switch from a client will be sent to all other clients. Enable is the default. *Disable* – If you disable this parameter, any routes that are received by the selected switch from a client will not be sent to all other clients. ***Note**: Disable this parameter if all clients are fully meshed.*

4. Choose Apply.

Configuring BGP Neighbors and Assigning Route Maps

In addition to configuring BGP neighbors, you must assign route filters to these BGP nodes. Route maps control the flow of route updates. You use a route filter to selectively accept, reject (or hide), or advertise. See Chapter 11, “Configuring Route Policies” for details on defining route maps.

To configure a BGP neighbor for a switch:

1. From the network map select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All BGP ⇒ Set All BGP Neighbors. The Set All BGP Neighbors dialog box appears (see Figure 8-6).

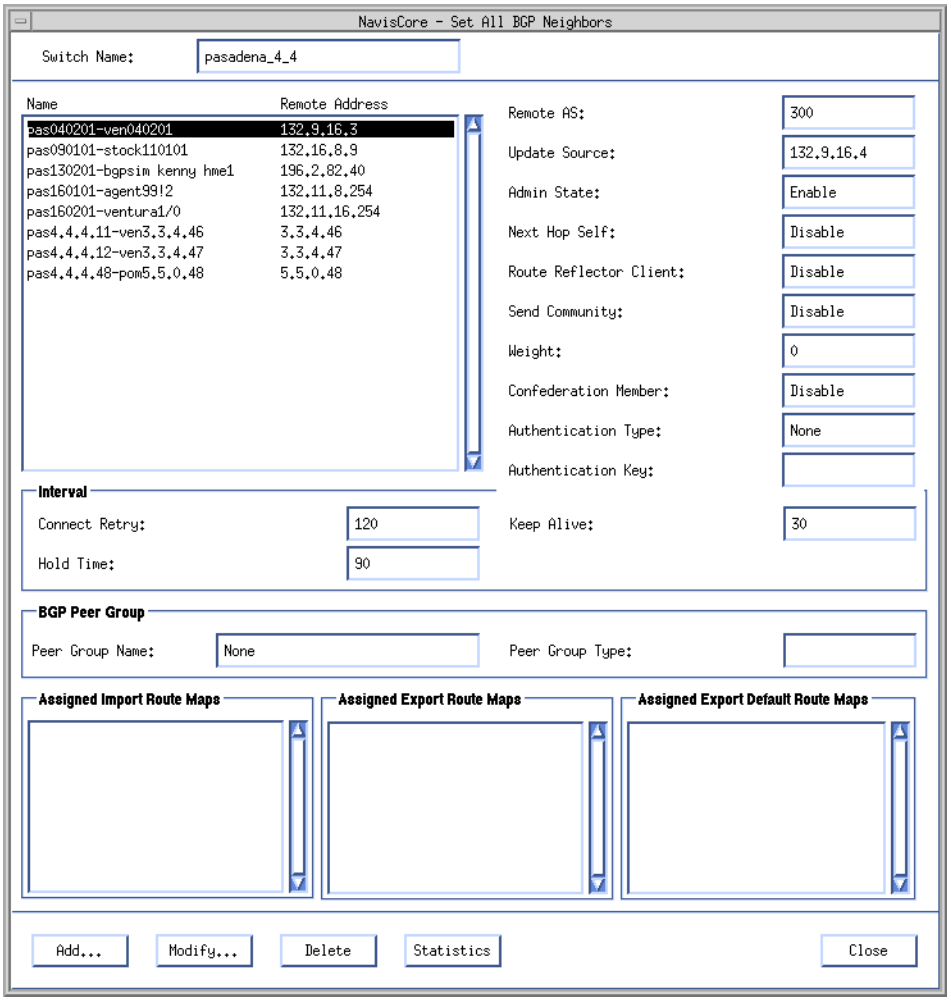

NavisCore - Set All BGP Neighbors

Switch Name: pasadena_4_4

Name	Remote Address
pas040201-ven040201	132.9.16.3
pas090101-stock110101	132.16.8.9
pas130201-bgpsim kenny hme1	196.2.82.40
pas160101-agent99!2	132.11.8.254
pas160201-ventura1/0	132.11.16.254
pas4.4.4.11-ven3.3.4.46	3.3.4.46
pas4.4.4.12-ven3.3.4.47	3.3.4.47
pas4.4.4.48-pom5.5.0.48	5.5.0.48

Remote AS: 300
Update Source: 132.9.16.4
Admin State: Enable
Next Hop Self: Disable
Route Reflector Client: Disable
Send Community: Disable
Weight: 0
Confederation Member: Disable
Authentication Type: None
Authentication Key:

Interval
Connect Retry: 120
Hold Time: 90
Keep Alive: 30

BGP Peer Group
Peer Group Name: None
Peer Group Type:

Assigned Import Route Maps
Assigned Export Route Maps
Assigned Export Default Route Maps
Add...
Modify...
Delete
Statistics
Close

To display the parameters for a listed route map, double-click on the route map.

Figure 8-6. Set All BGP Neighbors Dialog Box

The Set All BGP Neighbors dialog box displays the following buttons:

Table 8-3. Set All BGP Neighbors Buttons

Button	Function
Add	Enables you to add a BGP neighbor.
Modify	Enables you to modify a BGP neighbor.
Delete	Enables you to delete a BGP neighbor.
Statistics	Use the Statistics option to display BGP peer connection statistics. For more information, see the *NavisCore Diagnostics Guide.*

3. Choose Add.

If you selected a switch that has an AS of zero, the following error message appears:

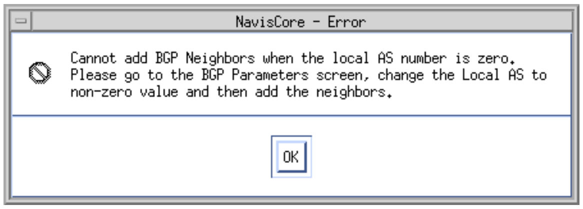

NavisCore - Error

Cannot add BGP Neighbors when the local AS number is zero. Please go to the BGP Parameters screen, change the Local AS to non-zero value and then add the neighbors.

Figure 8-7. BGP Neighbor Error Message

Go to “Configuring BGP Switch Parameters” on page 8-10 and change the switch’s AS number to a non-zero value.

If you selected a switch that has a non-zero AS value, the Add BGP Neighbor dialog box appears (see Figure 8-8).

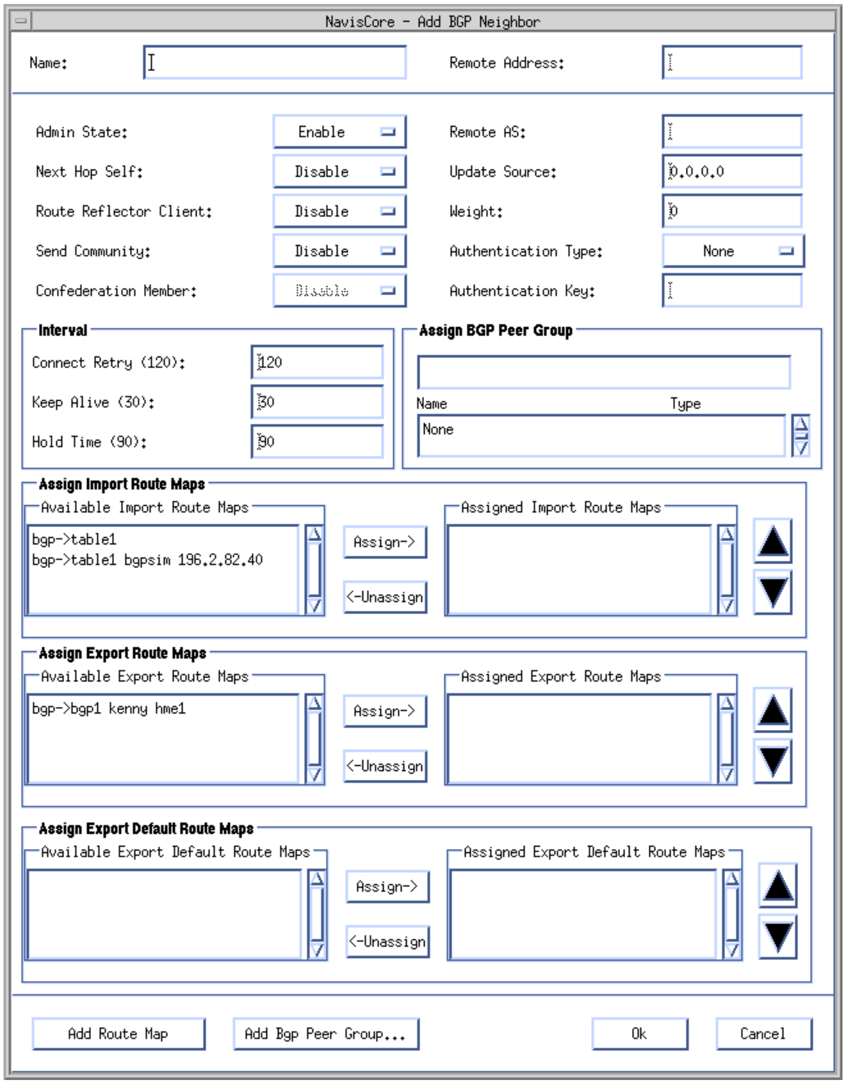
NavisCore - Add BGP Neighbor
Interval
Assign BGP Peer Group
Assign Import Route Maps
Available Import Route Maps
Assigned Import Route Maps
Assign->
<-Unassign
Assign Export Route Maps
Available Export Route Maps
Assigned Export Route Maps
Assign Export Default Route Maps
Available Export Default Route Maps
Assigned Export Default Route Maps
Add Route Map
Add Bgp Peer Group...
Ok
Cancel

Name:

Remote Address:

Admin State: Enable

Next Hop Self: Disable

Route Reflector Client: Disable

Send Community: Disable

Confederation Member: Disable

Remote AS:

Update Source: 0.0.0.0

Weight: 0

Authentication Type: None

Authentication Key:

Connect Retry (120): 120

Keep Alive (30): 30

Hold Time (90): 90

Name	Type
None	

Available Import Route Maps:
bgp->table1
bgp->table1 bgpsim 196.2.82.40

Available Export Route Maps:
bgp->bgp1 kenny hme1

Use the arrow buttons to specify the sequence that IP Navigator uses to filter routing information.

Figure 8-8. Add BGP Neighbor Dialog Box

4. Specify the values as listed in Table 8-4.

Table 8-4. BGP Neighbor Fields

Field	Action/Description
Name	Enter the name of the BGP neighbor.
Remote Address	Enter the IP address of the BGP neighbor.
Admin State	Select one of the following options: *Enable* – (default) Activates the connection between the selected switch and this BGP neighbor. *Disable* – Deactivates the connection between the selected switch and this BGP neighbor.

Table 8-4. BGP Neighbor Fields (Continued)

Field	Action/Description
Remote AS	Enter a value between 1 and 65535. This value is the neighbor's remote AS number.
Next Hop Self	Select one of the following options: *Enable* – For IBGP peers, enabling this parameter forces BGP to advertise the local address of the BGP connection as the next hop. For EBGP peers, BGP always advertises the local address as the next hop; therefore you do not need to enable next hop self for EBGP peers. *Disable* – (default) Disabling this parameter allows BGP to determine the next hop.
Update Source	Enter a valid IP address for the Update Source address, which is the source address for the BGP TCP connection.
Route Reflector Client	Select one of the following options: *Enable* – If you enable this parameter, the selected switch's neighbor is defined as a route reflector client, implicitly making the selected switch a route reflector. *Disable* – (default) If you disable this parameter, the selected switch's neighbor is not defined as a route reflector client. In addition, if you disable this parameter on all of the selected switch's BGP neighbors, the selected switch is not defined as a route reflector. However, if the route reflector client is enabled on at least one BGP neighbor, the selected switch is still considered a route reflector.
Weight	Enter a value between 0 and 65535 (the default value is zero). This parameter represents the path weight (received by the neighbor) that is applied to every EBGP route. IP Navigator applies the weight value to EBGP routes only. It does not use the weight value for IBGP routes.
Send Community	Select one of the following options: *Enable* – Enables you to send community attributes of all updates to this BGP neighbor. A community is a group of destinations that share some common property. A community is not restricted to one network or autonomous system; it has no physical boundaries. You use community attributes to simplify routing policies by identifying routes based on the logical property rather than IP prefix or AS number. *Disable* – (default) Disables the sending of community attributes of all updates to this BGP neighbor.
Authentication Type	Select one of the following options: *None* – (default) No authentication is used. *MD5* – Select MD5 if you want to use MD5 authentication to protect BGP peer sessions against the introduction of spoofed TCP segments into a TCP connection stream. For more information on how MD5 authentication works, see RFC 1321 (The MD5 Message-Digest Algorithm). In addition, see RFC 2385 (Protection of BGP Sessions via the TCP Signature Option) for more information on how MD5 authentication is used to protect BGP TCP sessions. If you select MD5, make sure you specify a key in the Authentication Key field.

Table 8-4. BGP Neighbor Fields (Continued)

Field	Action/Description
Confederation Member	Select one of the following options: *Enable* – Includes this BGP neighbor as a confederation member. Before you can configure a BGP confederation, ensure that you have entered a valid Confederation ID on the Set BGP Attributes dialog box (see Figure 8-5 on page 8-10). *Disable* – (default) Specifies that this BGP neighbor is not part of a confederation.
Authentication Key	If you selected MD5 as an authentication type, specify an alphanumeric string that will act as the MD5 authentication key.
Interval (The default value for each interval field is in parentheses.)	
Connect Retry (120)	Enter a value between 1 and 65535 (the default is 120). This parameter is the time, in seconds, that BGP waits before it tries to connect to this neighbor. The number of connection retries due to errors are generated with no regard to this value. The initial value is 60 seconds, which is doubled for each retry after that.
Keep Alive(30)	Enter a value between 0 and 21845 (the default is 30). This parameter is the time, in seconds, between consecutive keep alive messages sent to this neighbor. This event occurs after a connection is established. Keep alive messages are sent periodically between BGP neighbors to ensure that the connection is still alive.
Hold Time(90)	Enter either a value of 0, or a range of 3 to 65535 (the default is 90). The value 0 indicates not to use hold time with this neighbor. This parameter represents the time, in seconds, BGP holds before considering the connection to be down if messages are not received from this neighbor.
Assign BGP Peer Group	Enables you to assign a BGP peer group. See "Configuring BGP Peer Groups" on page 8-23 for more information.
Assign Import Route Maps	
Available Import Route Maps	The import route maps that are available for assignment to this BGP neighbor. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, "Configuring Route Policies." To display the parameters for any listed route map, double-click on the map.
Assigned Import Route Maps	The import route maps that are assigned to this BGP neighbor. All incoming routes on this BGP neighbor are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.

Table 8-4. BGP Neighbor Fields (Continued)

Field	Action/Description
Assign Export Route Map	
Available Export Route Maps	The export route maps that are available for assignment to this BGP neighbor. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, “Configuring Route Policies.” To display the parameters for any listed route map, double-click on the map.
Assigned Export Route Maps	The export route maps that are assigned to this BGP neighbor. All outgoing routes on this BGP neighbor are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.
Assign Export Default Route Maps	
Available Export Default Route Maps	The export default route maps that are available for assignment to this BGP neighbor. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, “Configuring Route Policies.” To display the parameters for any listed route map, double-click on the map.
Assigned Export Default Route Maps	The export default route maps that are assigned to this BGP neighbor. All outgoing routes on this BGP neighbor are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.

5. In the Available Import Route Maps List box, specify the import route map and choose assign.

   To remove an import route map from the list, select the import route map from the Assigned Import Route Maps List box and choose Unassign.

6. In the Available Export Route Maps List box, specify the export route map and choose assign.

   To remove an export route map from the list, select the export route map from the Assigned Export Route Maps List box and choose Unassign.

7. In the Available Export Default Route Maps List box, specify the export default route map and choose assign.

   To remove an export default route map from the list, select the export default route map from the Assigned Export Default Route Maps list box and choose Unassign.

8. Choose OK.

9. To add an additional network access list, choose Add Route Map. See "Adding Route Maps" on page 11-18 for more information.

   To add a BGP peer group, choose Add Bgp Peer Group. See "Configuring BGP Peer Groups" on page 8-23 for more information.

10. In the Set All BGP Neighbors dialog box, choose Close.

Configuring BGP Aggregates

To configure BGP aggregates:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, choose Lucent IP Parameters ⇒ Set All BGP ⇒ Set All BGP Aggregates. The Set All BGP Aggregates dialog box appears (see Figure 8-9).

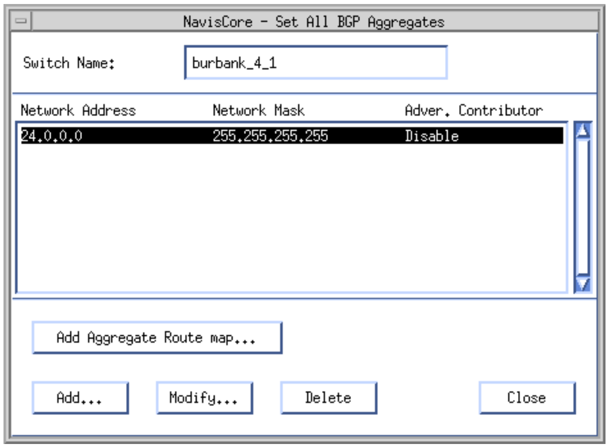
NavisCore - Set All BGP Aggregates
Add Aggregate Route map...
Add...
Modify...
Delete
Close

Switch Name: burbank_4_1

Network Address	Network Mask	Adver. Contributor
24.0.0.0	255.255.255.255	Disable

Figure 8-9. Set All BGP Aggregates Dialog Box

The Set All BGP Aggregates dialog box displays the following buttons.

Table 8-5. Set All BGP Aggregates Buttons

Button	Function
Add Aggregate Route Map	Enables you to add an Aggregate Route Map. For more information, see Table 11-33 on page 11-55.
Add	Enables you to add a BGP aggregate.
Modify	Enables you to modify a BGP aggregate.
Delete	Enables you to delete a BGP aggregate.

3. At the Set All BGP Aggregates dialog box, choose Add. The Add BGP Aggregates dialog box appears (see Figure 8-10).

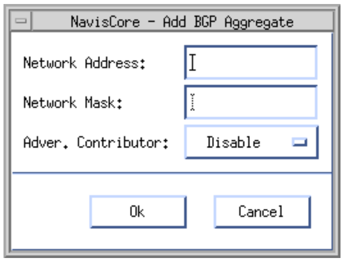
NavisCore - Add BGP Aggregate

Network Address:

Network Mask:

Adver. Contributor: Disable

Figure 8-10. Add BGP Aggregate Dialog Box

4. Specify the values described in Table 8-6.

Table 8-6. BGP Aggregate Fields

Field	Action/Description
Network Address	This parameter is the aggregate network IP address.
Network Mask	This parameter is the aggregate network mask.
Adver. Contributor	Select one of the following options: *Enable* – Enabling this parameter allows you to advertise components of the aggregate network. *Disable* – Disabling this parameter enables you to stop advertising components of the aggregate network.

5. Choose OK.

6. At the Set All BGP Aggregates dialog box, choose Close.

Configuring BGP Peer Groups

This section provides instructions on how to configure peer groups. Before you configure peer groups, note the following:

- The peer group configuration process involves associating a group of BGP neighbors with one or more route maps. See "Configuring BGP Neighbors and Assigning Route Maps" on page 8-14 for more information on configuring BGP neighbors. See Chapter 11, "Configuring Route Policies" for more information on configuring route maps.
- All members of the same peer group must be the same BGP node type (EBGP, IBGP client, or IBGP non-client).
- Route maps in a peer group can be sequenced. This is no different than sequencing the route maps that you assign directly to a BGP neighbor.
- Some neighbors in a peer group may have slightly different route map requirements than other neighbors. You can add additional route maps to the neighbor that complements the set of route maps already assigned to the peer group.
- For each neighbor in the peer group, IP Navigator evaluates the route maps defined specifically for that neighbor first, then it evaluates the peer group route maps.
- A switch can support a maximum of 10 different peer groups and a maximum of 50 associations between peer groups and route maps.

To configure BGP peer groups:

1. From the network map, select the appropriate switch icon.

2. From the Administer menu, select Ascend IP Parameters ⇒ Set All BGP ⇒ Set All BGP Peer Groups. The Set All BGP Peer Groups dialog box appears (see Figure 8-11).

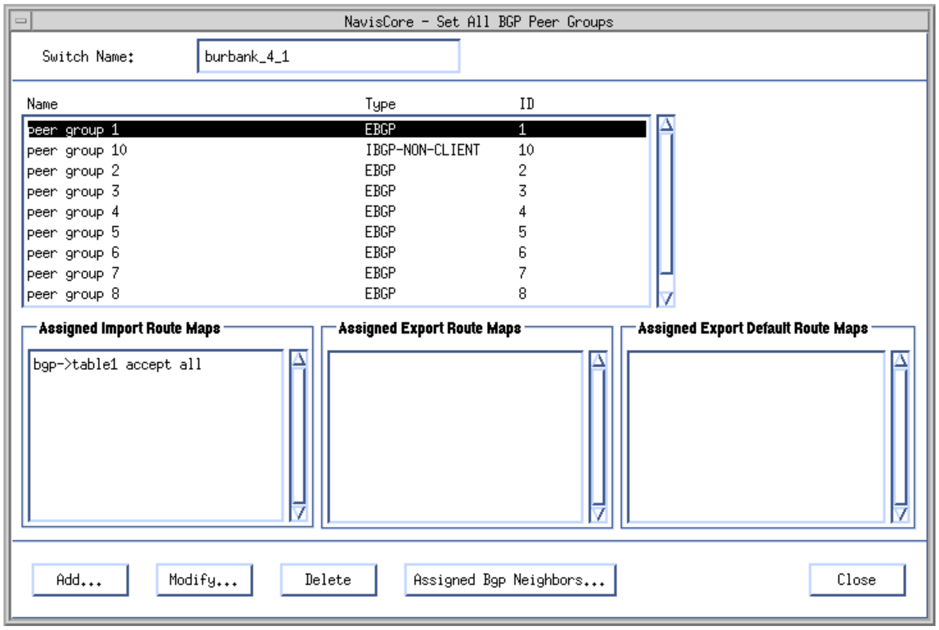

NavisCore - Set All BGP Peer Groups

Switch Name: burbank_4_1

Name	Type	ID
peer group 1	EBGP	1
peer group 10	IBGP-NON-CLIENT	10
peer group 2	EBGP	2
peer group 3	EBGP	3
peer group 4	EBGP	4
peer group 5	EBGP	5
peer group 6	EBGP	6
peer group 7	EBGP	7
peer group 8	EBGP	8

Assigned Import Route Maps: bgp->table1 accept all

Assigned Export Route Maps
Assigned Export Default Route Maps
Add...
Modify...
Delete
Assigned Bgp Neighbors...
Close

Figure 8-11. Set All BGP Peer Groups Dialog Box

The Set All BGP Peer Groups dialog box displays all of the peer groups in the network. You can select a peer group, and then modify its configuration, delete it, or display its assigned BGP neighbors. For a description of the fields on this dialog box, see Table 8-9 on page 8-26. Table 8-7 describes the buttons on the dialog box.

Table 8-7. Set All BGP Peer Groups Dialog Box Buttons

Button	Function
Add	Enables you to add a BGP peer group.
Modify	Enables you to modify the selected BGP peer group.
Delete	Enables you to delete the selected BGP peer group.
Assigned BGP Neighbors	Enables you to display the BGP neighbors assigned to the selected peer group.

3. Choose Add. The Add BGP Peer Group Dialog Box appears (see Figure 8-12).

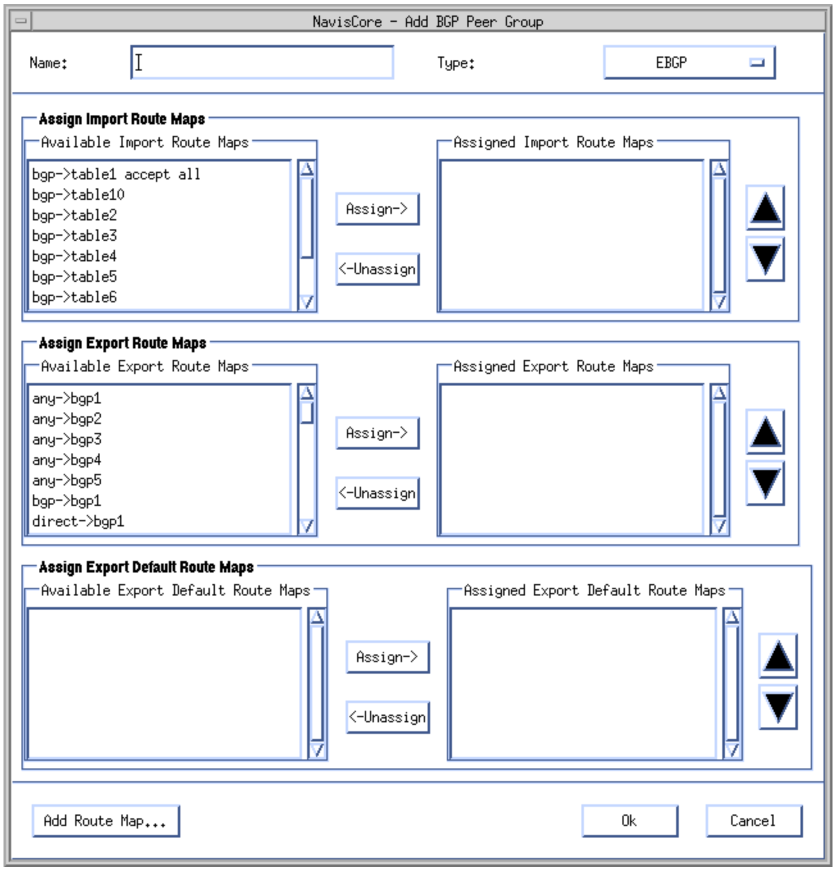

NavisCore - Add BGP Peer Group

Name:

Type: EBGP

Assign Import Route Maps

Available Import Route Maps

bgp->table1 accept all
bgp->table10
bgp->table2
bgp->table3
bgp->table4
bgp->table5
bgp->table6

Assigned Import Route Maps

Assign Export Route Maps

Available Export Route Maps

any->bgp1
any->bgp2
any->bgp3
any->bgp4
any->bgp5
bgp->bgp1
direct->bgp1

Assigned Export Route Maps

Assign Export Default Route Maps

Available Export Default Route Maps

Assigned Export Default Route Maps

Figure 8-12. Add BGP Peer Group Dialog Box

For a description of the fields on this dialog box, see Table 8-9 on page 8-26. Table 8-8 describes the buttons on the dialog box.

Table 8-8. Add BGP Peer Group Dialog Box Buttons

Button	Function
Add Route Map	Adds a route map. See Chapter 11, “Configuring Route Policies” for more information on adding route maps.
OK	Puts your changes into effect.
Cancel	Returns you to the previous dialog box without putting your changes into effect.

4. Complete the fields described in Table 8-9.

Table 8-9. Add BGP Peer Group Fields

Field	Action/Description
Name	Specify a unique alphanumeric name for the peer group.
Type	Specify one of the following BGP neighbor types: *EBGP* – (default) BGP neighbors that are members of this peer group must all be Exterior Border Gateway Protocol (EBGP) peers. *IBGP-CLIENT* – BGP neighbors that are members of this peer group must all be Interior Border Gateway Protocol (IBGP) clients in the same AS. They must also be configured as peers of a route reflector and part of its cluster. *IBGP-NON-CLIENT* – BGP neighbors that are members of this peer group must all be IBGP non-clients in the same AS. They must also be configured as peers of a route reflector, but not part of its cluster.
Assign Import Route Maps	
Available Import Route Maps	The import route maps that are available for assignment to this BGP peer group. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, “Configuring Route Policies.” To display the parameters for any listed route map, double-click on the map.
Assigned Import Route Maps	The import route maps that are assigned to this BGP peer group. All incoming routes on the members of the BGP peer group are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.

Table 8-9. Add BGP Peer Group Fields (Continued)

Field	Action/Description
Assign Export Route Map	
Available Export Route Maps	The export route maps that are available for assignment to this BGP peer group. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, "Configuring Route Policies." To display the parameters for any listed route map, double-click on the map.
Assigned Export Route Maps	The export route maps that are assigned to this BGP peer group. All outgoing routes on the members of this BGP peer group are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.
Assign Export Default Route Maps	
Available Export Default Route Maps	The export default route maps that are available for assignment to this BGP peer group. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, "Configuring Route Policies." To display the parameters for any listed route map, double-click on the map.
Assigned Export Default Route Maps	The export default route maps that are assigned to this BGP peer group. All outgoing routes on the members of this BGP peer group are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.

5. Choose OK. The Set All BGP Peer Groups dialog box appears (see Figure 8-11 on page 8-24).

Assigning BGP Neighbors to Peer Groups

You can assign a BGP neighbor to a peer group when you add or modify the neighbor. To assign a BGP neighbor to a peer group:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All BGP ⇒ Set All BGP Neighbors. The Set All BGP Neighbors dialog box appears (see Figure 8-6 on page 8-14).
3. The action you perform at this point depends on whether you want to add a BGP neighbor or modify an existing neighbor:
   - To add a new BGP neighbor, choose Add. The Add BGP Neighbor dialog box appears (see Figure 8-8 on page 8-16).
   - To modify a BGP neighbor, select a BGP neighbor from the list of neighbors and choose Modify. The Modify BGP Neighbor dialog box appears. This dialog box is similar to the Add BGP Neighbor dialog box (see Figure 8-8 on page 8-16).
4. Select the peer group to which you are assigning the BGP neighbor. NavisCore displays an error message if you select a peer group that is incompatible with the neighbor. For example, if the peer group type is configured as EBGP, but the neighbor you want to assign is an IBGP client, NavisCore will display an error.

   If no peer groups are defined, choose Add Bgp Peer Group to add a peer group. The Add BGP Peer Group dialog box appears (see Figure 8-12 on page 8-25). After you complete the fields on this dialog box, you can return to the Add BGP Neighbor dialog box to assign the neighbor to a peer group. See "Configuring BGP Peer Groups" on page 8-23 for more information on adding peer groups.
5. Complete the other fields on the Add BGP Neighbor dialog box if you are adding a new peer group. See "Configuring BGP Neighbors and Assigning Route Maps" on page 8-14 for more information.
6. Choose OK.

Configuring BGP Route Dampening

You configure IP Navigator to suppress propagation of unstable BGP routes between a destination (D) and an EBGP peer (P). For each route, IP Navigator maintains an *instability metric*, which provides a way to measure the degree of a route's instability. Whenever P deletes or changes its route to D, IP Navigator increments the associated instability metric. In other words, the more frequently a route changes, the route's degree of instability increases and, therefore, the instability metric increases.

The instability metric decays exponentially over time based on a configurable *half-life time*. You configure the decay rates differently when D is reachable or unreachable.

When a route's instability metric exceeds a specified *upper threshold*, IP Navigator suppresses advertisement of the route even if the route is up. A route's instability metric continues to increase even after the route is suppressed. However, the metric will not exceed a configurable *ceiling*, which effectively limits the time it takes to readvertise the route regardless of the route's prior history. IP Navigator reuses the route only when the instability metric goes below the configurable *lower threshold*.

> IP Navigator does not apply route dampening to routes that it learns through IBGP. This restriction prevents the creation of forwarding loops, and prevents IBGP peers from enforcing a higher penalty for routes that are external to the AS than for routes that are internal to the AS.

To configure BGP route dampening:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Ascend IP Parameters ⇒ Set All BGP ⇒ Set All BGP Parameters. The Set BGP dialog box appears (see Figure 8-5 on page 8-10).
3. Choose Route Dampening. The Set BGP Route Dampening dialog box appears (see Figure 8-13).

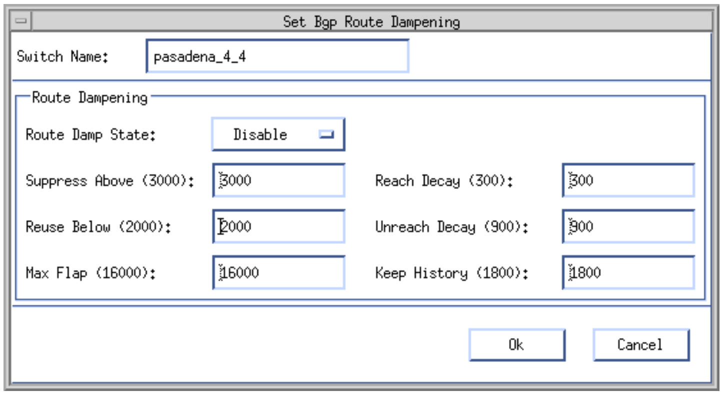

Switch Name: pasadena_4_4

Route Dampening

Route Damp State: Disable

Suppress Above (3000): 3000
Reach Decay (300): 300
Reuse Below (2000): 2000
Unreach Decay (900): 900
Max Flap (16000): 16000
Keep History (1800): 1800

Figure 8-13. Set BGP Route Dampening Dialog Box

4. Complete the fields described in Table 8-10.

Table 8-10. Set BGP Route Dampening Fields

Field	Action/Description
Route Damp State	Specify one of the following BGP route dampening states: *Enable* – Suppress unstable routes based on the criteria you specify in the fields below. *Disable* – (default) Do not suppress unstable routes.
Suppress Above (*Upper Threshold*)	Specify the upper threshold of a route's instability metric that, if exceeded, causes the route to be suppressed. The value can range from 1 to 20000. The default is 3000. For example, if you set this field to 3000 (the default), any route with an instability metric greater than 3000 is suppressed.
Reuse Below (*Lower Threshold*)	Specify the lower threshold of a route's instability metric, below which a suppressed route is re-used, as long as the route is up. The value can range from 300 to 20000. The default is 2000. For example, if you set this field to 2000 (the default), any route with an instability metric less than 2000 is re-used.
Max Flap (*Ceiling*)	Specify the maximum instability metric (that is, the highest instability metric that any route can have). This value imposes a ceiling for the instability metric, determining the longest period of time that a route can be suppressed. The value can range from 1 to 20000. The default is 16000. The Max Flap value must be greater than the Suppress Above value.
Reach Decay (*Half-Life Time for Reachable Routes*)	Specify the number of seconds that must elapse before the instability metric for a route to a reachable destination decreases to half of its current value. The value can range from 300 to 1800. The default value is 300. For example, if you specify a Reach Decay value of 500, then 500 seconds must elapse before the instability metric for a route to a reachable destination decreases to half of its current value.
Unreach Decay (*Half-Life Time for Unreachable Routes*)	Specify the number of seconds that must elapse before the instability metric for a route to an unreachable destination decreases to half of its current value. The value can range from 300 to 1800. The default value is 900. For example, if you specify a Reach Decay value of 1000, then 1000 seconds must elapse before the instability metric for a route to an unreachable destination decreases to half of its current value.
Keep History	Specify the number of seconds that IP Navigator maintains a history of (that is, keeps track of) any route's instability. The value can range from 60 seconds to 1800 seconds. The default is 1800 seconds.

5. Choose OK. The Set BGP Dialog Box (see Figure 8-5 on page 8-10) appears.

Configuring IP Loopback Addresses

The Set IP Loopback Address function enables you to establish an IP loopback address that is not associated with any physical port. Because the loopback address is independent of a physical interface, the status of the physical link does not affect the IP loopback address. If you use an IP loopback address as a BGP neighbor address, you ensure that the BGP connection will not go down.

To set an IP loopback address:

1. From the network map, select the appropriate switch icon.

2. From the Administer menu, choose Lucent IP Parameters ⇒ Set All IP Loopback Addresses. The Set All IP Loopback Addresses dialog box appears (see Figure 8-14).

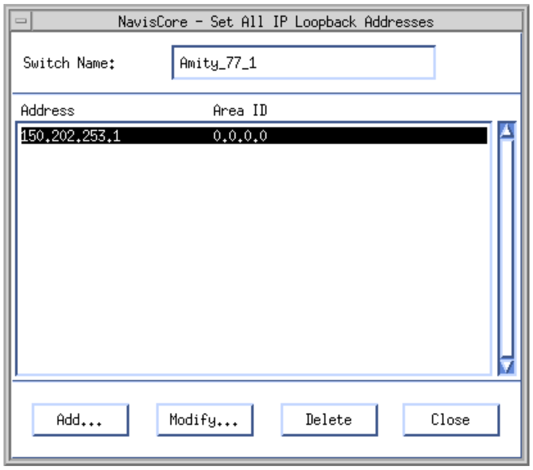

NavisCore - Set All IP Loopback Addresses

Switch Name: Amity_77_1

Address	Area ID
150.202.253.1	0.0.0.0

Figure 8-14. Set All IP Loopback Addresses Dialog Box

3. Choose Add. The Add IP Loopback Address dialog box appears (see Figure 8-15).

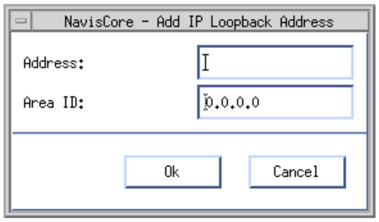

NavisCore - Add IP Loopback Address

Address:

Area ID: 0.0.0.0

Figure 8-15. Add IP Loopback Address Dialog Box

4. Enter the IP address.

5. Enter the Area ID.

6. Choose OK. The Set All IP Loopback Addresses dialog box reappears and the new IP loopback address is included in the list.

9

Configuring IP OSPF and VNN OSPF

Lucent switches can perform routing tasks using IP Navigator and Virtual Network Navigator (VNN) routing software. Both IP Navigator and VNN support the Open Shortest Path First (OSPF) routing protocol. In previous releases of Lucent switch software, VNN and IP Navigator shared the same OSPF software component to perform their OSPF routing tasks. However, in this release, VNN and IP Navigator have their own OSPF software components, or *instances*. These instances are called *VNN OSPF* and *IP OSPF*.

Because VNN and IP Navigator now have their own OSPF instances, you can configure Lucent switches so that VNN and IP Navigator have separate OSPF views of the network. For example, you can configure separate areas for VNN OSPF and IP OSPF, and you can configure the network so that some trunks are hidden from IP OSPF but visible to VNN OSPF.

► Although IP Navigator is not supported on GX 550 switches, GX 550 switches that run switch software version 01.05.xx.xx support separate instances of OSPF for IP and VNN. With the exception of the information that applies to configuring IP OSPF on IP logical ports, the information in this chapter applies to GX 550 switches. IP logical ports are not supported on GX 550 switches.

This chapter discusses:

- An overview of OSPF, IP Navigator, and VNN
- Planning separate IP OSPF and VNN OSPF network views
- Configuring IP OSPF area IDs and VNN OSPF area IDs for trunks
- Configuring IP OSPF
- Configuring VNN OSPF
- Configuring multiple IP OSPF areas and VNN OSPF areas
- Configuring route maps using IP OSPF and VNN OSPF as route map sources and destinations.

About OSPF

OSPF is a link-state routing protocol that works within an Autonomous System (AS). Within an AS, each OSPF router maintains an identical link-state database that describes the network topology. This link-state database allows a router to calculate a routing table by constructing a shortest-path tree. Both IP Navigator and VNN support the OSPF Version 2 protocol, as described in RFC 2328.

The OSPF protocol has the following advantages over the Routing Information Protocol (RIP):

Authentication — Provides security. Only an authorized router can generate route updates to other routers.

Type of Service (TOS) — Enables your network to make routing decisions based on the quality of service required by a host application.

Areas — Restricts flooding to configured areas, thereby reducing the database size.

OSPF Link-State Database

OSPF uses a link-state database to describe the networks and gateways within an Autonomous System. OSPF uses a *flooding mechanism* to distribute and synchronize the link-state database between routers. When a network’s topology changes, each router receives one or more link-state advertisements (LSAs), which contain topology information. Each router stores LSAs in its link-state database.

Link-state databases include:

- Known router addresses
- Known links and their associated costs
- Known network addresses

Routers use the link-state database and Dijkstra’s algorithm (algorithm used to calculate best routes) to determine the best route.

Designated Routers and OSPF Relationships

Designated Routers are responsible for sending copies of the link-state database to routers in the network. When new routers send Hello packets to the Designated Router, the Designated Router responds with an acknowledgment message. The new router then sends a database description packet requesting a copy of the link-state database. The Designated Router responds by sending a database description packet that contains a copy of the link-state database to the new router and updates other nodes in the network about the new router/switch.

In addition, Designated Routers:

- Monitor the health of adjacent routers
- Establish adjacencies with non-designated routers

A Backup Designated Router is typically defined in case the Designated Router goes down. The Backup Designated Router keeps track of the same information as the Designated Router, but keeps silent. If the backup detects a failure of the Designated Router, it immediately becomes active.

OSPF Flooding Controls

Flooding is a reliable way to send link-state advertisements (LSAs) to ensure that all routers in an area have identical link-state databases. Multiple copies of the message travel through the network, ensuring that one message will arrive safely at each node. However, flooding causes significant network traffic. To reduce network traffic, OSPF implements the following flooding controls:

- The Designated Router is the only router that can generate link-state updates. This control reduces the number of copies created.
- Before forwarding OSPF link-state updates, the Designated Router checks its own link-state database to see if the update was received. If it was, the copy is discarded.
- OSPF supports areas where flooding is restricted. Smaller areas mean fewer copies of a message and less traffic.
- Each node acknowledges each LSA update that it receives from the Designated Router.

OSPF Areas

Link-state database size increases in proportion to network size. This can limit the scalability of the network for the following reasons:

- Increased memory space is consumed
- Route table generation becomes more processor-intensive
- It takes longer to:
  - Calculate link costs for more links
  - Calculate the spanning tree for a large network
  - Generate large routing tables required by large networks

To address large link-state databases, OSPF uses *areas*. An area is a group of OSPF routers that exchanges topology information. Designated Routers only send LSAs to routers that are part of the same area. If an Autonomous System (AS) has one area, all routers in the Autonomous System receive LSAs. However, if the Autonomous System is divided into many areas, LSAs only go to the appropriate areas, thereby minimizing traffic and the link-state database size. The Autonomous System works like a collection of smaller networks.

Because of flooding controls, the topology of one area is unknown by routers in another area. This means a router knows nothing of network topology outside its own area. Each area has a unique link-state database, and all routers in the given area should have the same database.

If an Autonomous System consists of more than one area, a backbone area (Area 0) must be created for OSPF to distribute consistent routing information throughout the Autonomous System. The backbone area allows information from one area to be advertised in other areas.

Figure 9-1 illustrates the concept of areas.

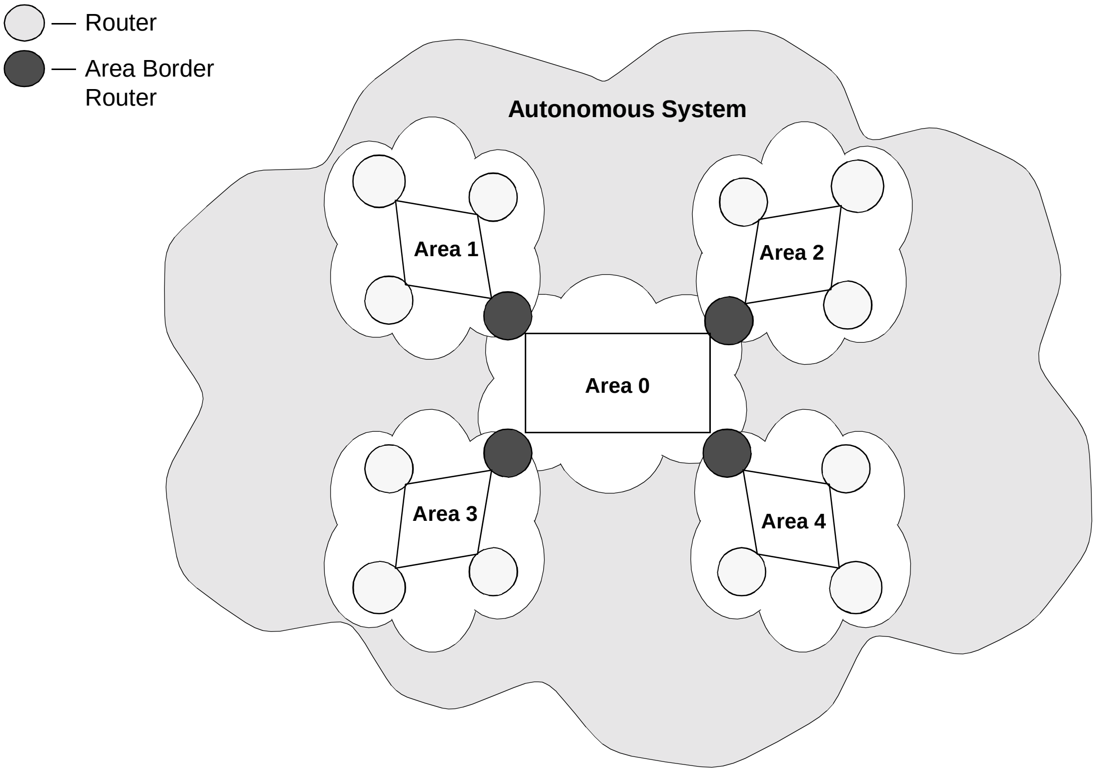

Figure 9-1. OSPF Areas

In Figure 9-1, the Autonomous System is divided into five areas. Each area represents smaller networks within the Autonomous System, and maintains separate link-state databases. Area 0 is the backbone, connecting all areas within the Autonomous System.

OSPF requires the backbone to be contiguous to all areas in the Autonomous System. All non-backbone areas (that is, areas other than area 0) must connect to the backbone area directly or through a virtual link.

Packets are forwarded between areas, from the source area, through the backbone, and then into the destination area. Lucent’s OSPF area implementation assigns each trunk to a specific area. This provides maximum flexibility in setting area boundaries and changing area boundaries in the future.

OSPF Routing and Router Classifications

There are three types of routing:

- Intra-area routing — Routing within an area
- Inter-area routing — Routing between areas
- AS-external routing — Routing between the OSPF Autonomous System and other Autonomous Systems

These types of routing are performed by different classifications of routers, including:

Internal routers — Routers that are directly connected and belong to the same area. In addition, routers with interfaces connected only to the backbone are classified as internal routers.

Area border routers (ABRs) — Routers with links to more than one area. ABRs maintain separate link-state databases of each area to which they belong.

Backbone routers — Routers with an interface to the backbone. A backbone is either an area border router or an internal router.

AS boundary routers — Routers that connect an OSPF Autonomous System to a region that uses a different routing protocol. AS boundary routers may be internal routers, area border routers, or backbone routers.

Figure 9-2 shows an example of OSPF routing and router classifications.

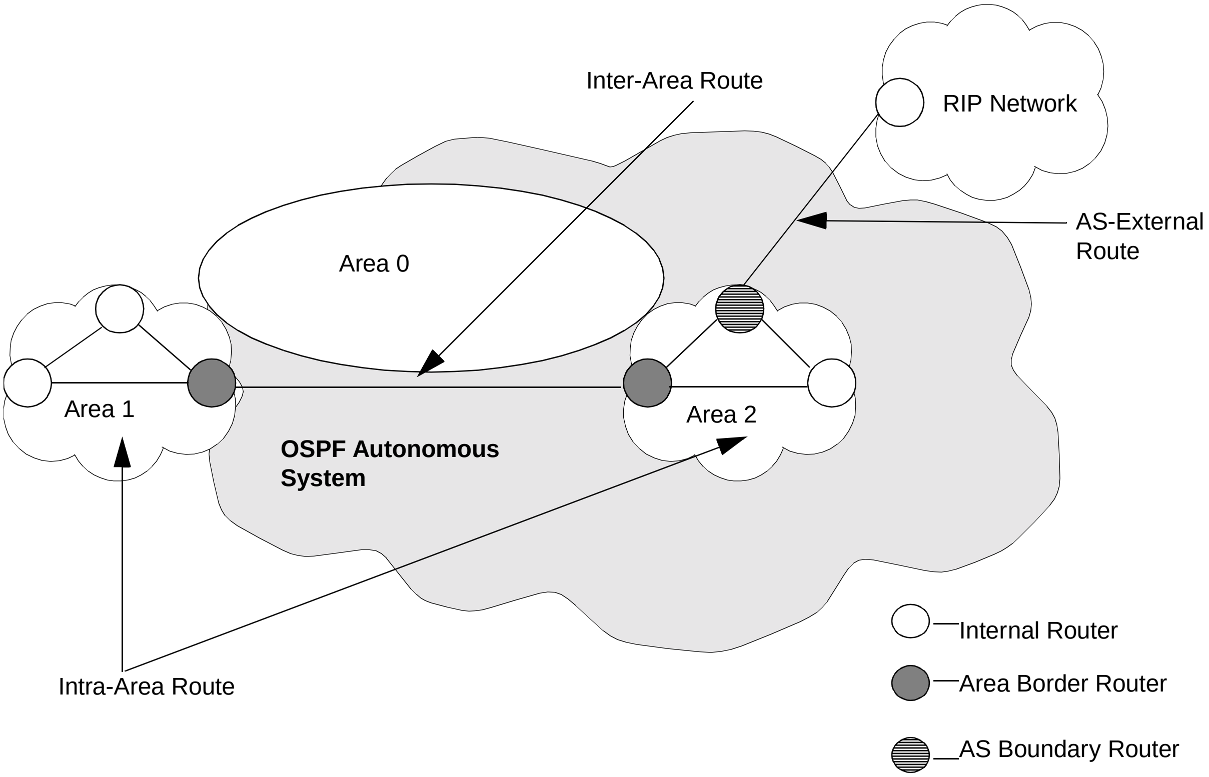

Figure 9-2. Router Classifications

OSPF Area Aggregates

Area aggregates consolidate multiple routes (or IP addresses) within an area (or areas) into one single LSA. This consolidation enables one advertisement representing a range of IP addresses within an area (or areas) to be broadcast.

Area aggregates:

- Reduce the size of the OSPF routing table
- Provide better control over network instabilities
- Provide a better mechanism to summarize route updates across areas
- Reduce memory requirements for link-state databases
- Reduce the cost of route calculation
- Have a maximum area size of 400 switches and routers, or 1,000 interfaces

Although address aggregation is not required, it results in fewer summary LSAs. Fewer summary LSAs reduce the size of the routing table and OSPF link-state databases. See "OSPF Summary LSAs" for more information on summary LSAs.

OSPF Summary LSAs

IP addressing information is advertised across area boundaries in OSPF summary LSAs. Each summary LSA advertises a single range of IP addresses. The IP address ranges are configured in the ABRs. See "OSPF Routing and Router Classifications" on page 9-6 for more information on ABRs.

For example, Area 1 is assigned a subnet of 107.109.220.0/*24*. The number 24 specifies a subnet mask of 24, so all IP addresses in the range 107.109.220.1-254 are sent as a single OSPF summary LSA.

OSPF Virtual Links

OSPF requires that all OSPF areas be connected to the OSPF backbone area. In addition to direct connections, you can use virtual links to logically connect physically separate portions of a network to the backbone. OSPF uses virtual links for the following purposes:

- To connect areas that are not physically connected to the backbone.
- To preserve the continuity of the OSPF backbone.

The two endpoints of a virtual link are ABRs. Figure 9-3 illustrates a network that connects Area 0.0.0.4 to the backbone through Area 0.0.0.3 by using a virtual link from Switch 7 to Switch 4.

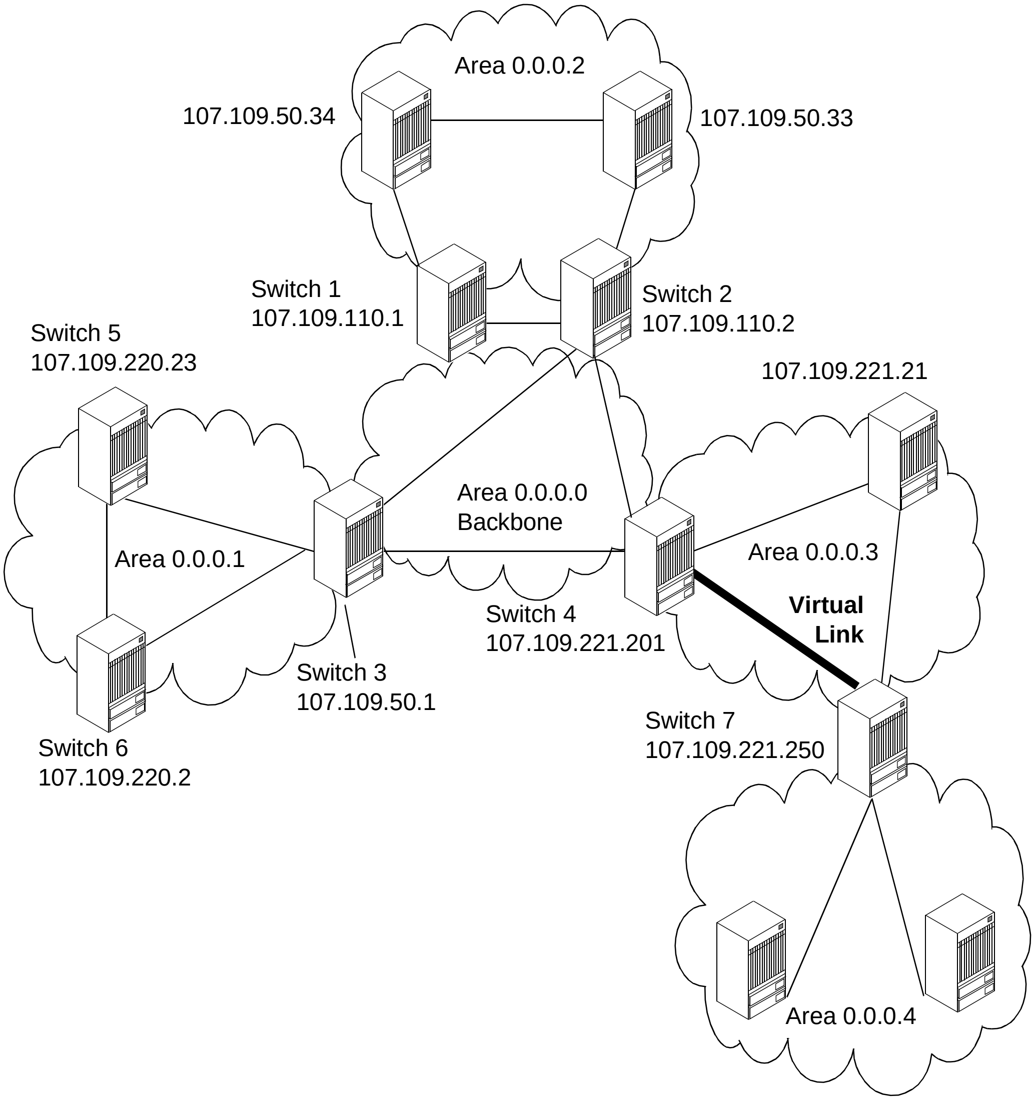

Figure 9-3. OSPF Area Configuration Example

Figure 9-3 shows a single OSPF backbone area with an Area ID of 0.0.0.0. All three non-backbone areas (0.0.0.1, 0.0.0.2, and 0.0.0.3) are directly connected to the OSPF 0.0.0.0 backbone area. Area 0.0.0.4 is connected to the backbone through the virtual link that is configured in Area 0.0.0.3.

About Clustering

Clustering is a way of grouping VNN OSPF areas into subareas (the IP instance of OSPF does not support clustering). Clustering enables you to use set increments (a set of three bits of the internal IP address to assign a cluster address between 000 and 111, or 0 and 7) of the host ID address in different VNN OSPF areas, while performing route aggregation at the Area Border Router. A cluster forms a subset of an OSPF area. A cluster enables additional address aggregation at the ABR and reduces the size of the IP routing table, link-state database, and the number of summary LSAs.

Use clustering only if you deploy new nodes with the same subnet addresses into multiple VNN OSPF areas (for example, due to a lack of IP addresses).

▶ Do not configure clustering if you implement OSPF areas using B-STDX switch software versions prior to version 5.0.

Version 2.3 and later NMS versions allow you to define an IP address subnet as part of a cluster, define a cluster ID, and designate a switch as part of a cluster at switch deployment.

You assign a cluster ID to the IP address to be clustered. The cluster ID specifies the upper three bits of the host ID. As you add switches in that cluster ID, the switch number/host ID in the IP address increments according to the cluster ID. Table 9-1 shows the cluster ID IP-address range using 107.109.50.x as the default IP address.

Table 9-1. Cluster ID and IP Addresses Example

Cluster ID	IP Address Range
0	107.109.50.1 - 107.109.50.30
1	107.109.50.33 - 107.109.50.62
2	107.109.50.65 - 107.109.50.94
3	107.109.50.97 - 107.109.50.126
4	107.109.50.129 - 107.109.50.158
5	107.109.50.161 - 107.109.50.190
6	107.109.50.193 - 107.109.50.222
7	107.109.50.225 - 107.109.50.254

VNN and IP Navigator Overview

Before you can start to configure separate network views for VNN OSPF and IP OSPF, it is helpful to review how each routing method works separately, and understand how both routing methods work together.

VNN

VNN performs the following functions:

- Routes ATM and Frame Relay PVC and SVC traffic to other Lucent switches through the Lucent network. To route ATM and Frame Relay between Lucent switches, VNN uses OSPF with some proprietary extensions.
- Routes SMDS datagrams.
- Routes switch management traffic (e.g., SNMP) through the Lucent network. Management connections such as management DLCIs and management VPI/VCIs are handled by VNN.
- Sets up Multipoint-to-Point Tunnel (MPT) label switched paths (LSPs) for use by IP Navigator and selects the paths that they use. VNN determines when to add new leaves to each MPT LSP and the path that each MPT LSP uses.
- Distributes topological data and other data from the switch processor (CP or SP) to the Input/Output Processors (IOPs) so that the IOPs can perform VC calculations.

IP Navigator

IP Navigator performs the following functions:

- Determines whether IP packets are routed at Layer 3 (IP) or switched at Layer 2 (using MPT LSPs). When the IP path is through an IP logical port, IP Navigator routes the packet. When the IP path traverses trunks, IP Navigator forwards the packet via an MPT LSP (which is set up by VNN).
- Maintains the IP routing table using IP OSPF, RIP, BGP, and static routing. Note that IP OSPF is based on the OSPF standard while VNN OSPF implements some proprietary extensions to standard OSPF.
- Distributes IP routing information between other routing protocols. For BGP, IP Navigator distributes information on IBGP next hops.

How VNN OSPF and IP OSPF View the Network

Separate instances of OSPF for VNN and IP Navigator have the following benefits:

Create a simplified topological view for IP routers from other vendors — IP routers from other vendors cannot support OSPF areas as large as the areas that Lucent switches can support. By supporting separate topological views for VNN OSPF and IP OSPF, you can configure large areas for VNN OSPF, thereby optimizing virtual circuit paths, and create smaller OSPF areas for IP OSPF to accommodate IP routers from other vendors.

Hide management addresses from IP routers — By configuring different network views for VNN OSPF and IP OSPF, you can prevent the management IP addresses of switches from being advertised to IP routers. Unless you explicitly allow the management addresses to be advertised to IP routers (through a route map), management addresses are advertised by VNN OSPF only.

Separate OSPF views provide you with a flexible way to configure the network. For example, you can divide the network into three areas for IP OSPF, but have only one area for VNN OSPF. This approach would have the following benefits:

- Allow VNN to use optimal paths for virtual circuits.
- Reduce the size of IP areas, thereby reducing the size of the OSPF databases on IP routers that forward packets through the Lucent network.
- Hide management addresses from the IP routers.

When you initially upgrade a Lucent switch to this release of IP Navigator, VNN OSPF and IP OSPF have identical OSPF network views. You have the option to continue with this configuration, or create different views for VNN OSPF and IP OSPF.

Figure 9-4 shows a Lucent network in which all switches are running a release of Lucent switch software prior to this release. The Lucent network is a single-area AS. Notice that, on Switch1, VNN and IP Navigator have a common OSPF network view of the Lucent network through a common OSPF instance.

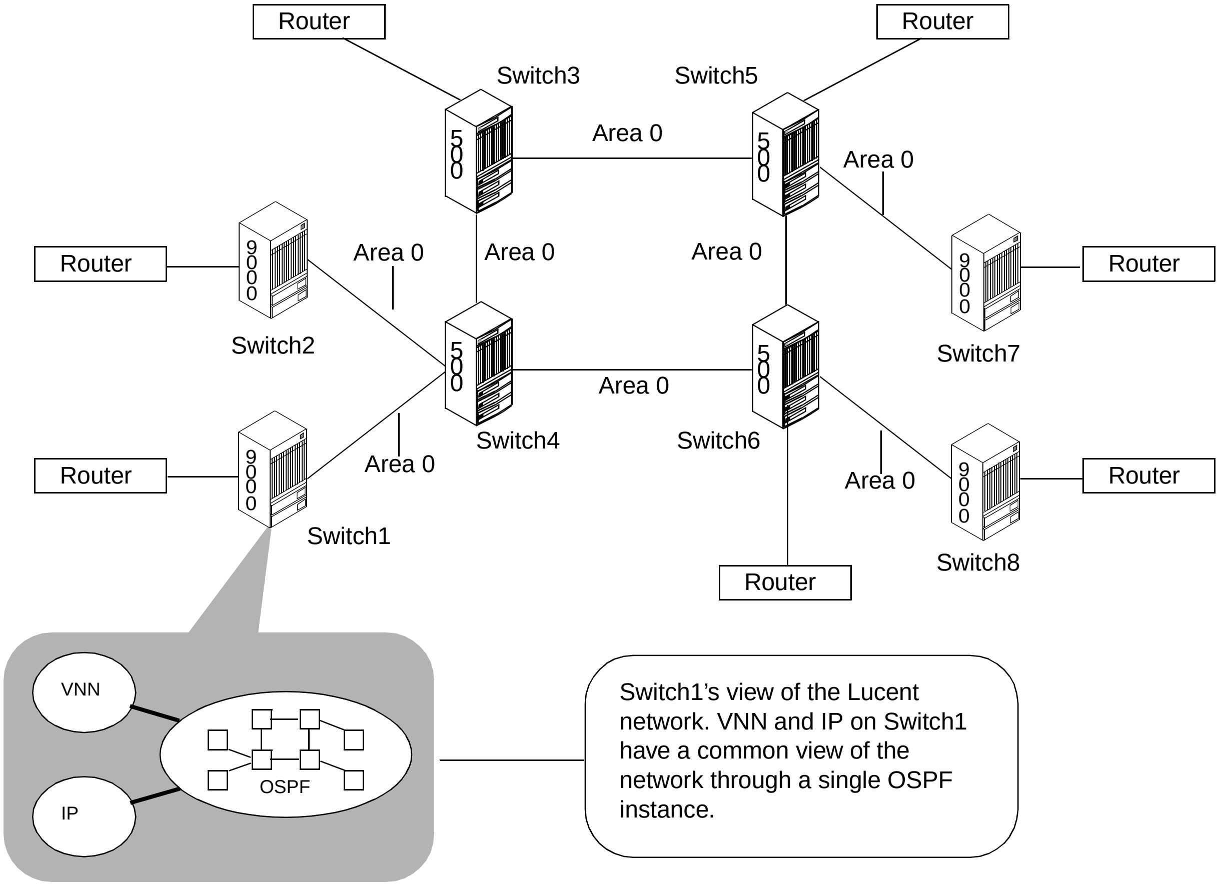

Switch1's view of the Lucent network. VNN and IP on Switch1 have a common view of the network through a single OSPF instance.

Figure 9-4. Common OSPF Network View Through a Single OSPF Instance

In Figure 9-5, all the switches in the network are upgraded to this release of Lucent switch software, creating separate instances of OSPF (VNN OSPF and IP OSPF). In both the VNN OSPF and IP OSPF network views, the Lucent network remains a single-area AS. However, in the IP OSPF view in Figure 9-5, the network manager has hidden part of the Lucent network (the trunks indicated by dashed lines) from the router connected to Switch1. The hidden trunks are part of the VNN OSPF view only.

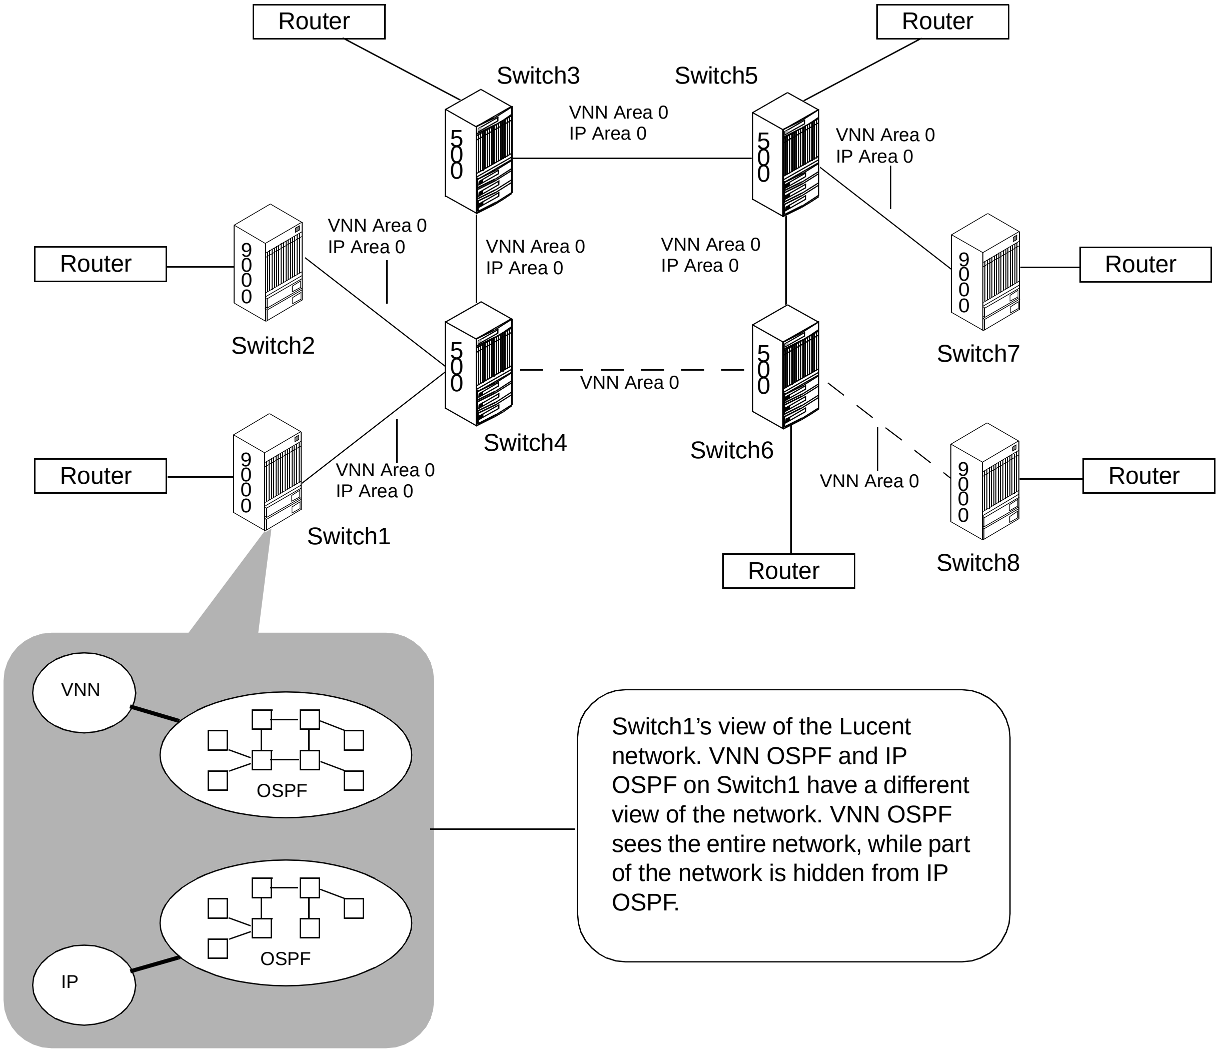

Switch1's view of the Lucent network. VNN OSPF and IP OSPF on Switch1 have a different view of the network. VNN OSPF sees the entire network, while part of the network is hidden from IP OSPF.

Figure 9-5. Different OSPF Network Views

If required, the network manager could have configured multiple areas for either VNN OSPF, IP OSPF, or both. Figure 9-6 shows a Lucent network in which VNN OSPF has only one area, but IP OSPF has three areas.

► You can configure an IP OSPF router ID on any switch connected to a trunk that is a member of an IP OSPF area. See "Planning Router IDs" on page 9-22 for more information.

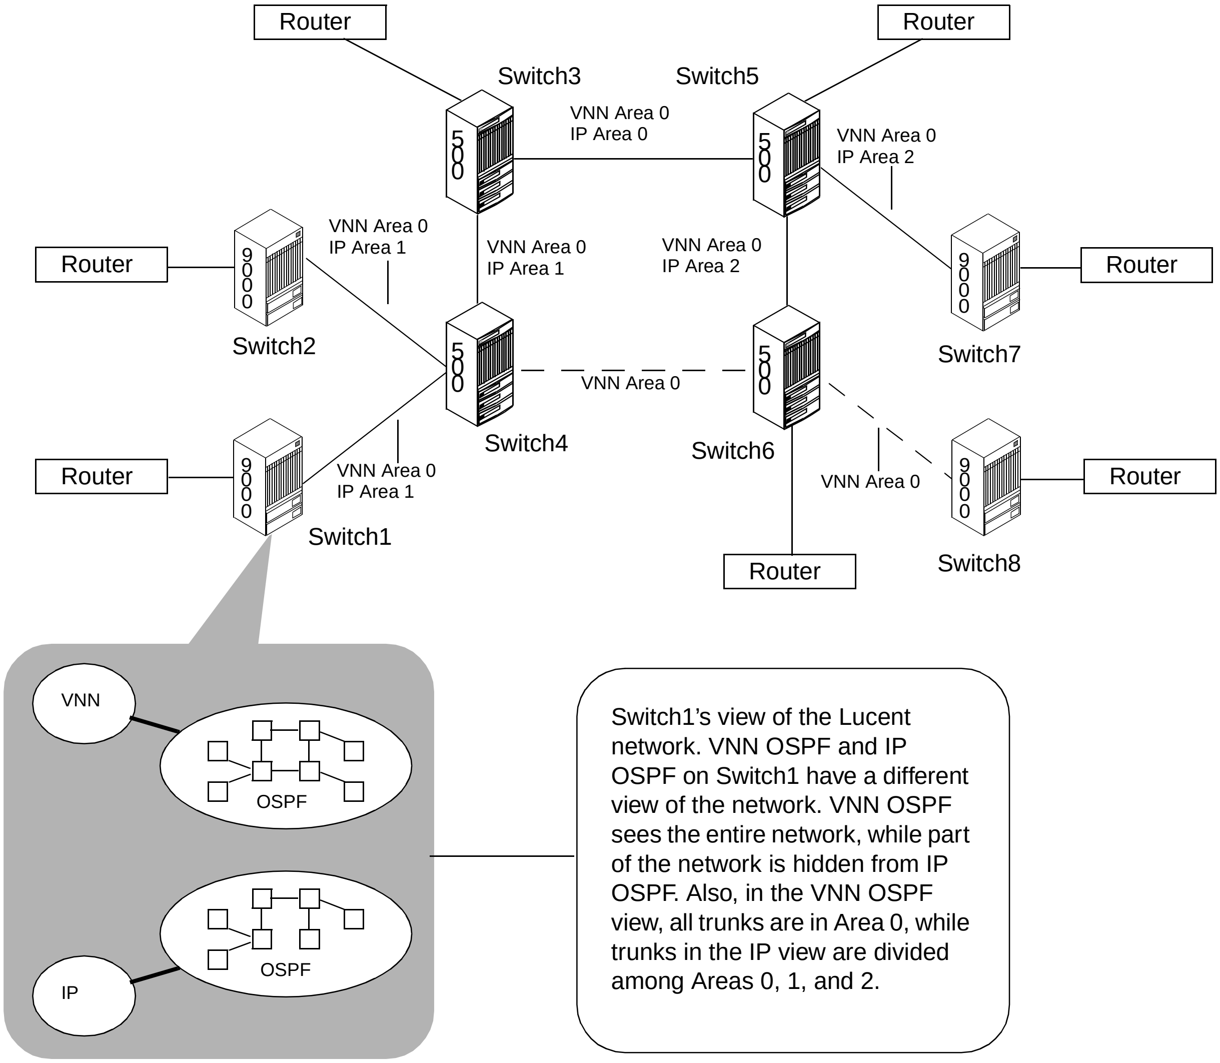

Switch1's view of the Lucent network. VNN OSPF and IP OSPF on Switch1 have a different view of the network. VNN OSPF sees the entire network, while part of the network is hidden from IP OSPF. Also, in the VNN OSPF view, all trunks are in Area 0, while trunks in the IP view are divided among Areas 0, 1, and 2.

Figure 9-6. Different Area Configurations for VNN OSPF and IP OSPF

Planning Separate OSPF Network Views

Before you start to configure separate network views for VNN OSPF and IP OSPF, perform the following tasks:

- Consider network size and types of network traffic
- Plan your trunk configuration
- Plan for configuring VNN OSPF loopbacks, area aggregates, and virtual links
- Plan for configuring router IDs for IP OSPF
- Plan route maps

Considering Network Size and Types of Network Traffic

To develop an overall strategy for configuring separate network views for VNN OSPF and IP OSPF, you need to take into account:

- Types of network traffic
- Network size
- Topology

IP Navigator handles IP traffic while VNN handles non-IP traffic, which includes management traffic and ATM and Frame Relay PVC and SVC traffic.

If your network is large and handles a mix of IP, management, ATM, and Frame Relay traffic, consider configuring one or two large areas for VNN OSPF and several smaller areas for IP OSPF. The benefits of this approach are:

- Large OSPF areas enable VNN to create optimal paths for ATM and Frame Relay virtual circuits. In addition, you can include management trunks in the VNN areas only, hiding them from third-party IP routers outside the Lucent network.
- Small OSPF areas enable IP Navigator to accommodate third-party router limitations.

Planning Trunk Configuration

Your trunk configuration determines:

- VNN OSPF areas
- IP OSPF areas

All trunks are in VNN OSPF areas. In order to add a trunk, you must configure a VNN OSPF area ID for the trunk.

Optionally, you can enable IP and configure an IP OSPF area ID for a trunk so that a trunk can belong to:

- Both a VNN OSPF area and an IP OSPF area
- A VNN OSPF area only

A trunk cannot belong to an IP OSPF area only.

Figure 9-7 shows a Lucent network in which some trunks belong to both VNN OSPF areas and IP OSPF areas (trunks 1-6), and other trunks belong to VNN OSPF areas only (trunks 7 and 8). Note that trunks 7 and 8 (shown by the dashed lines) are best suited for carrying management traffic because they are the most secure (that is, they cannot be seen by third-party IP routers outside the Lucent network).

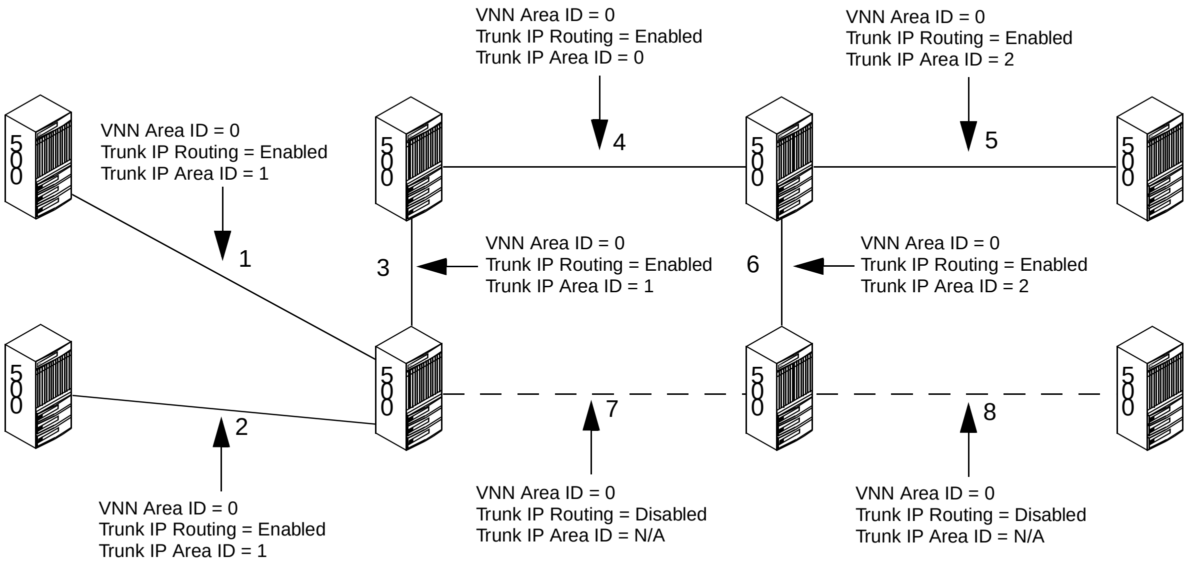

Figure 9-7. Trunk VNN OSPF and IP OSPF Area Membership

Configuration parameters on the Add Trunk and Modify Trunk dialog boxes allow you to enable IP routing and configure VNN OSPF and IP OSPF area IDs on trunks. When you enable IP routing for a trunk, NavisCore automatically configures the IP OSPF interface for each trunk.

NavisCore uses the following trunk configuration parameters to configure the IP OSPF interface for each trunk endpoint:

Trunk IP Area ID — NavisCore sets the OSPF Area ID to the value of the Trunk IP Area ID field.

TOS Zero Metric (End 1) — NavisCore sets the TOS 0 metric of the first trunk endpoint's IP OSPF interface to the value in this field.

TOS Zero Metric (End 2) — NavisCore sets the TOS 0 metric of the second trunk endpoint's IP OSPF interface to the value in this field.

For more information on configuring trunks for separate OSPF network views, see "Configuring Trunks for IP OSPF and VNN OSPF" on page 9-25.

Planning VNN OSPF and IP OSPF Loopback Addresses, Area Aggregates, and Virtual Links

▶ This section applies only if you configure multiple VNN OSPF areas or IP OSPF areas.

Because VNN OSPF and IP OSPF have separate network views, you must configure the following parameters separately if the network supports multiple IP OSPF areas or multiple VNN OSPF areas:

Loopback Addresses — Loopback addresses act as endpoint addresses for virtual links. A loopback address is not associated with any physical port.

Area Aggregates — Aggregation is the process of advertising a range of IP addresses by advertising only a single IP address prefix, which is specified by an IP address/IP network mask pair. Addresses can be aggregated at area borders. This practice improves route scaling by reducing the size of the topology database and the routing table. You configure aggregates in the Area Border Router (ABR).

Virtual Links — OSPF requires that all areas be attached to the backbone area. However, areas do not have to be physically attached to the backbone. Instead, you can configure virtual links to logically attach an area to the backbone. Before you can configure a virtual link you must know:

- The non-backbone transit area for the virtual link.
- The router IDs for the two endpoint switches. For VNN OSPF, the router ID is the same value as the switch ID. IP OSPF has its own router ID. The endpoint switches are the ABRs that border on the transit area.

In addition, you must ensure that both switches have OSPF interfaces in the transit area.

See "Configuring Multiple IP OSPF and VNN OSPF Areas" on page 9-54 for more information on configuration considerations for multiple OSPF areas. For more conceptual information on loopback addresses, area aggregates, and virtual links, see "About OSPF" on page 9-2.

Figure 9-8 and Figure 9-9 show the same physical network, but from different OSPF perspectives. Figure 9-8 shows the VNN OSPF view of the network and Figure 9-9 shows the IP OSPF view of the network. In Figure 9-8, two VNN OSPF virtual links are configured, but in Figure 9-9 only one IP OSPF virtual link is configured because IP OSPF has a different view of the network.

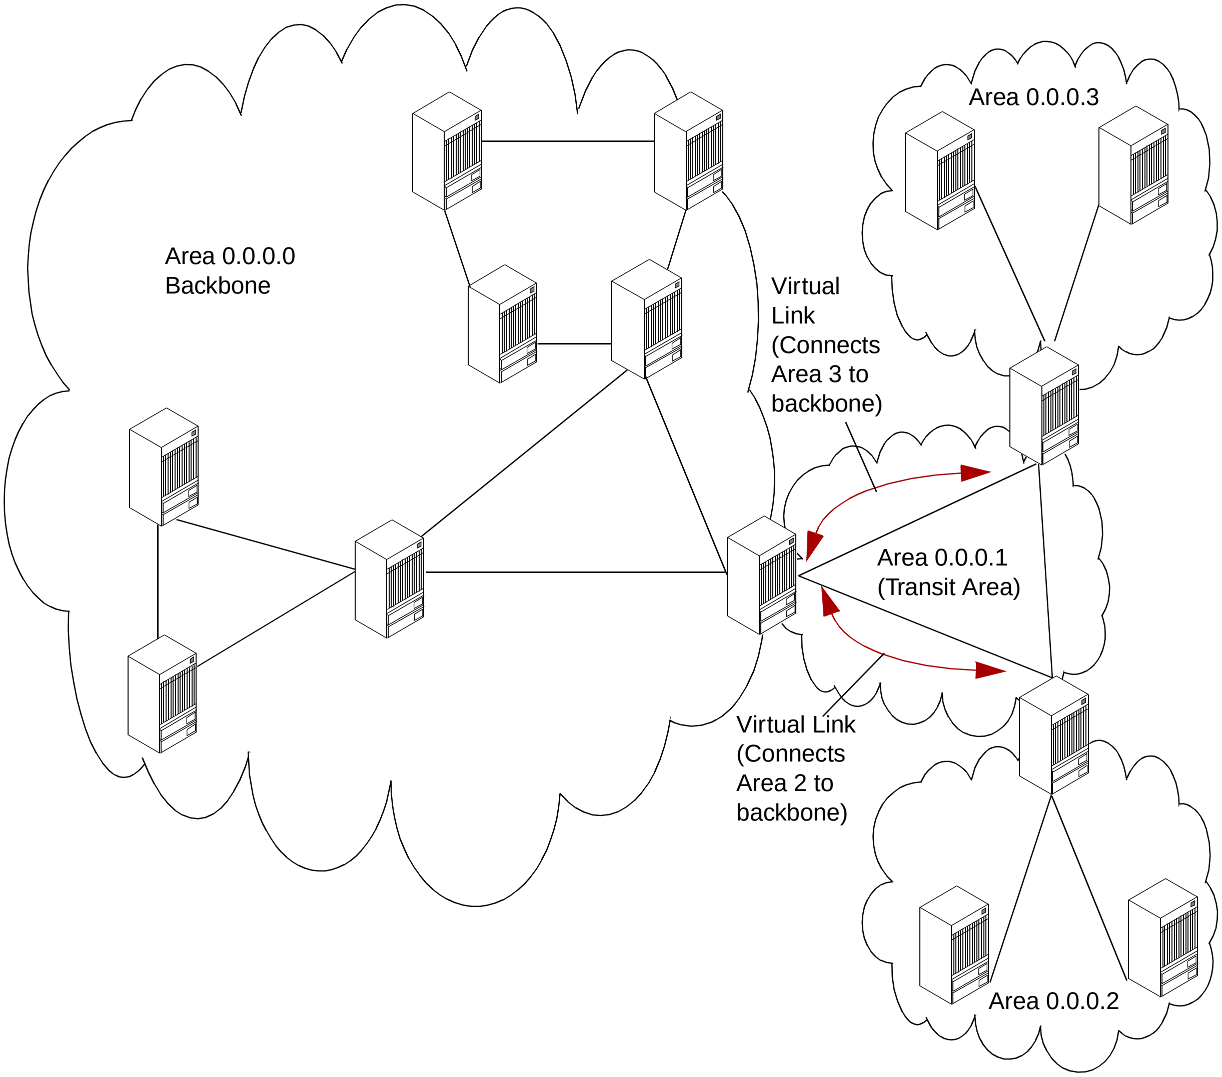

Figure 9-8. Virtual Links for VNN OSPF Network View

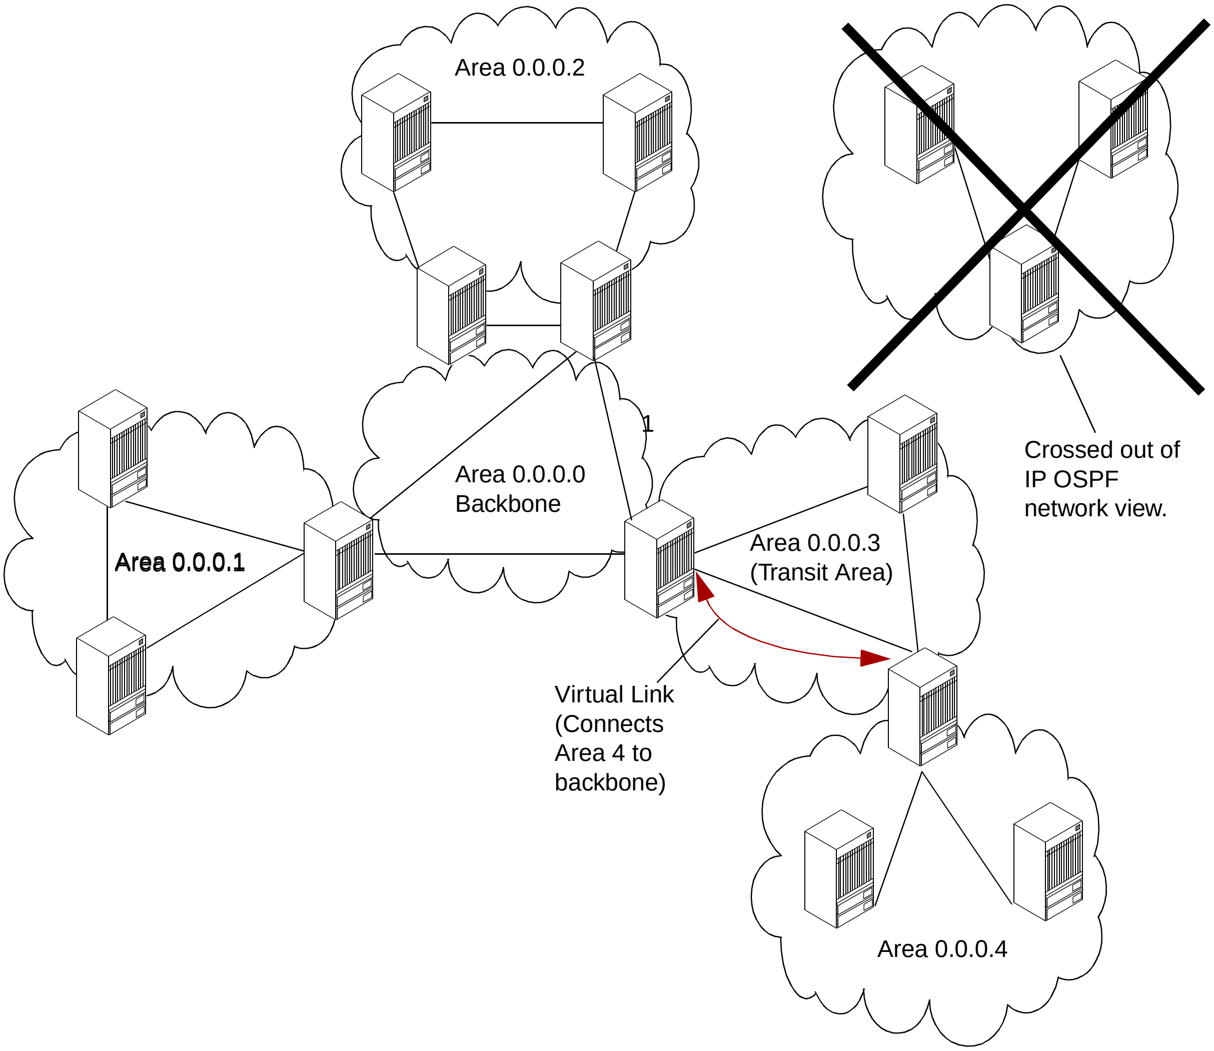

Figure 9-9. Virtual Link for IP OSPF Network View

Planning Router IDs

Switches have two OSPF router IDs: one for VNN OSPF and one for IP OSPF. Whereas VNN OSPF uses the internal switch ID as the router ID, IP OSPF has its own router ID.

Switch IDs are configured when the switch is installed. In prior releases of IP Navigator, both VNN and IP Navigator used the switch ID as the router ID. Once you upgrade to this release of IP Navigator, only VNN OSPF continues to use the switch ID as the router ID.

▶ You do not have to configure an IP OSPF router ID on switches that have IP logical ports or IP Server logical ports (and associated IP interfaces and IP OSPF interfaces) configured. These switches automatically choose the IP address of one of their IP interfaces to be the IP OSPF router ID. If a switch has multiple IP interfaces, the switch chooses the highest IP address to be the IP OSPF router ID.

The automatically chosen IP OSPF router ID applies to the public IP VPN only. You must have an IP OSPF router ID for the public IP VPN in order for IP traffic to be forwarded properly. If necessary, you may manually override the automatically chosen IP OSPF router ID by configuring an IP OSPF router ID of your own choosing for the public IP VPN. You can configure IP OSPF router IDs for private IP VPNs, but this is not required. See Chapter 16, “Configuring IP Virtual Private Networks” for more information on IP VPNs.

For an IP OSPF router ID to be selected automatically, it is not sufficient to configure just IP logical ports (or IP Server logical ports) and IP interfaces. You must also configure at least one IP OSPF interface.

You must manually configure an IP OSPF router ID on each switch (within the context of the public IP VPN) that meets the following conditions:

- No IP logical port or IP Server logical port (and associated IP interfaces) configured on the switch
- Forwards IP traffic on IP trunk interfaces (that is, trunks assigned IP OSPF area IDs)

Switches that meet these conditions cannot automatically choose an IP OSPF router ID. You must configure one manually.

For example, in Figure 9-10, all switches, with one exception, have an IP logical port configured. Each of these switches does not require you to configure an IP OSPF router ID. The single exception is Switch4, which requires an IP OSPF router ID because it is connected to IP OSPF trunks (and therefore has to route IP traffic), but does not have an IP logical port or IP Server logical port. You would configure an IP OSPF router ID for Switch4 in the context of the public IP VPN.

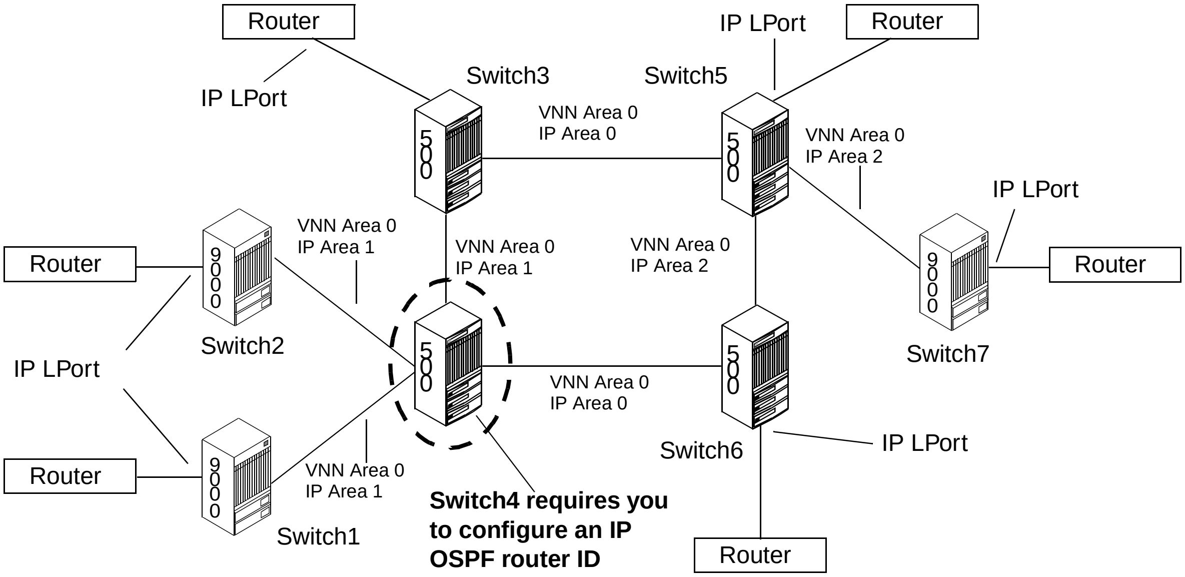

Figure 9-10. IP OSPF Router ID Requirements

For more information on configuring the router ID for IP OSPF, see "Configuring IP OSPF Router IDs" on page 9-42.

Planning Route Maps

Although VNN OSPF and IP OSPF have separate OSPF databases, they share a common IP routing table. VNN OSPF is the source of management routes (e.g., routes associated with switch internal IP addresses and routes to the NMS). IP routing protocols (such as IP OSPF) are the source of IP user routes.

You can configure route maps to:

- Leak routes learned through VNN OSPF into IP Navigator. For example. if you want third-party routers to use management routes, you can leak these routes into IP OSPF.
- Leak routes learned through IP OSPF into VNN OSPF.

Switches perform some automatic route leaking for you. For example, when some switches are upgraded to support separate network views for VNN OSPF and IP OSPF, and other switches are not upgraded, routes are automatically leaked as follows:

- Routes learned by IP OSPF are leaked into VNN OSPF.
- Non-management routes learned by VNN OSPF are leaked into IP OSPF. Management routes remain hidden from IP OSPF.

Only the switch with the highest router ID is responsible for leaking routes.

Important Upgrade Considerations

This section lists some important considerations you should take into account before you upgrade a switch to this release of IP Navigator.

IP OSPF Router IDs

After you upgrade any switch to this release, make sure you configure a router ID for IP OSPF on the switch if the switch meets the following criteria:

- The switch does not have an IP logical port or IP Server logical port (and associated IP interfaces) configured
- The switch is required to route IP traffic over IP trunk interfaces

See “Planning Router IDs” on page 9-22 and “Configuring IP OSPF Router IDs” on page 9-42 for more information.

Virtual Links

When you upgrade a switch to this release, virtual links that you configured in previous releases become VNN OSPF virtual links. They do not become IP OSPF virtual links. If you want IP OSPF virtual links, you must configure them.

Performance

As you upgrade switches to this release of IP Navigator, network routing may not be optimal. Optimal routing performance returns as soon as the whole network is upgraded or a routing policy is configured. See “Planning Route Maps” on page 9-23 for more information on this process.

Viewing Routing Information

Following upgrade, VNN OSPF routing information is labeled “VNN” in the routing table. IP OSPF routing information is labeled “OSPF.”

Configuring Trunks for IP OSPF and VNN OSPF

When you add or modify a trunk, you can enable or disable IP routing for the trunk and assign the trunk to different areas (one for VNN OSPF and one for IP OSPF).

To add a new trunk or modify an existing trunk to support different network views for VNN OSPF and IP OSPF:

1. Follow the instructions for adding or modifying trunks in the *NavisCore ATM Configuration Guide* or the *NavisCore Frame Relay Configuration Guide*.
2. At the Add Trunk dialog box (see Figure 9-11) or Modify Trunk dialog box (which is similar to the Add Trunk dialog box), perform the following actions to configure separate network views for VNN OSPF and IP OSPF:
   - Enter the VNN OSPF Area ID in the Area ID field.
   - Optionally, enter the IP OSPF Area ID in the Trunk Ip Area ID field.
   - By default, IP routing is enabled for the trunk. To disable IP routing, specify Disable in the Trunk IP routing field. For example, you would specify Disable if the trunk handles management traffic and you want to hide the trunk from third-party routers outside the Lucent network.
   - Configure the values for Admin Cost and TOS 0 Metric, where Admin Cost is used by VNN OSPF and TOS 0 Metric is used by IP OSPF.
3. Complete all the procedures for adding or modifying a trunk described in the *NavisCore ATM Configuration Guide* or the *NavisCore Frame Relay Configuration Guide.*

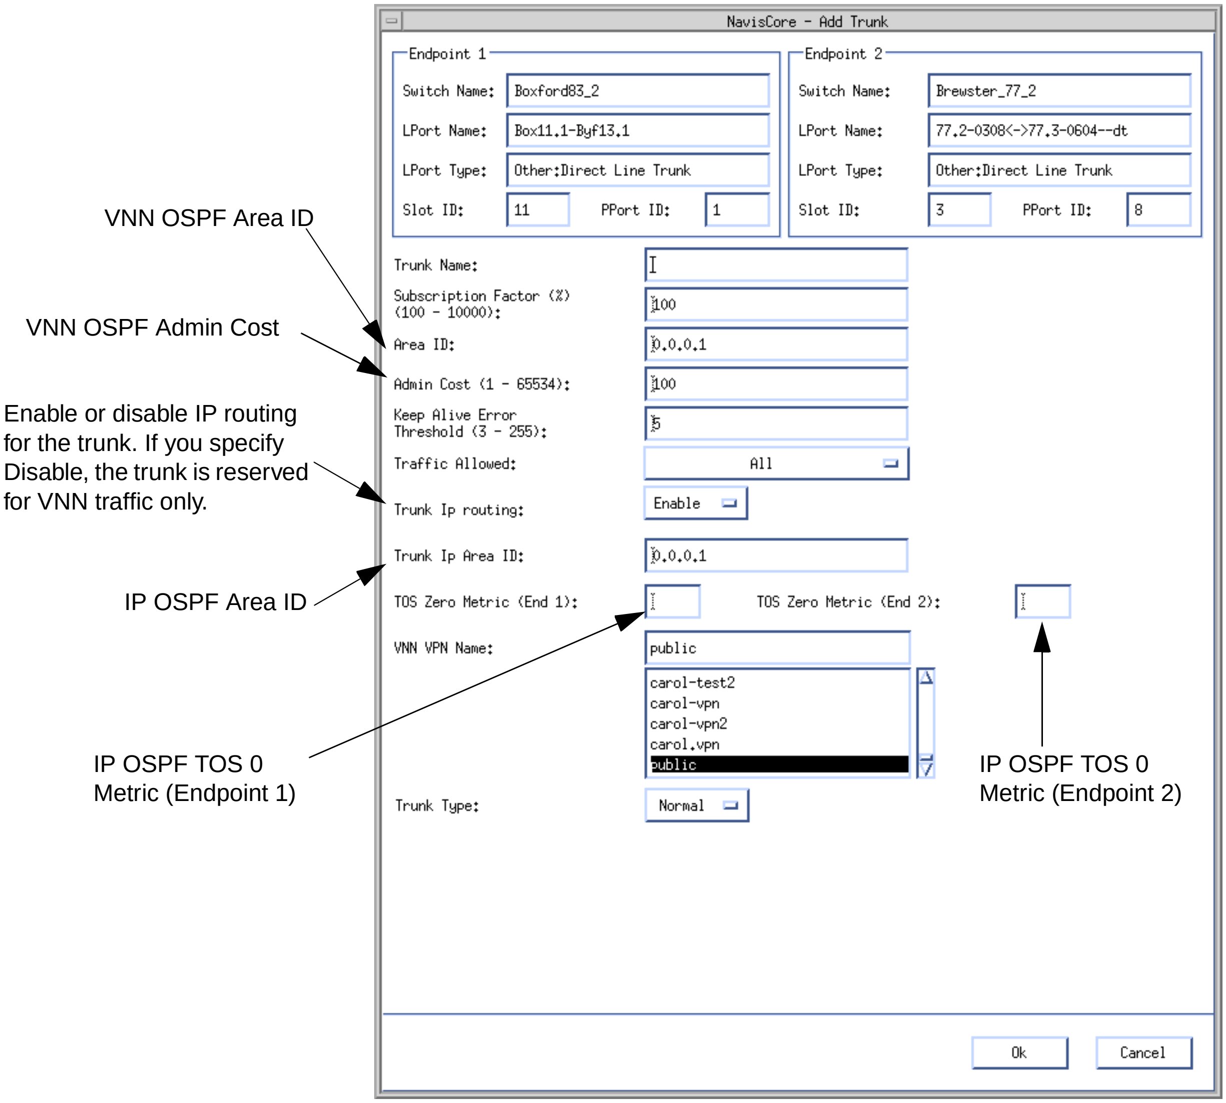

NavisCore - Add Trunk

Endpoint 1

Switch Name: Boxford83_2

LPort Name: Box11.1-Byf13.1

LPort Type: Other:Direct Line Trunk

Slot ID: 11 PPort ID: 1

Endpoint 2

Switch Name: Brewster_77_2

LPort Name: 77.2-0308<->77.3-0604--dt

LPort Type: Other:Direct Line Trunk

Slot ID: 3 PPort ID: 8

Trunk Name:

Subscription Factor (%) (100 - 10000): 100

Area ID: 0.0.0.1

Admin Cost (1 - 65534): 100

Keep Alive Error Threshold (3 - 255): 5

Traffic Allowed: All

Trunk Ip routing: Enable

Trunk Ip Area ID: 0.0.0.1

TOS Zero Metric (End 1):

TOS Zero Metric (End 2):

VNN VPN Name: public

carol-test2
carol-vpn
carol-vpn2
carol.vpn
public

Trunk Type: Normal

VNN OSPF Area ID
VNN OSPF Admin Cost
Enable or disable IP routing for the trunk. If you specify Disable, the trunk is reserved for VNN traffic only.
IP OSPF Area ID
IP OSPF TOS 0 Metric (Endpoint 1)
IP OSPF TOS 0 Metric (Endpoint 2)
Ok
Cancel

Figure 9-11. Add Trunk Dialog Box

If you disable IP routing for a trunk, NavisCore automatically sets the value of TOS Zero Metric (End 1) and the value of TOS Zero Metric (End 2) to zero. *You should write down these values before you disable IP routing, since you must reconfigure these values if you re-enable the trunk.*

When you enable IP routing for a trunk, NavisCore automatically configures an IP OSPF interface for each trunk logical port endpoint. NavisCore uses the following trunk configuration parameters to configure the IP OSPF interface for each trunk endpoint:

Trunk IP Area ID — NavisCore sets the IP OSPF Area ID to the value of this field.

TOS Zero Metric (End 1) — Enter a value between 1 and 65535. This value specifies the TOS 0 metric for the first trunk endpoint, and represents the cost of using this path.

TOS Zero Metric (End 2) — Enter a value between 1 and 65535. This value specifies the TOS 0 metric for the second trunk endpoint.

> ▶ The TOS 0 Metric should be the same for both trunk endpoints; however, they do not have to be the same. For example, you may want traffic in the forward direction to use a different path than traffic in the reverse direction. Be careful when configuring different TOS 0 Metric values for the trunk endpoints. Otherwise, routing problems may result.

Figure 9-12 shows a trunk configured for IP routing.

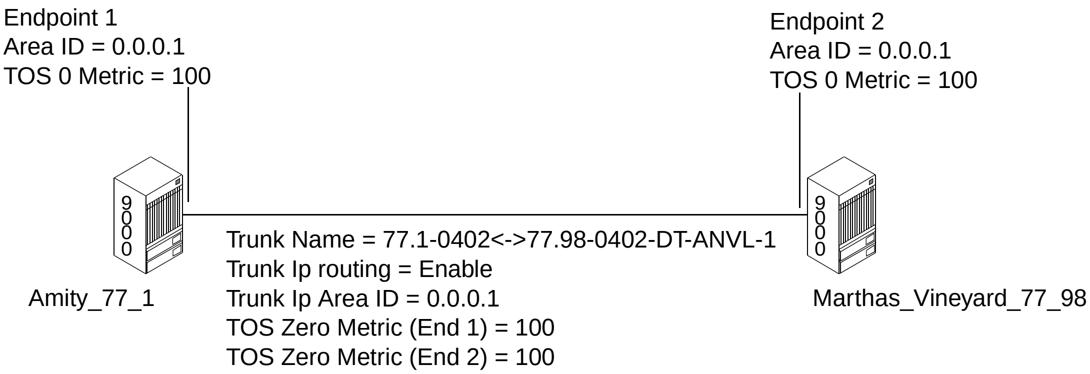

Endpoint 1
Area ID = 0.0.0.1
TOS 0 Metric = 100

Endpoint 2
Area ID = 0.0.0.1
TOS 0 Metric = 100

Trunk Name = 77.1-0402<->77.98-0402-DT-ANVL-1
Trunk Ip routing = Enable
Trunk Ip Area ID = 0.0.0.1
TOS Zero Metric (End 1) = 100
TOS Zero Metric (End 2) = 100

Figure 9-12. Trunk Configured for IP Routing

Figure 9-13 shows the Modify Trunk dialog box for the trunk in Figure 9-12. Note the matching parameters.

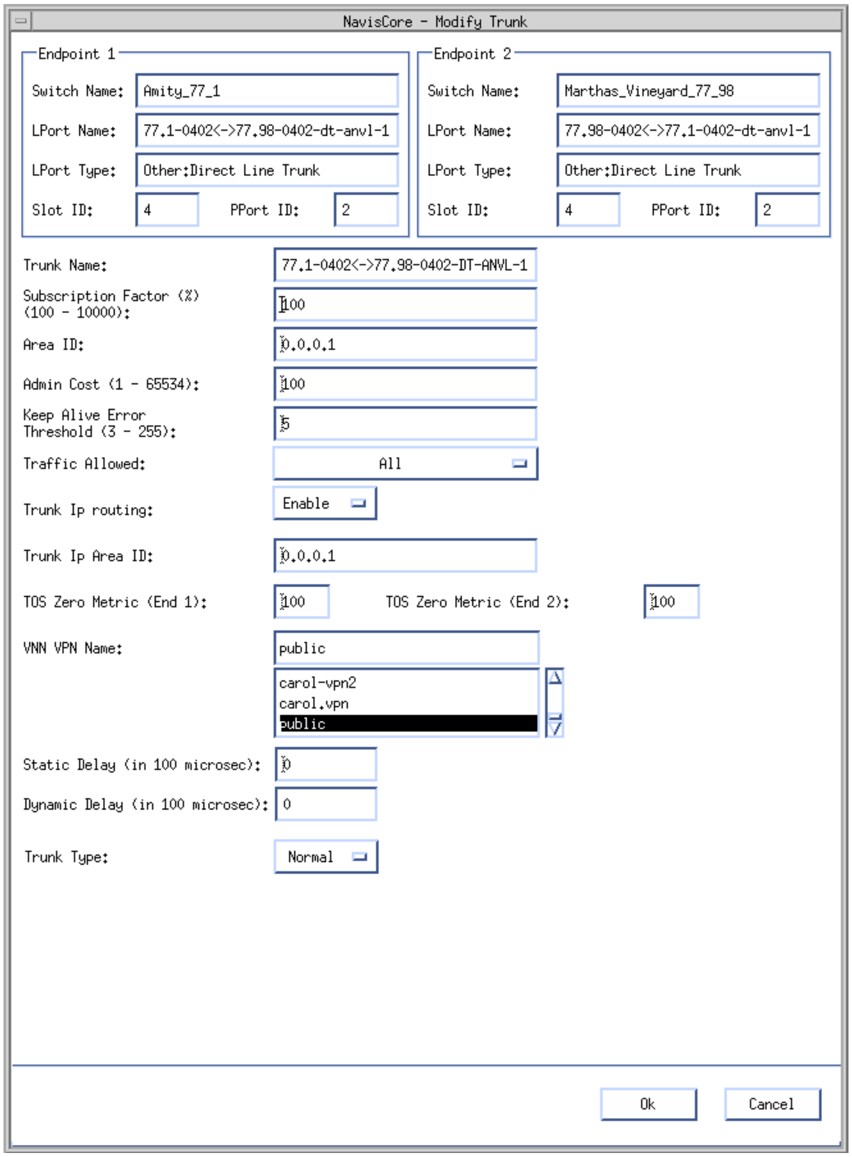

Endpoint 1

Switch Name: Amity_77_1
LPort Name: 77.1-0402<->77.98-0402-dt-anvl-1
LPort Type: Other:Direct Line Trunk
Slot ID: 4
PPort ID: 2

Endpoint 2

Switch Name: Marthas_Vineyard_77_98
LPort Name: 77.98-0402<->77.1-0402-dt-anvl-1
LPort Type: Other:Direct Line Trunk
Slot ID: 4
PPort ID: 2

Trunk Name: 77.1-0402<->77.98-0402-DT-ANVL-1
Subscription Factor (%) (100 - 10000): 100
Area ID: 0.0.0.1
Admin Cost (1 - 65534): 100
Keep Alive Error Threshold (3 - 255): 5
Traffic Allowed: All
Trunk Ip routing: Enable
Trunk Ip Area ID: 0.0.0.1
TOS Zero Metric (End 1): 100
TOS Zero Metric (End 2): 100
VNN VPN Name: public
carol-vpn2
carol.vpn
public
Static Delay (in 100 microsec): 0
Dynamic Delay (in 100 microsec): 0
Trunk Type: Normal

Figure 9-13. Modify Trunk Dialog Box

Configuring IP OSPF

This section describes how to set IP OSPF parameters and includes the following tasks:

- Configuring IP OSPF parameters on the logical port
- Configuring IP OSPF neighbors
- Configuring IP OSPF area aggregates
- Configuring IP OSPF virtual links
- Configuring IP OSPF route maps
- Configuring IP OSPF router IDs
- Configuring multiple IP OSPF authentication keys

In addition, this section also discusses equal-cost multipath (ECMP) routing for IP OSPF.

Configuring IP OSPF on the IP Logical Port

To configure IP OSPF on the IP logical port:

1. Enable the logical port for IP services as described in "Configuring IP Logical Ports" on page 3-5.
2. Choose Add OSPF from the Set IP Interface Addresses dialog box (see Figure 3-6 on page 3-14). The Add OSPF Interface dialog box appears (see Figure 9-14).

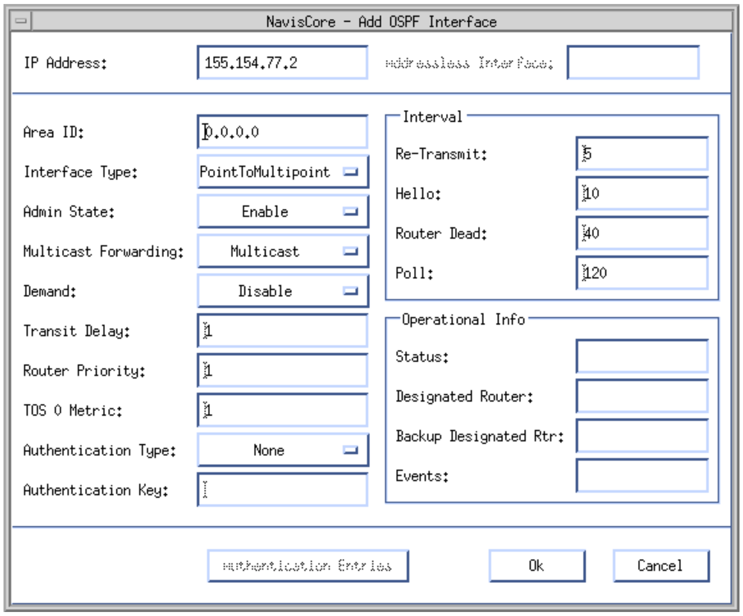

IP Address: 155.154.77.2
Addressless Interface:
Area ID: 0.0.0.0
Interface Type: PointToMultipoint
Admin State: Enable
Multicast Forwarding: Multicast
Demand: Disable
Transit Delay: 1
Router Priority: 1
TOS 0 Metric: 1
Authentication Type: None
Authentication Key:
Re-Transmit: 5
Hello: 10
Router Dead: 40
Poll: 120
Status:
Designated Router:
Backup Designated Rtr:
Events:

Figure 9-14. Add OSPF Interface Dialog Box

The Add OSPF Interface dialog box displays the following buttons:

Table 9-2. Add OSPF Interface Buttons

Button	Function
Authentication Entries	Enables you to add multiple authentication keys for the OSPF interface. See "Configuring Multiple IP OSPF Authentication Keys" on page 9-43 for more information.
OK	Enables you to put your changes into effect.
Cancel	Enables you to exit the dialog box without putting your changes into effect.

3. Complete the fields described in Table 9-3.

Table 9-3. Add OSPF Interface Fields

Field	Action/Description
IP Address	Displays the name assigned to the IP unicast address, with which this IP interface will communicate.
Addressless Interface	Enter the addressless interface. If the interface has an IP address, the value is 0.0.0.0. If the interface is addressless, the value is the logical port number or interface number.
Area ID	Enter the area ID (x.x.x.x) for the area in which you want to locate this interface. Area 0.0.0.0 is the network backbone area. Areas are collections of networks, hosts, and routers. The area ID identifies the area.
Interface Type	Select one of the following options: *Broadcast* – (default for Ethernet and IP VPN cloud interfaces) A broadcast network supports many routers and has a Designated Router that addresses a single physical message to all attached routers. The Hello protocol dynamically discovers neighboring routers on these networks. *NBMA* – (default for Frame Relay and ATM interfaces) A non-broadcast multi-access (NBMA) network supports many routers, but does not have broadcast capability. This type of network requires full-mesh connectivity. *Point-to-Multipoint* – A point-to-multipoint network supports multiple router connections, which are treated like point-to-point connections. The IP addresses of the remote router's interfaces are advertised. *Point-to-Point* – (default for PPP interfaces) A point-to-point network joins two routers together. The IP address of the neighboring router's interface is advertised. Hello packets are sent to the neighbor at regular intervals based on the value that you specify for the *Hello Interval* parameter. Note that this selection may not be available, depending on the type of data link interface. For example, this selection is not available for ATM and Frame Relay interfaces.
Admin State	Select one of the following options: *Enable* – (default) This parameter allows this interface to communicate using IP OSPF. In addition, this interface can send or receive Hello packets. *Disable* – This parameter prevents this interface from communicating using IP OSPF. In addition, this interface cannot send or receive Hello packets.
Multicast Forwarding	See "Configuring MOSPF" on page 15-43 for more information.
Demand	Not supported in this release.
Transit Delay	Enter a value betweeen 1 and 3600 (the default value is 1). This value is the estimated number of seconds it takes to transmit a link-state update packet over this interface.

Table 9-3. Add OSPF Interface Fields (Continued)

Field	Action/Description
Router Priority	Enter a value between 0 and 255 (the default value is 1). This number identifies the priority of the router associated with this logical port and is used to elect the Designated Routers and Backup Designated Routers. The router with the highest priority is considered the Designated Router. A value of 0 indicates the router is not eligible to be the designated or Backup Designated Router. If all routers have the same priority, the router ID is used to determine the Designated Router.
TOS 0 (Zero) Metric	Enter a value between 1 and 65535 (the default value is 1). This value specifies the type of service cost. The lowest TOS 0 has the highest priority for routing.
Authentication Type	Specify the type of authentication that OSPF uses as a security measure to ensure that this logical port and router exchange information with correct neighbors. Options include: *None* – (default) Specifies that no authentication is performed. *Simple Password* – Specifies a simple password authentication method that includes a password in all OSPF messages on an interface-by-interface basis. When a router receives a message on an interface that uses simple password authentication, the router checks the incoming OSPF message to see if the password is included in the message. If the password is correct, the message is processed normally. If the password is not part of the incoming message, the message is ignored and dropped. *MD5* – Use MD5 authentication to verify a key that is appended to the end of an IP OSPF protocol packet. For more information on how MD5 authentication works, see RFC 1321 (The MD5 Message-Digest Algorithm). In addition to RFC 1321, RFC 2178 (OSPF Version 2) provides information on how MD5 authentication is used with IP OSPF.
Authentication Key	Enter an authentication password if you specified either *Simple* or *MD5* as the authentication type. This value is not required if you specified *None* as the authentication type.
Interval	
Re-Transmit	Enter a value between 0 and 3600 (the default value is 5 seconds). This value specifies the time to wait before resending a packet if no acknowledgment is received.
Hello	Enter a value between 1 and 65535 (the default value is 10 seconds). Specifies the number of seconds between router Hello messages. This parameter controls the frequency of router Hello messages on an interface.
Router Dead	Enter a value greater than or equal to 0 (the default value is 40 seconds). This value is a multiple of the Hello interval. For example, if the Hello interval is set to 10, the router dead interval should be configured at 40. This parameter is the number of seconds a router waits to hear a Hello message from a neighbor before the router declares the neighbor unreachable. The value that you specify can affect OSPF operation. If the interval is too short, neighbors are considered unreachable when they are available. If the interval is too long, routers that are unreachable are not identified soon enough to reroute data properly.

Table 9-3. Add OSPF Interface Fields (Continued)

Field	Action/Description
Poll	Enter a value greater than or equal to 0 (the default value for this field is 120). Specifies the time, in seconds, between Hello packets sent to an inactive non-broadcast multi-access (NBMA) neighbor.
Operational Info (All values are read-only)	
Status	Displays the status of OSPF communication. Options for Point-to-Point, Point-to-Multipoint, Broadcast, and Virtual link networks are: *Up* – Indicates the network interface is operational. *Point-to-Point* – Indicates the interface is at the highest level of connection. In this state, the interface is operational and connects either to a physical point-to-point network or to a virtual link. Upon entering this state, the router attempts to form an adjacency with the neighboring router. Hello packets are sent to the neighbor at regular intervals based on the value that you specify for the *Hello Interval* parameter. *Init* – In this state, the neighbor sees a Hello packet. However, bidirectional communication has not been established with the neighbor. All neighbors in this state are listed in the Hello packets sent from the associated interface. *Down* – Indicates the interface is not usable. No protocol traffic will be sent or received on this interface. Options for an NBMA network are: *Loopback* – In this state, the router's interface to the network is "looped back." The interface may be looped back in hardware and software. While in loopback, the interface is not available for regular traffic data traffic. *Waiting* – In this state, the router tries to determine the Backup Designated Router's identity. To do this, the router monitors received Hello packets. The router cannot elect a Designated Router or Backup Designated Router until it leaves the waiting state. This prevents any unnecessary changes to the Backup Designated Router. *Designated Router* – In this state, the router is the Designated Router on the attached network. Adjacencies are established to all other routers attached to the network. The router must also originate network link advertisements for the network node. The advertisement provides link information to all routers (including the Designated Router itself) attached to the network. *Backup Designated Router* – In this state, the router is the Backup Designated Router on the attached network. When the present Designated Router fails, this router takes over. The router establishes adjacencies to all other routers attached to the network. *Other* – In this state, the router forms adjacencies to both the Designated Router and the Backup Designated Router.
Designated Router	Displays the 32-bit IP address of the Designated Router for this network as seen by the advertising router. An IP address of 0.0.0.0 indicates that a Designated Router has not been specified for this network. If all routers have the same priority, the router ID is used to specify the Designated Router.

Table 9-3. Add OSPF Interface Fields (Continued)

Field	Action/Description
Backup Designated Rtr	Displays the 32-bit IP address of the Backup Designated Router for this network as seen by the advertising router. An IP address of 0.0.0.0 indicates that a Backup Designated Router has not been specified for this network.
Events	Displays the number of times this OSPF interface changed its state, or the number of times an error occurred.

4. When you are done setting parameters, choose OK.

Configuring IP OSPF Neighbors

> You do not have to define IP OSPF neighbors that are reached on IP OSPF interfaces. OSPF automatically discovers its neighbors through Hello packets.
>
> If you manually define IP OSPF neighbors for an IP OSPF interface configured with an NBMA interface type where IARP is disabled, only those neighbors are in effect. Dynamic discovery is disabled on that interface.

To define an IP OSPF neighbor:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Neighbors. The Set All OSPF Neighbors dialog box appears (see Figure 9-15).

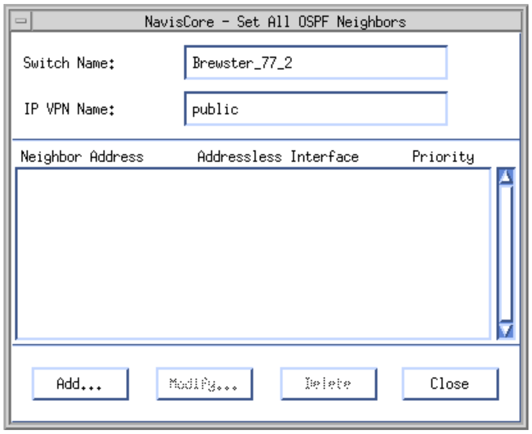

Switch Name: Brewster_77_2

IP VPN Name: public

Figure 9-15. Set All OSPF Neighbors Dialog Box

3. Choose Add. The Add OSPF Neighbor dialog box appears (see Figure 9-16).

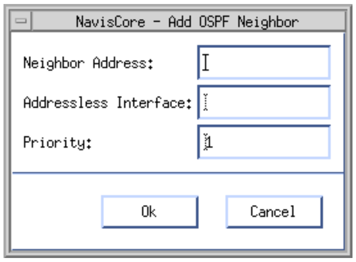

NavisCore - Add OSPF Neighbor

Neighbor Address:

Addressless Interface:

Priority: 1

Figure 9-16. Add OSPF Neighbor Dialog Box

4. Complete the fields described in Table 9-4.

Table 9-4. Add OSPF Neighbor Fields

Field	Action/Description
Neighbor Address	Enter the IP address this neighbor uses in its IP source address. On addressless links, the address is not 0.0.0.0 but the address of another of the neighbor's interfaces.
Addressless Interface	Enter the addressless interface. If the interface has an IP address, the value is 0.0.0.0. If the interface is addressless, the value is the logical port number or interface number.
Priority	Enter a value between 0 and 255 (the default value is 1). The neighbor with the highest priority is the Designated Router. This field only applies to NBMA and broadcast networks. The value zero signifies the neighbor cannot be the Designated Router on this network.

5. When you are done setting parameters, choose OK.

6. At the Set All OSPF Neighbors dialog box, choose Close.

Configuring IP OSPF Area Aggregates

To define an OSPF area aggregate:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Area Aggregates. The Set All OSPF Area Aggregates dialog box appears (see Figure 9-17).

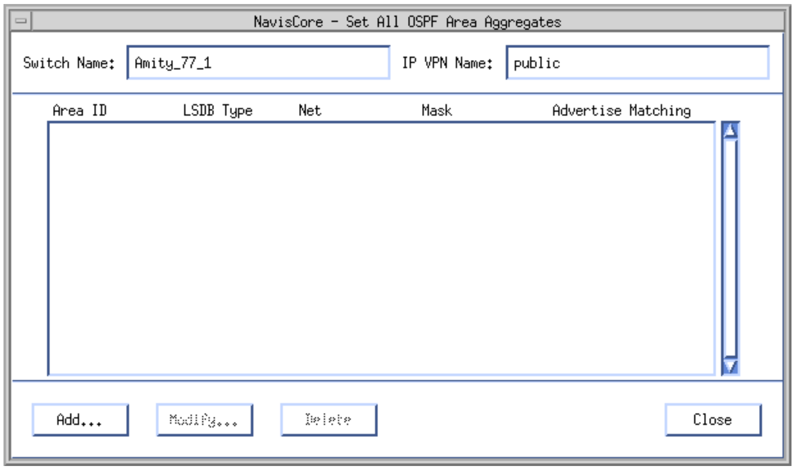

Switch Name: Amity_77_1
IP VPN Name: public

Figure 9-17. Set All OSPF Area Aggregates Dialog Box

3. Choose Add. The Add OSPF Area Aggregate dialog box appears (see Figure 9-18).

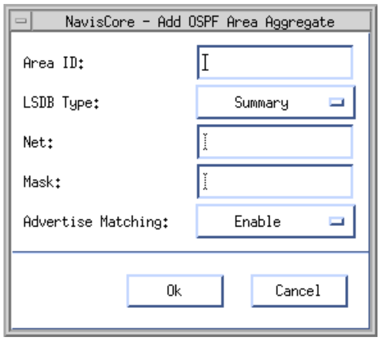

Area ID:
LSDB Type: Summary
Net:
Mask:
Advertise Matching: Enable

Figure 9-18. Add OSPF Area Aggregate Dialog Box

4. Complete the fields described in Table 9-5.

Table 9-5. Add OSPF Area Aggregate Fields

Field	Action/Description
Area ID	Enter the area ID (x.x.x.x) in which you want to locate the node. Area 0.0.0.0 is the network backbone. Areas are collections of networks, hosts, and routers. The area ID identifies the area.
LSDB Type	Specify the link state database type to which this address aggregate applies. Options include: *Summary* – (default) Area border routers generate summary link advertisements, which describe inter-area routes (routes between areas) to networks. *NSSA External* – Not So Stubby Area external (NSSA) link advertisements allow an AS border router within a stub area and the routers within that area to learn about the external networks accessible through the AS border router in the area.
Net	Enter the IP address of the net or subnet, indicated by the range.
Mask	Enter the subnet mask that pertains to the net or subnet.
Advertise Matching	Select one of the following options: *Enable* – (default) If you enable this parameter, you “leak” the net/mask you specified for the given area. *Disable* – If you disable this parameter, you hide the net/mask you specified for the given area.

5. When you are done setting parameters, choose OK.

6. At the Set All OSPF Area Aggregates dialog box, choose Close.

Configuring IP OSPF Virtual Links

To define an OSPF virtual link:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Virtual Links. The Set All OSPF Virtual Links dialog box appears (see Figure 9-19).

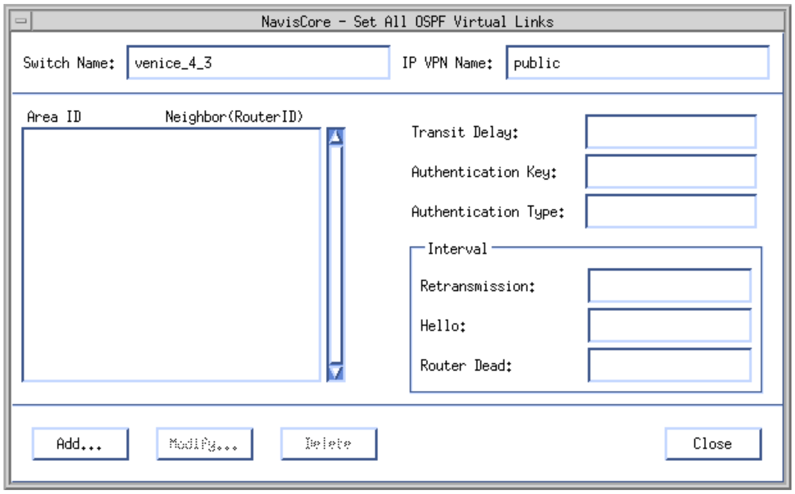

Switch Name: venice_4_3
IP VPN Name: public

Figure 9-19. Set All OSPF Virtual Links Dialog Box

3. Choose Add. The Add OSPF Virtual Link dialog box appears (see Figure 9-20).

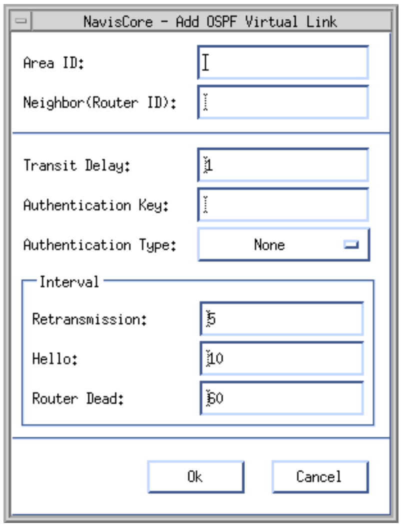

Transit Delay: 1
Authentication Type: None
Retransmission: 5
Hello: 10
Router Dead: 60

Figure 9-20. Add OSPF Virtual Link Dialog Box

4. Complete the fields described in Table 9-6.

Table 9-6. OSPF Virtual Link Fields

Field	Action/Description
Area ID	Enter the area ID (*x.x.x.x*) of the transit area, which is the non-backbone area that the virtual link traverses to connect to the backbone area. This ID cannot be 0.0.0.0 (the Area ID of the backbone area). Areas are collections of networks, hosts, and routers. The area ID identifies the area.
Neighbor (Router ID)	Enter the Router ID (internal IP address) of the switch (that is, the neighbor) on the other end of the virtual link. The router ID is configured when the switch is installed. To determine the internal IP address, access the switch console and issue the **`show ip ospf statistics`** command. In the command output, the internal IP address appears in the OSPF Router ID field. For example: **OSPF Router ID: 150.202.77.2**
Transit Delay	Enter a value between 0 and 3600 (the default value is 1). This field specifies the estimated number of seconds it takes to transmit a link-state update packet over this interface.
Authentication Key	Enter an authentication password in this field if you specify either *Simple* or *MD5* as the authentication type. This value is not required if you specify *None* as the authentication type.
Authentication Type	Specify the type of authentication that OSPF uses as a security measure to ensure that this logical port and router exchange information with correct neighbors. Options include: *None* – (default) Specifies that no authentication is performed. *Simple Password* – Specifies a simple password authentication method that includes a password in all OSPF messages on an interface-by-interface basis. When a router receives a message on an interface that uses simple password authentication, the router checks the incoming OSPF message to see if the password is included in the message. If the password is correct, the message is processed normally. If the password is not part of the incoming message, the message is ignored and dropped. *MD5* – Use MD5 authentication to verify a key that is appended to the end of an IP OSPF protocol packet. For more information on how MD5 authentication works, see RFC 1321 (The MD5 Message-Digest Algorithm). In addition to RFC 1321, RFC 2178 (OSPF Version 2) provides information on how MD5 authentication is used with IP OSPF.
Interval	
Retransmission	Enter a value between 0 and 3600 (the default value is 5 seconds) This field specifies the time to wait before resending a packet if no acknowledgment is received.

Table 9-6. OSPF Virtual Link Fields (Continued)

Field	Action/Description
Hello	Enter a value between 1 and 65535 (the default value is 10 seconds). This field specifies the number of seconds between router Hello messages and controls the frequency of router Hello messages on an interface. The neighbor switch must use the value you enter here.
Router Dead	Enter a value greater than or equal to 0 (the default value for this field is 40 seconds). This value is a multiple of the Hello interval. For example, if the Hello interval is set to 10, the router dead interval should be configured at 20, 30, 40, etc. Specify this parameter if you have bad connections or if a link in the network is down. The neighbor switch must use the value you enter here. This parameter is the number of seconds a router waits to hear a Hello message from a neighbor before the router declares the neighbor “down.” The value that you specify can affect OSPF operation. If the interval is too short, neighbors are considered down when they are reachable. If set for too long, routers that are really down are not considered down soon enough to properly reroute data.

5. Choose OK.
6. At the Set All OSPF Virtual Links dialog box, choose Close.

Configuring IP OSPF Route Maps

Chapter 11, “Configuring Route Policies” and “Configuring Route Maps for IP OSPF and VNN OSPF” on page 9-57 provide detailed information about all types of route maps (including OSPF route maps) that you can configure using IP Navigator. See Chapter 11 and “Configuring Route Maps for IP OSPF and VNN OSPF” before you begin any route map configuration.

To configure an IP OSPF route map from the IP OSPF parameter menu:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Route Maps. The Set All OSPF Route Maps dialog box appears (see Figure 9-21).

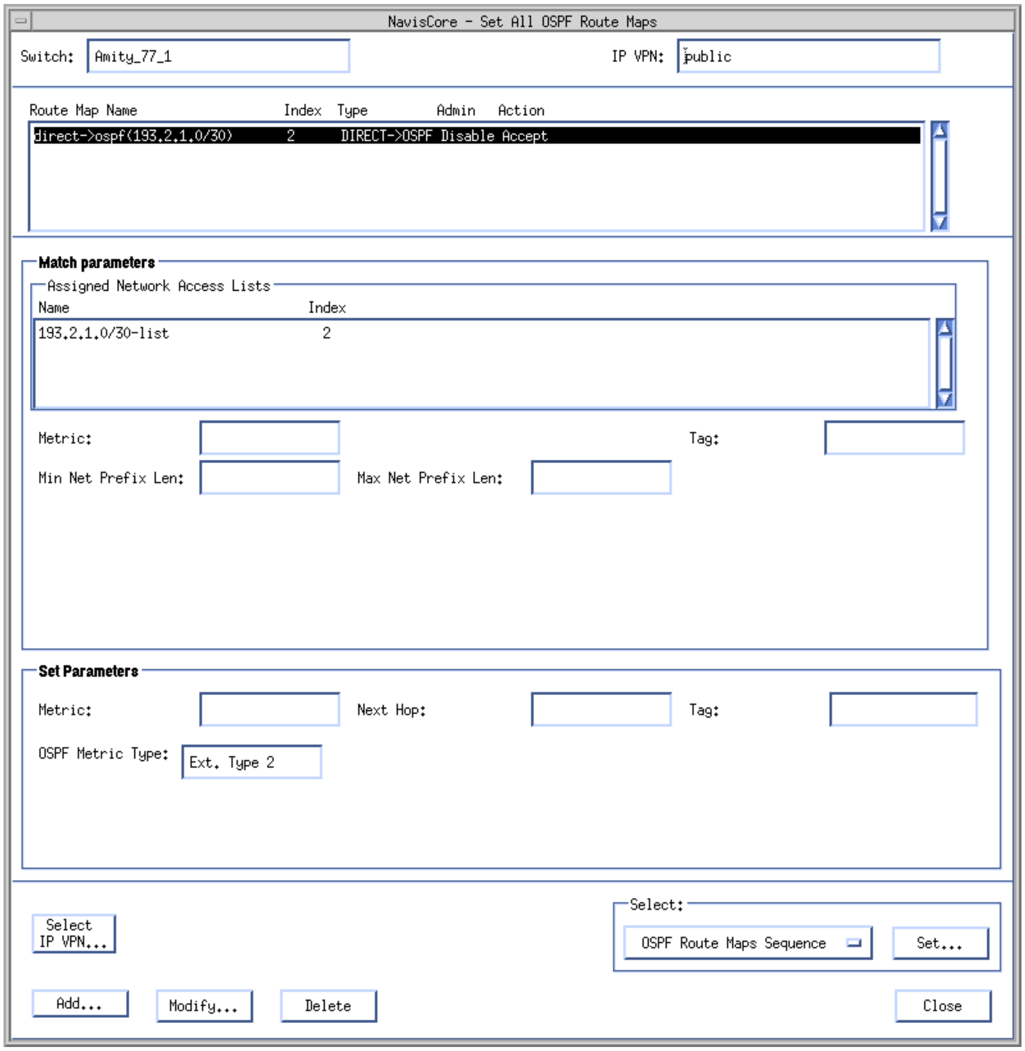

NavisCore - Set All OSPF Route Maps

Switch: Amity_77_1   IP VPN: public

Route Map Name	Index	Type	Admin	Action
direct->ospf(193.2.1.0/30)	2	DIRECT->OSPF	Disable	Accept

Match parameters

Assigned Network Access Lists

Name	Index
193.2.1.0/30-list	2

Metric: Tag:

Min Net Prefix Len: Max Net Prefix Len:

Set Parameters

Metric: Next Hop: Tag:

OSPF Metric Type: Ext. Type 2

Select: OSPF Route Maps Sequence

Use the OSPF Route Maps Sequence option and choose Set to display a dialog box that enables you to order the route map sequence.

Figure 9-21. Set All OSPF Route Maps

Configuring IP OSPF Router IDs

> ▶ Before you configure IP OSPF router IDs, make sure you read the section "Planning Router IDs" on page 9-22.

To configure the router ID for IP OSPF:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All OSPF ⇒ Set OSPF Router ID. The Set OSPF Router ID dialog box (see Figure 9-22) appears.

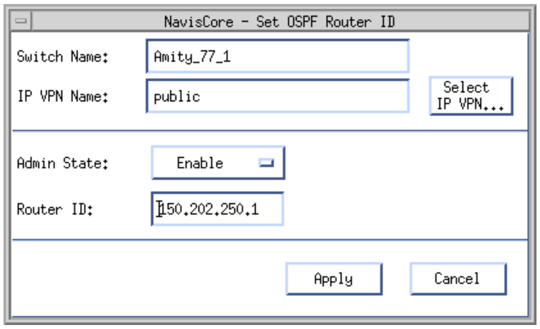
NavisCore - Set OSPF Router ID

Switch Name: Amity_77_1

IP VPN Name: public

Admin State: Enable

Router ID: 150.202.250.1

Figure 9-22. Set OSPF Router ID Dialog Box

2. Complete the following fields:

Admin State — Specify Enable (the default) to enable the router ID. Specify Disable to disable the router ID. If you specify Disable, IP Navigator will not function properly on the switch.

Router ID — Enter the router ID for the switch. This ID is a 32-bit IP address (for example, 155.10.10.10) that uniquely identifies the switch within the Autonomous System (AS). The router ID can be the same as an IP address assigned to any of the switch's IP interfaces. Consider using the smallest IP interface address as the router ID.

3. Choose Apply when you are done.

Note that you can configure an IP OSPF router ID on an IP VPN-by-IP VPN basis. You can choose the Select IP VPN button to access a specific IP VPN. See Chapter 16, "Configuring IP Virtual Private Networks" for more information on managing IP VPNs.

Configuring Multiple IP OSPF Authentication Keys

You can configure multiple authentication keys for an IP OSPF interface. You can then activate these keys at different times, making it extremely difficult for network intruders to compromise your security.

About Multiple Authentication Keys

Each authentication key has four time constants:

Start Accept — The date and time when the IP OSPF interface starts accepting packets with the given key.

Stop Accept — The date and time when the IP OSPF interface stops accepting packets with the given key.

Start Generate — The date and time when the IP OSPF interface starts to use the key to generate packets.

Stop Generate — The date and time when the IP OSPF interface stops using the key to generate packets.

For a given authentication key, the time constants should be synchronized as follows:

- The Start Accept date and time should be less than the Start Generate date and time. This guarantees that the IP OSPF interface is ready to accept packets with the given authentication key before they are generated.
- The Stop Generate date and time should be less than the Stop Accept date and time. This guarantees that the generation of packets with the given authentication key stops before the IP OSPF interface is ready to reject them.

Make sure that the same time constant values are configured on all network nodes for a given authentication key. The Start Accept day and time you configure on one node should be the same as the Start Accept day and time you configure on another node, and so on.

When a new authentication key replaces an old one, the Start Generate day and time for the new key must be less than or equal to the Stop Generate day and time of the old key. This helps to avoid the creation of a window of time during which network traffic is interrupted. The interruption results from the following:

- Packets with the old key are not generated
- Packets with the new key are generated but not accepted

Figure 9-23 shows two authentication keys and their respective time constants that span a one-year period.

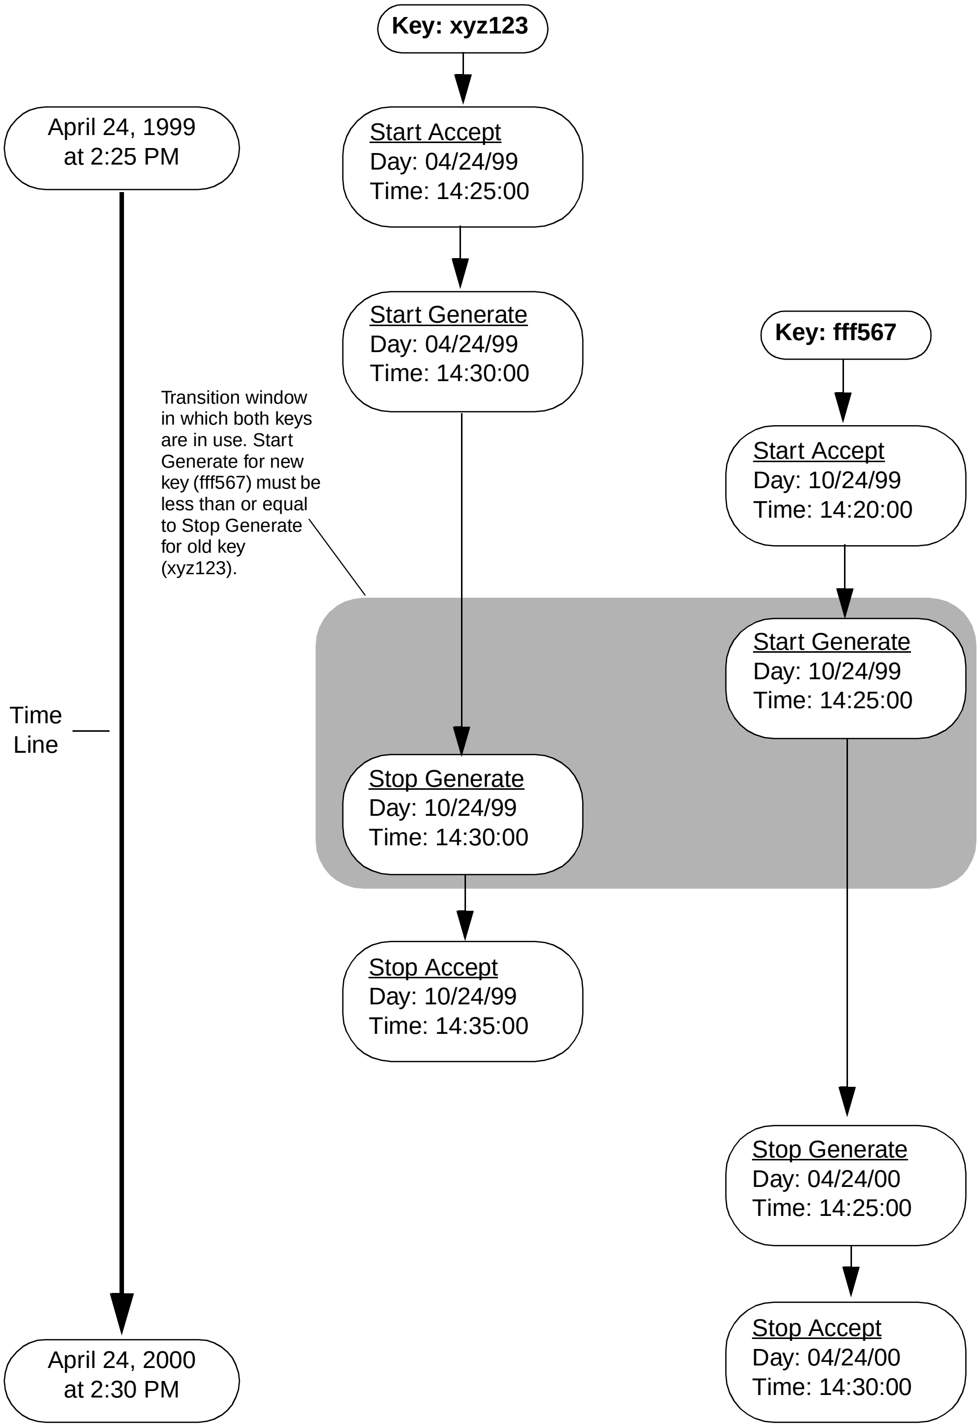

Figure 9-23. Multiple Authentication Keys for an IP OSPF Interface

Configuring Authentication Keys

Before you attempt to configure multiple authentication keys, make sure that you have added an OSPF interface and specified an Authentication Type of MD5 or Simple Password.

To configure an authentication key:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears (see Figure 3-2 on page 3-6).
3. Select the IP logical port associated with the OSPF interface for which you want to configure MD5 authentication or simple password authentication.
4. Choose IP Parameters. The Set IP Parameters dialog box appears (see Figure 3-3 on page 3-7).
5. Select IP Interface.
6. Choose Go. The Set IP Interface Addresses dialog box appears (see Figure 3-6 on page 3-14).
7. Choose Modify OSPF. The Modify OSPF Interface dialog box appears (see Figure 9-24).

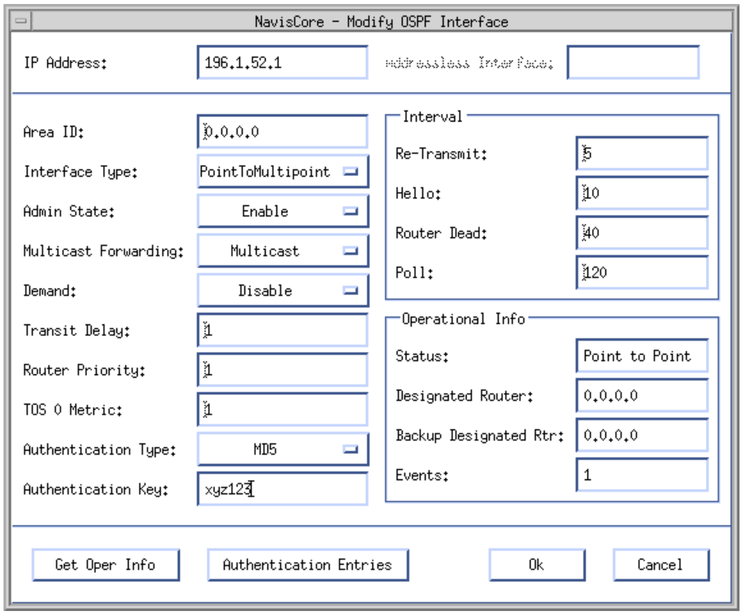

NavisCore - Modify OSPF Interface

IP Address: 196.1.52.1
Addressless Interface:

Area ID: 0.0.0.0
Interface Type: PointToMultipoint
Admin State: Enable
Multicast Forwarding: Multicast
Demand: Disable
Transit Delay: 1
Router Priority: 1
TOS 0 Metric: 1
Authentication Type: MD5
Authentication Key: xyz123

Interval
Re-Transmit: 5
Hello: 10
Router Dead: 40
Poll: 120

Operational Info
Status: Point to Point
Designated Router: 0.0.0.0
Backup Designated Rtr: 0.0.0.0
Events: 1

Figure 9-24. Modify OSPF Interface Dialog Box

8. Choose Authentication Entries. The Set OSPF Authentication Entries dialog box appears (see Figure 9-25).

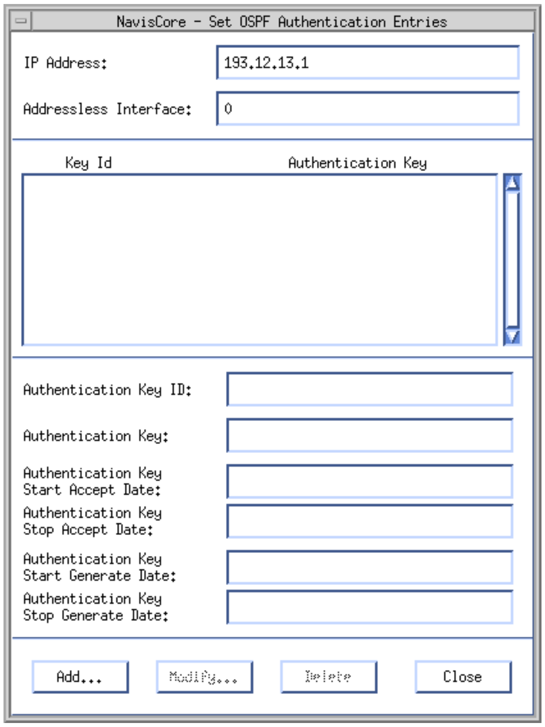

NavisCore - Set OSPF Authentication Entries

IP Address: 193.12.13.1

Addressless Interface: 0

Key Id
Authentication Key
Authentication Key ID:
Authentication Key:
Authentication Key Start Accept Date:
Authentication Key Stop Accept Date:
Authentication Key Start Generate Date:
Authentication Key Stop Generate Date:
Add...
Modify...
Delete
Close

Figure 9-25. Set OSPF Authentication Entries Dialog Box

The Set OSPF Authentication Entries dialog box displays all of the authentication keys you have configured for the OSPF interface, and their associated time constants. The fields on the dialog box are described in Table 9-8.

The buttons on the dialog box are described in Table 9-7.

Table 9-7. Set OSPF Authentication Entries Buttons

Button	Function
Add	Adds an authentication key entry.
Modify	Modifies a selected authentication key.
Delete	Deletes a selected authentication key.

9. Choose Add. The Add OSPF Authentication Entry dialog box appears (see Figure 9-26).

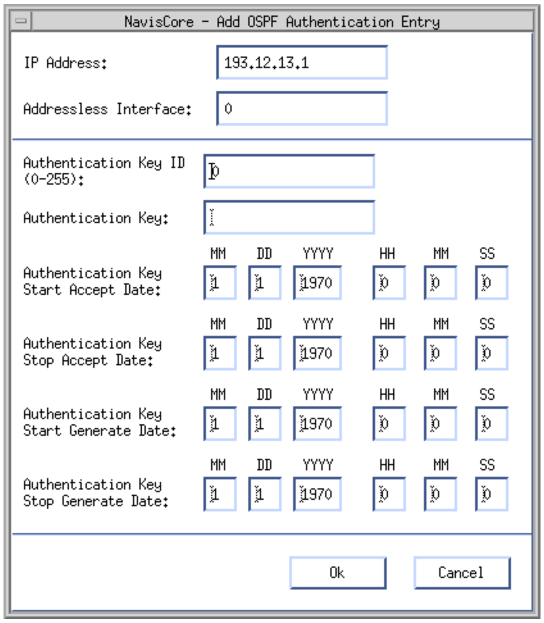

NavisCore - Add OSPF Authentication Entry

IP Address: 193.12.13.1

Addressless Interface: 0

Authentication Key ID (0-255): 0

Authentication Key:

	MM	DD	YYYY	HH	MM	SS
Authentication Key Start Accept Date:	1	1	1970	0	0	0
Authentication Key Stop Accept Date:	1	1	1970	0	0	0
Authentication Key Start Generate Date:	1	1	1970	0	0	0
Authentication Key Stop Generate Date:	1	1	1970	0	0	0

Figure 9-26. Add OSPF Authentication Entry Dialog Box

10. Complete the fields in Table 9-8.

Table 9-8. Add OSPF Authentication Entry Fields

Field	Action/Description
Authentication Key ID	Specify a numeric authentication key ID, from 0 to 255. Each authentication key is identified by a unique ID. The default is 0.
Authentication Key	Specify an alphanumeric string that will act as the authentication key. The sting can be up to 16 characters long.
Authentication Key Start Accept Date	Specify the date and time that the OSPF interface accepts packets with the specified key. The Start Accept date and time should be less than the Start Generate date and time. This guarantees that the OSPF interface is ready to accept packets with the given authentication key before they are generated. Specify date and time as follows: • Specify the day in the DD field. The default is 1 (the first day of the month). • Specify the year in the YYYY field. You must specify all 4 digits of the year (for example, 2000). The default is 1970. • Specify the hour, minute, and second on the specified date in the HH, MM, and SS fields respectively using military-style time (for example, 2:05 PM is 14:05:00). The default for all of these fields in 0. If you leave these fields unchanged, the key goes into effect at midnight on the specified date.

Table 9-8. Add OSPF Authentication Entry Fields (Continued)

Field	Action/Description
Authentication Key Stop Accept Date	Specify the date and time that the OSPF interface stops accepting packets with the specified key. See the description of the Authentication Key Start Accept Date for information on how to specify date and time.
Authentication Key Start Generate Date	Specify the date and time that the OSPF interface generates packets with the specified key. The Start Generate date and time should be greater than the Start Accept date and time. This guarantees that the OSPF interface is ready to accept packets with the given authentication key before they are generated. See the description of the Authentication Key Start Accept Date for information on how to specify date and time.
Authentication Key Stop Generate Date	Specify the date and time that the OSPF interface stops generating packets with the specified key. The Stop Generate date and time should be less than the Stop Accept date and time. This guarantees that the generation of packets with the given authentication key stops before the OSPF interface is ready to reject them. See the description of the Authentication Key Start Accept Date for information on how to specify date and time.

11. Choose OK.

About Equal-cost Multipath Routing for IP OSPF

This release supports equal-cost multipath (ECMP) routing, which load balances IP traffic as many as four routes of equal cost to the same destination. If more than four routes of equal cost to the same destination exist, ECMP uses the four most recent routes. This feature applies to routes learned through BGP, RIP, OSPF, and manually configured static routes.

IP OSPF automatically calculates multiple equal-cost paths to the same destination IP prefix. No manual configuration is required.

If multiple next hops to a router that advertises a destination IP prefix exist, IP traffic is load balanced across all of these next hops if the following conditions are met:

- The paths associated with the next hops have the same administrative cost.
- The administrative cost is the lowest administrative cost out of all the paths that can be used to reach the destination IP prefix.

Instead of next hops, the equal-cost paths could also be associated with multiple circuits (label switched paths) if the following conditions are met:

- Multiple egress switches to the destination IP prefix exist
- Label switched paths to the egress switches exist

Configuring VNN OSPF

This section describes how to configure loopback addresses, area aggregates, and virtual links for VNN OSPF.

Configuring VNN Loopback Addresses

To configure loopback addresses for VNN:

1. From the Administer menu, select Lucent Parameters ⇒ Set All VNN ⇒ Set All VNN Loopbacks. The Set All VNN Loopback Addresses dialog box (see Figure 9-27) appears.

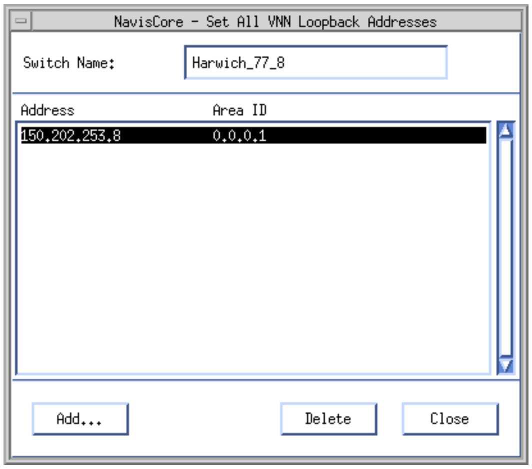

NavisCore - Set All VNN Loopback Addresses

Switch Name: Harwich_77_8

Address	Area ID
150.202.253.8	0.0.0.1

Figure 9-27. Set All VNN Loopback Addresses Dialog Box

2. Choose Add. The Add VNN Loopback Address dialog box (see Figure 9-28) appears.

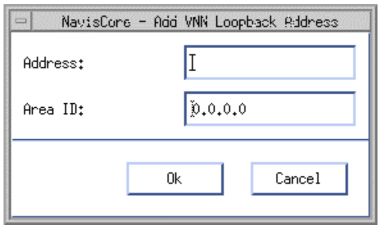

NavisCore - Add VNN Loopback Address

Address:

Area ID: 0.0.0.0

Figure 9-28. Add VNN Loopback Address Dialog Box

3. Enter the loopback IP address (for example, 152.148.30.5).
4. Enter the Area ID (for example, 0.0.0.2).
5. Choose OK. The Set All VNN Loopback Addresses dialog box appears, displaying the new loopback address.

Configuring VNN Area Aggregates

To configure area aggregates for VNN:

1. From the Administer menu, select Lucent Parameters ⇒ Set All VNN ⇒ Set All VNN Area Aggregates. The Set All VNN Area Aggregates dialog box (see Figure 9-29) appears.

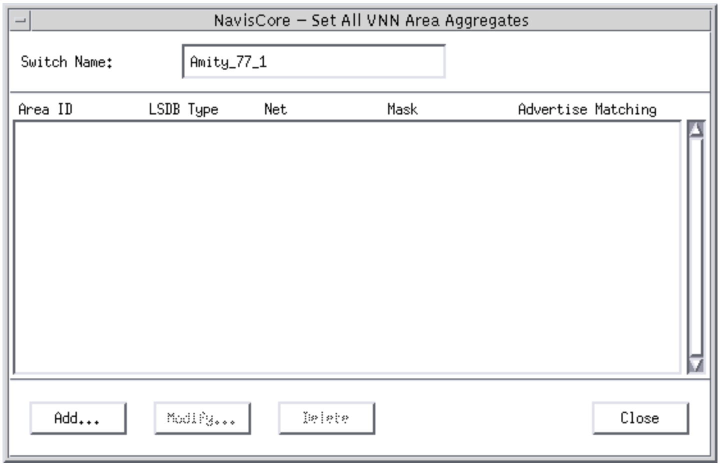

NavisCore – Set All VNN Area Aggregates

Switch Name: Amity_77_1

Figure 9-29. Set All VNN Area Aggregates Dialog Box

2. Choose Add. The Add VNN Area Aggregate dialog box (see Figure 9-30) appears.

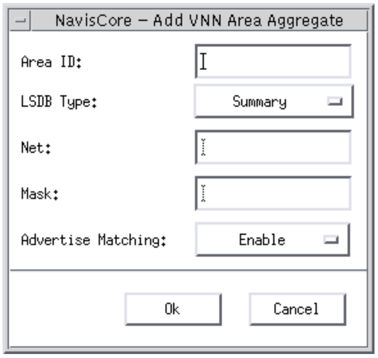

NavisCore – Add VNN Area Aggregate

Area ID:

LSDB Type: Summary

Net:

Mask:

Advertise Matching: Enable

Figure 9-30. Add VNN Area Aggregate Dialog Box

3. Complete the fields described in Table 9-9.

Table 9-9. Add VNN OSPF Area Aggregate Fields

Field	Action/Description
Area ID	Enter the ID (x.x.x.x) of the area in which the IP address range is located. Area 0.0.0.0 is the network backbone. Areas are collections of networks, hosts, and routers. The area ID identifies the area.
LSDB Type	Specify the link state database type to which this address aggregate applies. The only option is as follows: *Summary* – (default) Area border routers generate summary link advertisements, which describe inter-area routes (routes between areas) to networks. Note that *NSSA External* is not a supported option for VNN OSPF area aggregates. It is only supported for IP OSPF area aggregates.
Net	Enter the IP address of the network or subnetwork that encompasses the range of addresses you want to advertise.
Mask	Enter the subnet mask that pertains to the net or subnet.
Advertise Matching	Select one of the following options: *Enable* – (default) If you enable this parameter, you “leak” the net/mask you specified for the given area, making it available to the rest of the network. *Disable* – If you disable this parameter, you hide the net/mask you specified for the given area.

4. When you are done setting parameters, choose OK.

Configuring VNN Virtual Links

To configure virtual links for VNN:

1. From the Administer menu, select Lucent Parameters ⇒ Set All VNN ⇒ Set All VNN Virtual Links. The Set All VNN Virtual Links dialog box (see Figure 9-31) appears.

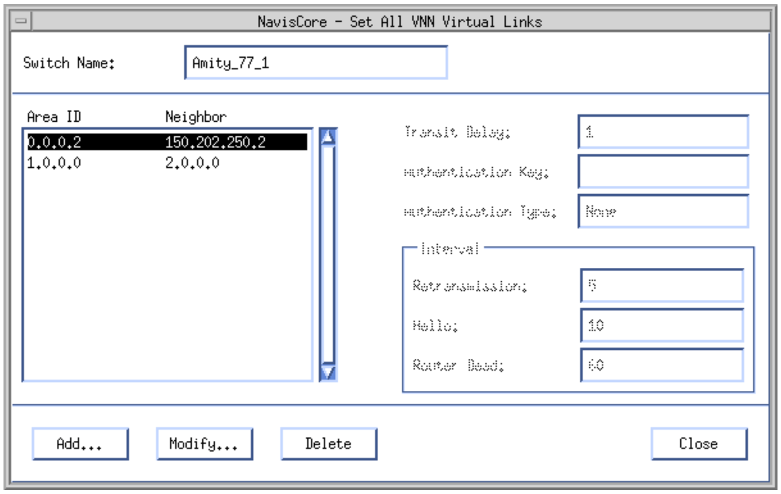

NavisCore - Set All VNN Virtual Links

Switch Name: Amity_77_1

Area ID	Neighbor
0.0.0.2	150.202.250.2
1.0.0.0	2.0.0.0

Transit Delay: 1

Authentication Key:

Authentication Type: None

Interval

Retransmission: 5

Hello: 10

Router Dead: 60

Figure 9-31. Set All VNN Virtual Links Dialog Box

2. Choose Add. The Add VNN Virtual Link dialog box appears (see Figure 9-32).

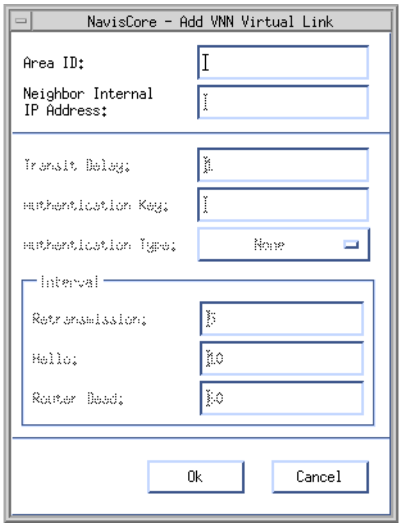

NavisCore - Add VNN Virtual Link

Area ID:

Neighbor Internal IP Address:

Transit Delay: 1

Authentication Key:

Authentication Type: None

Interval

Retransmission: 5

Hello: 10

Router Dead: 60

Figure 9-32. Add VNN Virtual Link Dialog Box

3. Complete the fields described in Table 9-10.

Table 9-10. VNN OSPF Virtual Link Fields

Field	Action/Description
Area ID	Enter the area ID (*x.x.x.x*) of the transit area, which is the non-backbone area that the virtual link traverses to connect to the backbone area. This ID cannot be 0.0.0.0 (the Area ID of the backbone area). Areas are collections of networks, hosts, and routers. The area ID identifies the area.
Neighbor Internal IP Address	Enter the internal IP address of the switch (that is, the neighbor) on the other end of the virtual link. The internal IP address is configured when the switch is installed. To determine the internal IP address, access the switch console and issue the **`show system`** command. In the command output, the internal IP address appears in the Internal IP Addr field. For example: **Internal IP Addr: 150.202.77.2** In this example, the internal IP address is 150.202.77.2.

4. Choose OK.
5. At the Set All VNN Virtual Links dialog box, choose Close.

Configuring Multiple IP OSPF and VNN OSPF Areas

Configuring multiple IP OSPF areas and VNN OSPF areas consists of the following procedures:

- OSPF Area Configuration
- Virtual Link Configuration
- Address Aggregation

See “Steps for Configuring Multiple OSPF Areas” on page 9-54 for more information about each of these procedures.

▶

> Area 0 is the OSPF backbone area. Areas do not have to be physically attached to the backbone. Instead, virtual links can be configured to logically attach an area to the backbone.
>
> Every area border router must be connected to the backbone area. You use area 0 trunks or configured virtual links to connect each area border router to the backbone area.

Steps for Configuring Multiple OSPF Areas

The following sections outline each of the steps for OSPF Area configuration.

Prerequisites for Multiple OSPF Areas

To configure multiple IP OSPF areas, the IP logical ports and associated IP interfaces must be defined. See “Configuring IP Logical Ports” on page 3-5 for details. To configure multiple VNN OSPF areas, no special prerequisites must be met.

Configuration Recommendations

Be aware of the following recommendations when configuring multiple IP OSPF areas or multiple VNN OSPF areas:

- *Do not make areas too small.*
  - Area boundaries may be difficult to configure, and can cause sub-optimal routing for both IP and circuits.
  - If every switch is an area border router, there will be no improvements to the route scaling.
- *Plan ahead when assigning areas.* Modification of the trunk Area ID is a service-affecting procedure that causes the trunk to bounce.
- *Do not unnecessarily aggregate IP addresses.* Modifying switch IP addresses is time-consuming and should not be done for the sole purpose of aggregation.

IP OSPF Area Configuration

To configure multiple IP OSPF areas:

1. Set the IP OSPF Area ID of all trunks. From the Administer menu, select Lucent Parameters ⇒ Set All Trunks. The NMS displays the NavisCore Set All Trunks dialog box. See "Configuring Trunks for IP OSPF and VNN OSPF" on page 9-25 for more information.
2. Set the Area ID of the IP logical ports:
   a. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears.
   b. Choose IP Parameters. The Set IP Parameters dialog box appears.
   c. Choose IP Interface. The Set IP Interface Addresses dialog box appears.
   d. Choose Add OSPF. The Add OSPF Interface dialog box appears. See "Configuring IP OSPF on the IP Logical Port" on page 9-30 for more information about how to complete this dialog box.
3. Set the Area ID of the loopback addresses. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP Loopback Addresses. The Set All IP Loopback Addresses dialog box appears. See "Configuring IP Loopback Addresses" on page 8-31 for more information.

VNN OSPF Area Configuration

To configure multiple VNN OSPF areas:

1. Set the VNN OSPF Area ID of all trunks. From the Administer menu, select Lucent Parameters ⇒ Set All Trunks. The NMS displays the NavisCore Set All Trunks dialog box. See "Configuring Trunks for IP OSPF and VNN OSPF" on page 9-25 for more information.
2. Set the Area ID of the Network Service Access Points (NSAPs). From the Administer menu, select Lucent Parameters ⇒ Set All SVC Parameters ⇒ Set All Node Prefixes. The Set All Node Prefixes dialog box appears. See the *NavisCore Frame Relay Configuration Guide* or the *NavisCore ATM Configuration Guide* for more information.

> The Area ID of the switch IP address is set automatically to one of the following:
>
> – Area 1, if it exists.
>
> – If Area 1 does not exist, the Area ID is set to the ID of the switch with the lowest Area ID.

3. Set the Area ID of the loopback addresses. From the Administer menu, select Lucent Parameters ⇒ Set All VNN ⇒ Set All VNN Loopbacks. See "Configuring VNN Loopback Addresses" on page 9-49 for more information.

Virtual Link Configuration

Areas do not have to be physically attached to the backbone. Instead, you can configure virtual links to logically attach an area to the backbone. Before you can configure a virtual link, you must know:

- The non-backbone transit area for the virtual link.
- The router IDs for the two endpoint switches. (The router ID of each switch endpoint is the same value as the switch ID).

In addition, you must ensure that both switches have IP addresses in the transit area. Add loopback addresses if necessary.

To add virtual links for IP OSPF and VNN OSPF, see the following sections:

- "Configuring IP OSPF Virtual Links" on page 9-38
- "Configuring VNN Virtual Links" on page 9-52

Address Aggregation

Aggregation is the process of advertising a single address prefix (rather than advertising multiple, more specific prefixes). Addresses can be aggregated at area borders. This practice further improves route scaling by reducing the size of the link-state database and the routing table.

You configure aggregates in the area border router.

VNN OSPF Address Aggregation

You can configure VNN OSPF area aggregates. See "Configuring VNN Area Aggregates" on page 9-50 for more information.

You can also configure an aggregate for an NSAP. From the Administer menu, select Lucent Parameters ⇒ Set All SVC Parameters ⇒ Set All Node Prefixes. The Set All Node Prefixes dialog box appears. See the *NavisCore Frame Relay Configuration Guide* and the *NavisCore ATM Configuration Guide* for more information.

IP OSPF Address Aggregation

You can configure IP OSPF area aggregates. See "Configuring IP OSPF Area Aggregates" on page 9-36 for more information.

Configuring Route Maps for IP OSPF and VNN OSPF

Because of the separation of VNN OSPF and IP OSPF, you can now use VNN OSPF as a route map source and destination. For example, you can configure a route map to advertise management routes (which are learned through VNN OSPF) to BGP, RIP, or IP OSPF.

▶ VNN OSPF is identified as “VNN” in the route map dialog boxes, and IP OSPF is identified as “OSPF.”

As a source, VNN OSPF can be used in the following combinations:

- VNN ⇒ BGP
- VNN ⇒ RIP
- VNN ⇒ OSPF

As a destination, VNN OSPF can be used in the following combinations:

- BGP ⇒ VNN
- RIP ⇒ VNN
- OSPF ⇒ VNN
- Static ⇒ VNN
- Direct ⇒ VNN
- Any ⇒ VNN
- Aggregate ⇒ VNN

For more information on configuring route maps, see Chapter 11, “Configuring Route Policies.”

10

Configuring Static Routes

This chapter describes how to configure static routes.

About Static Routes

You configure static routes manually only if they are reachable. Static routes do not disappear from the IP routing table and will always be advertised. However, static routes do not respond to network topology changes. The only way a static route can change is if the network administrator changes them. In addition, static routes provide redundancy if a primary connection fails.

Configuring a Static Route

To configure a static route:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Static Routes. The Set All Static Routes dialog box appears (see Figure 10-1).

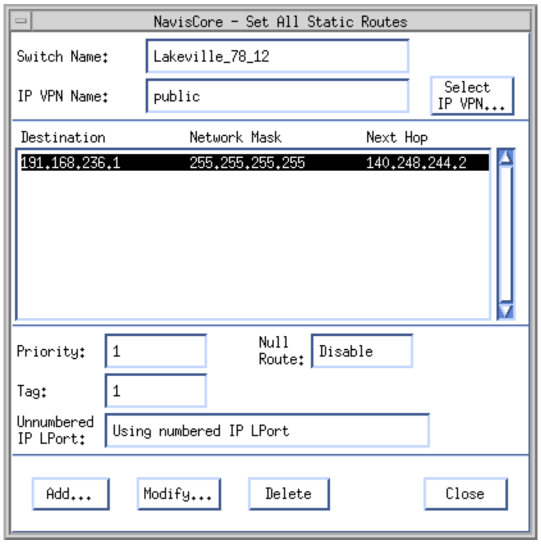

NavisCore - Set All Static Routes

Switch Name: Lakeville_78_12

IP VPN Name: public

Destination	Network Mask	Next Hop
191.168.236.1	255.255.255.255	140.248.244.2

Priority: 1

Null Route: Disable

Tag: 1

Unnumbered IP LPort: Using numbered IP LPort

Figure 10-1. Set All Static Routes Dialog Box

The Set All Static Routes dialog box displays the following buttons.

Table 10-1. Set All Static Routes Buttons

Button	Function
Select IP VPN	Enables you to select an IP VPN. Once you select the IP VPN, all static routes you configure are for the selected IP VPN only. By default, all static routes you configure are public network resources. See "Selecting the IP VPN" on page 16-33 for more information on selecting an IP VPN.
Add	Enables you to add a static route.
Modify	Enables you to modify a static route.
Delete	Enables you to delete a static route.

> If you are creating a static route for an IP VPN, make sure you are in the context of that IP VPN. For example, if you are adding a static route for an IP VPN called "IPVPN1," make sure you are in the context of IPVPN1. To enter the context of an IP VPN, choose the Select IP VPN button and select the appropriate IP VPN.

3. Choose Add. The Set Static Route dialog box appears (see Figure 10-2).

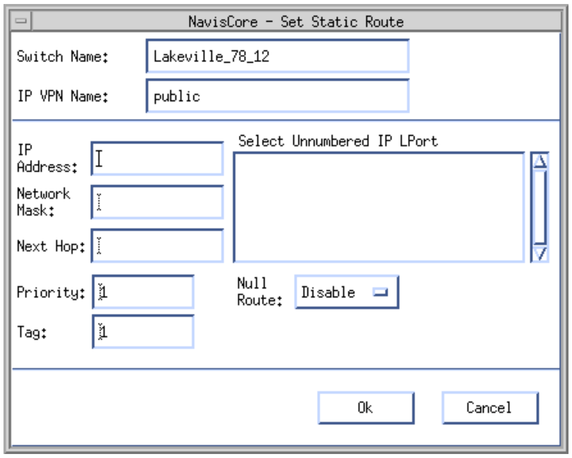
NavisCore - Set Static Route

Switch Name: Lakeville_78_12

IP VPN Name: public

IP Address:

Network Mask:

Next Hop:

Priority: 1

Tag: 1

Null Route: Disable

Figure 10-2. Set Static Route Dialog Box

4. Complete the fields described in Table 10-2.

Table 10-2. Static Route Fields

Field	Description
IP Address	Enter the IP address of the destination network.
Network Mask	Enter the network mask.
Next Hop	Enter the IP address of the next hop. The next hop field is disabled if you: – Selected an unnumbered IP logical port (see “Select Unnumbered IP LPort”), or – Enabled null route (see “Null Route”)
Priority	Enter a value from 1 to 20 to specify the static route priority. The highest number is the preferred priority. The priority of the static route is in relation to other route protocols. If you assign equal priority to multiple static routes to the same destination, IP traffic is load balanced across those routes. (This is called equal-cost multipath routing.) Keep in mind that routes with a higher priority will still take precedence over routes with a lower priority. If you want to always load balance traffic across multiple static routes to a given destination, do the following: – Assign these routes the same priority. – Make sure that routes with a higher priority do not exist.
Tag	Enter the tag value, which you use to group multiple static route entries together.
Null Route	Select one of the following: *Enable* – If you enable this parameter, packets destined for this network will be discarded. In addition, the next hop is disabled. *Disable* – If you disable this parameter, packets destined for this network will be forwarded.
Select Unnumbered IP LPort	Select an unnumbered IP logical port to set up a static route to an IP interface that is not part of a subnet and does not have a specific address. Instead, the unnumbered IP logical port uses the router ID as its address.

5. Choose OK.
6. At the Static Route dialog box, choose Close.

Configuring Route Policies

This chapter describes the following configuration tasks:

- Adding a Network Filter
- Adding a Network Access List
- Adding a Route Map

You perform these tasks in order to set *route policies* for your network. These policies control the flow of routing information.

The *route map* provides the primary mechanism for setting route policies. The purpose of a route map is to control and modify routing information and to define the parameters that your system uses to redistribute routes between routing domains. Route maps are used to alter route parameters that are then stored in the routing table, or sent via routing updates to other routers.

You can optionally define the following components for use in a route map:

- Network filters
- Network access lists

After you define the route map, you must assign it to a neighbor (in the case of BGP or RIP). If you have multiple route maps for the same neighbor, you can specify the order in which IP Navigator uses these maps.

The following sections define the concepts for using network filters, network access lists, and route maps. In addition, these sections describe how to use route maps to redistribute routes between routing domains.

About Network Filters

Network filters control the flow of route distribution. You can use a network filter to select routes that will be accepted or rejected by route maps. The specified filters must be used in a network access list and then applied to route maps.

When you create a network filter, you specify the following information:

- A network address
- A network mask value
- Coverage (inclusive or exact)

The network address and network mask value identify the route. The coverage specifies the type of access. *Inclusive* filters allow access to all networks that match the specified network address (including addresses that may be more specific such as a subnetwork address). *Exact* filters allow access only to the network that is specified in the network address.

A network filter is an optional component of a route map; however, if you want to use one or more network filters, you must include the filter in an access list and then include the access list in a route map. The route map must then be assigned to the appropriate neighbor or interface. A network filter by itself cannot be applied to a route map, neighbor, or interface.

About Access Lists

A network filter access list is an object that contains a set of unique network filters. Up to 300 network filters can be included in an access list.

You can create an empty network access list and later add defined network filters to the list. You use network access lists to logically group network filters.

A network access list is an optional component of a route map; however, if you want to use one or more network access lists, you must include the list in a route map and then assign the map to the appropriate neighbor or interface. A network access list by itself cannot be applied to a neighbor or interface.

About Route Maps

Route maps enable you to specify the direction of route traffic based on the source of the traffic or a combination of both the traffic source and destination. You can enable or disable a route map as required by setting the Admin Status value for the route map.

When you create a route map, you specify two routing protocols: a From Protocol and a To Protocol. The route map specifies how routes are redistributed from one routing protocol to another. This is done between two different protocols as well as within the same protocol (for example, from BGP to BGP). The route maps are also used to selectively accept routes from a particular routing protocol into the router's main routing table.

In addition, you can optionally specify the following values as route map match or set parameters:

- Metric value
- Tag value
- Next hop address
- Autonomous System path values (BGP only)
- Community values (BGP only)
- OSPF route type

Route Map From and To Choices

Each time you define a route map, you must specify a From and a To choice to specify the two protocols used for route redistribution. The protocols that you specify govern the direction (import or export) as well as the set of affected routes.

The To choices that you can select vary depending upon the previously selected From choices. For example, the routing table option can only be selected if the From choice protocol is BGP, RIP, MOSPF, and DVMRP.

Table 11-1 lists all of the route map From and To choices.

Table 11-1. Route Map From and To Choices

Choice	From	To
BGP	Yes	Yes
OSPF (IP OSPF instance)	Yes	Yes
RIP	Yes	Yes
Static	Yes	No
Routing Table	No	Yes
Direct	Yes	No
Aggregate	Yes	No
Any	Yes	No
VNN (VNN OSPF instance)	Yes	Yes
MOSPF	Yes	No
DVMRP	Yes	Yes

The Any option enables you to select routes from the routing table regardless of the origin protocol. For example, you could select a specific route from the routing table and then advertise that route to BGP. The protocol used to transport the route to the routing table is not important.

Determining if a Route Map is for Import or Export

The protocol that you select for the To choice specifies whether a map is an import or export map as follows:

- Route maps that use a To choice of BGP, OSPF, or RIP are automatically created as export route maps.
- Route maps that use a To choice of Routing Table are created as import route maps.
- All route selections for a route map that uses a To choice of "Routing Table" are performed before IP Navigator adds the routes to the routing table.

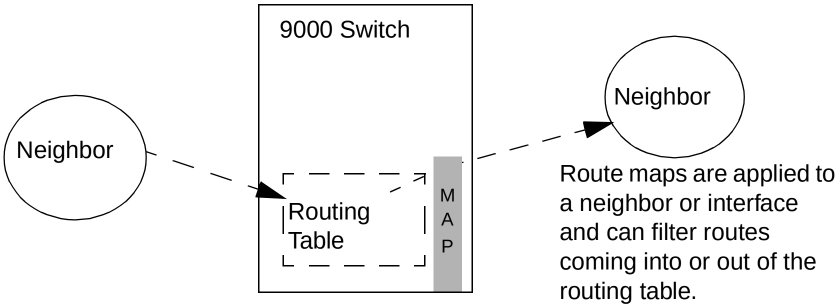

Route maps are applied to a neighbor or interface and can filter routes coming into or out of the routing table.

Figure 11-1. Using Route Maps to Filter Routes

Route Map Guidelines

Route maps are required if you want to accomplish any of the following tasks:

- Route filtering
- Route redistribution
- Altering route parameters such as metric, next hop, tag, and BGP path attributes.

See Figure 11-2 and the sections that follow for a description of the guidelines for route map use.

When are Route Maps Not Used?

You cannot use a route map to specify a routing policy for the following pairs:

OSPF to Routing Table — IP Navigator always adds IP OSPF routes to the routing table. For this reason, you cannot use a route map to determine the acceptance or rejection of specific routes between IP OSPF and the routing table.

VNN to Routing Table — IP Navigator always adds VNN OSPF routes to the routing table. For this reason, you cannot use a route map to determine the acceptance or rejection of specific routes between VNN OSPF and the routing table.

OSPF to OSPF — IP Navigator always advertises IP OSPF routes to the IP OSPF routing domain. Link state protocols assume that all routes share the same information. For this reason, you cannot use a route map to determine the acceptance or rejection of specific routes being sent to an IP OSPF neighbor.

VNN to VNN — IP Navigator always advertises VNN OSPF routes to the VNN OSPF routing domain. Link state protocols assume that all routes share the same information. For this reason, you cannot use a route map to determine the acceptance or rejection of specific routes being sent to a VNN OSPF neighbor.

Figure 11-2 illustrates the logical flow of routing information through the switch and where route maps can optionally be applied in this flow.

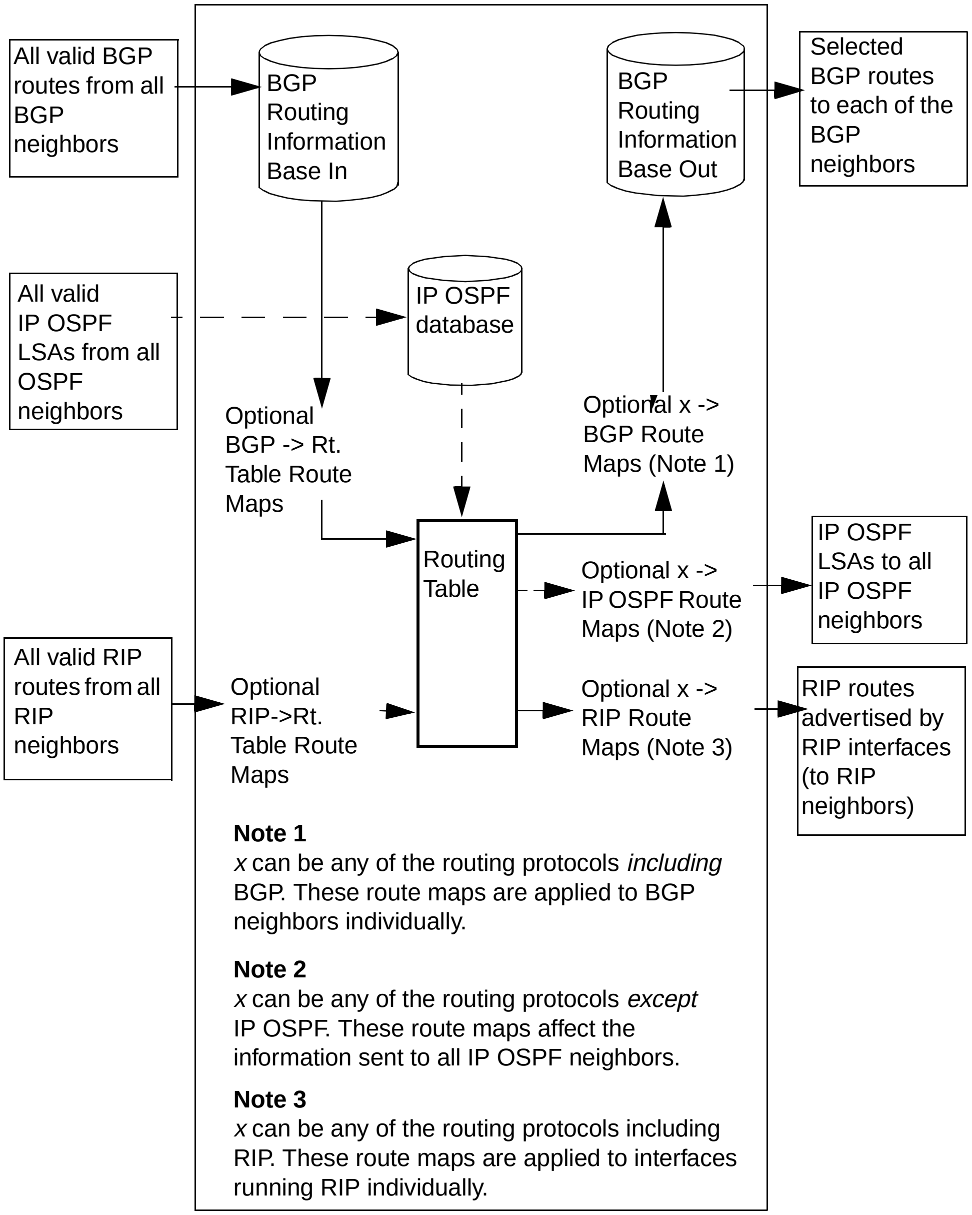
All valid BGP routes from all BGP neighbors
BGP Routing Information Base In
BGP Routing Information Base Out
Selected BGP routes to each of the BGP neighbors
All valid IP OSPF LSAs from all OSPF neighbors
IP OSPF database
Optional BGP -> Rt. Table Route Maps
Optional x -> BGP Route Maps (Note 1)
Routing Table
IP OSPF LSAs to all IP OSPF neighbors
Optional x -> IP OSPF Route Maps (Note 2)
All valid RIP routes from all RIP neighbors
Optional RIP->Rt. Table Route Maps
Optional x -> RIP Route Maps (Note 3)
RIP routes advertised by RIP interfaces (to RIP neighbors)
Note 1
x can be any of the routing protocols including BGP. These route maps are applied to BGP neighbors individually.
Note 2
x can be any of the routing protocols except IP OSPF. These route maps affect the information sent to all IP OSPF neighbors.
Note 3
x can be any of the routing protocols including RIP. These route maps are applied to interfaces running RIP individually.

Figure 11-2. Flow of Routing Information Through the Switch

What Happens if You Do Not Use a Route Map?

If you do not use a route map for route filtering or route redistribution, the following import and export operations occur by default:

- Routes from all protocols, except for EBGP, are imported into the routing table by default.
- EBGP routes are not imported into the routing table by default for security reasons. You must specify a route map and optionally specify an access list containing any EBGP routes that you may want to import into the routing table.
- All RIP routes are exported to any RIP interface addresses that are configured for the IP interface.

Protocol Pairs That Do Not Require Route Maps

Route maps are not required for each of the following protocol pairs:

- IBGP Peer ⇒ Routing Table
- BGP ⇒ BGP
- RIP ⇒ Routing Table
- RIP ⇒ RIP

Protocol Pairs That Require Route Maps

Route maps are required for each of the following protocol pairs:

Table 11-2. Protocol Pair Route Map Requirements

Protocol Pair	Description
Static ⇒ OSPF Direct ⇒ OSPF BGP ⇒ OSPF RIP ⇒ OSPF VNN ⇒ OSPF	Route maps are required to advertise any Static, Direct, BGP, RIP, and VNN OSPF routes into the IP OSPF routing domain. By default, IP Navigator does not advertise Static, direct, BGP, RIP, or VNN OSPF routes into the IP OSPF routing domain. ***Note***: Switches perform some automatic route leaking for you between VNN OSPF and IP OSPF. For example, in a Lucent network, when some switches are upgraded to support separate OSPF network views for VNN and IP Navigator, and other switches are not upgraded, routes are automatically leaked as follows: • Routes learned by IP OSPF are leaked into VNN OSPF. • Non-management routes learned by VNN OSPF are leaked into IP OSPF. Management routes remain hidden from IP OSPF.
Static ⇒ VNN Direct ⇒ VNN BGP ⇒ VNN RIP ⇒ VNN OSPF ⇒ VNN	Route maps are required in order to advertise any Static, Direct, BGP, RIP, and IP OSPF routes into the VNN OSPF routing domain. By default, IP Navigator does not advertise Static, direct, BGP, RIP, or IP OSPF routes into the VNN OSPF routing domain. ***Note***: Switches perform some automatic route leaking for you between VNN OSPF and IP OSPF. For example, in a Lucent network, when some switches are upgraded to support separate OSPF network views for VNN and IP Navigator, and other switches are not upgraded, routes are automatically leaked as follows: • Routes learned by IP OSPF are leaked into VNN OSPF. • Non-management routes learned by VNN OSPF are leaked into IP OSPF. Management routes remain hidden from IP OSPF.
Static ⇒ RIP Direct ⇒ RIP BGP ⇒ RIP OSPF ⇒ RIP VNN ⇒ RIP	Route maps are required to advertise any Static, Direct, BGP, VNN OSPF, and IP OSPF routes into the RIP routing domain. By default, IP Navigator does not advertise Static, Direct, BGP, VNN OSPF, and IP OSPF routes into the RIP routing domain.
Static ⇒ BGP Direct ⇒ BGP OSPF ⇒ BGP RIP ⇒ BGP VNN ⇒ BGP	Route maps are required to advertise any Static, Direct, BGP, RIP, and OSPF routes into the BGP routing domain. By default, IP Navigator does not advertise Static, Direct, BGP, RIP, VNN OSPF, and IP OSPF routes into the BGP routing domain.
MOSPF ⇒ DVMRP	Route maps are required to export any routes learned by MOSPF into the DVMRP routing domain.

Table 11-2. Protocol Pair Route Map Requirements (Continued)

Protocol Pair	Description
BGP ⇒ Routing Table	Route maps are required to install any routes advertised by neighboring EBGP peers into the main routing table. By default, IP Navigator does not install EBGP routes into the main routing table. IBGP routes are installed into the routing table even if there are no route maps.

► IP Navigator applies multiple route maps using first match logic. This means that, as each route map is applied, any matching route entries are accepted or rejected immediately. *Subsequent route maps cannot consider the route entries that were already accepted or rejected.* For this reason, you should arrange the sequence of multiple route maps so that the *more specific matches are first in the list*.

► Route maps that use a To choice of VNN OSPF or DVMRP are automatically created as export route maps.

Steps For Configuring a Route Map

To configure a route map, use the following steps:

1. *(Optional)* Define the network filters depending on your system's needs. See "Adding a Network Filter" on page 11-13 for more information.
2. *(Optional)* Use the defined network filters to create the network access lists. See "Adding a Network Access List" on page 11-15 for more information.
3. Specify the routing policies that define the match parameters to be used to filter routes and the set parameters for all selected routes. See "Adding Route Maps" on page 11-18 for more information.
4. Assign the route map to a BGP neighbor or a RIP interface. You assign route maps to BGP interfaces using the Add BGP Neighbor dialog box or the Modify BGP Neighbor dialog box. See Chapter 8, "Configuring BGP Parameters" for more information about accessing the BGP functions. You assign route maps to RIP interfaces using the Add RIP Interface dialog box or the Modify RIP Interface dialog box. See Chapter 7, "Configuring RIP" for more information about accessing the RIP functions.
5. If you have multiple route maps, use the arrow buttons (see Figure 11-3) on the Add/Modify BGP Neighbor and Add/Modify RIP Interface dialog boxes to specify the order in which IP Navigator uses the assigned route maps. Route maps filter routes on the interface in the order in which they are specified on these dialog boxes. Route maps should be ordered from *most specific* to least *specific*.

▶ Route maps that have a To protocol of OSPF are global and for this reason do not need to be assigned to an OSPF interface. IP Navigator uses this type of route map as soon as you create the map.

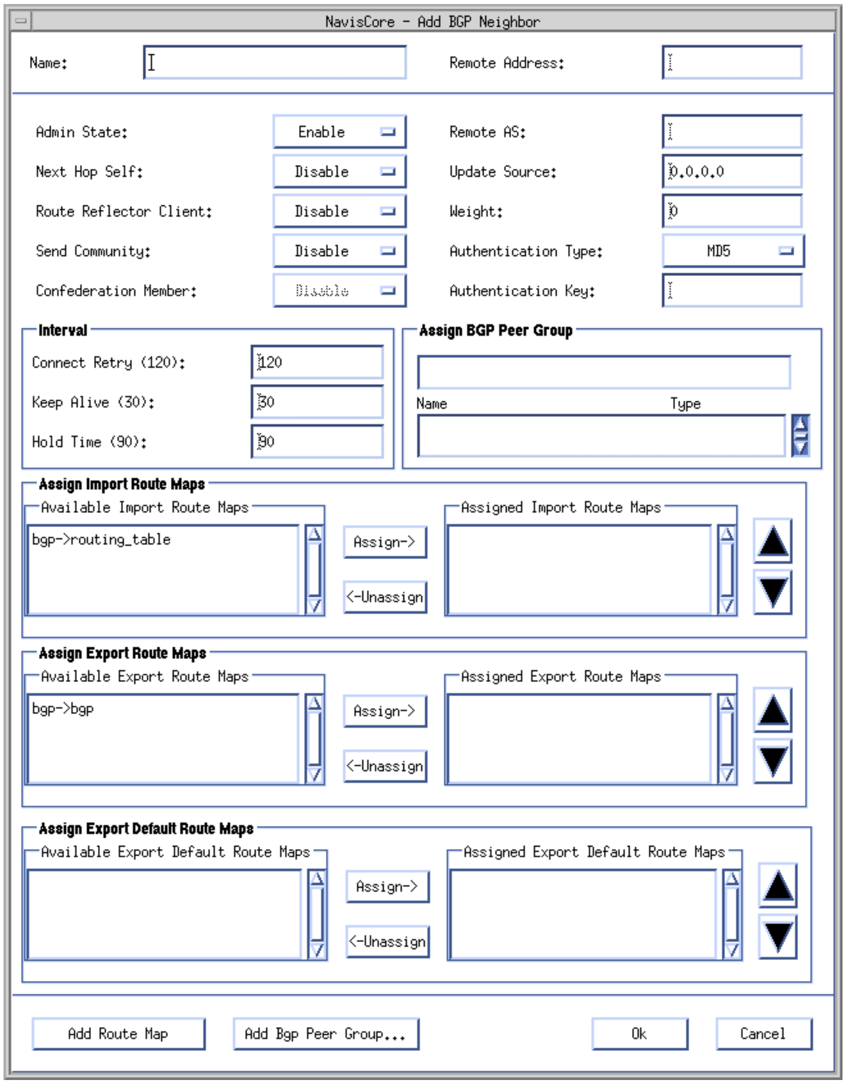

NavisCore - Add BGP Neighbor

Name:

Remote Address:

Admin State: Enable

Next Hop Self: Disable

Route Reflector Client: Disable

Send Community: Disable

Confederation Member: Disable

Remote AS:

Update Source: 0.0.0.0

Weight: 0

Authentication Type: MD5

Authentication Key:

Interval

Connect Retry (120): 120

Keep Alive (30): 30

Hold Time (90): 90

Assign BGP Peer Group

Name Type

Assign Import Route Maps

Available Import Route Maps: bgp->routing_table

Assigned Import Route Maps:

Assign Export Route Maps

Available Export Route Maps: bgp->bgp

Assigned Export Route Maps:

Assign Export Default Route Maps

Available Export Default Route Maps:

Assigned Export Default Route Maps:

Use the arrow buttons to specify the sequence that IP Navigator uses to filter routing information.

Figure 11-3. Using the Arrow Buttons to Sequence Route Maps

Adding a Network Filter

To add a network filter:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Route Policies ⇒ Set All Network Filters. The Set All Network Filters dialog box appears (see Figure 11-4).

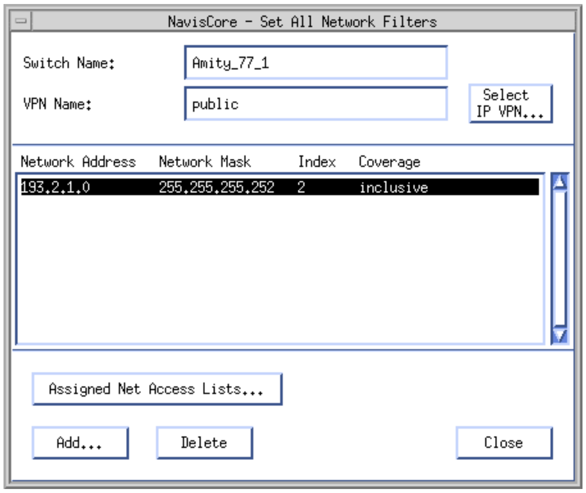

Switch Name: Amity_77_1
VPN Name: public

Network Address	Network Mask	Index	Coverage
193.2.1.0	255.255.255.252	2	inclusive

Figure 11-4. Set All Network Filters Dialog Box

Table 11-3 describes each of the Set All Network Filters buttons.

Table 11-3. Set All Network Filters Buttons

Button	Function
Select IP VPN	Allows you to select the IP VPN to which you want to assign the network filters. By selecting an IP VPN, you enter the context of that VPN. For more information on selecting an IP VPN, see "Selecting the IP VPN" on page 16-33.
Assigned Net Access Lists	Displays any network access lists that use the selected filter.
Add	Displays the Add Network Filter dialog box to enable you to add a network filter.
Delete	Displays the Delete Network Filter dialog box to enable you to delete a network filter.

3. Choose Add. The Add Network Filter dialog box appears (see Figure 11-5).

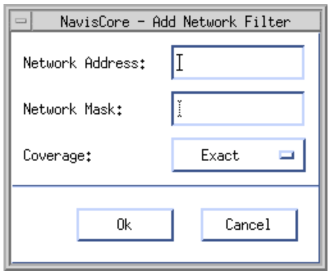

Figure 11-5. Add Network Filter Dialog Box

4. Specify the necessary network filter values listed in Table 11-4.

Table 11-4. Network Filter Fields

Field	Action/Description
Network Address	Specify the network address for this filter. For example, 0.0.0.0 specifies all network addresses.
Network Mask	Specify the network mask for this filter.
Coverage	Specify *inclusive* to allow all networks that match the specified network address (including addresses that may be more specific such as subnetwork addresses*)*. Specify *exact* to allow only the network that is specified in the network address and the network mask.

Adding a Network Access List

A network access list enables you to logically group a set of network filters. To add a network access list:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Route Policies ⇒ Set All Network Access Lists. The Set All Network Access Lists dialog box appears (see Figure 11-6).

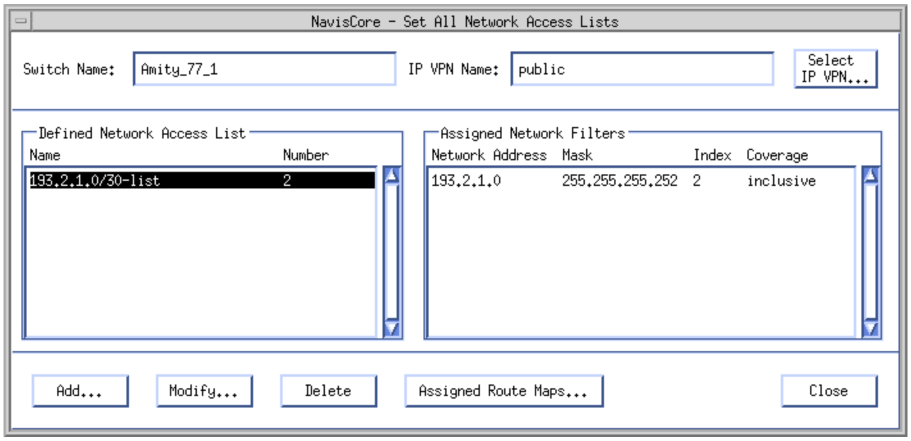

NavisCore - Set All Network Access Lists

Switch Name: Amity_77_1

IP VPN Name: public

Defined Network Access List

Name	Number
193.2.1.0/30-list	2

Assigned Network Filters

Network Address	Mask	Index	Coverage
193.2.1.0	255.255.255.252	2	inclusive

Figure 11-6. Set All Network Access Lists Dialog Box

Table 11-5 describes each of the Set All Network Access Lists buttons.

Table 11-5. Set All Network Access List Buttons

Button	Function
Select IP VPN	Allows you to select the IP VPN to which you want to assign the network access list. By selecting an IP VPN, you enter the context of that VPN. For more information on selecting an IP VPN, see "Selecting the IP VPN" on page 16-33.
Add	Displays the Add Network Access List dialog box to enable you to add a network access list.
Modify	Displays the Modify Network Access List dialog box to enable you to modify a selected network access list.
Delete	Displays the Delete Network Access List dialog box to enable you to delete a selected network access list.
Assigned Route Maps	Displays route maps that use a selected network access list.

3. Choose Add. The Add Network Access List dialog box appears (see Figure 11-7).

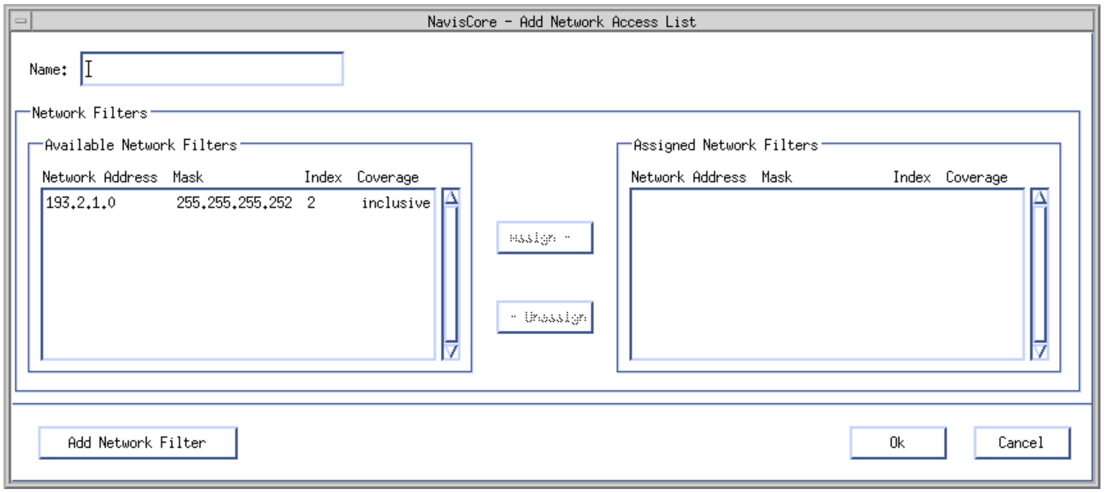

Network Address	Mask	Index	Coverage
193.2.1.0	255.255.255.252	2	inclusive

Figure 11-7. Add Network Access List Dialog Box

4. Specify a unique network access list name.

5. Use the Assign and Unassign buttons to specify the network filters that you want to include in the network access list. See Table 11-6 on page 11-17 for a description of each of the fields on the Add Network Access List dialog box.

6. To add a filter to the list of Available Network Filters, choose Add Network Filter to display the Add Network Filter dialog box shown in Figure 11-5 on page 11-14. Any filters that you add are included in either the list of available network filters or the list of assigned network filters.

7. Choose OK after the Assigned Network Filters list includes all of the filters that you want to use in the network access list. One network access list can include up to 300 network filters.

Table 11-6. Network Access List Fields

Field	Action/Description
Name	Specify a unique network access list name.
Available Network Filters	A list of filters that you can add to the network access list.
Network Address	The network address for the filter.
Mask	The network mask for the filter.
Index	The index field is generated by NavisCore and is unique within the switch. This field is for internal system use only and cannot be modified.
Coverage	*Inclusive* allows all networks that match the specified network address (including addresses that may be more specific). *Exact* allows only the network that is specified in the network address.
Assigned Network Filters	A list of network filters that are currently included in the network access list. Up to 300 filters can be included in the access list.
Network Address	The network address for the filter.
Mask	The network mask for the filter.
Index	The index field is generated by NavisCore and is unique within the switch. This field is for internal system use only and cannot be modified.
Coverage	*Inclusive* allows all networks that match the specified network address (including addresses that may be more specific). *Exact* allows only the network that is specified in the network address.

Adding Route Maps

To add a route map:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set All Route Policies ⇒ Set All Route Maps. The Set All Route Maps dialog box appears (see Figure 11-8).

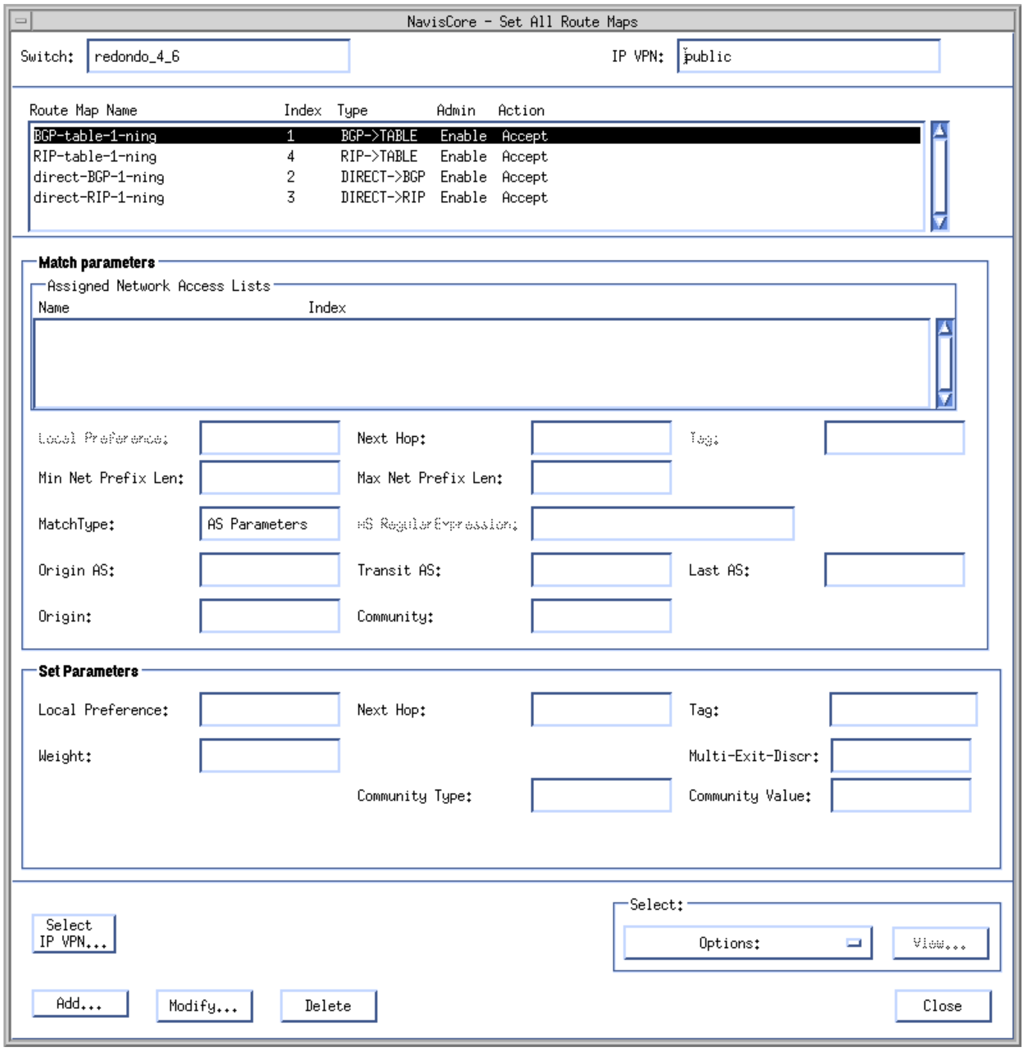

NavisCore - Set All Route Maps

Switch: redondo_4_6

IP VPN: public

Route Map Name	Index	Type	Admin	Action
BGP-table-1-ning	1	BGP->TABLE	Enable	Accept
RIP-table-1-ning	4	RIP->TABLE	Enable	Accept
direct-BGP-1-ning	2	DIRECT->BGP	Enable	Accept
direct-RIP-1-ning	3	DIRECT->RIP	Enable	Accept

MatchType: AS Parameters

Figure 11-8. Set All Route Maps Dialog Box

Table 11-7 describes each of the Set All Route Maps buttons. Table 11-8 on page 11-20 describes the fields at the top of the Set All Route Maps dialog box.

The Match Parameters and Set Parameters on the Set All Route Map dialog box vary depending on the type of route map that you are defining. See Table 11-11 on page 11-23 for a reference to the section of this chapter that describes the Match and Set parameters for each route map type.

Table 11-7. Set All Route Maps Buttons

Button	Function
Select IP VPN	Allows you to select the IP VPN to which you want to assign the network access list. By selecting an IP VPN, you enter the context of that VPN. For more information on selecting an IP VPN, see "Selecting the IP VPN" on page 16-33.
Add	Displays the Add Route Map dialog box to enable you to add a route policy.
Modify	Displays the Modify Route Map dialog box to enable you to modify a route policy.
Delete	Displays the Delete Route Map dialog box to enable you to delete a route policy.
Options	Use the Select: Options button to select one of the following options. **Assigned BGP Neighbors** — Lists all BGP neighbors that use a selected route map. **Assigned BGP Peer Groups –** Displays the assigned BGP peer groups. **Assigned RIP Interfaces** — Lists all RIP interfaces that use a selected route map. **BGP Neighbors** — Displays the Set All BGP Neighbors dialog box to enable you to assign a route map to a BGP neighbor. **OSPF Route Maps Sequence** — Displays the Change the Order of OSPF Route Maps dialog box to enable you to change the sequence of assigned route maps. **Show Bgp Route Dampening** – Displays the BGP route dampening configuration.

Table 11-8. Set All Route Maps Common Values

Field	Action/Description
Switch Name	Displays the name of the currently selected switch.
Route Map Name	A name that uniquely identifies the route map.
Index	The index field is generated by NavisCore and is unique within the switch. This field is for internal system use only and cannot be modified.
Type	Displays the From protocol and To protocol that identify the route distribution type. See Table 11-11 on page 11-23 for a list of route distribution types and a reference to the section of this chapter that describes how to redistribute routes between various routing protocols.
Admin	Specify Enable or Disable. *Enable* indicates that the route map is administratively enabled and can be used. *Disable* indicates that the route map is administratively disabled and cannot be used.
Action	Specify Accept, Deny, or Originate Default. *Accept* – indicates that all routes that match the specified Match parameters are accepted. *Deny* – indicates that all routes that match the specified Match parameters are denied. *Originate Default* – indicates that you can specify the match parameters that define where to send a default route heading. This option is used for the following types of route maps: BGP to BGP, ANY to BGP, or RIP to RIP.

3. Choose Add. The Add Route Map dialog box appears (see Figure 11-9).

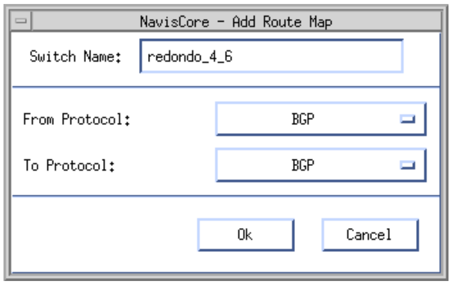
NavisCore - Add Route Map

Switch Name: redondo_4_6

From Protocol: BGP

To Protocol: BGP

Figure 11-9. Add Route Map Dialog Box

4. Specify the values listed in Table 11-9.

Table 11-9. Route Map Descriptions

Field	Action/Description
Switch Name	Displays the name of the currently selected switch.
From Protocol	Specify one of the following values: BGP, OSPF, RIP, STATIC, Direct, Aggregate, ANY, VNN, MOSPF, or DVMRP.
To Protocol	Specify one of the following values: BGP, OSPF, RIP, Routing Table, VNN, or DVMRP. The routing table option can only be selected if the From protocol is BGP, RIP, MOSPF, or DVMRP.

▶ If you configure a route map and specify ANY or DIRECT as the From Protocol, make sure that you also configure an access list that selects only those routes that you want to include as export routes.

5. Choose OK. The system displays a dialog box similar to the one shown in Figure 11-10.

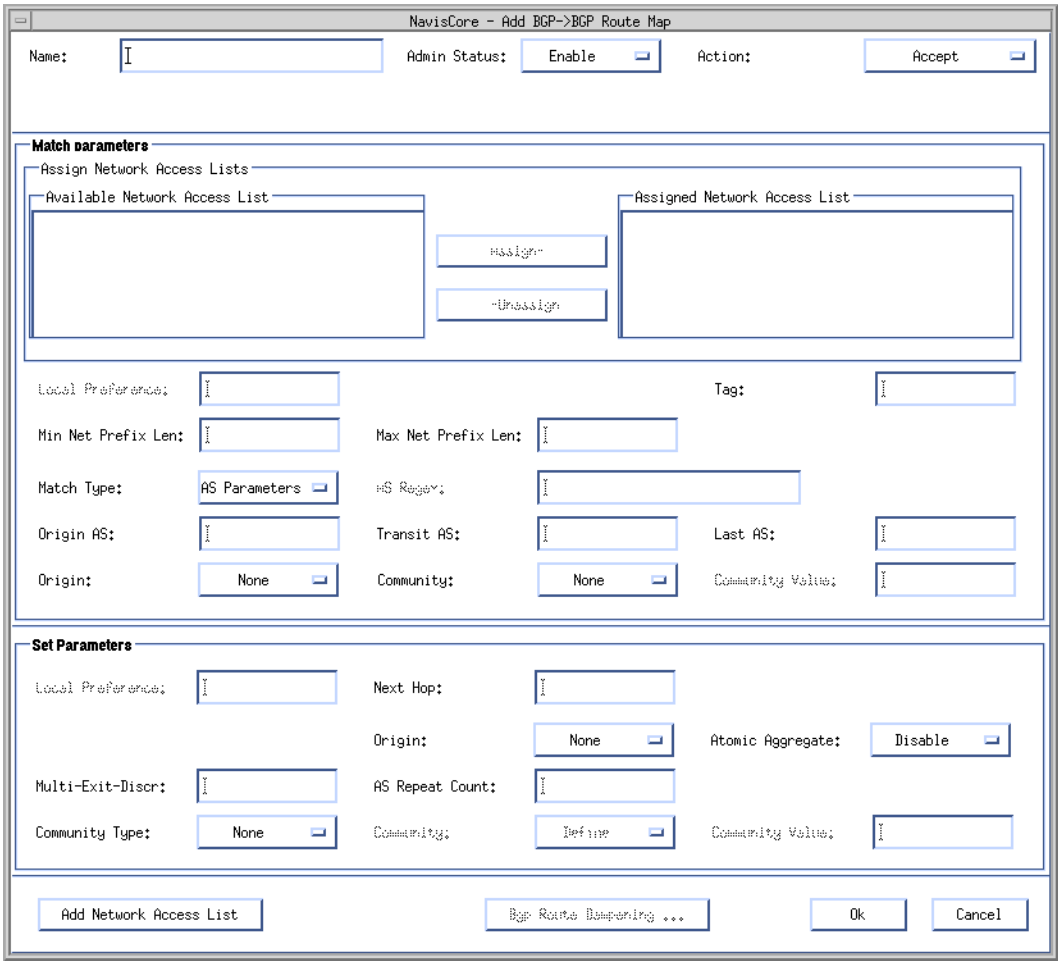

Admin Status: Enable

Action: Accept

Match Type: AS Parameters

Origin: None

Community: None

Origin: None

Atomic Aggregate: Disable

Community Type: None

Community: Define

Figure 11-10. Second Add Route Map Dialog Box

6. Specify the Route Map Name, Admin Status, and Action values as described in Table 11-10.

7. Specify the necessary match and set parameters for this route map. If you need to add an access list to the route map, choose Add Access Lists. Instructions for adding access lists start on page 11-15.

   The Match parameters and Set parameters on the Add Route Map dialog box vary depending on the type of route map that you are defining. See Table 11-11 for a reference to the section of this chapter that describes the Match and Set parameters for each route map type.

▶ All of the Match and Set Parameter fields described in Table 11-12 through Table 11-31 are optional. You can specify a routing policy that uses no match and no set values.

Table 11-10. Add Route Map Fields

Field	Action/Description
Route Map Name	Specify a unique name to identify the route map.
Admin Status	Specify Enable or Disable. *Enable* indicates that the route map is administratively enabled and can be used. *Disable* indicates that the route map is administratively disabled and cannot be used.
Action	Specify Accept, Deny, or Originate Default. *Accept* – Indicates that all routes that match the specified Match parameters are accepted. *Deny* – Indicates that all routes that match the specified Match parameters are denied. *Originate Default* – Indicates that you can specify the match parameters that define where to send a default route heading. This option is used for the following types of route maps: BGP to BGP, ANY to BGP, or RIP to RIP.

Table 11-11. Match and Set Parameter Descriptions

Route Map Type	See...
BGP to BGP	Table 11-12 on page 11-25
BGP to OSPF	Table 11-13 on page 11-28
BGP to RIP	Table 11-14 on page 11-30
BGP to Routing Table	Table 11-15 on page 11-32
BGP to VNN	Table 11-16 on page 11-34
OSPF to BGP	Table 11-17 on page 11-36
OSPF to RIP	Table 11-18 on page 11-38
OSPF to VNN	Table 11-19 on page 11-39
RIP to RIP	Table 11-20 on page 11-40
RIP to BGP	Table 11-21 on page 11-41
RIP to OSPF	Table 11-22 on page 11-43
RIP to Routing Table	Table 11-23 on page 11-44
RIP to VNN	Table 11-24 on page 11-45
Static to BGP	Table 11-25 on page 11-46

Table 11-11. Match and Set Parameter Descriptions (Continued)

Route Map Type	See...
Static to OSPF	Table 11-26 on page 11-47
Static to RIP	Table 11-27 on page 11-48
Static to VNN	Table 11-28 on page 11-49
Any or Direct to BGP	Table 11-29 on page 11-50
Any or Direct to OSPF	Table 11-30 on page 11-52
Any or Direct to RIP	Table 11-31 on page 11-53
Any or Direct to VNN	Table 11-32 on page 11-54
Aggregate to BGP	Table 11-33 on page 11-55
Aggregate to VNN	Table 11-34 on page 11-56
VNN to BGP	Table 11-35 on page 11-57
VNN to OSPF	Table 11-36 on page 11-59
VNN to RIP	Table 11-37 on page 11-60
MOSPF to Routing Table	Table 11-38 on page 11-61
MOSPF to DVMRP	Table 11-39 on page 11-62
DVMRP to Routing Table	Table 11-40 on page 11-62
DVMRP to DVMRP	Table 11-41 on page 11-63

▶ If you configure a route map and specify ANY or DIRECT as the From protocol, make sure that you also configure an access list that selects *only those routes* that you want to include as export routes.

Table 11-12. BGP to BGP Match and Set Parameters

Field	Description
Match Parameters	BGP routes can be distributed to BGP based on matches to the following parameters. *Any fields that you do not plan to use as a match parameter should be left blank.*
Assign Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	The route tag value. Tag values are used to further identify a route. Only routes matching the specified tag value are selected.
Match Type	Specify one of the following match types: *AS Parameters* – (default) Enables you to specify a match based on Origin AS, Transit AS, and Last AS. *AS Regex* – Enables you to specify a match based on a regular expression. If you specify AS Regex, the Origin AS, Transit AS, and Last AS fields are grayed out and the AS Regex field is no longer grayed out.
Origin AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the Autonomous System (AS) where the route originated. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
Transit AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
Last AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the last Autonomous System (AS) in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
AS Regex (*Appears only if AS Regex is the specified match type*)	Specify a regular expression. BGP will perform a sub-string match based on the regular expression you specify. If you want an exact match on an AS, do not use regular expressions. Instead, specify AS Parameters as the match type and specify an AS number in the Transit AS field. This is a more efficient method of processing exact matches than regular expressions. For more information on using regular expressions, see "Using Regular Expressions" on page 11-64.

Table 11-12. BGP to BGP Match and Set Parameters (Continued)

Field	Description
Origin	Specify one of the following values to indicate the BGP origin code for use as a match parameter: *IGP, EGP, Incomplete, None.*
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be used as a match parameter. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise, or Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank. No default is used if the field is left blank.*
Origin	Specify one of the following values to indicate the BGP origin code for use as a match parameter: *IGP, EGP, Incomplete, None.*
Atomic Aggregate	Specify *Enable* or *Disable* to indicate whether or not the atomic aggregate attribute should be set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that identifies the next hop to reach a network.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.
Community Type	Specify one of the following values to identify the community type: *Replacement* – Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification will occur.

Table 11-12. BGP to BGP Match and Set Parameters (Continued)

Field	Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* –Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise, or Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

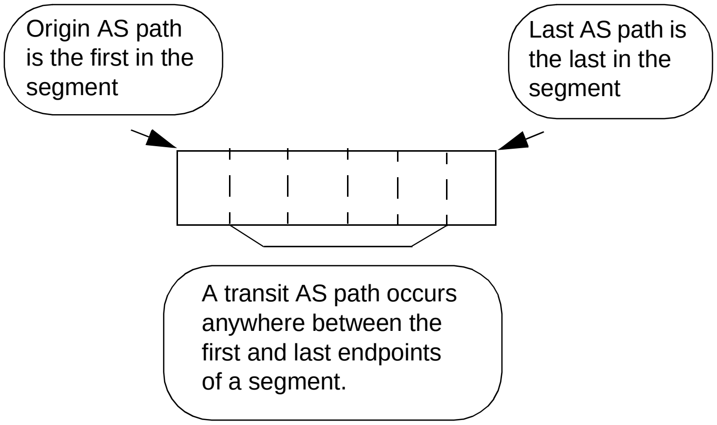
Origin AS path is the first in the segment
Last AS path is the last in the segment
A transit AS path occurs anywhere between the first and last endpoints of a segment.

Figure 11-11. Origin, Transit, and Last AS Paths

Table 11-13. BGP to OSPF Match and Set Parameters

Field	Description
Match Parameters	BGP routes can be distributed to IP OSPF based on matches to the following parameters. Only routes that match the specified parameters are selected for the Set operations. *Any fields that you do not plan to use as a match parameter should be left blank.*
Assign Network Access Lists	Use the *Assign* and *Unassign* options to specify network access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value. Tag values are used to further identify a route. Only routes matching the specified tag value will be selected.
Match Type	Specify one of the following match types: *AS Parameters* – (default) Enables you to specify a match based on Origin AS, Transit AS, and Last AS. *AS Regex* – Enables you to specify a match based on a regular expression. If you specify AS Regex, the Origin AS, Transit AS, and Last AS fields are grayed out and the AS Regex field is no longer grayed out.
Origin AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the Autonomous System (AS) where the route originated. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
Transit AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
Last AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the last Autonomous System (AS) in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.

Table 11-13. BGP to OSPF Match and Set Parameters (Continued)

Field	Description
AS Regex (*Appears only if AS Regex is the specified match type*)	Specify a regular expression. BGP will perform a sub-string match based on the regular expression you specify. If you want an exact match on an AS, do not use regular expressions. Instead, specify AS Parameters as the match type and specify an AS number in the Transit AS field. This is a more efficient method of processing exact matches than regular expressions. For more information on using regular expressions, see "Using Regular Expressions" on page 11-64.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, or None.*
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be used as a match parameter. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: No Export, No Advertise, or Local AS. If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-13. BGP to OSPF Match and Set Parameters (Continued)

Field	Description
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	Sets the IP OSPF route metric to the specified metric value. If you leave this field blank, a default metric from the routing table is used.
Tag	Sets the IP OSPF route tag value to the specified value. If you leave this field blank, a default tag from the routing table is used.
OSPF Metric Type	Specify *External-type-1* or *External-type-2.* If you leave this field blank, External-type-2 is used as the default.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default of 0 is used.

Table 11-14. BGP to RIP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from BGP to RIP are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify network access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Match Type	Specify one of the following match types: *AS Parameters* – (default) Enables you to specify a match based on Origin AS, Transit AS, and Last AS. *AS Regex* – Enables you to specify a match based on a regular expression. If you specify AS Regex, the Origin AS, Transit AS, and Last AS fields are grayed out and the AS Regex field is no longer grayed out.
Origin AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the Autonomous System (AS) where the route originated. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.

Table 11-14. BGP to RIP Match and Set Parameters (Continued)

Field	Description
Transit AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
Last AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the last Autonomous System (AS) in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
AS Regex (*Appears only if AS Regex is the specified match type*)	Specify a regular expression. BGP will perform a sub-string match based on the regular expression you specify. If you want an exact match on an AS, do not use regular expressions. Instead, specify AS Parameters as the match type and specify an AS number in the Transit AS field. This is a more efficient method of processing exact matches than regular expressions. For more information on using regular expressions, see "Using Regular Expressions" on page 11-64.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP*, *EGP*, *Incomplete*, or *None.*
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be used as a match parameter. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export*, *No Advertise*, or *Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	Sets the RIP route metric to the specified metric value. If you leave this field blank, a default metric from the routing table is used.
Tag	Sets the route tag field for the route. If you leave this field blank, a default tag from the routing table is used.

Table 11-14. BGP to RIP Match and Set Parameters (Continued)

Field	Description
Next Hop	The IP address that specifies the next hop to reach a network. If you leave this field blank, a default of 0 is used.

Table 11-15. BGP to Routing Table Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from BGP to the routing table is based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the Assign and Unassign options to specify network access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Next Hop	Specify the IP address that identifies the next hop to reach a network. Only routes that match this next hop value are selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Match Type	Specify one of the following match types: *AS Parameters* – (default) Enables you to specify a match based on Origin AS, Transit AS, and Last AS. *AS Regex* – Enables you to specify a match based on a regular expression. If you specify AS Regex, the Origin AS, Transit AS, and Last AS fields are grayed out and the AS Regex field is no longer grayed out.
Origin AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the Autonomous System (AS) where the route originated. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
Transit AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
Last AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the last Autonomous System (AS) in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.

Table 11-15. BGP to Routing Table Match and Set Parameters (Continued)

Field	Description
AS Regex (*Appears only if AS Regex is the specified match type*)	Specify a regular expression. BGP will perform a sub-string match based on the regular expression you specify. If you want an exact match on an AS, do not use regular expressions. Instead, specify AS Parameters as the match type and specify an AS number in the Transit AS field. This is a more efficient method of processing exact matches than regular expressions. For more information on using regular expressions, see "Using Regular Expressions" on page 11-64.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete*, or *None.*
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise, or Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes for the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Tag	Sets the route tag field for the route.
Weight	A weight value that is assigned to a route. This value is used only for routes from EBGP peers. The weight value is not used for routes from IBGP peers.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.

Table 11-15. BGP to Routing Table Match and Set Parameters (Continued)

Field	Description
Next Hop	Specify the IP address that identifies the next hop to reach a network.
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs.
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: No Export, No Advertise, or Local AS. If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-16. BGP to VNN Match and Set Parameters

Field	Description
Match Parameters	BGP routes can be distributed to VNN OSPF based on matches to the following parameters. Only routes that match the specified parameters are selected for the Set operations. *Any fields that you do not plan to use as a match parameter should be left blank.*
Assign Network Access Lists	Use the *Assign* and *Unassign* options to specify network access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value. Tag values are used to further identify a route. Only routes matching the specified tag value will be selected.

Table 11-16. BGP to VNN Match and Set Parameters (Continued)

Field	Description
Match Type	Specify one of the following match types: *AS Parameters* – (default) Enables you to specify a match based on Origin AS, Transit AS, and Last AS. *AS Regex* – Enables you to specify a match based on a regular expression. If you specify AS Regex, the Origin AS, Transit AS, and Last AS fields are grayed out and the AS Regex field is no longer grayed out.
Origin AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the Autonomous System (AS) where the route originated. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
Transit AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
Last AS (*Appears only if AS Parameters is the specified match type*)	Specify a match parameter for the last Autonomous System (AS) in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See Figure 11-11 on page 11-27 for an illustration of origin, transit, and last AS paths.
AS Regex (*Appears only if AS Regex is the specified match type*)	Specify a regular expression. BGP will perform a sub-string match based on the regular expression you specify. If you want an exact match on an AS, do not use regular expressions. Instead, specify AS Parameters as the match type and specify an AS number in the Transit AS field. This is a more efficient method of processing exact matches than regular expressions. For more information on using regular expressions, see “Using Regular Expressions” on page 11-64.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, or None.*
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.

Table 11-16. BGP to VNN Match and Set Parameters (Continued)

Field	Description
Community Value	If you chose *Define* for the Community field, specify the new community number that will be used as a match parameter. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: No Export, No Advertise, or Local AS. If you chose *None* for the Community field, the field is grayed out to indicate that it is not used.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	Sets the VNN OSPF route metric to the specified metric value. If you leave this field blank, a default metric from the routing table is used.
Tag	Sets the VNN OSPF route tag value to the specified value. If you leave this field blank, a default tag from the routing table is used.
VNN Metric Type	Specify *External-type-1* or *External-type-2*. If you leave this field blank, External-type-2 is used as the default.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default of 0 is used.

Table 11-17. OSPF to BGP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from an IP OSPF domain into BGP is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The IP OSPF cost. If you leave this field blank, a default value from the routing table is used.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the IP OSPF route tag value to be used as the match parameter. Only routes matching this value are selected.
OSPF Route Type	Specify one of the following OSPF Metric Type values: *Intra*, *Internal*, *External-1*, *External-2*, or *None*.

Table 11-17. OSPF to BGP Match and Set Parameters (Continued)

Field	Description
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes for the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, Do not set.*
Atomic Aggregate	Specify *Enable* or *Disable* to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that identifies the next hop to reach a network.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs.
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise, or Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-18. OSPF to RIP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from IP OSPF to RIP are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The IP OSPF cost. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the IP OSPF route tag value to be used as the match parameter. Only routes matching this value are selected.
OSPF Route Type	Specify one of the following IP OSPF Metric Type values: *Intra, Internal, External-1, External-2,* or *None.*
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The RIP metric. If you leave this field blank, a default metric from the routing table is used.
Tag	The route tag field for the route that you want to set. If you leave this field blank, a default tag from the routing table is used.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-19. OSPF to VNN Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from an IP OSPF domain into VNN OSPF is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The IP OSPF cost. If you leave this field blank, a default value from the routing table is used.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the IP OSPF route tag value to be used as the match parameter. Only routes matching this value are selected.
OSPF Route Type	Specify one of the following IP OSPF Route Type values: *Intra, Internal, External-1, External-2,* or *None.*
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Metric	The VNN OSPF cost. If you leave this field blank, a default value from the routing table is used.
Tag	The tag to be set in the redistributed routes to VNN OSPF. If you leave this field blank, a default tag from the routing table is used.
VNN Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*. If you leave this field blank, a value of External Type 2 is used.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-20. RIP to RIP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from RIP or RIP version 2 to RIP or RIP version 2 are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The RIP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The RIP metric. If you leave this field blank, a default metric from the routing table is used.
Tag	The route tag field for the route that you want to set. If you leave this field blank, a default tag from the routing table is used.
Next Hop	An IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-21. RIP to BGP Match and Set Parameters

Field	Description
Match Parameters	Routes from RIP and RIP version 2 can be redistributed into a BGP domain based on matches to one or more of the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The RIP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes for the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the origin of the route: *IGP, EGP, Incomplete*, or *None.*
Atomic Aggregate	Specify *Enable* or *Disable* to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that identifies the next hop to reach a network.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.

Table 11-21. RIP to BGP Match and Set Parameters (Continued)

Field	Description
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs.
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: *No Export, No Advertise, or Local AS.* *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise,* or *Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-22. RIP to OSPF Match and Set Parameters

Field	Description
Match Parameters	Routes from RIP and RIP version 2 can be redistributed into an IP OSPF domain based on matches to one or more of the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The RIP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The IP OSPF cost. If you leave this field blank, a default value from the routing table is used.
Tag	The tag to be set in the redistributed routes to IP OSPF. If you leave this field blank, a default tag from the routing table is used.
OSPF Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*. If you leave this field blank, a value of External Type 2 is used.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-23. RIP to Routing Table Match and Set Parameters

Field	Description
Match Parameters	Routes from RIP and RIP version 2 can be redistributed into an OSPF domain based on matches to one or more of the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The RIP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Next Hop	Specify the IP address that identifies the next hop to reach a network. Only routes that match this next hop value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The RIP metric. If you leave this field blank, a default metric from the routing table is used.
Tag	The tag to be set in the redistributed routes. If you leave this field blank, a default tag from the routing table is used.

Table 11-24. RIP to VNN Match and Set Parameters

Field	Description
Match Parameters	Routes from RIP and RIP version 2 can be redistributed into a VNN OSPF domain based on matches to one or more of the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The RIP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The VNN OSPF cost. If you leave this field blank, a default value from the routing table is used.
Tag	The tag to be set in the redistributed routes to VNN OSPF. If you leave this field blank, a default tag from the routing table is used.
VNN Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*. If you leave this field blank, a value of External Type 2 is used.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-25. Static to BGP Match and Set Parameters

Field	Description
Match Parameters	Static routes can be distributed to BGP based on matches to the following parameters. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the Assign and Unassign options to specify access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes to the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the origin of the route: *IGP, EGP, Incomplete,* or *None.*
Atomic Aggregate	Specify *Enable* or *Disable* to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that identifies the next hop to reach a network. The next hop value is set to this value on all selected routes.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs.

Table 11-25. Static to BGP Match and Set Parameters (Continued)

Field	Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise,* or *Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-26. Static to OSPF Match and Set Parameters

Field	Description
Match Parameters	Static routes can be distributed to IP OSPF based on matches to the following lists. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The IP OSPF cost. If no IP OSPF metric is specified, a default metric value from the routing table is used.
Tag	The tag to be set in the redistributed routes to IP OSPF. If none is specified, then a default tag value from the routing table is used.
OSPF Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*.

Table 11-26. Static to OSPF Match and Set Parameters (Continued)

Field	Description
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-27. Static to RIP Match and Set Parameters

Field	Description
Match Parameters	Static routes can be distributed to RIP based on matches to the following lists. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	Sets the metric value on all selected routes to the specified metric value. If you leave this field blank, a default metric value from the routing table is used.
Tag	Sets the tag value on all selected routes to the specified tag value. If you leave this field blank, a default tag value from the routing table is used.
Next Hop	An IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-28. Static to VNN Match and Set Parameters

Field	Description
Match Parameters	Static routes can be distributed to VNN OSPF based on matches to the following lists. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The VNN OSPF cost. If no VNN OSPF metric is specified, a default metric value from the routing table is used.
Tag	The tag to be set in the redistributed routes to VNN OSPF. If none is specified, then a default tag value from the routing table is used.
VNN Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-29. Any or Direct to BGP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from a Direct or Any domain into BGP is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The OSPF cost. Only routes matching this value are selected. This parameter is not used for Direct to BGP route maps.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the OSPF route tag value to be used as the match parameter. Only routes matching this value are selected. This parameter is not used for Direct to BGP route maps.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as a set parameter should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes to the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, Do not set.*
Atomic Aggregate	Specify Enable or Disable to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	An IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.

Table 11-29. Any or Direct to BGP Match and Set Parameters (Continued)

Field	Description
Community Type	Specify one of the following values. *Replacement* –Assigns a new community number to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs. All selected routes are set to the value that you specify.
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise,* or *Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-30. Any or Direct to OSPF Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from a Direct or Any domain into BGP is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The IP OSPF cost. Only routes matching this value are selected. This parameter is not used for Direct to OSPF route maps.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected. This parameter is not used for Direct to OSPF route maps.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The IP OSPF cost. If you leave this field blank, a default value from the routing table is used.
Tag	The tag to be set in the redistributed routes to IP OSPF. If no value is specified, a tag value from the routing table is used.
OSPF Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*. The IP OSPF Metric Type on selected routes is set to the specified value. A default value of *External Type 2* is used if no value is specified.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-31. Any or Direct to RIP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from a Direct or Any domain into RIP is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The metric value that is used as a match parameter. Only routes matching this value are selected. This parameter is not used for Direct to RIP route maps.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected. This parameter is not used for Direct to RIP route maps.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The RIP metric. If no RIP metric is specified, the value in the routing table is used.
Tag	The tag to be set in the redistributed routes to RIP. If no value is specified, a tag value from the routing table is used.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-32. Any or Direct to VNN Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from a Direct or Any domain into VNN OSPF is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The VNN OSPF cost. Only routes matching this value are selected. This parameter is not used for Direct to VNN OSPF route maps.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected. This parameter is not used for Direct to VNN OSPF route maps.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The VNN OSPF cost. If you leave this field blank, a default value from the routing table is used.
Tag	The tag to be set in the redistributed routes to VNN OSPF. If no value is specified, a tag value from the routing table is used.
VNN Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*. The VNN Metric Type on selected routes is set to the specified value. A default value of *External Type 2* is used if no value is specified.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-33. Aggregate to BGP Match and Set Parameters

Field	Action/Description
Match Parameters	There are no match parameters for an Aggregate to BGP route map.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as a set parameter should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes to the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, Do not set.*
Atomic Aggregate	Specify Enable or Disable to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that identifies the next hop to reach a network. Only routes that match this next hop value are selected.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number is assigned to replace the old value. *Additive* – Adds a community to an existing community. *None* – No community modification occurs. All selected routes are set to the value that you specify.

Table 11-33. Aggregate to BGP Match and Set Parameters (Continued)

Field	Action/Description
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: *No Export, No Advertise,* or *Local AS.* *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: *No Export, No Advertise, or Local AS.* If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-34. Aggregate to VNN Match and Set Parameters

Field	Description
Match Parameters	There are no match parameters for an Aggregate to VNN route map.
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as a set parameter should be left blank.* No default is used if the field is left blank.
Metric	The VNN OSPF cost. If no OSPF metric is specified, a default metric value from the routing table is used.
Tag	The tag to be set in the redistributed routes to VNN OSPF. If none is specified, then a default tag value from the routing table is used.
VNN Metric Type	Specify one of the following values: *External Type 1* or *External Type 2.*
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-35. VNN to BGP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from a VNN OSPF domain into BGP is based on matching the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The VNN OSPF cost. If you leave this field blank, a default value from the routing table is used.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the VNN OSPF route tag value to be used as the match parameter. Only routes matching this value are selected.
VNN Route Type	Specify one of the following VNN OSPF Route Type values: *Intra, Internal, External-1, External-2,* or *None.*
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *Any fields that you do not plan to use as set parameters should be left blank.* No default is used if the field is left blank.
Local Preference	The value that you specify is used as the local preference value for all selected routes. Local preference indicates a degree of preference given to a route to compare it with other routes for the same destination. A higher local preference value indicates a preferred route. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.
Origin	Specify one of the following values to indicate the BGP origin code: *IGP, EGP, Incomplete, Do not set.*
Atomic Aggregate	Specify *Enable* or *Disable* to indicate whether or not the atomic aggregate attribute is set as an indication of information loss.
Multi-Exit-Discr	The multi-exit-discriminator (MED) value. This value indicates the preferred path into an AS that has multiple entry points. Lower MED values indicate the preferred path. For example, a route with a MED value of 120 would be preferred over a route with a MED value of 200.
Next Hop	Specify the IP address that identifies the next hop to reach a network.
AS Repeat Count	A multiple number of the local AS number prepended to the existing segment. This number is the total number of times that IP Navigator adds the local AS to the AS path.

Table 11-35. VNN to BGP Match and Set Parameters (Continued)

Field	Description
Community Type	Specify one of the following values. *Replacement* – Assigns a new community number to replace the old value. Additive – Adds a community to an existing community. *None* – No community modification occurs.
Community	Specify one of the following values to identify the community: *Define* – Indicates that you will specify a user-defined community in the Community Value field. *Well Known* – Indicates that you will specify one of the following three reserved community values in the Community Value field: No Export, No Advertise, or Local AS. *None* – Indicates that no community value will be specified.
Community Value	If you chose *Define* for the Community field, specify the new community number that will be assigned to selected routes. If you chose *Well Known* for the Community field, specify one of the following three reserved community values: No Export, No Advertise, or Local AS. If you chose *None* for the Community field, this field is grayed out to indicate that it is not used.

Table 11-36. VNN to OSPF Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from VNN OSPF to IP OSPF are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The VNN OSPF cost. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the VNN OSPF route tag value to be used as the match parameter. Only routes matching this value are selected.
VNN Route Type	Specify one of the following VNN Route Type values: *Intra, Internal, External-1, External-2,* or *None.*
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The IP OSPF cost. If you leave this field blank, a default value from the routing table is used.
Tag	The tag to be set in the redistributed routes to IP OSPF. If you leave this field blank, a default tag from the routing table is used.
OSPF Metric Type	Specify one of the following values: *External Type 1* or *External Type 2*. If you leave this field blank, a value of External Type 2 is used.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-37. VNN to RIP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from VNN OSPF to RIP are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The VNN OSPF cost. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the VNN OSPF route tag value to be used as the match parameter. Only routes matching this value are selected.
VNN Route Type	Specify one of the following VNN Route Type values: *Intra, Internal, External-1, External-2,* or *None.*
Set Parameters	The following parameters are set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The RIP metric. If you leave this field blank, a default metric from the routing table is used.
Tag	The route tag field for the route that you want to set. If you leave this field blank, a default tag from the routing table is used.
Next Hop	The IP address that identifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-38. MOSPF to Routing Table Match and Set Parameters

Field	Description
Match Parameters	Routes from an MOSPF multicast network can be redistributed into a routing table based on matches to one or more of the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The MOSPF metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Tag	Specify the route tag value to be used as the match parameter. Tag values are used to further identify a route. Only routes matching this value are selected.
Next Hop	Specify the IP address that identifies the next hop to reach a network. Only routes that match this next hop value are selected.
Set Parameters	The following parameter is set on all selected routes. Routes are selected if they match the specified match parameters. *If you leave any of the following parameters blank, the system uses a default value.*
Metric	The routing table metric. If you leave this field blank, a default metric from the routing table is used.

Table 11-39. MOSPF to DVMRP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from an MOSPF multicast network to a DVMRP multicast network are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The MOSPF metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Set Parameters	Not applicable.

Table 11-40. DVMRP to Routing Table Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from DVMRP to the routing table are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The DVMRP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Set Parameters	Not applicable.

Table 11-41. DVMRP to DVMRP Match and Set Parameters

Field	Description
Match Parameters	The redistribution of routes from DVMRP multicast network to DVMRP multicast network are based on matches to the following objects. *Any fields that you do not plan to use as a match parameter should be left blank.*
Network Access Lists	Use the *Assign* and *Unassign* options to specify access lists as necessary.
Metric	The DVMRP metric value that is used as a match parameter. Only routes matching this value are selected.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length less than this value are not selected.
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length greater than this value are not selected.
Set Parameters	Not applicable.

Using Regular Expressions

A *regular expression* is a text string that describes a set of strings, which means that it can be used for pattern matching. For example, a regular expression “r” matches a string “s” if the string “s” is in the set of strings described by “r.”

Regular expressions are common in UNIX. If you have experience with UNIX programs that make use of regular expressions (e.g., *regexp*), you should be able to quickly grasp Lucent’s implementation of regular expressions. The manual page for regexp provides a lot of useful information on regular expressions. You can view this page by typing **`man regexp`** at the UNIX prompt.

Regular expression strings consist of characters and operators. Operators are special characters that specify the number of characters to match. Table 11-42 describes some commonly used regular expression operators.

Table 11-42. Commonly Used Regular Expression Operators

Operator	Description
.	Match any character in the string.
*	Match zero or more of the preceding characters in the string.
+	Match one or more of the preceding characters in the string.
?	Match zero or one of the preceding characters in the string.
^	Match from the beginning of the string.
$	Match at the end of the string.
|	Match the character that immediately precedes the operator and the character that immediately follows the operator.
[ ]	Match the characters enclosed in the brackets.
–	Match a range of characters.

For example, consider the following AS paths:

AS Path A — 32245 32246 56734 12356

AS Path B — 32245 32246 25348 19234 13456

AS Path C — 56743 41759 13456

AS Path D — 13456

The regular expressions in Table 11-43 illustrate how you can filter BGP control traffic from these autonomous systems.

Table 11-43. Regular Expression Examples

Regular Expression	Description	Matches AS Path...
^32	Match all paths that begin with “32”.	A, B
100$	Match all paths that end with “100”.	No matches
56$	Match all paths that end with “56”.	A, B, C, D
^13456$	Match all paths that begin and end with “13456”.	D
32+	Match all paths that have one or more consecutive twos.	A, B
^[5-7]+	Match all paths that begin with the numbers five, six, or seven.	C
34	Match all paths that contain “34”.	A, B, D
.*	Match all paths.	A, B, C, D
\<41759\>	Match all paths that contain “41759”.	C
^$	Match all paths that contain an empty string.	No matches.

You use regular expressions in route maps in which BGP is the source. See “Adding Route Maps” on page 11-18 for more information on route maps. In that section, the following tables provide information on using regular expressions:

- Table 11-12 on page 11-25
- Table 11-13 on page 11-28
- Table 11-14 on page 11-30
- Table 11-15 on page 11-32
- Table 11-16 on page 11-34

▶ Avoid using regular expressions for exact matches on AS numbers. Instead, specify AS Parameters as the match type and specify a transit AS in the route map configuration. This is a more efficient method of processing exact matches than regular expressions.

12

Configuring Label Switched Paths

This chapter provides an overview of label switched paths (LSPs) and describes the following configuration tasks:

- Enabling MPT LSPs and multicast LSPs
- Configuring point-to-point LSPs
- Defining a point-to-point connection path
- Displaying the operational status for point-to-point connection paths
- Configuring LSPs over OPTimum cell trunks

What Are Label Switched Paths?

A *label* is a temporary identifier that specifies a forwarding destination; unlike an IP address, a label is an arbitrary value that has no significance outside the switch network. *Label switching* is an advanced form of packet forwarding that replaces address-match routing with a more efficient forwarding algorithm. A *label switched path* (LSP) is an ATM or Frame Relay virtual circuit that uses labels to transport connectionless data, such as IP packets, across a switch network efficiently.

When an IP packet enters the Lucent switch network, a switch at the edge of the network reads the packet's IP header and encapsulates the packet with a label based on the packet's header contents. The edge switch then sends the packet through the core switch network, which can use the packet's label to identify the path the packet should take. Because each transit switch inside the network can forward the packet based solely on its label, rather than having to analyze the packet's IP header, do routing table lookups, and make routing decisions, label switched paths are much faster than hop-by-hop routing.

For example, Figure 12-1 illustrates a network with an ingress edge switch, several core switches inside the Lucent network cloud, and an egress edge switch. When the ingress switch receives an IP packet intended for Network B, the ingress switch reads the packet header and determines the packet should travel on path 111. It encapsulates the packet with a label and passes it to the switch network. The transit switches in the core network forward the packet on path 111 without reading the packet's header. When the packet reaches the egress switch, the egress switch strips off the label and routes the packet to Network B.

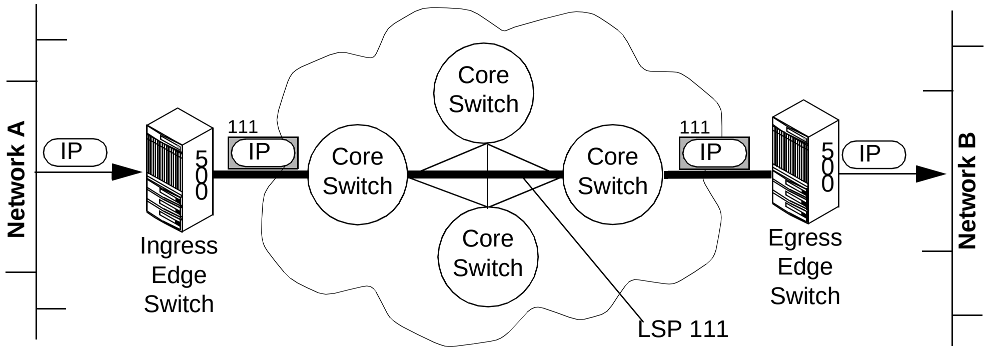

Figure 12-1. Ingress, Transit, and Egress Switches in an LSP

Lucent switches support three types of label switched paths:

- Multipoint-to-point LSPs (MPT LSPs) allow multiple leaf nodes to share the same circuit for transmission to a single destination (the root).
- Point-to-point LSPs allow a pair of nodes to share a point-to-point connection.
- Multicast LSPs allow a single root node to transmit IP multicast traffic to multiple leaf destinations.

Every switch that runs IP Navigator maintains an MPT LSP circuit network if you enable MPT LSPs on the switch and create at least one IP interface on the switch. (Note that CBX 500 switches also require a Frame card.)

Both the MPT LSP and multicast LSP networks are *rooted* at the switch. For this purpose, the switch maintains a root, which:

- Keeps track of MPT LSP and multicast LSP nodes
- Adds and deletes leaf nodes
- Keeps track of LSP circuits by periodically issuing keepalive messages

A root is a standard circuit endpoint that is created at initialization time on every CP card in the B-STDX 8000/9000 and every SP card in the CBX 500. All other nodes on this circuit network are considered leaves. As noted before, traffic flow occurs from the leaves to the root on MPT LSP, and from the root to the leaves on multicast LSPs. On point-to-point LSP connections, traffic is bi-directional, since each node configured on a point-to-point LSP connection acts as both a leaf and a root.

The following section discusses each type of LSP separately.

Multipoint-to-Point (MPT) LSPs

A multipoint-to-point LSP (formerly called a *reverse MPT*) is a unidirectional virtual circuit created automatically to route IP traffic to other Lucent switches. An MPT LSP lets all leaf nodes of a particular connectivity tree share the same virtual path when switching IP data to the root node of that tree. Data flow on an MPT LSP travels from the leaf switches to the root switch. Figure 12-2 shows a sample MPT LSP.

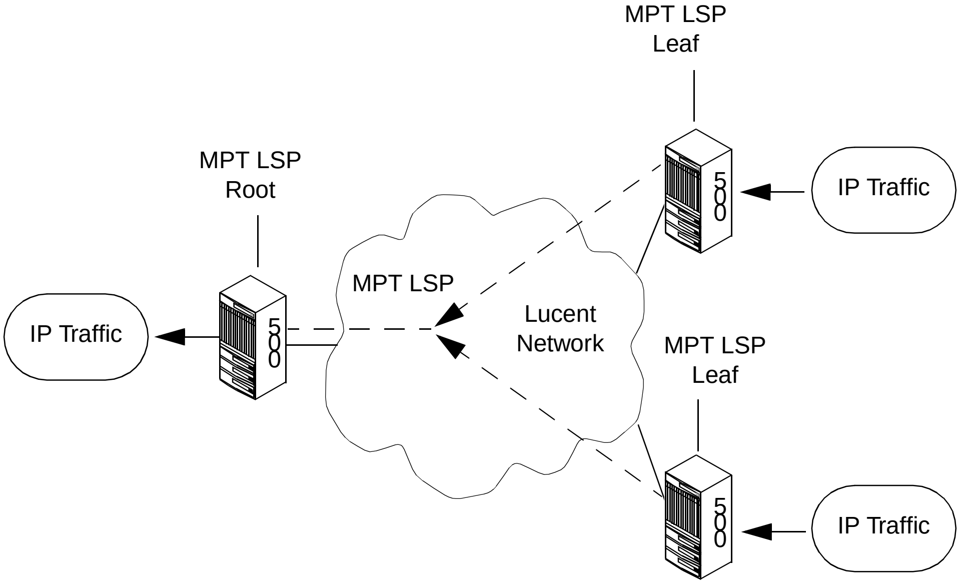

Figure 12-2. MPT Label Switched Path

Every switch in an IP Navigator network automatically establish MPT LSP circuits when they are initialized, and update the MPT LSP circuits when other switches join or leave the network. After a switch establishes an MPT LSP circuit from itself (root) to the leaf switches, it updates the MPT LSP circuits when other switches join or leave the network.

For example, Figure 12-3 illustrates a network that includes three nodes, each of which has established an MPT LSP to carry data from the leaf nodes to the root node. (The arrows in the figure indicate the data flow direction.) Node 1's MPT LSP lets ports on Nodes 2 and 3 forward packets through the Lucent switch network to Node 1. Similarly, the MPT LSPs that are rooted on Nodes 2 and 3 allow packets to be forwarded to them using labels.

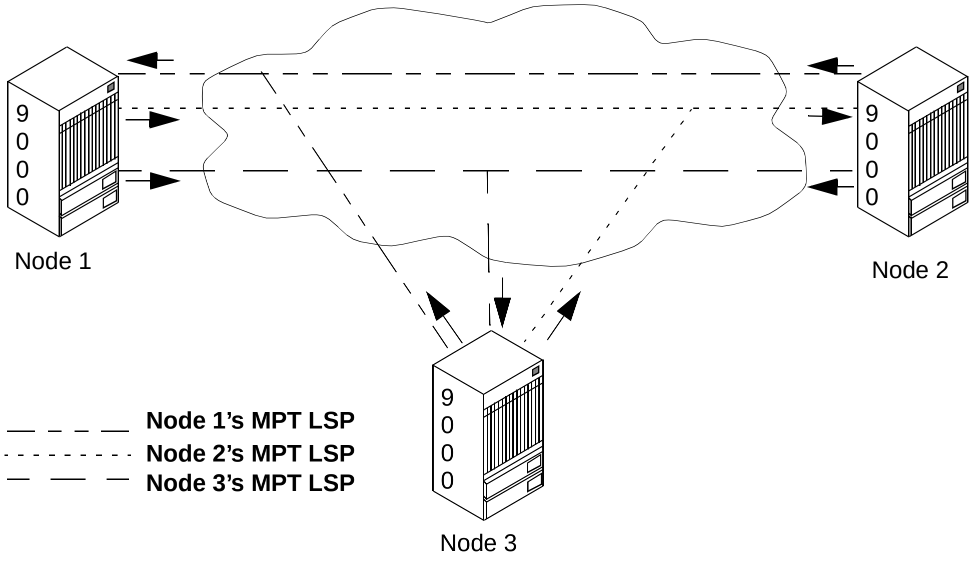

Figure 12-3. MPT LSP Network

MPT LSP Initialization

When OSPF finds a new node, it notifies the root LSP module, which in turn adds the node to a list of leaves. The list of leaves is updated based on OSPF notifications. Every 30 seconds, a grooming process scans the list, and LSPs may be rerouted based on current network conditions and the LSP list membership.

MPT LSP Requirements

The root for an MPT LSP is created at initialization time on every CP card in a B-STDX 8000/9000 and on every SP in a CBX 500. Roots track the state of other nodes or forwarding engines (FEs), and are responsible for adding and deleting nodes or FEs as well as keeping nodes or FEs alive.

MPT LSPs will run over direct trunk and Optimum Trunk links.

On each switch that acts as either the root or a leaf, MPT LSPs must be enabled and at least one IP interface must exist. For more information on enabling MPT LSPs, see "Enabling MPT LSPs and Multicast LSPs" on page 12-11. See Chapter 3, "Configuring IP Logical Ports and IP Servers," for more information on configuring IP interfaces.

Point-to-Point LSPs

Point-to-point LSPs are user-defined circuits for IP traffic between exactly two switches. Traffic on a point-to-point LSP connection is effectively bi-directional, as a point-to-point LSP consists of reciprocal unidirectional circuits.

A point-to-point LSP overrides an MPT LSP root-to-leaf connection. This feature enables you to specify the Quality of Service between two nodes. All traffic uses the point-to-point connection rather than the automatic connection.

Figure 12-4 shows a sample point-to-point LSP between two B-STDX 9000 switches (Node 1 and Node 2). To transport IP traffic between Node 1 and Node 2, these switches use the point-to-point LSP connection instead of the MPT LSP connection that is automatically created.

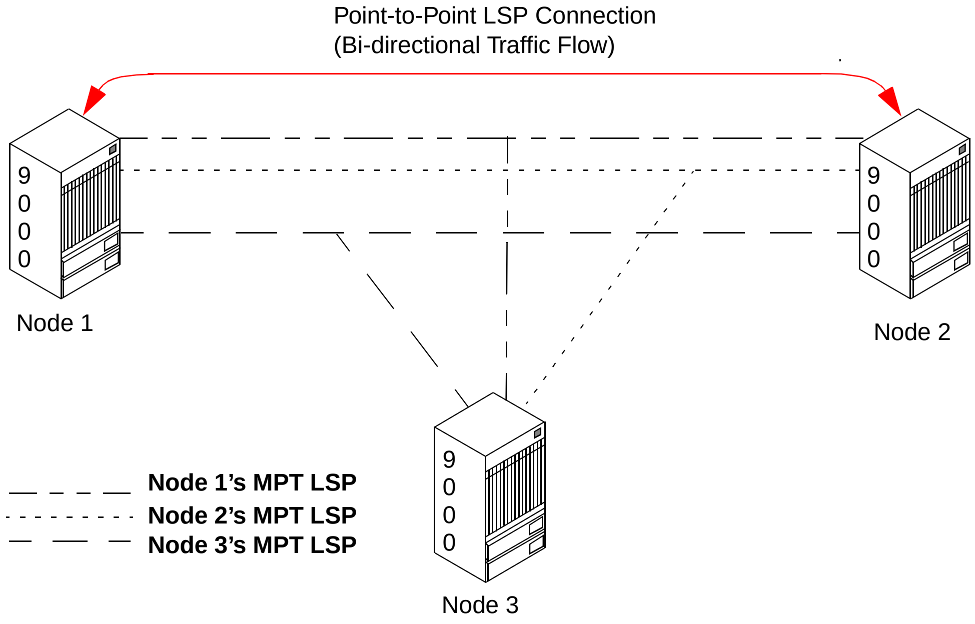

Figure 12-4. Point-to-Point LSP Connections

For each point-to-point LSP connection, you can assign preconfigured traffic descriptors that control the flow of traffic. See the *NavisCore ATM Configuration Guide* for more information on traffic descriptors.

You can define a point-to-point LSP connection between:

- Two B-STDX 8000/9000 switches
- A CBX 500 switch equipped with at least one Frame card (e.g., 6-port DS3 Frame card) and a B-STDX 8000/9000 switch
- Two CBX 500 switches, each equipped with at least one Frame card

When you configure point-to-point LSP connections between two switches, you can reserve them for use by either IP Virtual Private Networks (IP VPNs) or public traffic. IP VPNs use point-to-point LSPs to transport traffic through the Lucent network between *virtual routers* (VRs). A virtual router is an IP VPN's representative on a switch; it behaves very much like a physical router. All IP VPN traffic between the virtual routers on two switches travels through point-to-point LSP connection, if one is configured. For more information on IP VPNs, see Chapter 16, "Configuring IP Virtual Private Networks."

Typically, the point-to-point LSP connects switches at the network edge. Traffic is transported between the switches according to the traffic descriptors assigned to the point-to-point LSP connection.

An IP VPN can use:

- Point-to-point LSP connections configured explicitly for that IP VPN. Multiple point-to-point LSPs can connect two switches. However, for a single IP VPN, you can configure one (and only one) point-to-point LSP connection between two switches. For example, suppose that you have three IP VPNs that require point-to-point LSP connections between the same two switches. You can configure one point-to-point LSP connection for each VPN, but you cannot, for example, configure two point-to-point LSP connections for any of the VPNs.
- Point-to-point LSP connections configured for public use (available to all IP VPNs and public traffic).

> ▶ Only one point-to-point LSP connection may be configured for public use between the same two switches. The public point-to-point LSP connection is overridden by point-to-point LSP connections configured explicitly for IP VPNs.

- The default, best-effort MPT LSP, which is always available for public use. This MPT LSP is overridden by point-to-point LSPs configured for public use.

Figure 12-5 shows IP VPN traffic being transported between two edge switches via point-to-point LSP connections. Notice the following details:

- Both VPN1 and VPN2 have a point-to-point LSP connection configured explicitly for each of them.
- A public, configured point-to-point LSP connection is available for use by either VPN1 or VPN2.
- The public, default point-to-point LSP connection is available for use by either VPN1 or VPN2.

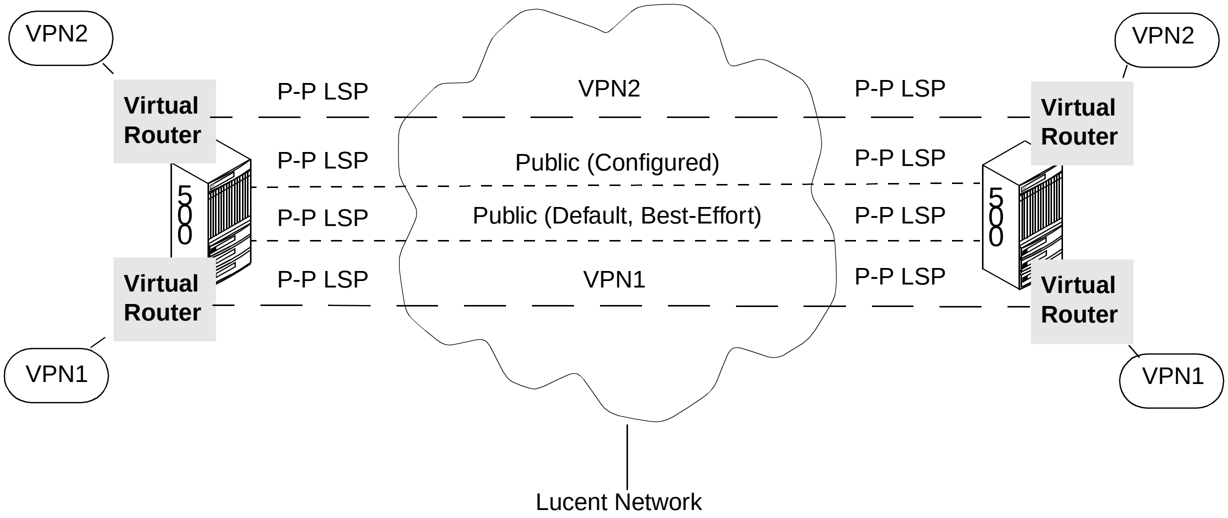

Figure 12-5. IP VPN Traffic over Point-to-Point LSP Connections

Multicast LSPs

IP multicasting is the transmission of an IP datagram from one host to a host group, a set of zero or more hosts identified by a single IP destination address. A multicast LSP (formerly called a *forward MPT*) is constructed as a tree for an ATM or Frame Relay multipoint virtual circuit, with the node constructing the multicast LSP serving as the root of the tree. A separate multicast LSP is constructed for each unique set of multicast group members for which the root switch wants to forward IP packets. Multicast LSPs transport only multicast traffic, and traffic flows from the root switch to the leaf switches.

Multicast LSPs provide a way to forward IP multicast traffic through the Lucent network. All of the switches that share a multicast LSP are in the same multicast host group.

▶ At this time, the Multicast Open Shortest Path First (MOSPF) multicast routing protocol creates multicast LSPs; but other applications, such as DVMRP and IP VPNs, actually use the multicast LSPs.

Figure 12-6 shows a sample multicast LSP.

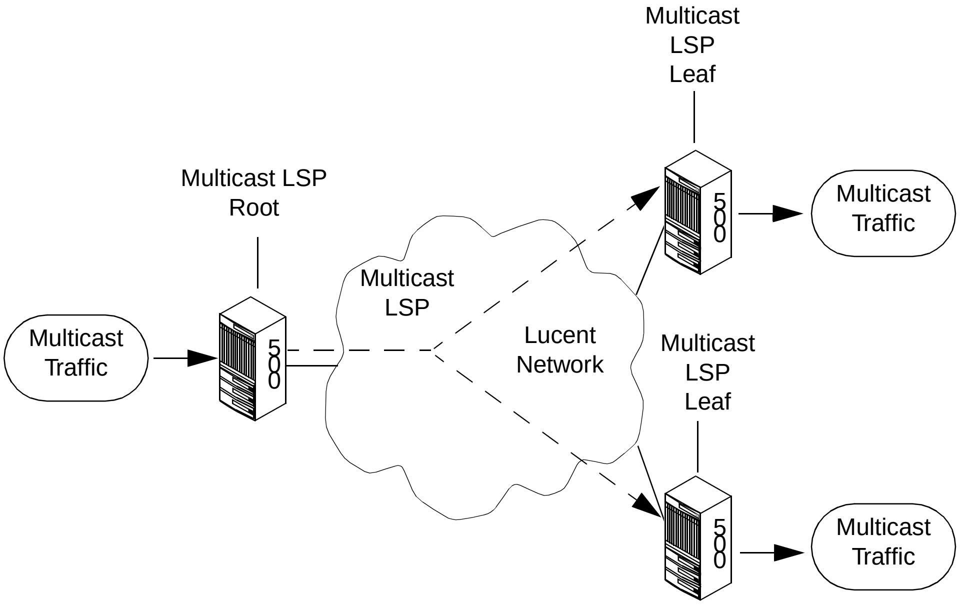

Figure 12-6. Sample Multicast LSP

Multicast LSP Initialization

When a root switch needs to send a message to a particular multicast group, it checks whether an appropriate multicast LSP already exists. If it does, the root switch uses this circuit to forward data to the multicast group members. If a multicast LSP is not already available, the root switch identifies the leaf switches in the multicast group and begins signalling to establish the virtual circuit to each leaf switch. Until all leaf nodes have confirmed the connection setup, the multicast LSP circuit is not used for data traffic, and packets are routed to the other group nodes.

You must enable multicast LSPs on all the switches in Figure 12-6. See “Enabling MPT LSPs and Multicast LSPs” on page 12-11 for details.

Processing LSPs

All LSPs on the CBX 500 are used to forward IP data over virtual paths (VPs) from one switch to another. LSPs are initiated at the SP of a CBX 500 in the same way that they are initiated at the CP on a B-STDX 8000/9000.

However, each leaf that is added to the LSP occurs:

- In the CP on a B-STDX 8000/9000.
- In the SP on a CBX 500. Each SP and each forwarding engine (FE) on a CBX 500 is added as a leaf of the LSP. FEs reassemble cells and perform IP lookups. There are two FEs in each of the following CBX 500 cards in Figure 12-7:
  - 6-port DS3 Frame card
  - 4-port Ethernet

Figure 12-7 illustrates LSP leaf occurrences in the CBX 500 and the B-STDX 8000/9000.

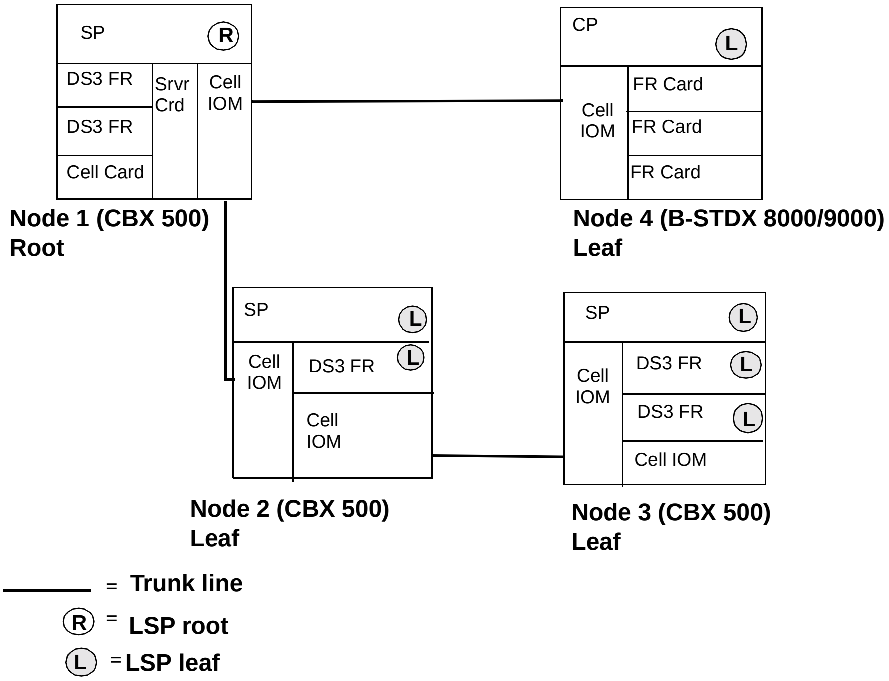

Figure 12-7. LSP Leaf Occurrences in the CBX 500 and B-STDX 8000/9000

LSPs and Switch Domains

There are two types of switch domains:

Cell Domain — Paths that traverse direct ATM trunks and ATM OPTimum trunks.

Frame Domain — Paths that traverse direct Frame trunks and Frame OPTimum trunks.

A switch can belong to multiple domains if the domains are adjacent. Switches that belong to multiple domains must reside at the border of these domains. In addition, these switches must perform additional protocol layer processing to determine routes across the different domains. The root maintains connections to each domain the switch belongs to.

The Virtual Network Navigator (VNN) OSPF instance determines how LSPs connect two switches in different domains. The following factors apply when determining LSPs:

- A switch that only belongs to one domain cannot add a switch from a different domain to its LSP. *LSPs are only established between switches in the same domain. To traverse different domains, a boundary switch that belongs to both domains must act as an intermediary. Each endpoint switch connects to the boundary switch via a separate LSP, and the boundary switch performs an IP lookup when routing traffic between the endpoints.*
- If the shortest path between two switches in the same domain traverses a different domain, the switches cannot add each other to their LSPs.

LSPs use cell and Frame domains to circumvent addressing limitations in switching ATM cells. *Whenever you cross boundaries between cell and Frame domains, an IP lookup is required.*

LSPs and VNN OSPF Areas

Point-to-point LSPs and multicast LSPs can traverse multiple VNN OSPF areas. However, MPT LSPs cannot traverse multiple VNN OSPF areas. Within a single VNN OSPF area, MPT LSPs and multicast LSPs are switched. However, as soon as traffic reaches an area boundary (that is, an area border router), a routing lookup must be made and the traffic switched on another multicast LSP or MPT LSP, if available.

▶ An MPT LSP can have multiple roots at the area border router. This eliminates the 2048-leaf limit for the aggregate of spanned areas.

Enabling MPT LSPs and Multicast LSPs

The Set IP Parameters dialog box allows you to enable MPT LSPs and multicast LSPs on a switch. By default, MPT LSPs are enabled for a switch.

The MPT LSPs value that you specify determines the use of MPT LSPs on the switch as follows:

- If the MPT LSPs value is set to Enable and no IP interfaces have been defined, the switch *does not* establish MPT LSPs unless the switch is a boundary switch.
- If the MPT LSPs value is set to Enable and IP interfaces have been defined, the switch *does* establish MPT LSPs as a means of forwarding IP traffic.
- If the MPT LSPs value is set to Disable and IP interfaces have been defined, the switch *does not* establish MPT LSPs to forward IP traffic**,** but instead uses a hop-by-hop forwarding method.
- If the MPT LSPs value is set to Disable and no IP interfaces have been defined, the switch *does not* establish MPT LSPs.

In order for the switch to process multicast LSPs, *the multicast LSPs value for the switch must be enabled*. Once you enable this setting, switches will create multicast LSPs when multicast traffic is received or IP VPNs are implemented. The multicast LSPs value is enabled by default.

To enable or disable MPT LSPs and multicast LSPs on a switch:

1. Select the switch on the network map.
2. From the Administer menu, choose Lucent IP Parameters ⇒ Set IP Parameters. The Set IP Parameters dialog box appears (see Figure 12-8).

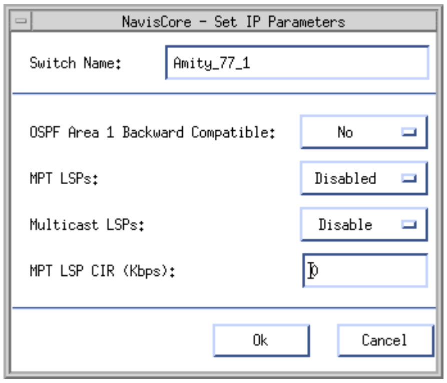

Switch Name: Amity_77_1

OSPF Area 1 Backward Compatible: No

MPT LSPs: Disabled

Multicast LSPs: Disable

MPT LSP CIR (Kbps): 0

Figure 12-8. Set IP Parameters Dialog Box

3. Specify information in the Set IP Parameters dialog box fields described in Table 12-1.

Table 12-1. Set IP Parameters Dialog Box: Field Descriptions

Field	Action/Description
Switch Name	Displays the name of the switch.
OSPF Area 1 Backward Compatible (*Applies to switches in VNN OSPF Area 1 only*)	Select either *Yes* or *No*. If you select *Yes*, the switch: – Can communicate with other Lucent switches in Area 1 running pre-5.0 switch software. – Can communicate with other Lucent switches in Area 1 running 5.0 switch software, which are set to *Yes* in this field. If you select *No*, the switch: – Cannot communicate with other Lucent switches in Area 1 running pre-5.0 switch software. – Can communicate with other Lucent switches in Area 1 running 5.0 switch software, which are set to *No* in this field. Note that this parameter applies to VNN OSPF only. It has nothing to do with IP OPSF. See Chapter 9, “Configuring IP OSPF and VNN OSPF” for more information on VNN OSPF and IP OSPF.
MPT LSPs	Set this value to Enable or Disable. To process MPT LSP connections, *the MPT LSPs value for the switch must be set to Enable*.
Multicast LSPs	Set this value to Enable or Disable. In order for the switch to process multicast LSPs, *the multicast LSPs value for the switch must be set to Enable*.
MPT LSP CIR (Kbps)	Enter the MPT LSP Committed Information Rate (CIR). The MPT LSP CIR specifies the rate in Kbps at which MPT LSPs transfer data, averaged over a minimum increment of time. In addition, this value reserves bandwidth for MPT LSPs, which the switch originates. This value *does not* apply to multicast LSPs. ***Note:*** *This value applies to all links in the MPT LSP.*

4. Choose OK.

Configuring Point-to-Point LSP Connections

The steps you perform to configure a point-to-point LSP connection depend on whether:

- You want to provision the point-to-point LSP for use by an IP VPN.
- You want to provision the point-to-point LSP for public use.

To configure a point-to-point LSP and provision it for use by an IP VPN, you must:

1. Verify that network-wide traffic descriptors have been configured, if you want to assign a traffic descriptor to the point-to-point LSP. Traffic descriptors consist of an ATM Quality of Service (QoS) class (for example, UBR/ABR or VBR) and associated descriptors (for example, peak cell rate). See "Verifying Network-Wide Traffic Descriptors" on page 12-13.
2. Select the IP VPN. See "Selecting the IP VPN" on page 16-33. Keep in mind that the Set All Point-to-Point LSP Connections dialog box (see Figure 12-9 on page 12-14) provides a Select IP VPN button, which allows you to select a VPN.
3. Configure one or more point-to-point LSP connections for the IP VPN. See "Configuring the Connections" on page 12-14.

To configure a point-to-point LSP connection and provision it for public use:

1. If you want to assign a traffic descriptor to the point-to-point LSP connection, verify that network-wide traffic descriptors have been configured. See "Verifying Network-Wide Traffic Descriptors" on page 12-13.
2. Configure the point-to-point LSP connection. Point-to-point LSP connections are reserved for use by the Public IP VPN unless you specify otherwise. See "Configuring the Connections" on page 12-14 for more information.

Verifying Network-Wide Traffic Descriptors

Before you assign traffic descriptors to point-to-point LSP connections:

1. Read about traffic descriptors in the *NavisCore ATM Configuration Guide*. Make sure that you use the ATM guide in conjunction with this guide when you assign traffic descriptors to point-to-point LSP connections, especially if the point-to-point LSP traverses ATM trunks.
2. Create network-wide traffic descriptors using the procedures described in the *NavisCore ATM Configuration Guide*.

▶ Traffic descriptors do not apply to Frame trunks. Instead, when the point-to-point LSP traverses Frame trunks, best-effort guarantees are made to meet the CIR.

Configuring the Connections

To configure a point-to-point LSP connection from the NavisCore menu:

1. Select Lucent IP Parameters ⇒ Set Point-to-Point LSP from the Administer menu. The Set All Point-to-Point LSP Connections dialog box appears (see Figure 12-9). Table 12-2 describes each of the Set All Point-to-Point LSP Connections buttons. Table 12-3 describes the Set All Point-to-Point LSP Connections dialog box fields.

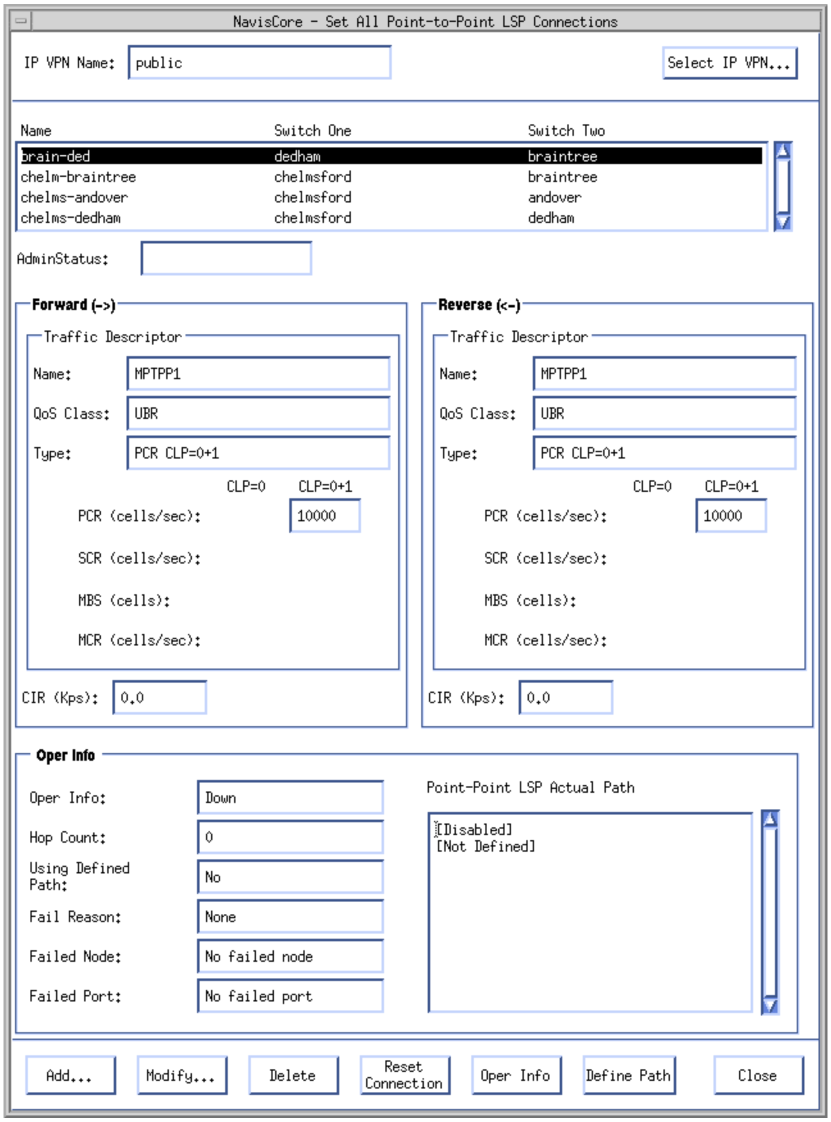

NavisCore - Set All Point-to-Point LSP Connections

IP VPN Name: public

Name	Switch One	Switch Two
brain-ded	dedham	braintree
chelm-braintree	chelmsford	braintree
chelms-andover	chelmsford	andover
chelms-dedham	chelmsford	dedham

AdminStatus:

Forward (->)

Traffic Descriptor

Name: MPTPP1
QoS Class: UBR
Type: PCR CLP=0+1

	CLP=0	CLP=0+1
PCR (cells/sec):		10000
SCR (cells/sec):		
MBS (cells):		
MCR (cells/sec):		

CIR (Kps): 0.0

Reverse (<-)

Traffic Descriptor

Name: MPTPP1
QoS Class: UBR
Type: PCR CLP=0+1

	CLP=0	CLP=0+1
PCR (cells/sec):		10000
SCR (cells/sec):		
MBS (cells):		
MCR (cells/sec):		

CIR (Kps): 0.0

Oper Info

Oper Info: Down
Hop Count: 0
Using Defined Path: No
Fail Reason: None
Failed Node: No failed node
Failed Port: No failed port

Point-Point LSP Actual Path

```
[Disabled]
[Not Defined]
```

Figure 12-9. Set All Point-to-Point LSP Connections Dialog Box

Table 12-2. Set All Point-to-Point LSP Connections Dialog Box: Buttons

Button	Function
Select IP VPN	Allows you to select an IP VPN. Once you select the IP VPN, all management tasks you perform apply only to that IP VPN. See "Selecting the IP VPN" on page 16-33 for more information on selecting an IP VPN.
Add	Enables you to add a point-to-point LSP connection. To add the connection, you specify a circuit name, the endpoints, traffic descriptors, and committed information rate (CIR).
Modify	Enables you to modify the parameters for a selected point-to-point LSP connection.
Delete	Deletes a selected point-to-point LSP connection.
Reset Connection	Enables you to toggle connection signalling on and off.
Oper Info	Displays operational information about a selected point-to-point LSP connection.
Define Path	Displays the Set Point-to-Point LSP Defined Path dialog box to enable you to define the path of a selected point-to-point LSP connection. See "Defining a Point-to-Point Connection Path" on page 12-18.

Table 12-3. Set All Point-to-Point LSP Connections Dialog Box: Fields

Field	Action/Description
Name/Switch One/ Switch Two	Displays all of the defined point-to-point LSP connections in your network.
Forward and Reverse CIR and Traffic Descriptors	
Name	Displays the name of the forward and reverse traffic descriptors.
QoS Class	Displays the Quality of Service (QoS) class for the forward and reverse traffic descriptors. See the *NavisCore ATM Configuration Guide* for more information on the QoS classes. Note that the CBR QoS class is not supported. Traffic descriptors are used to route traffic over ATM trunks only. For Frame trunks, best-effort attempts are made to meet the CIR.
Type	Displays the forward and reverse traffic descriptor type. See the *NavisCore ATM Configuration Guide* for more information.
PCR (cells/sec) (*ATM Only*)	Displays the Peak Cell Rate (PCR), in cells per second, for the traffic flow across the point-to-point LSP connection in the forward and reverse direction. See the *NavisCore ATM Configuration Guide* for more information.
SCR (cells/sec) (*ATM Only*)	Displays the Sustained Cell Rate (SCR), in cells per second, for the traffic flow across the point-to-point LSP connection in the forward and reverse direction. See the *NavisCore ATM Configuration Guide* for more information.

Table 12-3. Set All Point-to-Point LSP Connections Dialog Box: Fields (Continued)

Field	Action/Description
MBS (cells) (*ATM Only*)	Displays the Maximum Burst Size (MBS), in number of cells, for the traffic flow across the point-to-point LSP connection in the forward and reverse direction. See the *NavisCore ATM Configuration Guide* for more information.
MCR (cells/sec) (*ATM Only*)	Displays the Minimum Cell Rate (MCR), in cells per second, for the traffic flow across the point-to-point LSP connection in the forward and reverse direction. See the *NavisCore ATM Configuration Guide* for more information.
CIR (Kps)	Displays the forward and reverse CIR for the point-to-point LSP connection in kilobits per second (Kps). The CIR is the rate at which the network transfers data in the forward or reverse direction under normal conditions over Frame trunks. Normal conditions refer to a properly designed network with ample bandwidth and switch capacity. Over Frame trunks, the CIR is used by VNN OSPF to reserve bandwidth, and is not used for rate enforcement.
Oper Info	
Oper Info	Displays *Up* or *Down* to indicate the current operational status of a selected point-to-point LSP connection.
Hop Count	Displays the number of hops used in the path for a selected point-to-point LSP connection.
Using Defined Path	Displays one of the following values: *Yes* – Indicates that the point-to-point LSP connection uses a user-defined path. *No* – Indicates that the path uses the point-to-point LSP connection that was automatically defined by VNN.
Fail Reason	Displays the word *None* if no failure exists or displays the reason in the event of a failure. These fail reasons are reported by the switch to the NMS. Refer to "LSP Connection Failure Reasons" on page B-39 for a list of possible fail reasons.
Failed Node	Displays one of the following values: *No Failed Node* – Indicates that no failure exists. *A Node ID value* – Indicates a failure in the displayed node ID.
Failed Port	Displays one of the following values: *No failed port* – Indicates that no failure exists. *Logical Port Interface Number* – Indicates the number that identifies a logical port interface (that the point-to-point LSP connection is using to access the switch) in the event of a failure.
Point-Point LSP Actual Path	Displays the actual path for a selected point-to-point LSP connection.

2. Choose Add. The Add Point-to-Point LSP Connection dialog box appears (see Figure 12-10).

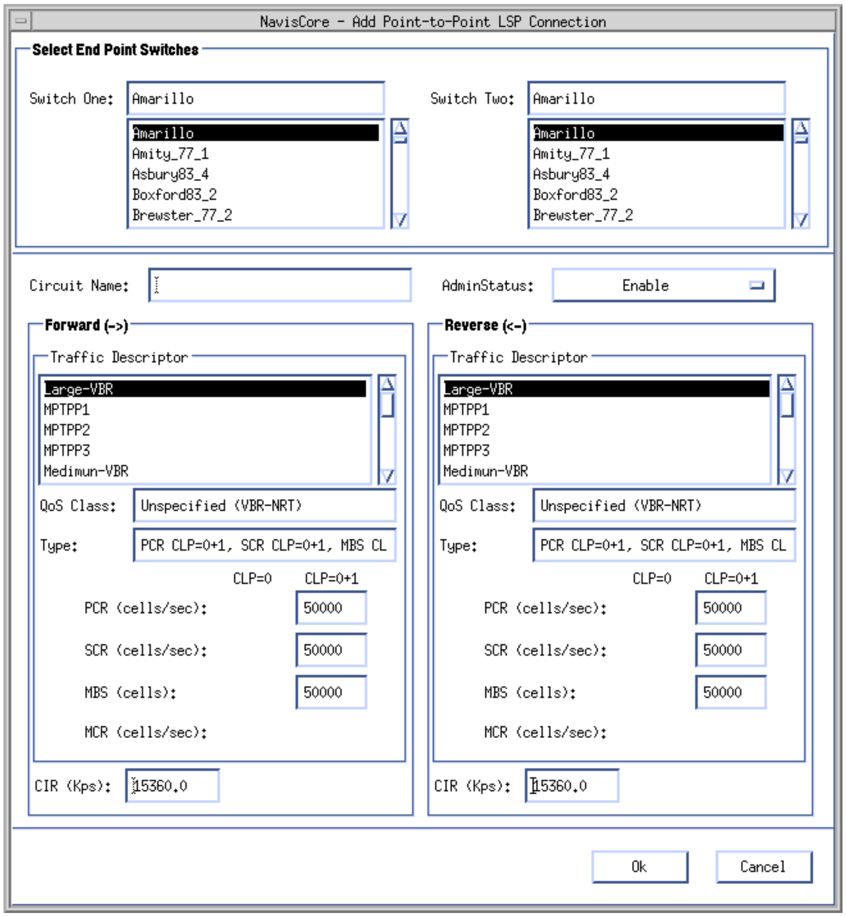

NavisCore - Add Point-to-Point LSP Connection

Select End Point Switches

Switch One: Amarillo

Switch Two: Amarillo

Circuit Name:

AdminStatus: Enable

Forward (->)

Traffic Descriptor: Large-VBR

QoS Class: Unspecified (VBR-NRT)

Type: PCR CLP=0+1, SCR CLP=0+1, MBS CL

	CLP=0	CLP=0+1
PCR (cells/sec):		50000
SCR (cells/sec):		50000
MBS (cells):		50000
MCR (cells/sec):		

CIR (Kps): 15360.0

Reverse (<-)

Traffic Descriptor: Large-VBR

QoS Class: Unspecified (VBR-NRT)

Type: PCR CLP=0+1, SCR CLP=0+1, MBS CL

	CLP=0	CLP=0+1
PCR (cells/sec):		50000
SCR (cells/sec):		50000
MBS (cells):		50000
MCR (cells/sec):		

CIR (Kps): 15360.0

Figure 12-10. Add Point-to-Point LSP Connection Dialog Box

3. Use the up and down arrows to select the two switches as the endpoints for this point-to-point LSP connection.

4. Specify the Circuit Name for this connection.

5. Select the Forward and Reverse Traffic Descriptors for the point-to-point LSP connection. The Forward Traffic Descriptors do not have to be the same as the Reverse Traffic Descriptors. See Table 12-3 on page 12-15 for more information on traffic descriptors.

6. Specify the Forward and Reverse Committed Information Rate (CIR). See Table 12-3 on page 12-15 for more information on CIR.

7. Choose OK. The Set All Point-to-Point LSP Connections dialog box reappears and the point-to-point LSP connection is included in the list of connections. VNN automatically defines the best path for a point-to-point LSP connection. However, you can use the Define Path option to create a user-defined path. See the following section, "Defining a Point-to-Point Connection Path," for details.

Defining a Point-to-Point Connection Path

VNN automatically uses the best path for a point-to-point LSP connection when you add the connection. However, you can use the Define Path option to create a user-defined path.

To define the path for a point-to-point LSP connection:

1. Select the connection from the Set All Point-to-Point LSP Connections dialog box (see Figure 12-9).
2. Choose Define Path. The Set Point-to-Point LSP Define Path dialog box appears.

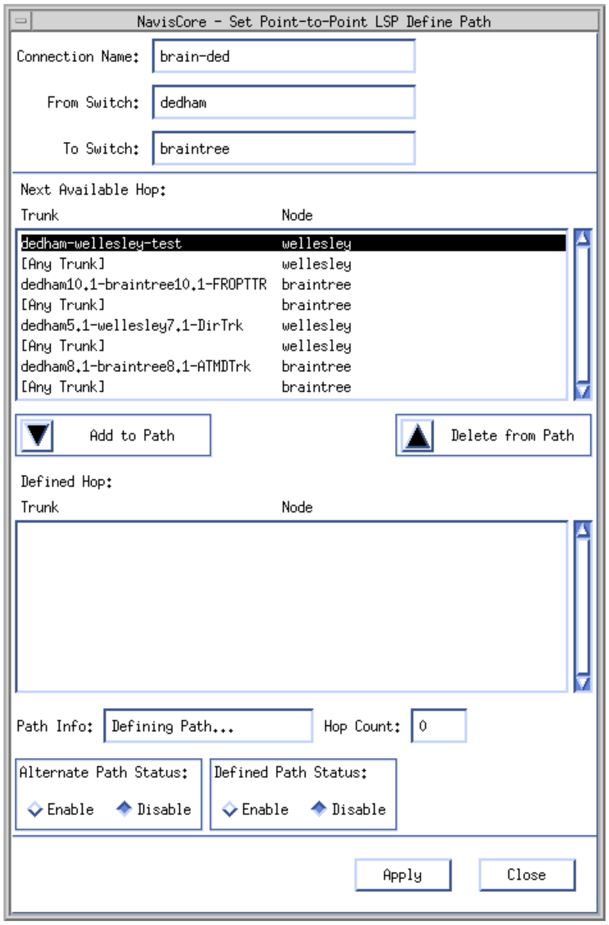

Connection Name: brain-ded

From Switch: dedham

To Switch: braintree

Next Available Hop:

Trunk	Node
dedham-wellesley-test	wellesley
[Any Trunk]	wellesley
dedham10.1-braintree10.1-FROPTTR	braintree
[Any Trunk]	braintree
dedham5.1-wellesley7.1-DirTrk	wellesley
[Any Trunk]	wellesley
dedham8.1-braintree8.1-ATMDTrk	braintree
[Any Trunk]	braintree

Defined Hop:

Trunk	Node

Path Info: Defining Path... Hop Count: 0

Alternate Path Status: ◇ Enable ◆ Disable

Defined Path Status: ◇ Enable ◆ Disable

Figure 12-11. Set Point-to-Point LSP Define Path Dialog Box

3. Use the Add to Path arrow to add a hop from the Next Available Hop list to the Defined Hop list. Use the Delete From Path arrow to delete a hop from the Defined Hop list.
4. Specify the necessary information for the Defined and Alternate Path status fields, as shown in Table 12-4.

Table 12-4. Set Point-to-Point LSP Defined Path Dialog Box: Fields

Field	Action/Description
Connection Name	Displays a unique name that identifies the point-to-point LSP connection.
From Switch	Displays a name that identifies the first switch endpoint of the point-to-point LSP connection.
To Switch	Displays a name that identifies the second switch endpoint of the point-to-point LSP connection.
Next Available Hop	Lists the trunks that are available for use in defining the path for this point-to-point LSP connection.
Defined Hop	Lists the trunks that you are using for this user-defined point-to-point LSP connection.
Path Info	Displays the status of the user-defined path.
Hop Count	Displays the number of hops used in the user-defined path for this point-to-point LSP connection.
Defined Path Status	Allows you to select Disable to indicate that the path is administratively down or Enable to indicate that the path is administratively up.
Alternate Path Status	Allows you to specify whether the alternate path (the path that VNN automatically defined before you created that user-defined path) is administratively up or down. • If you select Disable and the user-defined path for the point-to-point LSP connection fails, VNN will not use the alternate path. • If you select Enable, VNN will use the alternate path if the user-defined path fails.

Displaying Operational Status of Point-to-Point Paths

To display the operational status for a point-to-point LSP connection path:

1. Select the connection from the Set All Point-to-Point LSP Connections dialog box (see Figure 12-12).
2. Choose Oper Info. The system then polls the network and updates the information in the Oper Info portion of the dialog box if necessary.

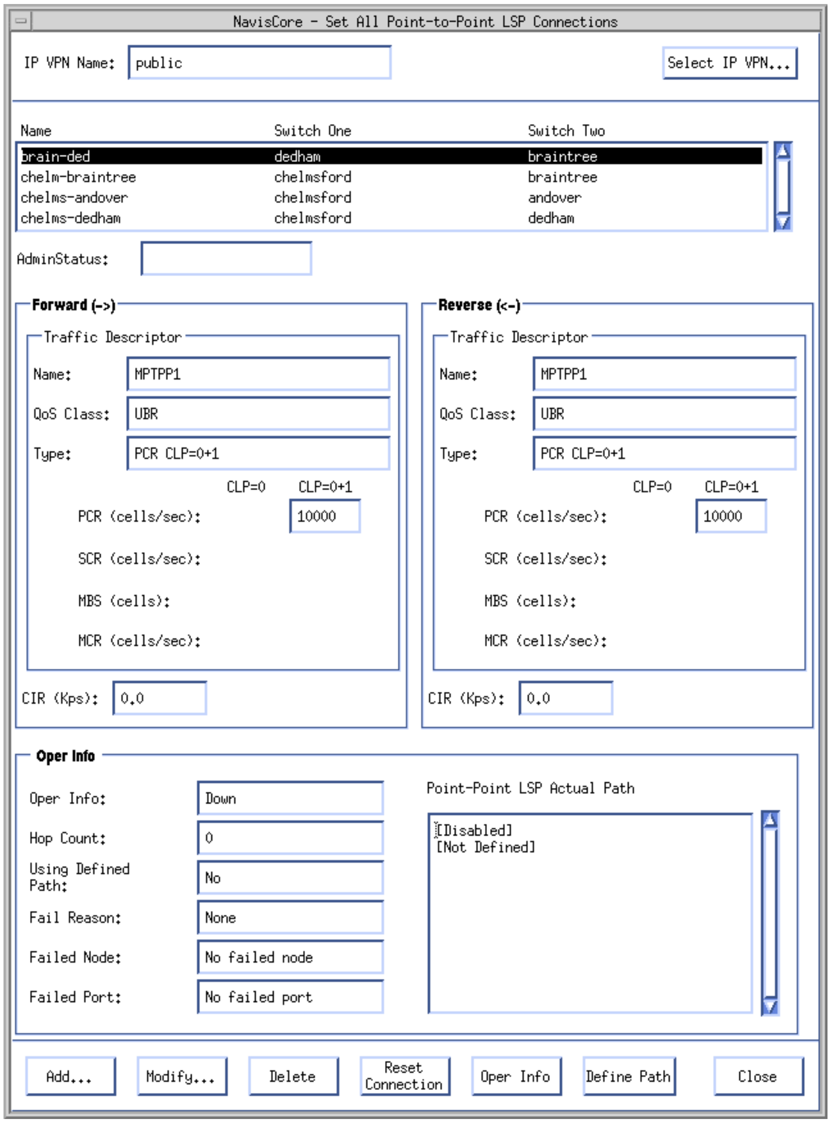

NavisCore - Set All Point-to-Point LSP Connections

IP VPN Name: public

Name	Switch One	Switch Two
brain-ded	dedham	braintree
chelm-braintree	chelmsford	braintree
chelms-andover	chelmsford	andover
chelms-dedham	chelmsford	dedham

AdminStatus:

Forward (->)

Traffic Descriptor

Name: MPTPP1

QoS Class: UBR

Type: PCR CLP=0+1

	CLP=0	CLP=0+1
PCR (cells/sec):		10000
SCR (cells/sec):		
MBS (cells):		
MCR (cells/sec):		

CIR (Kps): 0.0

Reverse (<-)

Traffic Descriptor

Name: MPTPP1

QoS Class: UBR

Type: PCR CLP=0+1

	CLP=0	CLP=0+1
PCR (cells/sec):		10000
SCR (cells/sec):		
MBS (cells):		
MCR (cells/sec):		

CIR (Kps): 0.0

Oper Info

Oper Info: Down

Hop Count: 0

Using Defined Path: No

Fail Reason: None

Failed Node: No failed node

Failed Port: No failed port

Point-Point LSP Actual Path

[Disabled]
[Not Defined]

Operational status displays for the selected connection

Figure 12-12. Displaying the Operational Status

Refer to "LSP Connection Failure Reasons" on page B-39 for a description of each of the possible failure reasons.

Configuring LSPs Over OPTimum Cell Trunks

An OPTimum cell trunk creates a switch-to-switch Lucent trunk through a Public Data Network (PDN). The Lucent OPTimum cell trunk feature allows private enterprise networks to purchase lower-cost, public-carrier services and configure OPTimum cell trunks between two Lucent switches. For more information about configuring an OPTimum cell trunk, see the *NavisCore ATM Configuration Guide*.

On CBX 500 switches, IP Navigator assigns a virtual path connection (VPC) to each MPT LSP or point-to-point LSP crossing an OPTimum trunk. These point-to-point VPCs carry MPT LSP and point-to-point LSP traffic in both directions and, therefore, reduce the number of paths required to interconnect switches in two given clusters.

On B-STDX 8000/9000 switches, IP Navigator assigns a Virtual Channel Connection (VCC) to each MPT LSP or point-to-point LSP crossing an OPTimum trunk.

Before IP Navigator can assign VPCs and VCCs, you must specify specific VPI values and ranges of VPI values for each logical port endpoint of the OPTimum trunk. You specify these values when you configure OPTimum Trunk VPI range attributes while adding an ATM OPTimum trunk logical port.

OPTimum Cell Trunk VPI Restrictions

The maximum allowable range of VPI values for an OPTimum cell trunk depends on whether the trunk’s feeder logical port is configured as UNI or NNI. (See the *NavisCore ATM Configuration Guide* for more information on feeder logical ports.) The values range from 1 through 255 for a UNI logical port and 1 through 4095 for an NNI logical port. Since you can specify more than one OPTimum trunk on the same physical link, make sure that the VPI value on each trunk does not exceed these limits.

► In rare cases, previous VPI range configurations for OPTimum cell trunk logical ports may no longer work once you upgrade to this release of NavisCore software and switch software. The upgrade procedure attempts to convert the old values and is successful in most cases. In the rare cases in which the upgrade procedure is not successful, you may have to reconfigure VPI ranges for these logical ports according to the rules and restrictions described in this section.

The following restrictions also apply:

- Specify a single VPI value for the default virtual channel connection (VCC). This connection is used for network management and virtual circuit control on the OPTimum trunk. This value must be outside the VPI value ranges you configure for transit MPT LSP, transit point-to-point LSP connections, and virtual UNI logical ports. It must also not conflict with the single VPI value you specify for the MPT LSP root.
- Specify a single VPI value for the MPT LSP root. If the switch where the OPTimum trunk logical port endpoint resides is also the root of an MPT LSP, a VPI is needed for the MPT LSP root. You can specify 0 if the switch is not an MPT LSP root. The value must be an even value between 2 and 30 (e.g., 4), and it must be outside the VPI value ranges you configure for transit MPT LSPs, transit point-to-point LSP connections, virtual UNI logical ports, and the default VCC on CBX 500 switches.
- Specify a range of VPI values for virtual UNI logical ports that use the OPTimum trunk. The range may be 0 if no PVC virtual paths traverse the trunk. This range cannot overlap the VPI value ranges you configure for transit MPT LSPs, transit point-to-point LSP connections, the MPT LSP root, and the default VCC on CBX 500 switches. See the *NavisCore ATM Configuration Guide* for more information on virtual UNI logical ports.
- Specify a range of VPI values for transit MPT LSPs. Transit MPT LSPs are root-to-leaf connections that traverse the trunk. This range cannot overlap the VPI value ranges you configure for virtual UNI logical ports, transit point-to-point LSP connections, the MPT LSP root, and the default VCC on CBX 500 switches.

  This range can be 0 if you do not want transit MPT LSPs to be established across the trunk, in which case the switches at the trunk endpoints must perform an IP routing table lookup to forward data. You may want to adopt this approach if you want to conserve VPIs.
- Specify a range of VPI values for transit point-to-point LSP connections. Transit point-to-point LSP connections are connections that traverse the trunk. This range cannot overlap the VPI value ranges you configure for virtual UNI logical ports, transit MPT LSPs, the MPT LSP root, and the default VCC on CBX 500 switches.

  This range can be 0 if you do not want transit point-to-point LSP connections to be established across the trunk, in which case the switches at the trunk endpoints must perform an IP routing table look-up in order to forward data. You may want to adopt this approach if you want to conserve VPIs.

► Switches that act as OPTimum cell trunk endpoints cannot also be point-to-point LSP connection endpoints. They may only act as transit switches for point-to-point LSP connections. If you configure a switch as both an OPTimum cell trunk endpoint and as a point-to-point LSP endpoint, the point-to-point LSP will fail.

It is important that ranges of VPI values do not overlap. For example, if you configure 2-to-30 as the range of VPIs for transit MPT LSPs, the range of VPIs for virtual UNI logical ports should be outside that range. Also, VPIs for virtual UNI logical ports should not fall in the transit point-to-point LSP connection range.

MPT LSP Traffic Forwarding Across OPTimum Cell Trunks

Table 12-5 summarizes how traffic is forwarded between switches at the leaves of MPT LSPs and switches at the root of MPT LSPs when the roots and switches are separated by an OPTimum cell trunk.

Table 12-5. MPT LSP Traffic Forwarding Across OPTimum Cell Trunks

Root	Leaf	How Data is Forwarded
B-STDX 8000/9000	B-STDX 8000/9000	Data is forwarded over the default VCC VPI. This value is configured using the Opt Trunk VPI Range fields in the Add Logical Port dialog box (see Figure 12-13 on page 12-24). This VPI may be different on the trunk endpoint switches if the intermediate ATM network that the trunk traverses remaps them as it sees fit.
CBX 500	CBX 500	Data is forwarded over a virtual path connection (VPC) that uses one of the following VPIs: – The MPT LSP VPI, which is configured using the Opt Trunk VPI Range fields in the Add Logical Port dialog box (see Figure 12-13 on page 12-24). – A VPI from the range of transit MPT LSP VPIs configured using the Opt Trunk VPI Range fields in the Add Logical Port dialog box (see Figure 12-13 on page 12-24). This VPI may be different on the trunk endpoint switches if the intermediate ATM network that the trunk traverses remaps them as it sees fit.
B-STDX 8000/9000	CBX 500	Data is forwarded over the default VCC VPI. This value is configured in the Opt Trunk VPI field on the Add Logical Port dialog box (see Figure 12-13 on page 12-24). This VPI may be different on the trunk endpoint switches if the intermediate ATM network that the trunk traverses remaps them as it sees fit.
CBX 500	B-STDX 8000/9000	Data is forwarded over a VPC VPI. Note that the VPIs configured on the CBX 500 should match the VPIs configured on nodes at the edge of the intermediate ATM network toward the B-STDX switch.

Configuring an ATM OPTimum Cell Trunk for LSP Traffic

Use the following steps to configure an ATM OPTimum cell trunk for MPT LSP and point-to-point LSP traffic:

1. Access the Add Logical Port dialog box. For complete details about how to access the Add Logical Port dialog box and information about provisioning OPTimum cell logical ports and trunks, see the *NavisCore ATM Configuration Guide.*

2. Access the Opt Trunk VPI Range attributes on the Set Attributes menu. NavisCore displays the dialog box shown in Figure 12-13.

   For complete details about how to access the Set Attributes menu, see the *NavisCore ATM Configuration Guide.*

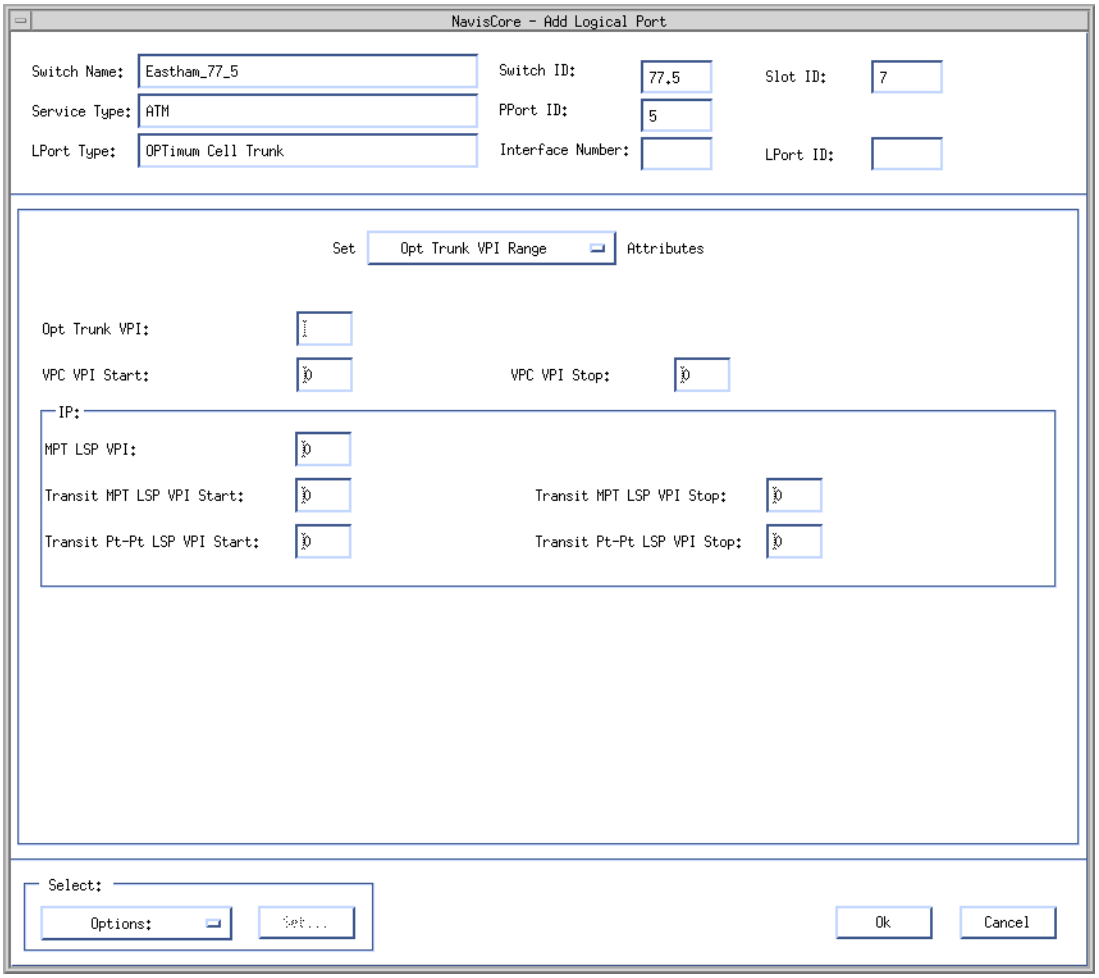

NavisCore - Add Logical Port

Switch Name: Eastham_77_5
Service Type: ATM
LPort Type: OPTimum Cell Trunk
Switch ID: 77.5
PPort ID: 5
Interface Number:
Slot ID: 7
LPort ID:

Set Opt Trunk VPI Range Attributes

Opt Trunk VPI:
VPC VPI Start: 0
VPC VPI Stop: 0

IP:
MPT LSP VPI: 0
Transit MPT LSP VPI Start: 0
Transit MPT LSP VPI Stop: 0
Transit Pt-Pt LSP VPI Start: 0
Transit Pt-Pt LSP VPI Stop: 0

Figure 12-13. Add Logical Port – Opt Trunk VPI Range Dialog Box

3. Specify the specific VPI values and value ranges in the fields described in Table 12-6.

Table 12-6. Opt Trunk VPI Range Attributes Fields

Field	Action/Description
Opt Trunk VPI	Specify the VPI value for the default virtual channel connection (VCC), which is used for network management and virtual circuit control on the OPTimum trunk. This value must be outside the VPI value ranges you configure for transit MPT LSPs, transit point-to-point LSP connections, and virtual UNIs. It must also not conflict with the MPT LSP VPI value.
VPC VPI Start (*OPTimum cell trunk logical port endpoints on CBX 500 and GX 550 switches only*)	Specify the first VPI value in the range of VPI values for virtual UNI logical ports that use the OPTimum trunk. For example, if the range is 155 to 255, you would specify 155 in this field. The default is 0. The range that you specify must not overlap the ranges that you specify for transit MPT LSPs and transit point-to-point LSP connections. It must also not conflict with the Opt Trunk VPI value and the MPT LSP VPI value. See the *NavisCore ATM Configuration Guide* for more information on virtual UNI logical ports.
VPC VPI Stop (*OPTimum cell trunk logical port endpoints on CBX 500 and GX 550 switches only*)	Specify the last VPI value in the range of VPI values for virtual UNI logical ports that use the OPTimum trunk. For example, if the range is 155 to 255, you would specify 255 in this field. The range that you specify must not overlap the ranges that you specify for transit MPT LSPs and transit point-to-point LSP connections. It must also not conflict with the Opt Trunk VPI value and the MPT LSP VPI value. See the *NavisCore ATM Configuration Guide* for more information on virtual UNI logical ports.
MPT LSP VPI (*OPTimum cell trunk logical port endpoints on CBX 500 and GX 550 switches only*)	Specify the VPI value for the MPT LSP root. If the switch where the OPTimum cell trunk logical port endpoint resides is also the root of an MPT LSP, a VPI is needed for the MPT LSP root. The value must be an even value between 2 and 30 (e.g., 4). The default is 0. The MPT LSP VPI value must be outside the VPI value ranges you configure for transit MPT LSPs, transit point-to-point LSP connections, and virtual UNI logical ports. It must also not conflict with the Opt Trunk VPI value.
Transit MPT LSP VPI Start (*OPTimum cell trunk logical port endpoints on CBX 500 and GX 550 switches only*)	Specify the first VPI value in the range of VPI values for transit MPT LSPs. The default is 0. The range that you specify must not overlap the ranges that you specify for virtual UNI logical ports and transit point-to-point LSP connections. It must also not conflict with the Opt Trunk VPI value and the MPT LSP VPI value.

Table 12-6. Opt Trunk VPI Range Attributes Fields (Continued)

Field	Action/Description
Transit MPT LSP VPI Stop	Specify the last VPI value in the range of VPI values for transit MPT LSPs. The default is 0. The range that you specify must not overlap the ranges that you specify for virtual UNI logical ports and transit point-to-point LSP connections. It must also not conflict with the Opt Trunk VPI value and the MPT LSP VPI value.
Transit Pt-Pt LSP VPI Start	Specify the first VPI value in the range of VPI values for transit point-to-point LSP connections. If you enter 0, the transit point-to-point LSP connection is disabled. The range that you specify must not overlap the ranges that you specify for virtual UNI logical ports and transit MPT LSPs. It must also not conflict with the Opt Trunk VPI value and the MPT LSP VPI value.
Transit Pt-Pt LSP VPI Stop	Specify the last VPI value in the range of VPI values for transit point-to-point LSP connections. If you enter 0, the transit point-to-point LSP connection is disabled. The range that you specify must not overlap the ranges that you specify for virtual UNI logical ports and transit MPT LSPs. It must also not conflict with the Opt Trunk VPI value and the MPT LSP VPI value.

4. Choose OK.

About Next Hop Resolution Protocol

This chapter provides:

- An overview of the Next Hop Resolution Protocol (NHRP)
- An overview of Lucent's implementation of NHRP

Overview of NHRP

NHRP is a Non-Broadcast Multiple Access (NBMA) address resolution protocol that, when supplied with a network-layer address (such as an IP address), provides a source station (for example, host or router) with the NBMA address of the "next hop" toward a destination station. (Examples of NBMA addresses are ATM End System Addresses [AESAs] and native E.164 addresses, which resemble telephone numbers.) Data packets do not have to pass through extra hops of routers when the source station and destination station belong to different Logical Internet Subnets (LIS) of the same logical NBMA.

▶ Although the NHRP protocol can support various network-layer and NBMA-layer address types, the Lucent IP Navigator implementation only issues and processes NHRP packets for IP (IPv4) over ATM. See "Lucent NHRP Implementation" on page 13-13 for more information on the Lucent NHRP implementation.

NHRP enables a shortcut SVC connection to be established between two points in the network so that they can exchange data packets without the services of intermediate routers. By supporting these types of connections, the underlying NBMA network facilities, such as end-to-end QoS guarantees, may be fully utilized when transmitting network-layer packets. For example, by using shortcut routing, ATM-provided QoS guarantees can be implemented without having to reassemble IP packets at each network-layer hop.

In the protocol's basic form, the “next hop” NBMA address to be resolved is that of the destination station itself when the destination station is directly connected to the logical NBMA subnetwork; otherwise, the “next hop” NBMA address to be resolved is that of the egress router for the NBMA subnetwork that is closest to the destination.

For example, in Figure 13-1, three CPE are directly connected to the logical NBMA network, and one CPE is not. When the three CPE that are directly connected to the logical NBMA network are the destination stations, their NBMA addresses are resolved. When the fourth CPE that is not part of the logical NBMA network is the destination station, the NBMA address of the switch to which the CPE is connected is the address that is resolved.

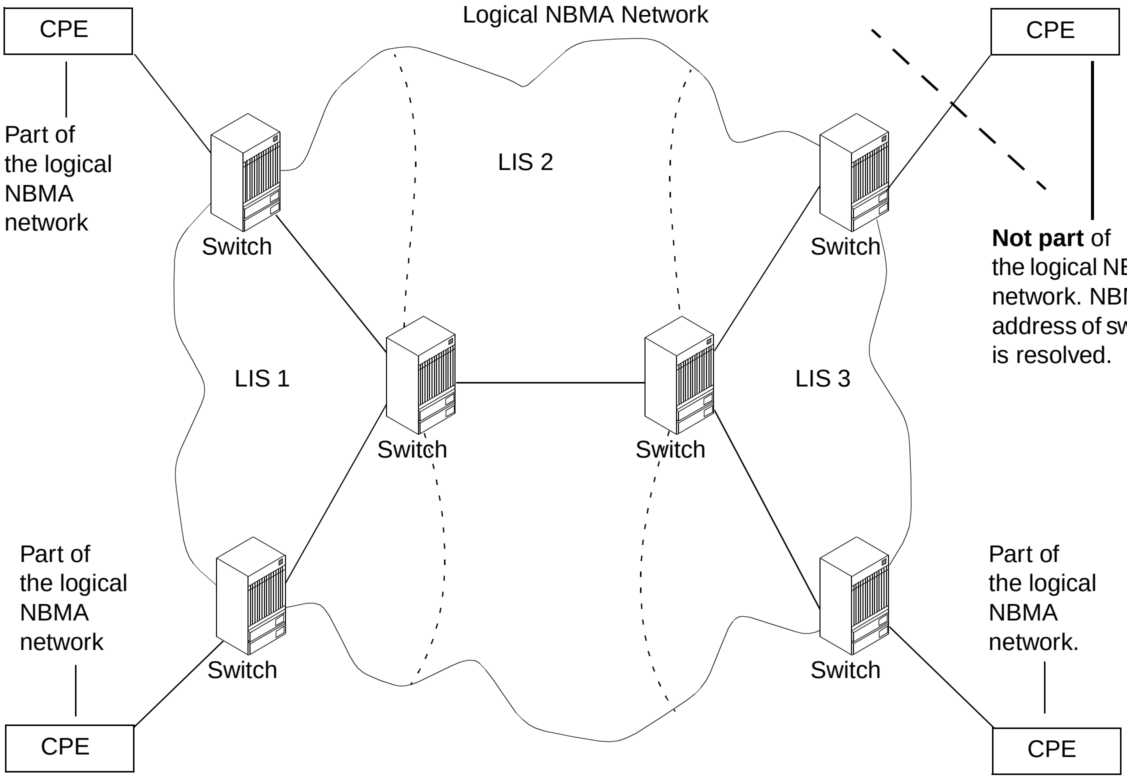
CPE
Logical NBMA Network
CPE
Part of the logical NBMA network
LIS 2
Switch
Switch
Not part of the logical NBMA network. NBMA address of switch is resolved.
LIS 1
LIS 3
Switch
Switch
Part of the logical NBMA network
Switch
Switch
Part of the logical NBMA network.
CPE
CPE

Figure 13-1. Sample Logical NBMA Network

NHRP Components

NHRP uses the client/server paradigm. It consists of Next Hop Server (NHS) components and Next Hop Client (NHC) components communicating via NHRP requests and replies in order to resolve NBMA addresses. The same node may act as both an NHS and an NHC. Figure 13-2 shows some sample NHCs and NHSs exchanging NHRP requests and replies.

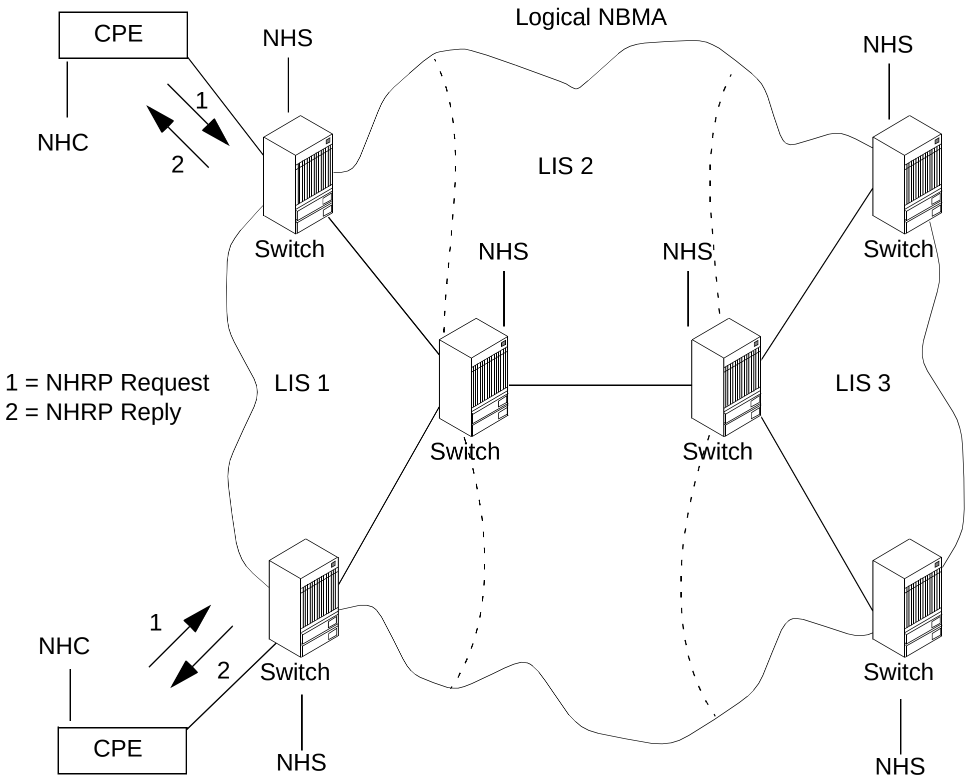

Figure 13-2. Sample NHCs and NHSs

NHRP uses network-layer routing in resolving destination addresses by using local routing tables to determine how to forward NHRP request and replies. Initiating an NHRP resolution request from an NHC to establish a shortcut is an application-dependent issue, and may depend on the cost trade-off between the data flow and shortcut setup factors of the environment; therefore, NHRP does not prohibit source stations from transmitting data packets to destinations using routing mechanisms. In fact, data packets may be transmitted using routing mechanisms while the NHRP shortcut is being established.

NHS Component

An NHS is an entity that performs NHRP services. It maintains the bindings of the network layer to NBMA layer addresses for its clients. When an NHS receives a request that it is unable to fulfill, it attempts to pass this request along the routed path towards the destination to the next NHS. In the protocol's basic form, NHSs communicate with each other simply by forwarding NHRP packets to the next network-layer hop; this next network layer hop is determined by performing a network layer routing lookup on the requested destination address. Therefore, a contiguous deployment of NHSs along the routed path towards a destination is an essential requirement for the NHRP protocol.

NHC Component

An NHC is an entity that initiates various types of NHRP requests. An NHC may send an NHRP Registration Request to an NHS to register its network layer to NBMA layer address mapping. An NHC may also initiate an NHRP Resolution Request to determine an NBMA address resolution of another destination. A proxy client is an entity that acts as an NHC for one or more sources of IP traffic. Examples of these sources include NHCs or IP traffic flows. See “NHS and NHC Support on Lucent Switches” on page 13-4 for a description of IP traffic flows.

NHS and NHC Support on Lucent Switches

Lucent switches may act as both:

- An NHS
- A proxy client

NHS Support

In many cases, you need to create only one NHS server instance, and map that instance to one or more IP logical ports or IP server logical ports. However, you may create multiple NHS instances on a Lucent switch, one per ATM or Frame Relay logical port (or IP server logical port) that also supports IP. For example, if the switch serves two ATM networks, you may want to maintain separate NHRP databases by creating two separate NHS server instances (one for each network).

▶ NHRP shortcuts cannot be established through Frame Relay logical ports. See “Frame Relay UNI and PPP Logical Ports” on page 13-22 for details.

You explicitly configure an NHS for an ATM or Frame Relay UNI logical port that serves NHCs (these ports must also be configured for IP). See “Placement of NHS-capable Switches and the Proxy Client” on page 13-16 for more information.

For trunk logical ports, you do not explicitly configure an NHS. Instead, a special NHS instance called the *default server* forwards NHRP traffic on these ports. See "Placement of NHS-capable Switches and the Proxy Client" on page 13-16 for more information.

Proxy Client Support

A proxy client can provision bandwidth and QoS guarantees for IP traffic flows. An IP flow is a set of packets that match a particular profile, called a *flow profile*, containing definitions of source/destination IP addresses with CIDR masks, IP protocol, source/destination protocol port numbers, and Type of Service (TOS) requirements. See "Guaranteeing QoS for IP Traffic Flows" on page 13-25 for more information on flow profiles.

Only one proxy client instance exists on a Lucent switch, but it can send NHRP Resolution Requests for multiple logical ports, as long as you enable proxy client functionality on those ports.

Creating Shortcuts

Once the NBMA address of the destination is known by a source (an NHC), an SVC may be created for a direct connection, thereby eliminating the need for subsequent data packets to traverse intermediate hops, as illustrated in Figure 13-3.

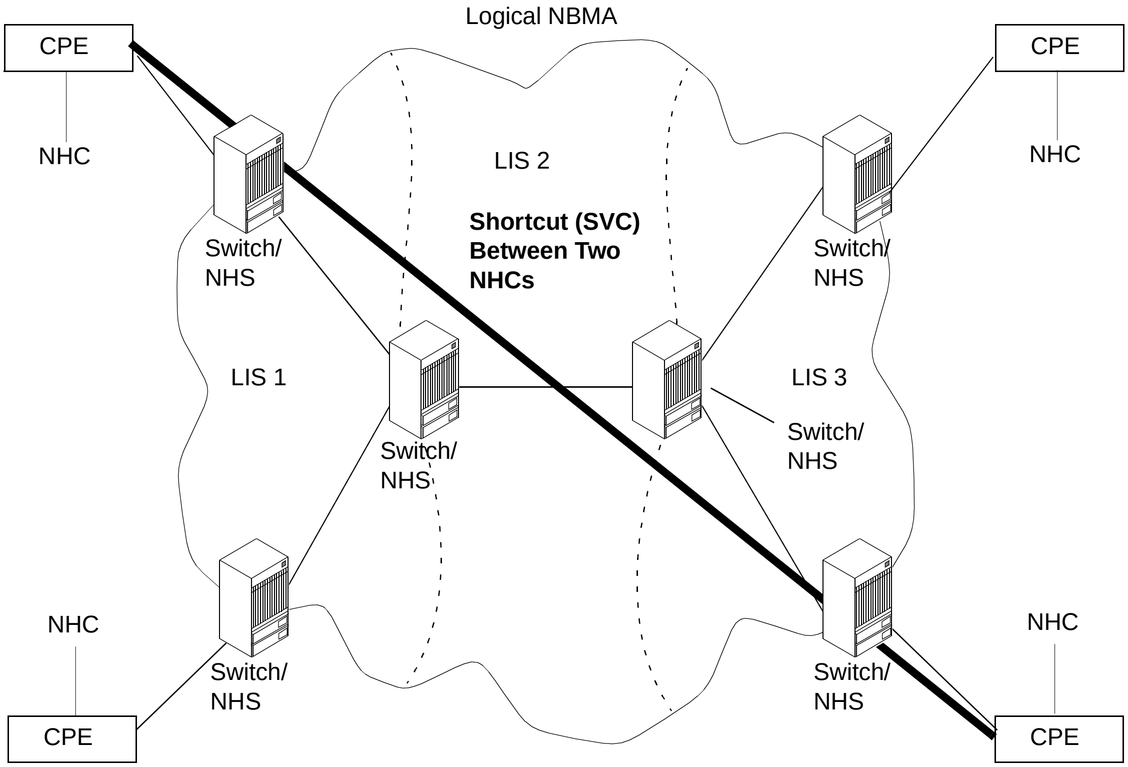

Figure 13-3. Shortcut Between Two NHCs

NHRP Protocol

The NHRP protocol supports four fundamental types of packet interactions:

- NHRP Registration Request/Reply
- NHRP Resolution Request/Reply
- NHRP Purge Request/Reply
- NHRP Error Indication

NHRP Registration Request/Reply

An NHC may first register its network-layer/NBMA address mapping with a serving NHS. Typically, the network-layer address is an IP address, though NHRP can potentially work with protocol suites other than TCP/IP. In most cases, the serving NHS is also the NHC’s default router, but the serving NHS may be a network node other than the default router.

To register with the NHS, the NHC sends an NHRP Registration Request to the NHS. If the NHC is configured with the serving NHS's address, the NHS's network-layer address may be used in the Destination Address of the NHRP Registration Request. Otherwise, the NHC's own network layer address is put in both the Destination and Source Address fields of the Request.

When an NHS receives an NHRP Registration Request from an NHC, the NHS should process the request if one of the following conditions is met:

- The Destination Address is equal to its own
- The Source and Destination Addresses are equal and the Destination Address represents the LIS that this NHS is serving

Otherwise, the NHS should forward the NHRP Registration Request along the routed path.

In response to an NHRP Registration Request, the serving NHS should send an NHRP Registration Reply over a direct connection to the requesting NHC. Each registration has a hold time value for which the registration is valid; therefore, the NHC should periodically transmit the NHRP Registration Requests to prevent the registration from timing out.

Figure 13-4 depicts the NHRP Registration process. NHC A and NHC B send NHRP Registration Requests to NHS A and NHS B respectively, and NHS A and NHS B respond to NHC A and NHC B with NHRP Registration Replies.

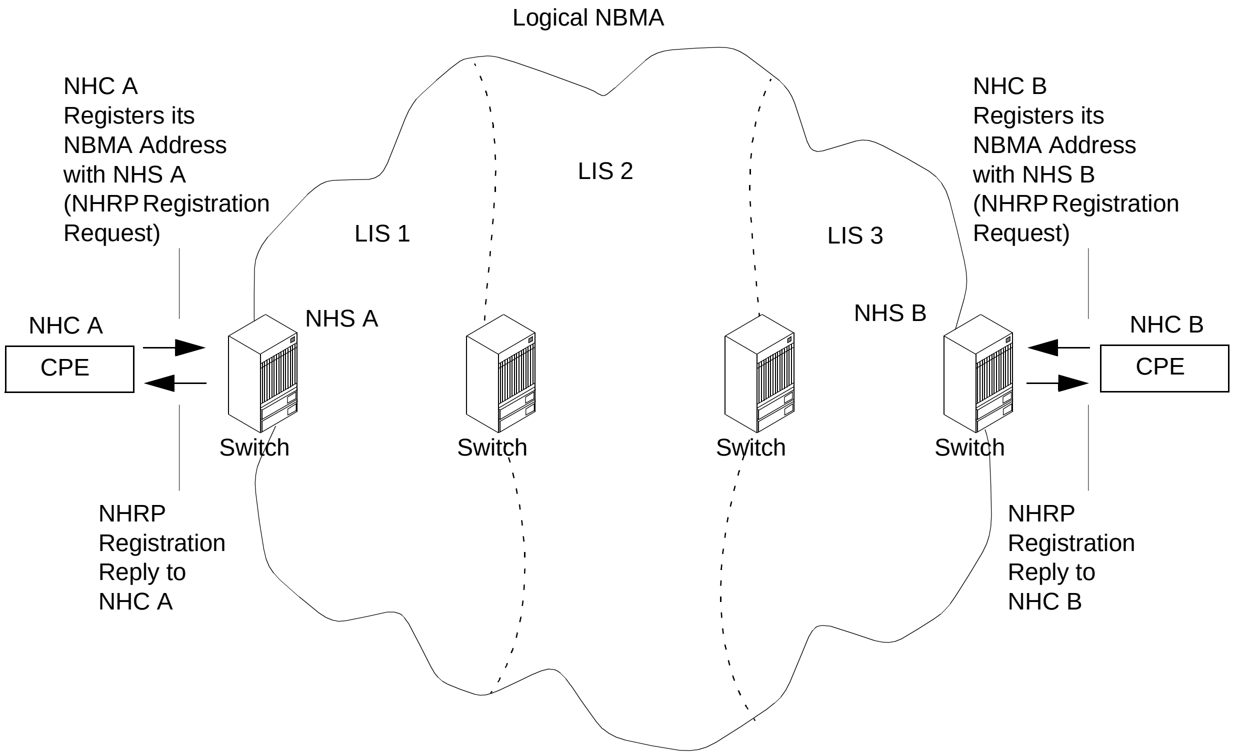

Figure 13-4. NHRP Registration Process

NHRP Resolution Request/Reply

When an NHC (or proxy client) attempts to resolve an NBMA address, it issues an NHRP Resolution Request toward the routed path.

The manner in which an NHC initiates an NHRP Resolution Request depends on the specific application. For example, an NHC might generate an NHRP Resolution Request when a data packet addressed to a particular network-layer destination (for which an NBMA address is unresolved) is forwarded from either a host or a transit router. This Resolution Request contains the destination's network layer address, the source's network layer address, and the source's NBMA layer address. While it awaits an NHRP Resolution Reply, the NHC may choose to perform one of the following actions:

- Drop the data packet
- Retain the data packet until the reply is received
- Forward the data packet along the routed path

When an NHS receives an NHRP Resolution Request, it initially determines whether the requested network-layer destination address matches any of its cached entries. If there is no match, the NHS performs a routing table lookup and forwards the NHRP Resolution Request along the routed path to the downstream NHS. If the NHS has a cached match, it generates a positive NHRP Resolution Reply, which is forwarded back along the routed path toward the requesting NHC.

Figure 13-5 and Figure 13-6 illustrate the NHRP resolution transaction. This scenario assumes the following conditions:

- NHC B already registered its NBMA address with NHS B
- All protocol operations are successful
- No synchronization problems have occurred

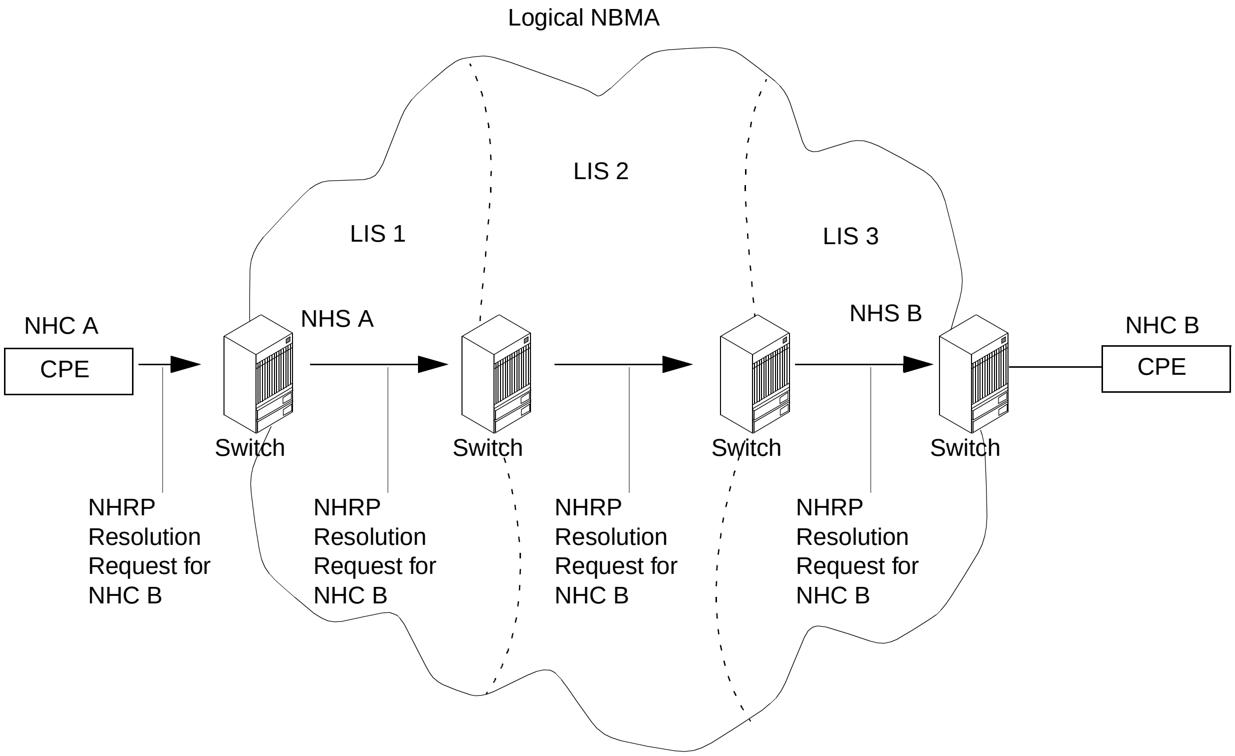
Logical NBMA
LIS 2
LIS 1
LIS 3
NHC A
CPE
NHS A
NHS B
NHC B
CPE
Switch
Switch
Switch
Switch
NHRP Resolution Request for NHC B
NHRP Resolution Request for NHC B
NHRP Resolution Request for NHC B
NHRP Resolution Request for NHC B

Figure 13-5. NHRP Resolution Request for NHC B's NBMA Address

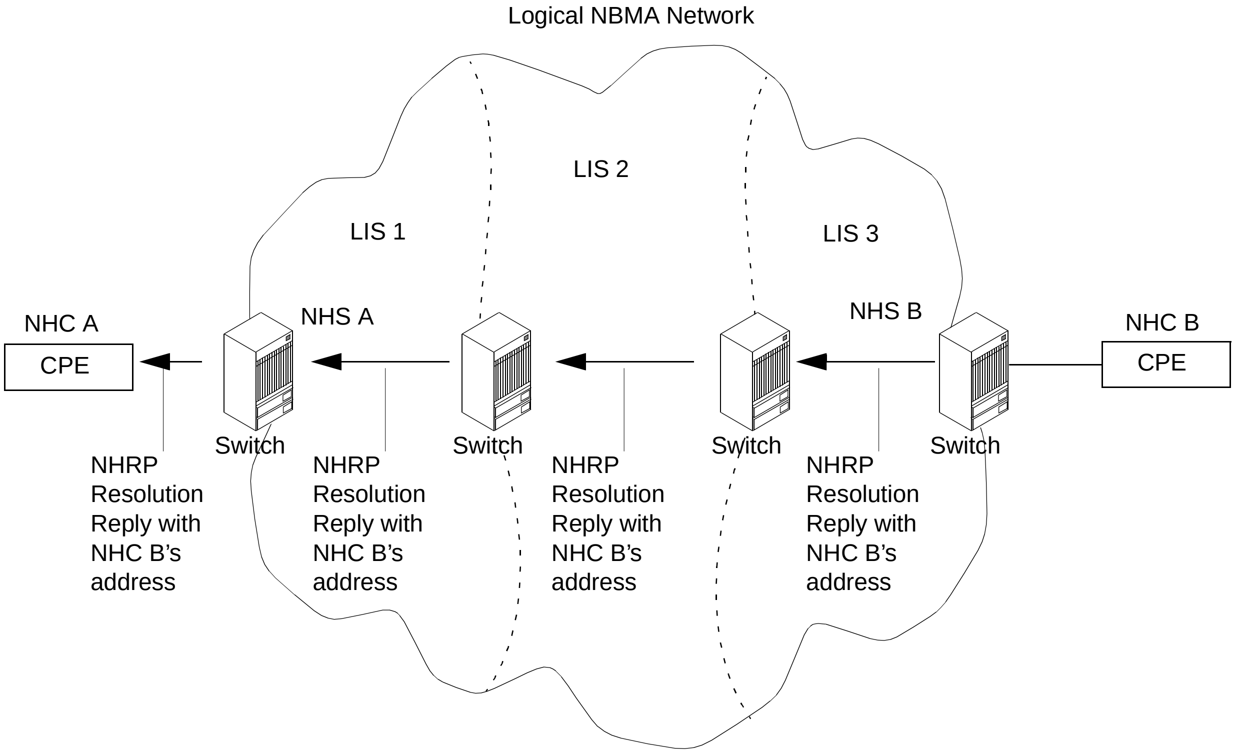
Logical NBMA Network
LIS 1
LIS 2
LIS 3
NHC A
CPE
NHS A
NHS B
NHC B
CPE
Switch
Switch
Switch
Switch
NHRP Resolution Reply with NHC B's address
NHRP Resolution Reply with NHC B's address
NHRP Resolution Reply with NHC B's address
NHRP Resolution Reply with NHC B's address

Figure 13-6. NHRP Resolution Reply with NHC B's NBMA Address

An NHS returns a negative NHRP Resolution Reply if no NHSs in the logical NBMA network can reply to the NHRP Resolution Request. This event could occur when an NHS is unable to resolve the request itself and one of the following conditions is met:

- The NHS is unable to forward the packet because there is no corresponding entry in the routing table
- The NHS is the egress node closest to the destination, but does not have the appropriate network-layer-to-NBMA mapping

NHRP Purge Request

The purpose of the NHRP Purge Request is to invalidate cached address resolution information before it expires due to the holding time. Either an NHS or an NHC may initiate the NHRP Purge Request.

The need to invalidate cached address resolution information might arise as a result of a change in the relationship between the NHC and serving NHS. If the NHC issues an NHRP Purge Request, the same destination addressing rules as discussed in "NHRP Resolution Request/Reply" on page 13-7 apply. If a serving NHS receives an NHRP Purge Request from its client, and the NHS previously responded to one or more NHCs with an authoritative NHRP Resolution Reply containing the purged information, the serving NHS must purge those NHCs of the information by sending a new NHRP Purge Request to each NHC.

▶ A requesting NHC sends an authoritative NHRP Resolution Request when it wants to receive an NHRP Resolution Reply from the *serving NHS* only. The serving NHS is the NHS with which the destination NHC has registered its NBMA address.

In response to an NHRP Purge Request, the receiving entity purges the requested information from its cache, if it exists. Then, the receiving entity sends an NHRP Purge Reply to the node identified by the source address contained in the NHRP Purge Request. When the NHC initiates an NHRP Purge Request, the NHS must send the NHRP Purge Reply over a direct connection between the serving NHS and the NHC. The entity that sends the NHRP Purge Request may retransmit it periodically until the request is acknowledged or the holding time for the purged address resolution information has expired.

NHRP Error Indication

The NHC or NHS sends an NHRP Error Indication packet to indicate an error in a received NHRP packet. Some error examples include:

- Unrecognized Extension
- Invalid Extension
- NHRP Loop Detected
- Protocol Address Unreachable
- Protocol Error
- SDU Size Exceeded
- Invalid Resolution Reply Received
- Hop Count Exceeded

When an NHC or NHS generates an NHRP Error Indication packet, the NHC or NHS discards the offending NHRP packet.

NHRP Extensions

To give you administrative flexibility, NHRP supports several optional extensions to the NHRP packets. An NHS must not change the order of the extensions. Lucent switches currently support the following extensions:

- Responder Address. See "Responder Address Extension" later in this section.
- Route Record. See "Route Record Extension" later in this section.

▶ Lucent switches do not currently support the Authentication extension.

The NHRP packet contains a compulsory flag which indicates whether it is mandatory for a receiver to process the extension. If this flag is set, but the responder is unable to process the extension, the responder sends an NHRP Error Indication. If the compulsory flag is cleared, then the responder can safely ignore the extension and can return the extension unchanged in the NHRP Reply packet.

If a transit NHS cannot process an extension with the compulsory flag set, it just forwards the packet including the extensions — it is not allowed to cache any resolution data from the NHRP Resolution Reply.

Responder Address Extension

The Responder Address Extension is used to determine the address of the NHRP responder. This address is not the same as the "next-hop" address in the event that a serving NHS responds to an NHC with a cached resolution address.

If an NHS generates a reply to a request containing this extension, the NHS also includes this extension with its own network-layer address in the reply. This extension might be useful in detecting routing loops.

Route Record Extension

Route record extensions — known as the *NHRP Forward Transit NHS Record Extension* and the *NHRP Reverse Transit NHS Record Extension* — contain a list of transit NHSs that were traversed by an NHRP packet. When a transit NHS receives a packet with one of these extensions, it appends its network-layer address to the extension, and updates the extension length field and the packet checksum field.

The responding NHS does not update the extension. These extensions may be useful for loop detection, diagnostic tracing, and subnetwork-layer filtering detection.

NHRP Feature Summary

In summary, NHRP has the following key features:

- Resolves “next hop” NBMA addresses, regardless of whether the destination station is in the same LIS as the source station.
- Avoids extra hops in an NBMA with multiple LISes.
- Provides several optional extensions.
- Deals with unidirectional data flows.
- Is not specific to a particular NBMA technology or network layer technology.
- Can be used in host-to-host, host-to-router, and router-to-router communications, but some additional efforts should be applied in router-to-router communications to avoid persistent routing loops.
- Uses the client/server paradigm consisting of Next Hop Server (NHS) components and Next Hop Client (NHC) components.
- Supports Proxy Next Hop Clients.
- Is not a routing protocol, but depends upon present and future routing protocols.
- Does not prohibit a source station from using conventional router mechanisms to transmit packets, instead of establishing NHRP shortcuts.

Lucent NHRP Implementation

This section provides notes on Lucent’s NHRP implementation and related configuration and management issues.

Configuration and Management Notes

This section provides notes on NHRP configuration and management issues you will encounter.

NHRP and IP VPNs

NHRP resources are public resources (that is, they are assigned to the public IP VPN). They cannot be reserved for exclusive use by private IP VPNs. The sources and destinations of customer data that traverse NHRP SVC shortcuts, as well as control traffic, must not be part of a private IP VPN. For example, if an NHRP traffic flow traverses an ATM PVC from the CPE to an ingress IP logical port, the VPI/VCI must not be assigned to a private IP VPN. Rather, it must be assigned to the IP VPN labeled as “public.” See Chapter 16, “Configuring IP Virtual Private Networks” for more information on IP VPNs.

SVC Node Prefixes

On each egress switch that processes NHRP traffic, you must configure an SVC node prefix. See “Configuring SVC Node Prefixes” on page 14-9 and the *NavisCore ATM Configuration Guide* for more information on configuring SVC node prefixes.

For example, in Figure 13-8 on page 13-17, you would configure an SVC node prefix on Switch 4, since Switch 4 is the egress switch. In Figure 13-9 on page 13-18, you would configure an SVC node prefix on Switch 1, since Switch 1 is the egress switch.

NHRP Logical Ports

You must add an NHRP logical port on an IP logical port (or IP Server logical port) that is associated with an NHS or a proxy client. See “Adding and Deleting NHRP Logical Ports” on page 14-10 for more information on adding NHRP logical ports.

Addressing

Although the NHRP protocol can support various network-layer and NBMA-layer address types, the Lucent IP Navigator implementation only issues and processes NHRP packets for IP (IPv4) over ATM, which includes the following address formats:

- Native E.164
- AESA (i.e., NSAP)
- E.164 with Network Service Access Point (NSAP)

► AESA addressing will be used in the Lucent switch for NSAP addressing.

See "About NBMA Addressing" on page 14-4 for more information on address formats.

MPT Aggregation Technology For Connecting Ingress and Egress Switches

B-STDX switches use proprietary features of the Lucent Virtual Network Navigator (VNN) to exchange NHRP packets in a Lucent network. To link all ingress and egress routers, the B-STDX switches pre-establish label switched paths (LSPs) through the use of Multipoint-to-Point Tunnel (MPT) aggregation technology. Instead of using traditional hop-by-hop IP routing to forward NHRP packets, the B-STDX switches use MPT LSPs as a best-effort message delivery mechanism.

In addition, both B-STDX and CBX switches use MPT LSPs for all data packets until the NHRP shortcut is established. Figure 13-7 shows an example of an MPT LSP with three switches that act as ingress and egress routers.

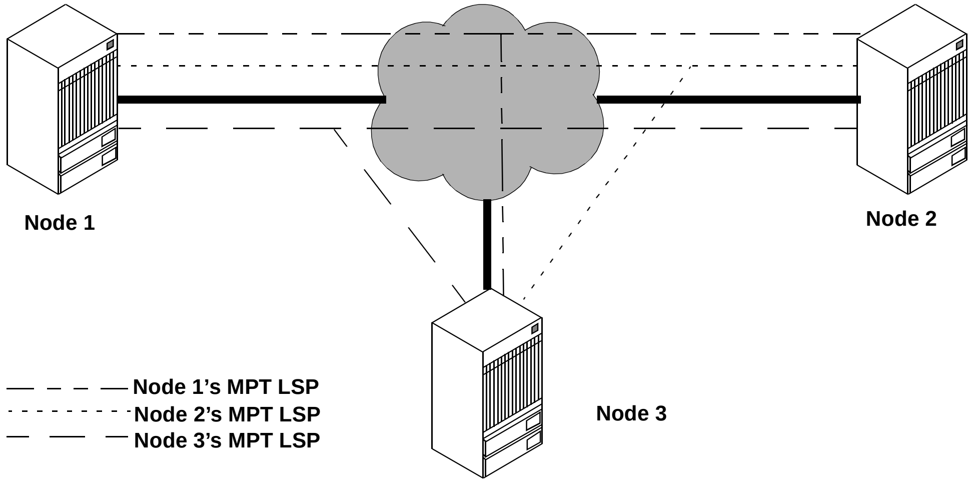

Figure 13-7. Sample MPT LSPs

In Figure 13-7, the solid lines represent trunks in the Lucent network. Node 1 establishes its MPT LSP first, followed by nodes 2 and 3. When node 1 establishes its MPT LSP, nodes 2 and 3 can forward packets back to node 1, thereby reversing the direction of the point-to-multipoint MPT LSP. The same idea applies to nodes 2 and 3 when they establish their MPT LSPs. The number of MPT LSPs is equal to the number of nodes.

See Chapter 12, “Configuring Label Switched Paths” for more information on MPT LSPs.

NHRP Packet Exchange Between CBX 500 Switches

CBX 500 switches use management PVCs (MPVCs) to exchange NHRP packets in a Lucent network. A CBX 500 can use an MPT LSP to send NHRP packets to a B-STDX 9000, but it cannot use an MPT LSP to receive NHRP packets from a B-STDX 9000. B-STDX 9000 switches send NHRP packets to CBX 500 switches hop-by-hop.

Proxy Clients and Packet Forwarding

In a Lucent switch, a proxy client will forward data packets toward the destination (possibly over MPT LSPs) while the shortcut is being established. See “NHS and NHC Support on Lucent Switches” on page 13-4 for a description of the proxy client.

Placement of NHS-capable Switches and the Proxy Client

It is important to configure NHSs and proxy clients in the right places in the network:

Configuring an NHS — Associate an NHS with each logical port that serves an NHC destination (that is, each logical port that responds to NHRP Registration Requests from the NHC destination). These logical ports are at the ingress/egress points of the Lucent network, lying along the routed paths to destinations. Do not forget to add an NHRP logical port to each IP logical port (or IP Server logical port) with which you associate an NHS. See "Adding and Deleting NHRP Logical Ports" on page 14-10 for more information on adding NHRP logical ports.

Configuring the proxy client — To send NHRP Resolution Requests on each ingress logical port where you want to filter network traffic, which is done through the use of flow profiles. Make sure that only one proxy client exists along the routed path between a traffic source and a destination. Make sure that you add an NHRP logical port to each IP logical port (or IP Server logical port) that supports a proxy client. See "Adding and Deleting NHRP Logical Ports" on page 14-10 for more information on adding NHRP logical ports.

You do not have to explicitly configure the NHS for trunk logical ports. A special NHS called the *default server* handles NHRP request and reply forwarding on all trunk logical ports (that is, logical ports that are internal to the Lucent network).

Placement of NHS-capable Switches

For a serving NHS to supply the ATM address of its client, NHS-capable routers/switches must be contiguously deployed at each hop along the routed path between the following nodes:

- The NHC requesting the address resolution
- The requested destination

In a Lucent network with established MPT LSPs, NHS-capable switches must exist at the ingress and egress of the network. Otherwise, the non-NHRP capable router/switch drops the NHRP packet.

Figure 13-8 and Figure 13-9 show NHS-capable switches serving two NHCs (NHC A and NHC B) at different ends of the Lucent network. In Figure 13-8, NHC A is the source of the NHRP Resolution Request, and, in Figure 13-9, NHC B is the source of the NHRP Resolution Request. In both figures, the ATM/IP UNI logical ports that interface with the NHCs are explicitly configured by the network manager as NHSs. This means that, on each switch, the network manager:

1. Created an NHS instance (see "Configuring Servers" on page 14-32 for more information).
2. Mapped that instance to the logical port (see "Configuring NHRP Logical Port Parameters" on page 14-59 for more information.).

The intermediate trunk logical ports have NHS capabilities automatically enabled by the default server.

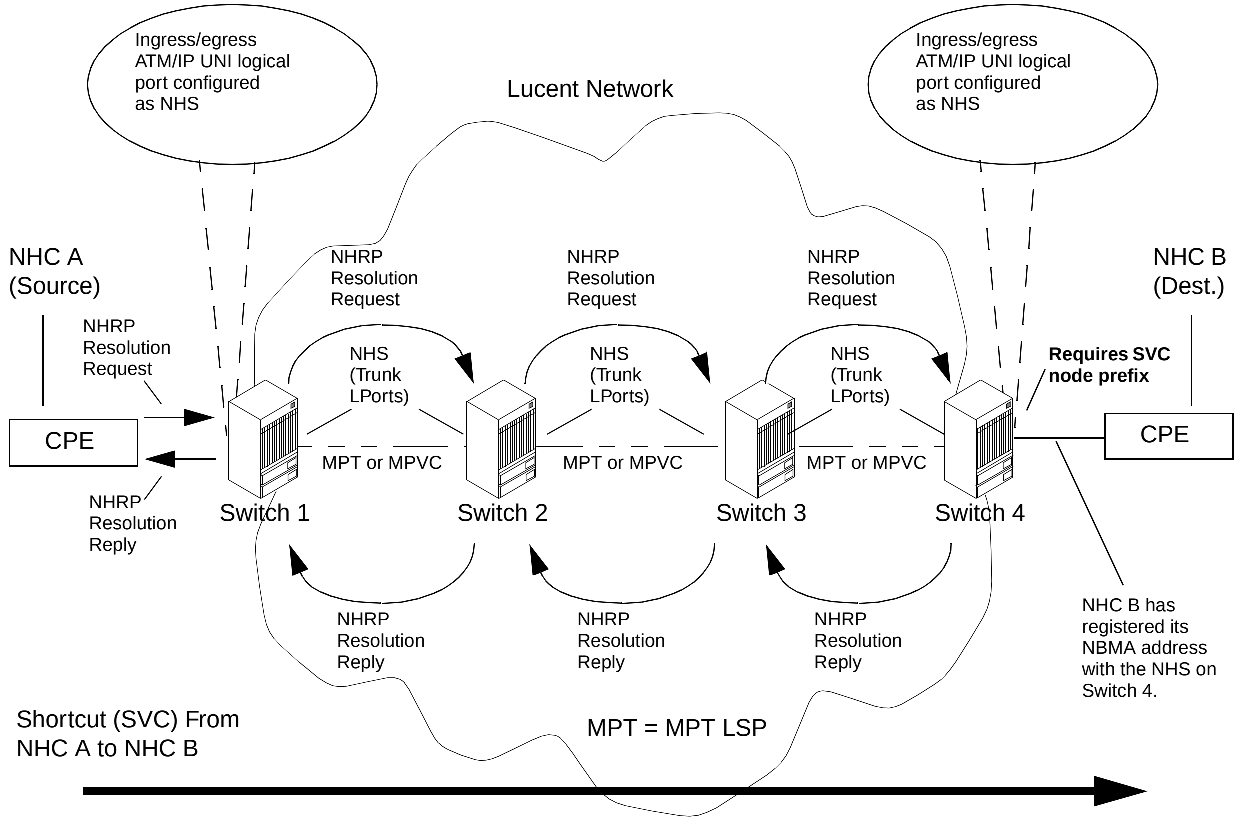

Figure 13-8. First Example of NHS-capable Switches

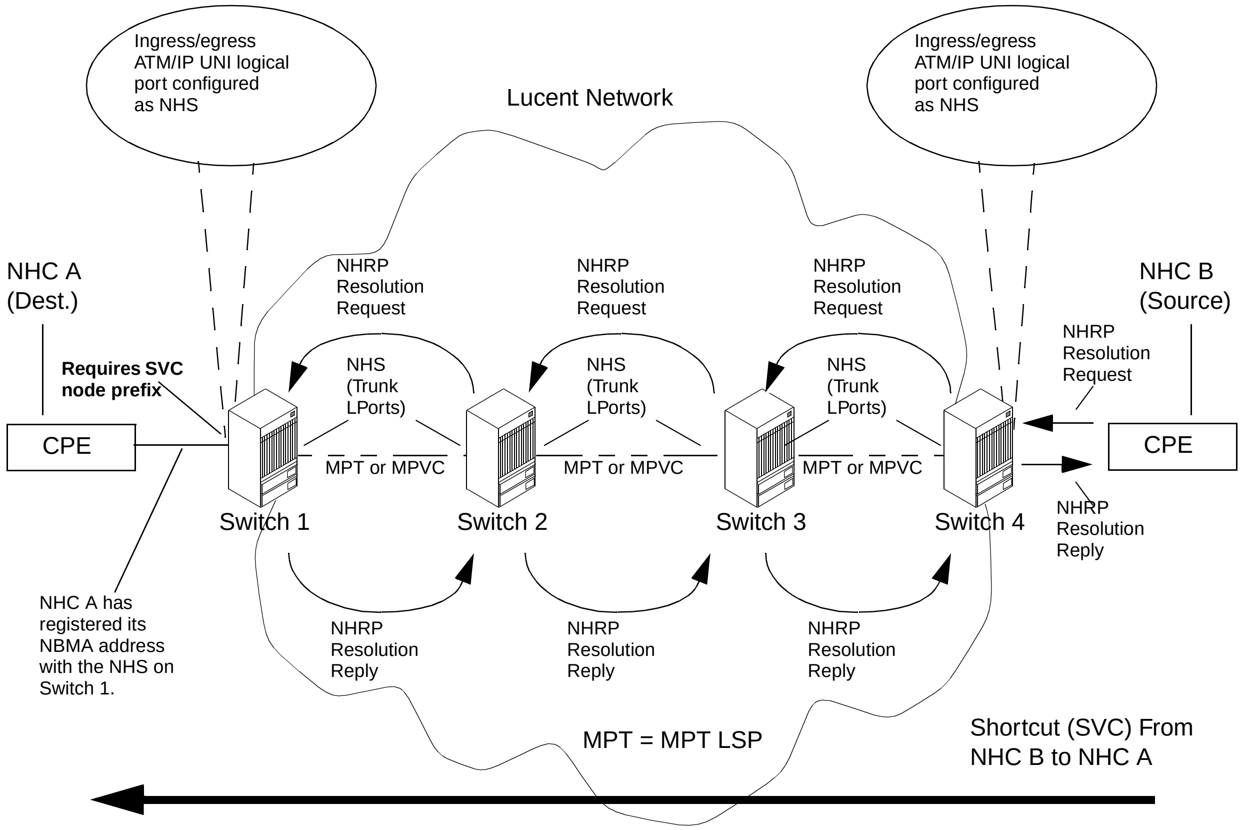

Figure 13-9. Second Example of NHS-capable Switches

Placement of Proxy Clients

The proxy client must be enabled on ingress IP logical ports where IP traffic flows enter the Lucent network. Once the proxy client detects a flow, it sends an NHRP Resolution Request to an NHS at an egress logical port. The proxy client must be the only node in the routed path that initiates NHRP Resolution Requests.

Figure 13-10 illustrates the proper placement of the proxy client. The figure assumes that IP traffic flows travel in one direction only — hence the presence of the proxy client on an ingress logical port at only one end of the network.

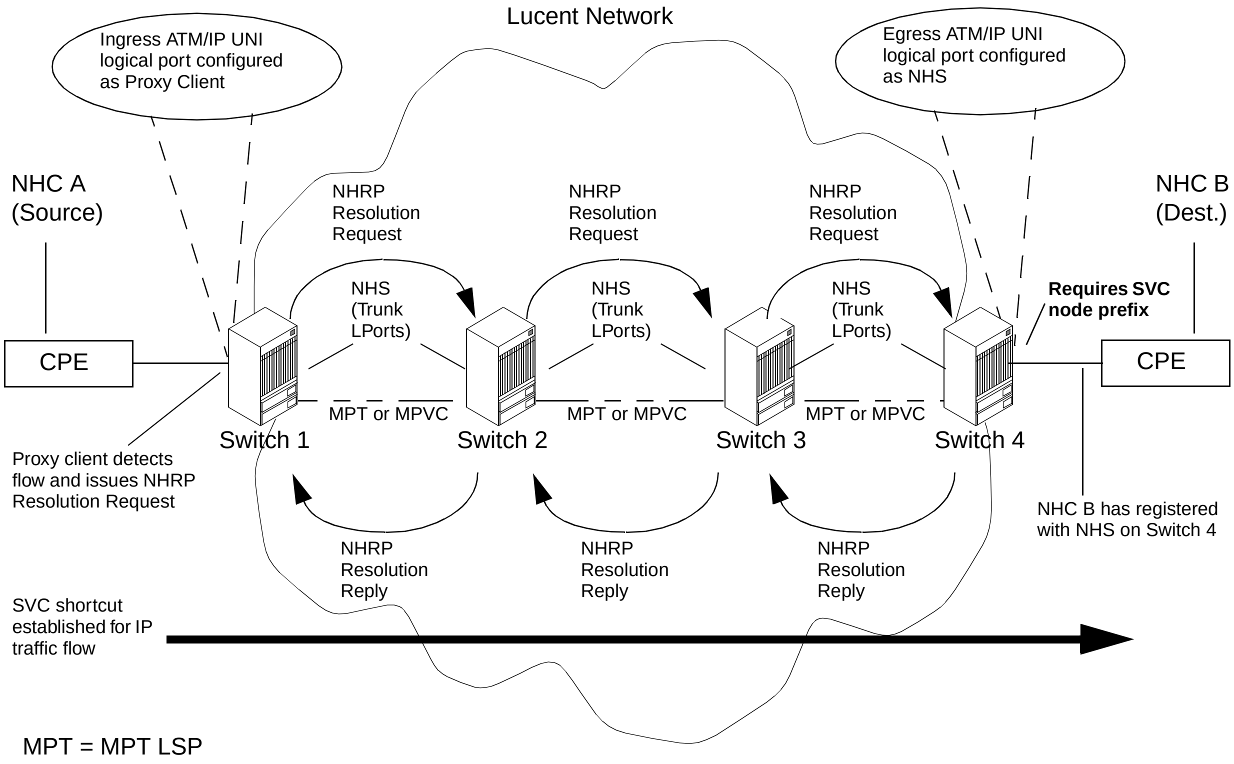

Figure 13-10. Sample Placement of Proxy Client for Uni-directional IP Traffic Flow

Figure 13-11 illustrates a network in which IP traffic flows in both directions — hence the presence of the proxy client on ingress logical ports at both ends of the network.

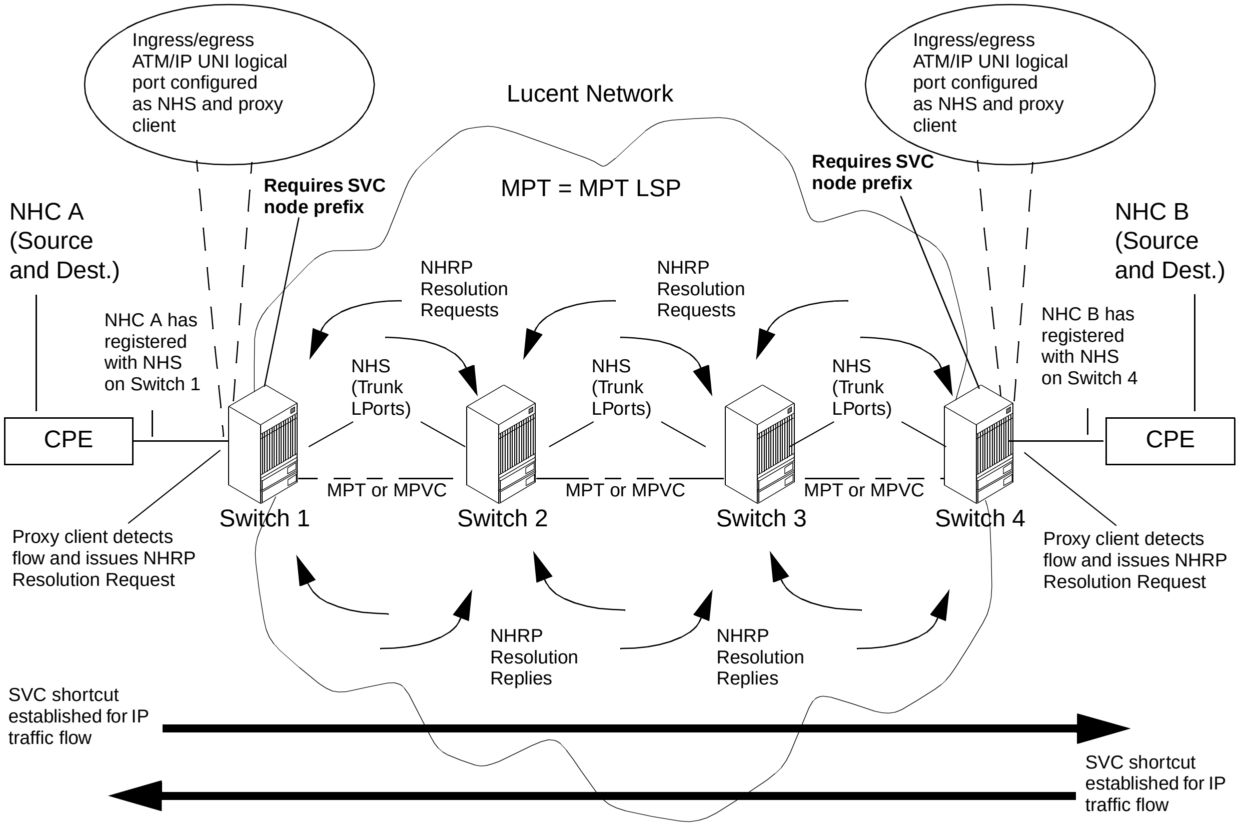
Ingress/egress ATM/IP UNI logical port configured as NHS and proxy client
Ingress/egress ATM/IP UNI logical port configured as NHS and proxy client
Lucent Network
MPT = MPT LSP
Requires SVC node prefix
Requires SVC node prefix
NHC A (Source and Dest.)
NHC B (Source and Dest.)
NHC A has registered with NHS on Switch 1
NHC B has registered with NHS on Switch 4
NHRP Resolution Requests
NHRP Resolution Requests
NHS (Trunk LPorts)
NHS (Trunk LPorts)
NHS (Trunk LPorts)
CPE
CPE
MPT or MPVC
MPT or MPVC
MPT or MPVC
Switch 1
Switch 2
Switch 3
Switch 4
Proxy client detects flow and issues NHRP Resolution Request
Proxy client detects flow and issues NHRP Resolution Request
NHRP Resolution Replies
NHRP Resolution Replies
SVC shortcut established for IP traffic flow
SVC shortcut established for IP traffic flow

Figure 13-11. Sample Placement of Proxy Clients for Bi-directional IP Traffic Flows

Multiple NHSs in a LIS

Lucent switches do not currently support the Server Cache Synchronization Protocol (SCSP), which is recommended to support synchronized caches when multiple, redundant NHSs exist in a LIS. If more than one serving NHS is desired in a LIS to avoid a single point of failure, the NHC must register with both NHSs, and transmit all subsequent registration refreshes and purge requests to both NHS destinations.

Enabling/Disabling NHRP Requests

The NHS serving a particular NHC must lie along the routed path towards the NHC destination. In practice, this means that all egress switches must double as NHSs serving the destination beyond them.

The Lucent switch allows you to enable/disable the generation of NHRP requests per ATM or Frame Relay UNI logical port. By default, these logical ports are configured so that they do not generate NHRP requests.

By default, the ability to forward NHRP requests and replies is enabled on trunk logical ports. You cannot disable this feature.

Connection Requirements for Ingress Logical Ports

NHRP Registration Reply and NHC-initiated Purge Requests must be sent over a PVC between an external NHC and a serving NHS that is associated with an ingress logical port. If a PVC does not exist, one should be created. For example, a PVC must connect the CPE in Figure 13-11 to the ingress logical ports on Switch 1 and Switch 4 (and these logical ports must have NHSs configured for them).

If a proxy client is enabled on an ingress logical port, make sure that CPE (for example, routers and remote access concentrators) can properly connect to the port.

SVC Connection and Termination

ATM link-layer connectivity must exist between the requesting node and the node identified by the NHRP-supplied ATM address. For example, in Figure 13-11, in order for one CPE to request the ATM address of the other, the two CPE must be able to communicate using ATM.

If the destination address identifies a node that is off the logical NBMA network, the border NHS must reply with its own ATM address to terminate the SVC. In a Lucent switch, the address used for SVC termination is automatically generated using:

- The 13-byte AESA prefix of the management logical port address of the B-STDX CP or the CBX 500 SP
- The 7-byte ESI derived from the IOM/IOP logical port address toward the destination

Therefore, in order to support an SVC termination from the ingress management logical port to the egress management logical port (and vice versa), you must configure the IP management logical port addresses on the CPs or SPs of both end point nodes.

No Resolution Cache Match

If an NHS receives an NHRP Resolution Request and finds no resolution cache match, the NHS applies the following rules:

- If the NHS reaches the next hop over a Frame Relay UNI logical port or a PPP logical port, the NHS terminates the request and responds with an NHRP Resolution Reply that contains its own ATM address. See "Frame Relay UNI and PPP Logical Ports" on page 13-22 for more information.
- The NHS forwards the NHRP Resolution Request if the NHS reaches the next hop over an ATM/Frame Relay OPTimum trunk, an ATM/Frame Relay Direct trunk, or an ATM interface (such as UNI, PNNI, IISP, or B-ICI).
- If the NHRP Resolution Request does not meet the criteria for termination or forwarding, the NHRP Resolution Request is negatively acknowledged by the NHS.

▶ See "Preventing Persistent Routing Loops" on page 13-34 for additional information on reasons for terminating and negatively acknowledging NHRP resolution requests.

Frame Relay UNI and PPP Logical Ports

Although you can configure an NHS on a Frame Relay UNI or PPP logical port, an SVC shortcut connection cannot be established through it. Instead, the SVC is terminated at the Frame Relay UNI or PPP logical port, and data is forwarded to the destination using IP routing (hop-by-hop).

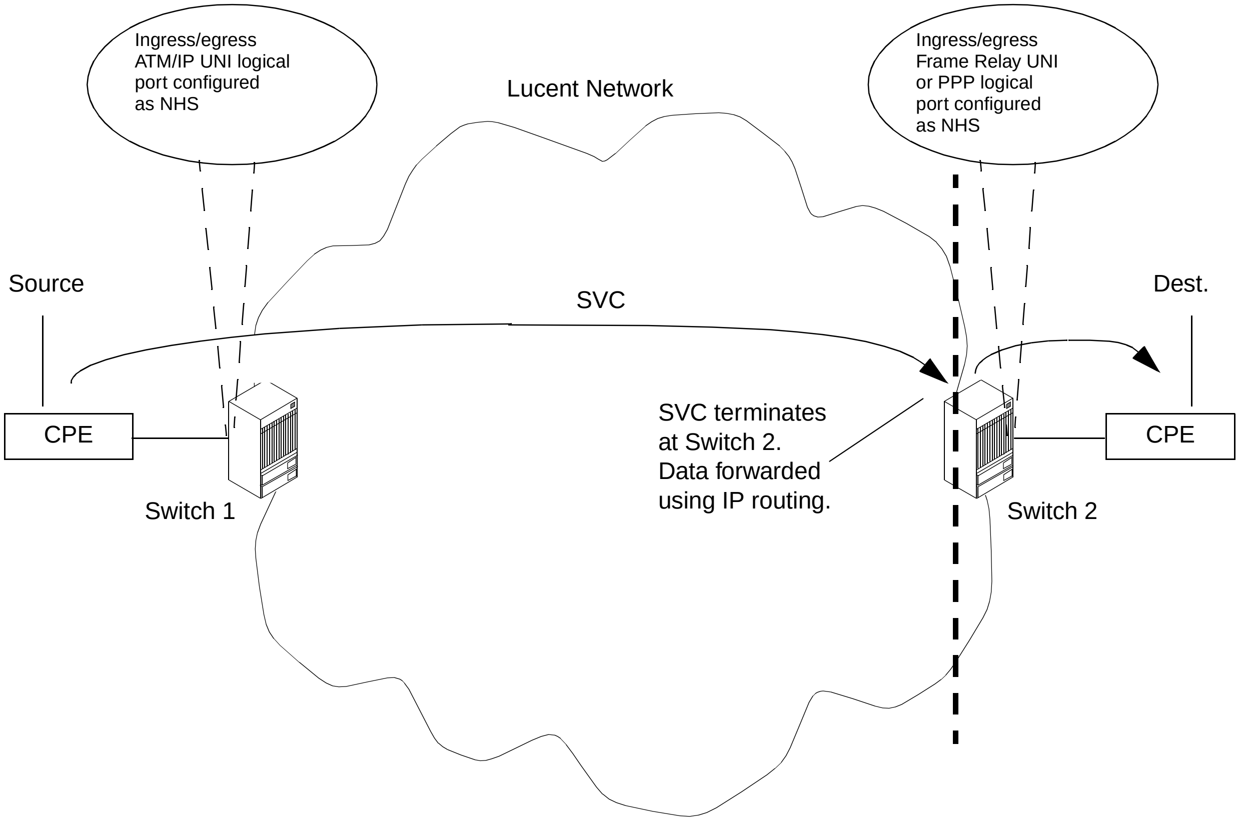

Figure 13-12. SVC Terminated at Frame Relay or PPP Logical Port

Switch 2 terminates the SVC with its own ATM address. See "SVC Connection and Termination" on page 13-21 for more information.

Extensions

In a Lucent switch, the NHRP Responder Address Extension and the Route Record Extensions are configurable options via the NMS and are disabled by default. However, if a Lucent switch receives an NHRP request with any of these extensions, the switch processes it appropriately.

> At this time, Lucent switches do not support Authentication Extension. If a Lucent switch receives this extension, the switch ignores the extension.

Domino Effect

When the proxy client on multiple switches along the path between sources and destinations initiates NHRP Resolution Requests, the NHRP domino effect can occur. The NHRP domino effect produces excessive NHRP traffic and/or the establishment of unnecessary SVCs.

As the network manager, you are responsible for solving this problem. On each switch, you are responsible for enabling proxy client functionality on a per logical port basis, guaranteeing that the proxy client is the only node in the routed path, from the source to the destination, responsible for initiating appropriate NHRP Resolution Requests. Typically, this node is the ingress switch.

In addition, NHRP Resolution Request initiation is prohibited from originating at trunk logical ports (that is, inside a Lucent network).

Queuing of NHRP Packets

NHRP packets are queued separately for each logical port on a CP or SP to prevent one logical port from starving out other logical ports.

Backward Feedback

The NHRP-established shortcuts are uni-directional. In order to support backward feedback (CCRMs and BCMs) from CBX 500 ports configured with the ATM Flow-Control Processor (ATM FCP), an NMS configuration option — the Bandwidth Reservation node parameter — is included on a per-node basis. If set, the backward traffic descriptor contains a QoS class of UBR with 34 cells per second bandwidth reservation. Otherwise, the reverse traffic descriptor is set to zero. See Table 14-4 on page 14-13 for a description of the Bandwidth Reservation parameter.

No Support for Non-authoritative Requests

All NHRP Resolution Replies from Lucent switches are authoritative. Lucent switches currently reply to non-authoritative requests only when they act as authoritative switches.

Guaranteeing QoS for IP Traffic Flows

In addition to supporting direct communications between two NBMA-connected nodes, the Lucent implementation of NHRP supports QoS guarantees for *IP traffic flows*. As described in "NHS and NHC Support on Lucent Switches" on page 13-4, an IP traffic flow is a set of packets that match a particular profile, called a *flow profile*, containing definitions of source/destination IP addresses, IP protocol, source/destination protocol port numbers, and Type of Service (TOS) requirements.

IP traffic flows traverse SVCs. The flow profile controls traffic flows in one direction: from the source to the destination. The flow profile is associated with an ATM QoS class (e.g., VBR-RT) and traffic descriptors which define service guarantees for the SVCs.

Figure 13-13 shows two IP traffic flows (and two corresponding SVCs) that match a single flow profile. Note that the 0.0.0.0 source address acts as a wildcard (that is, any source address is accepted as a match).

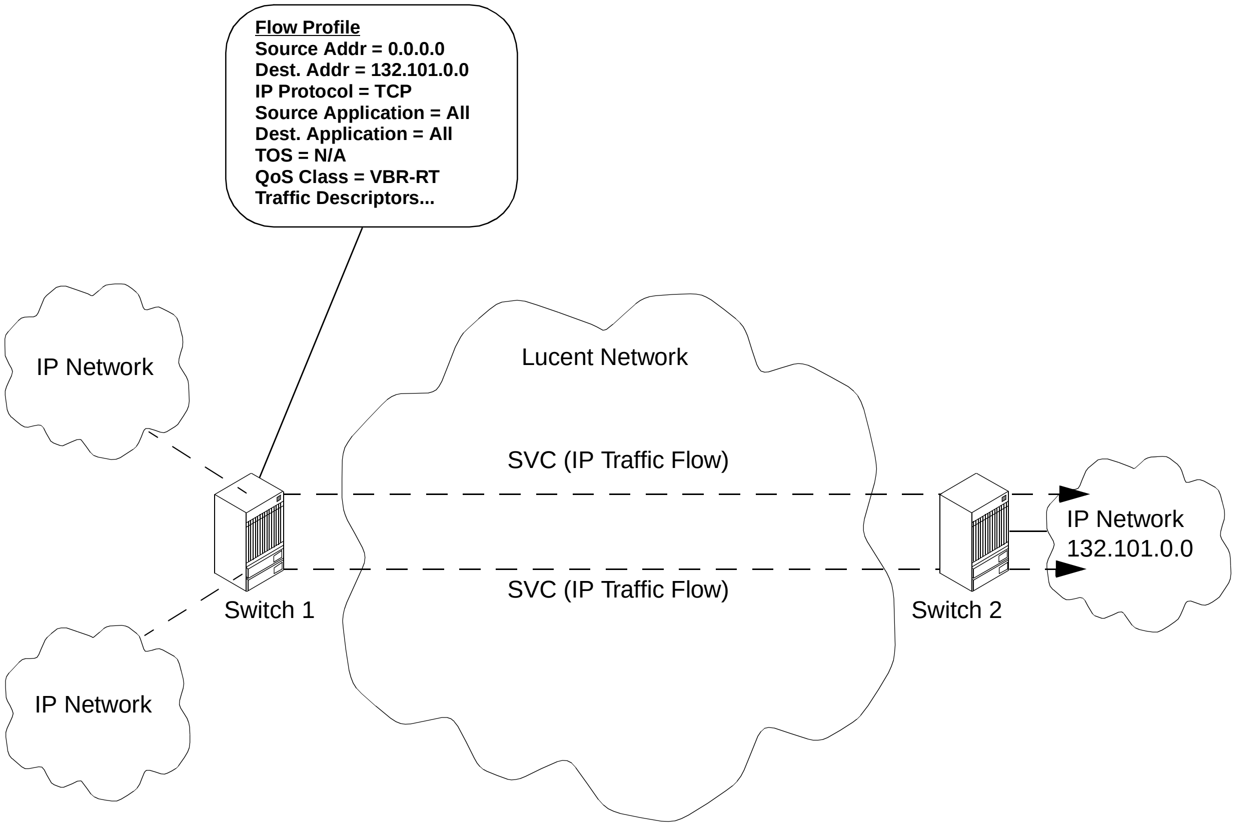

Figure 13-13. IP Traffic Flows

You can guarantee QoS on these SVCs through the use of ATM traffic descriptors (QoS parameters) mapped to flow profiles.

The best way to understand how flow profiles and traffic descriptors work is to understand the four possible roles of a Lucent switch:

- Serving NHS
- Ingress Transiting NHS
- Egress Transiting NHS
- Proxy Client

The roles apply to different scenarios, depending on the following factors:

- Whether the NHRP Resolution Request is initiated by an NHC off the Lucent network (for example, from customer premises) or by a proxy client on the Lucent network (that is, a Lucent switch)
- Whether the NHRP Resolution Reply is initiated from an NHS on or off the Lucent network

The following four sections discuss each of the Lucent switch roles. The section "Configuring Flow Profiles and Associated QoS Parameters" on page 13-32 discusses the Flow Profile and QoS Parameters in more detail.

Serving NHS

In this role, the Lucent switch acts as a serving NHS for the requested destination. When the switch receives an NHRP Resolution Request from either an NHC or a proxy client, it responds with an NHRP Resolution Reply.

The serving NHS's role is illustrated in Figure 13-14 (shown as the egress NHS).

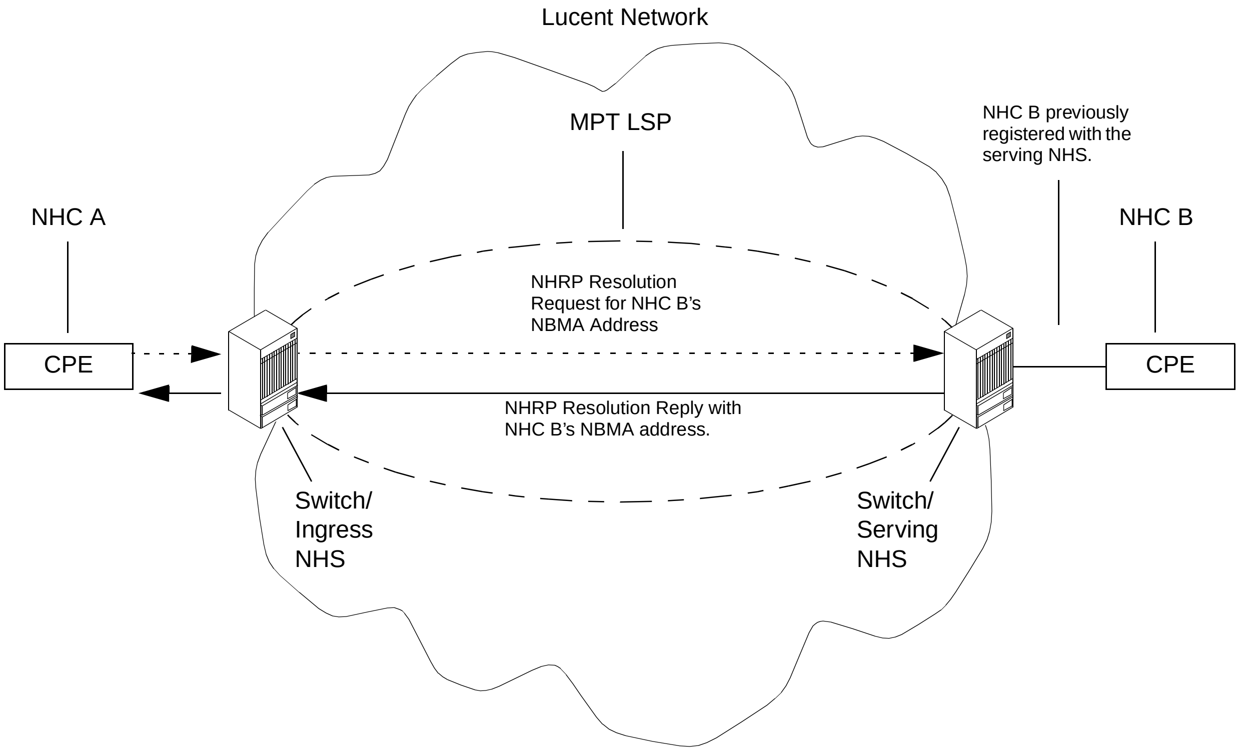
Lucent Network
MPT LSP
NHC B previously registered with the serving NHS.
NHC A
NHC B
NHRP Resolution Request for NHC B's NBMA Address
CPE
CPE
NHRP Resolution Reply with NHC B's NBMA address.
Switch/ Ingress NHS
Switch/ Serving NHS

Figure 13-14. Serving NHS

Ingress Transiting NHS

When the NHC that initiates the NHRP Resolution Request is off the Lucent network, the nearest ingress NHS that receives the NHRP Resolution Reply tries to match the destination address in the reply against the configured flow profiles (that is, the flow profiles associated with the same logical port with which the NHS is associated). These flow profiles are discussed in detail in "Configuring Flow Profiles and Associated QoS Parameters" on page 13-32.

As shown in Figure 13-15, the ingress NHS (considered the egress NHS for the NHRP Resolution Reply) receives the NHRP Resolution Reply that is addressed to the NHC.

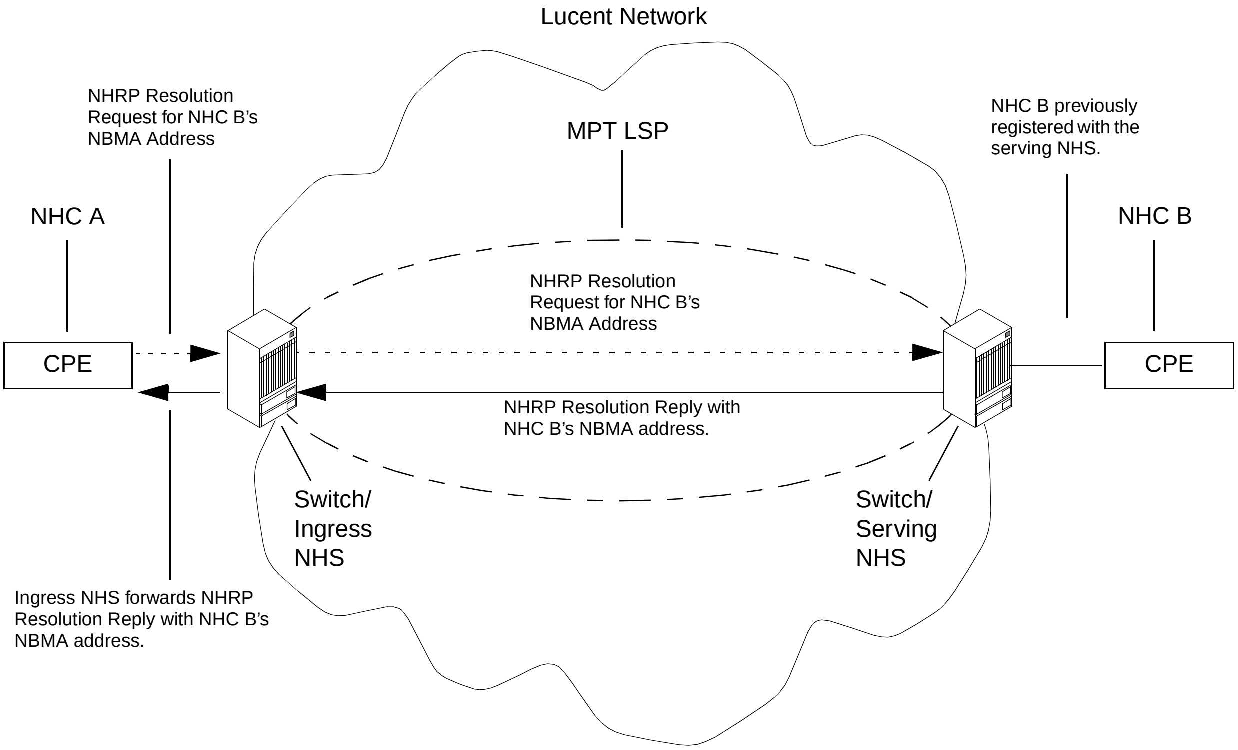
Lucent Network
NHRP Resolution Request for NHC B's NBMA Address
MPT LSP
NHC B previously registered with the serving NHS.
NHC A
NHC B
NHRP Resolution Request for NHC B's NBMA Address
CPE
CPE
NHRP Resolution Reply with NHC B's NBMA address.
Switch/ Ingress NHS
Switch/ Serving NHS
Ingress NHS forwards NHRP Resolution Reply with NHC B's NBMA address.

Figure 13-15. Ingress NHS Responsibilities

The Lucent NHS should have only one unique matching source/destination flow profile configured and saved per node from which it receives an NHRP Resolution Request outside the Lucent network. If configuration errors result in more than one profile, the NHS uses the first profile with a destination address that matches the destination address in the NHRP Resolution Reply. When choosing the first profile from a set of matching profiles, the NHS chooses the most specific profile first, then the second-most specific profile, and so on. See "Configuring Flow Profiles" on page 14-19 for more information.

Egress Transiting NHS

In the scenario shown in Figure 13-16, a Lucent switch that acts as a border NHS receives an NHRP Resolution Reply from an NHS off the Lucent network. The Lucent switch (acting as the border NHS) forwards it along.

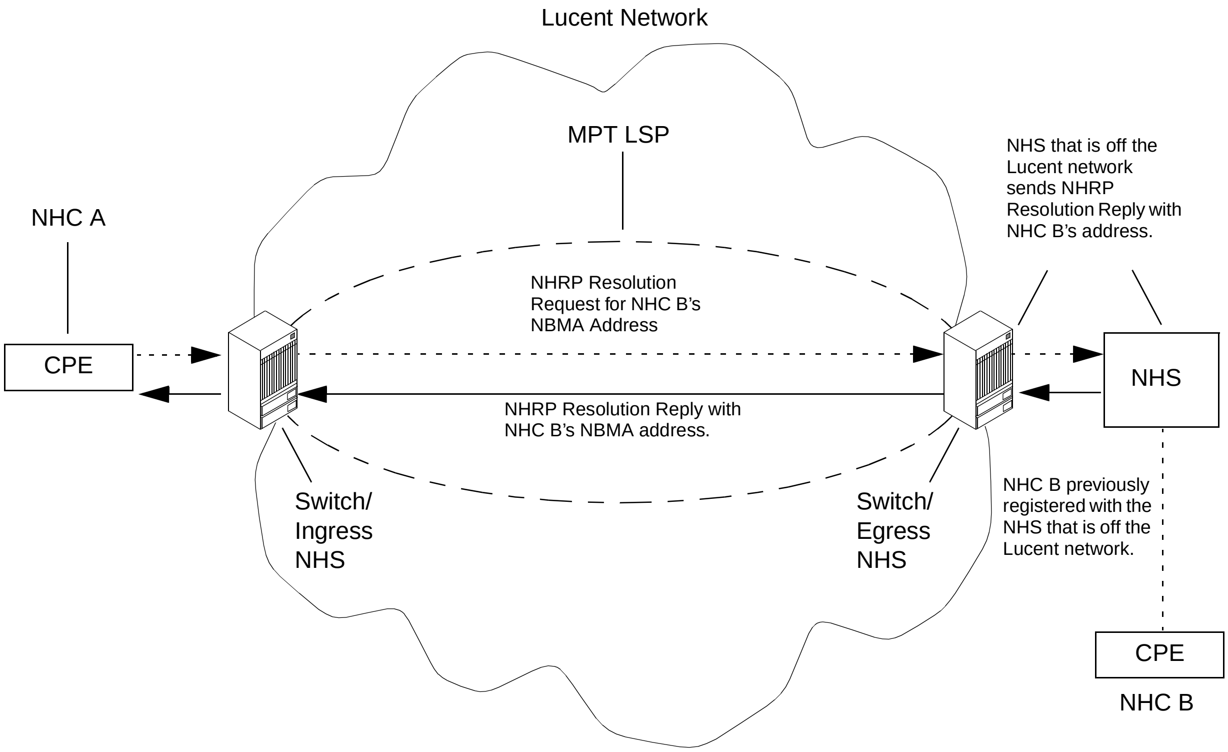

Figure 13-16. NHRP Resolution Reply Received from NHS Outside the Lucent Network

Proxy Client

Use of a proxy client in a Lucent switch provides you with the ability to provision bandwidth and QoS for IP flows. An IP flow is a set of packets that match a particular profile containing definitions of source/destination IP addresses, IP protocol, source/destination protocol port numbers, and Type of Service (TOS) requirements. Flow profiles are described in more detail in "Configuring Flow Profiles and Associated QoS Parameters" on page 13-32.

If the proxy client capability is enabled on a logical port (that is, the logical port can send NHRP Resolution Requests), the Lucent switch attempts to map IP flows into specific ATM SVCs. These SVCs have service characteristics that are based on provisioned information for the logical port from which the switch first detects the flow.

The proxy client that resides in the ingress switch sends an NHRP Resolution Request if all the following criteria are met:

- An IP packet arrives on a logical port for which NHRP Resolution Request initiation has been enabled
- The packet belongs to an IP flow with which a parameter profile is associated
- A corresponding shortcut connection does not already exist
- The number of SVCs on the logical port for that type of flow does not exceed the provisioned threshold
- The number of IP flows does not exceed the provisioned threshold
- The number of packets observed for the IP flow over the specified time interval exceeds the threshold

When the serving NHS receives the NHRP Resolution Request, it responds with the ATM address of the destination for the QoS shortcut. When the proxy client receives the NHRP Resolution Reply, it attempts to establish an SVC, and it uses the traffic descriptors configured for the flow profile (see Figure 13-17).

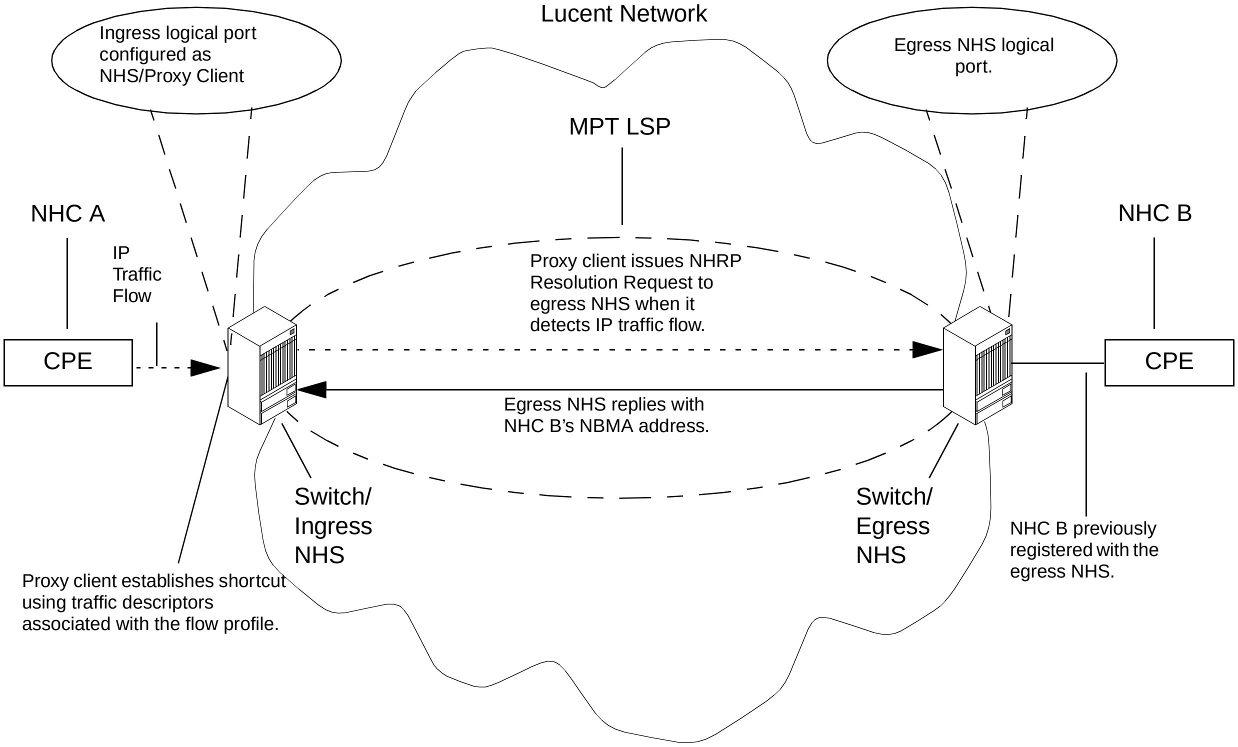
Lucent Network
Ingress logical port configured as NHS/Proxy Client
Egress NHS logical port.
MPT LSP
NHC A
NHC B
IP Traffic Flow
Proxy client issues NHRP Resolution Request to egress NHS when it detects IP traffic flow.
CPE
CPE
Egress NHS replies with NHC B's NBMA address.
Switch/ Ingress NHS
Switch/ Egress NHS
NHC B previously registered with the egress NHS.
Proxy client establishes shortcut using traffic descriptors associated with the flow profile.

Figure 13-17. Proxy Client Role

The source address of the SVC's termination is the address associated with the logical port on the IOM/IOP on which the proxy client detects the flow.

The transaction between the proxy client and the serving NHS may fail for one of the reasons described in Table 13-1.

Table 13-1. Fail Reasons

Fail Reason	Result
The proxy client does not receive an NHRP Resolution Reply within the specified timeout value.	The proxy client resends an NHRP Resolution Request according to the proxy client request retry and request backoff rules configured through the NMS. See "Configuring the Proxy Client" on page 14-47 for more information.
The proxy client receives an NHRP Resolution Reply negative acknowledgment due to the absence of an IP/ATM mapping.	The proxy client resends an NHRP Resolution Request according to the proxy client request retry and request backoff rules configured through the NMS. See "Configuring the Proxy Client" on page 14-47 for more information.
The resulting ATM call setup attempt fails.	The proxy client retries the SVC establishment according to the per-node retry and backoff rules configured through the NMS. See "Configuring Node Parameters" on page 14-12 for more information.
***Note**: If no retries are desired in any of the above cases, you may set the number of retries to 0. If the failure persists even after the proxy client performs the provisioned retries, the proxy client reinitiates flow detection and resends an NHRP Resolution Request once the trigger is satisfied again.*	

▶ No traps are generated for these errors; however, the failure described for the second fail reason is logged, if provisioned.

Once the SVC is established, the switch monitors flow abatement according to provisioned criteria (see "Configuring Flow Profiles" on page 14-19 for more information on this criteria). If the switch detects flow abatement, the switch removes the SVC. However, the proxy client retains the Destination Address and Traffic Descriptors in its resolution cache for the duration of the hold time. This retention expedites future shortcut establishment if the flow is rediscovered within the hold time. The switch should clear the resolution cache of an entry if the possibility of a routing loop arises for that entry.

Configuring Flow Profiles and Associated QoS Parameters

The IP flow profile is defined in terms of:

- IP source and destination addresses (with CIDR masking)
- IP protocol type
- Source and destination TCP/UDP-like protocol ports
- Type of Service (TOS) (e.g., voice)

The IP flow profile also includes the trigger criteria for flow onset and abatement. You may use wildcard values for any of these parameters (except trigger criteria) so that a packet only has to match the non-wildcard parameters to match the flow profile.

For example, a specific flow profile defines traffic flows for all TCP traffic that originates from a particular source, but you do not care about the specific destination. You could assign a wildcard to the destination address.

Although you may use wildcards for defining the destination addresses in flow profiles, the associated destination address for the flows themselves must be unambiguous for establishing an actual SVC. Therefore, several unambiguous flows (and subsequent SVCs) may correspond to a single flow profile. The switch tracks unambiguous flows separately.

You can configure QoS traffic descriptor parameters for the flow profiles. The traffic descriptors are used to establish the SVC shortcuts that carry the associated traffic flows. The traffic descriptors contain the ATM UNI 4.0 parameters with the exception of the following optional extended QOS parameters:

- Cell Loss Ratio (CLR)
- Cell Transfer Delay (CTD)
- Cell Delay Variation (CDV)

The ATM UNI 4.0 parameters are specified in the direction of data flow, from source to destination (that is, from the ingress logical port that has NHS and/or proxy client capabilities enabled to the egress NHS logical port).

When the full amount of bandwidth requested is not available on the route, you can also provision the ability to negotiate bandwidth parameters through UNI 4.0 signaling (the alternate/minimum traffic descriptors).

Once you define the flow profiles and the QoS traffic descriptor (TD) parameters, you associate them in any combination to ingress logical ports that have proxy client capabilities enabled. You may map one flow profile or TD to several logical ports. By the same token, you may map several corresponding flow profiles and TD pairs to a single logical port.

In addition, you may delete and modify these IP flow profiles or TDs and their logical port mappings. While a flow profile is in use, you may modify or delete its logical port mapping, but you may not modify or delete the flow profile itself. If you modify or delete a flow profile mapping after a shortcut has been established, the shortcut is removed immediately to avoid problems involved with sustaining the shortcut indefinitely.

Since switch system performance and memory utilization depend on the number of unambiguous flows that the switch detects, you can provision, per logical port, the upper limit on the number of simultaneously tracked flows that match a particular flow profile. You can also provision the following parameters:

- The maximum number of shortcuts for the flow profile
- The maximum number of traffic flows tracked per flow profile
- The maximum number of simultaneously tracked flows for all flow profiles associated with a logical port
- The maximum number of shortcuts established for all flow profiles associated with a logical port

During the time when the upper limit of flows for a flow profile is reached on the associated logical port, the switch does not detect a new unambiguous flow that matches an ambiguous flow profile. However, the list of unambiguous flows that the switch detects is dynamic. When an unambiguous flow does not meet its trigger criteria within a specified period of time (such as 2.5 times the detection period), the switch removes it from the detection list, thereby making room for others to be detected.

A particular packet may map to more than one of the defined flow profiles; therefore you are responsible for prioritizing the flow profiles. The traffic descriptor used for call setup is the one corresponding to the first (higher priority) flow profile that the packet matches.

See your switch Software Release Notice (SRN) for information on the maximum number of flow profile/logical port mappings, and other NHRP limits.

Interaction Between NHRP Shortcuts and Policy PVCs

You may have situations in which either an NHRP shortcut or a policy PVC can be used to reach a single destination. In this case, the NHRP shortcut takes precedence over the policy PVC. If no NHRP shortcut exists, NHRP passes traffic to the policy PVC. See “About Policy PVCs” on page 5-3 for more information on policy PVCs.

Preventing Persistent Routing Loops

You should be aware of a condition called a *persistent routing loop* that results in traffic traversing the network in a circular path until the Time-to-Live timer expires. This condition tends to appear in networks that use dynamic routing protocols such as NHRP.

The following scenario illustrates a typical cause of a persistent routing loop:

1. Multiple paths connect a source and destination. One path uses an NHRP shortcut; the other path uses an intermediate hop through a router. The path that uses the NHRP shortcut is the least-cost path and, therefore, is used most often.
2. Both paths intersect at a common intermediate node between the source and destination, merging into one path.
3. A link beyond the intersection breaks, creating the persistent routing loop.

Lucent switches can automatically prevent the creation of persistent routing loops within the Lucent network, as follows:

When a Lucent Switch Acts as a Proxy Client — In this situation, the switch monitors the MPT LSP status of the shortcuts it establishes to carry IP traffic flows. If the MPT LSP status of a shortcut changes, the switch tears down the shortcut and clears the associated resolution cache entry. If the traffic flow persists, the switch establishes a new shortcut immediately.

However, this solution automatically prevents the creation of persistent routing loops within the Lucent network only. It does not prevent the creation of these routing loops outside the Lucent network. To prevent routing loop creation outside the Lucent network, the NMS allows you to manually terminate shortcuts by setting a logical port configuration option. If you enable this option, and an egress NHS cannot resolve an NHRP Resolution Request from a proxy client, it terminates the request and sends an NHRP Resolution Reply with the NHS's own NBMA address, thereby terminating the shortcut.

When an NHC External to the Lucent Network Initiates the Shortcut — In this situation, when the destination is off the NBMA, the egress Lucent switch that acts as an NHS terminates the shortcut if the destination is directly connected to it (that is, there are no intermediate routers between the egress switch and the destination).

If the destination is not directly connected to the egress switch, no shortcut can be established between the external NHC and the destination in the first place. The switch will reply to the external NHC's NHRP Resolution Request with an error in the NHRP Resolution Reply. The external NHC will have to use a Lucent switch that acts as a proxy client. If you want to guarantee a redundant path to a destination, you may use the NMS to override this capability on a logical port basis.

14

Configuring Next Hop Resolution Protocol

This chapter provides instructions on how to configure Next Hop Resolution Protocol (NHRP).

Before You Begin

The NHRP configuration tasks, the sections that describe them, and the order in which you should perform them are listed in Figure 14-1.

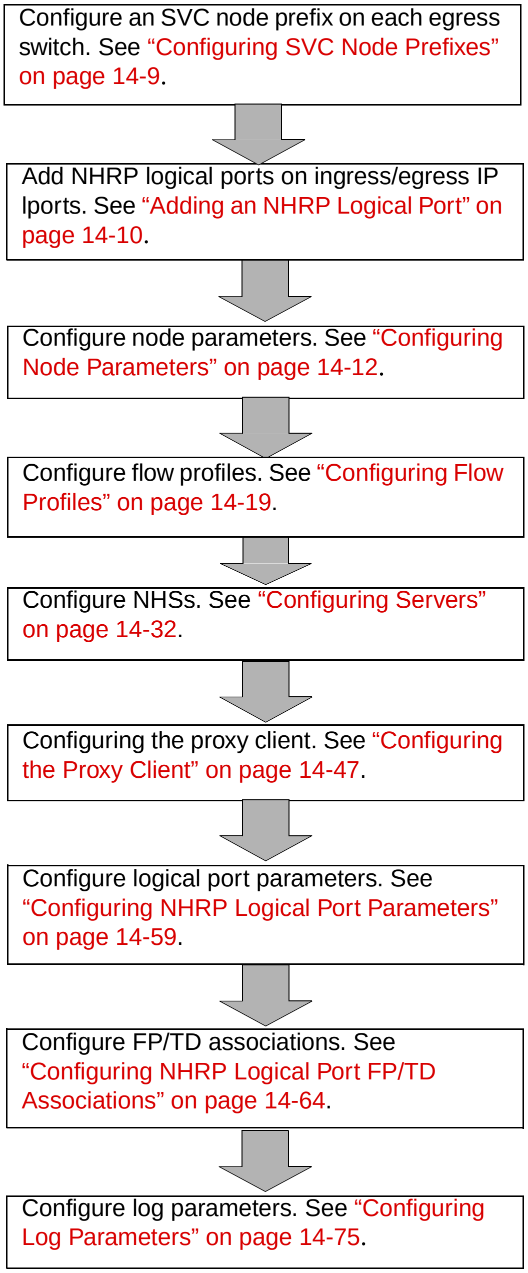

Configure an SVC node prefix on each egress switch. See "Configuring SVC Node Prefixes" on page 14-9.

Add NHRP logical ports on ingress/egress IP lports. See "Adding an NHRP Logical Port" on page 14-10.

Configure node parameters. See "Configuring Node Parameters" on page 14-12.

Configure flow profiles. See "Configuring Flow Profiles" on page 14-19.

Configure NHSs. See "Configuring Servers" on page 14-32.

Configuring the proxy client. See "Configuring the Proxy Client" on page 14-47.

Configure logical port parameters. See "Configuring NHRP Logical Port Parameters" on page 14-59.

Configure FP/TD associations. See "Configuring NHRP Logical Port FP/TD Associations" on page 14-64.

Configure log parameters. See "Configuring Log Parameters" on page 14-75.

Figure 14-1. NHRP Configuration Tasks

Before you perform these tasks, you should:

- Plan your NHRP network configuration
- Understand NBMA addressing

Planning Your NHRP Network Configuration

You should plan for configuring NHRP in your network as follows:

- Identify NHCs at customer premises (for example, CPE that act as NHCs). Make sure that each NHC has a proper PVC connection to the ingress/egress logical port on a Lucent switch that has an NHS configured.
- Configure an SVC node prefix on each egress switch. When you configure the SVC node prefix, enable Internal Management. See "Configuring SVC Node Prefixes" on page 14-9 for details.
- Add NHRP logical ports on all of the ingress/egress IP logical ports or IP Server logical ports that will handle NHRP traffic. See "Adding an NHRP Logical Port" on page 14-10 for details.
- Identify the switches that will act as proxy clients. The main purpose of the proxy client is to provide IP sources with provisioned bandwidth and QoS guarantees for IP flows (that is, the proxy client attempts to map IP flows into ATM SVCs that have service characteristics based on provisioned information for the logical port where the flow is first detected). You enable proxy client functionality at ingress logical ports. Make sure that CPE (for example, routers and remote access concentrators) can properly connect to the ingress ports.
- Identify the ingress/egress switches that will act as NHSs, and the ingress/egress logical ports on those switches that will need to have NHS capabilities enabled. Keep in mind that you can have multiple NHS instances on a switch, one per IP logical port.
- Identify all the necessary IP addresses and NBMA addresses you need to configure for NHSs, proxy clients, cache entries, and so on. It is suggested that you read "About NBMA Addressing" on page 14-4 before you configure servers, the proxy client, and associated cache entries.
- Verify that all network nodes (switches, routers, etc.) communicate successfully using IP.
- Verify that MPT LSPs and/or MPVCs connect ingress/egress switches within the Lucent network that act as NHSs and proxy clients. B-STDX switches that act as NHSs and proxy clients use MPT LSPs to exchange NHRP packets with each other, if they are available. CBX 500s use MPVCs to exchange NHRP packets with each other. CBX 500s use MPT LSPs to send NHRP packets to B-STDX switches; B-STDX switches send NHRP packets to CBX 500 switches hop-by-hop. If no MPT LSPs or MPVCs are available, switches forward NHRP packets hop-by-hop. Note that both B-STDX and CBX 500 switches use MPT LSPs to transmit data packets.
- Identify types of application-level traffic (for example, WWW, FTP) that traverse your network. Try to assess the relative importance of each type of traffic. This information will help you set up flow profiles and their associated traffic descriptors, and then map them to IP logical ports.

- Identify nodes and their associated addressing information that require you to configure static NHRP cache entries for either the proxy client or the NHS. Typically, nodes that require manually created cache entries do not support NHRP and, therefore, cannot participate in the NHRP address registration process. As an alternative to creating static NHRP cache entries, you can terminate SVCs to non-NHRP-compatible destinations at the management logical port of the egress switch. Traffic is then forwarded hop-by-hop using IP routing.
- Identify potential persistent routing loop problems by analyzing your network topology. See "Preventing Persistent Routing Loops" on page 13-34 for more information on persistent routing loops.

About NBMA Addressing

Before you configure servers or cache entries, you should understand Non-Broadcast Multiple Access (NBMA) addressing. Remember that network nodes are identified by both their IP addresses and their NBMA addresses, and the purpose of NHRP is to resolve IP addresses to NBMA addresses. When an IP address is resolved to an NBMA address, NHRP uses the NBMA address to establish an SVC shortcut toward the next hop.

When you configure a server or a proxy client, you must specify its IP address and its NBMA address. When you create a server cache entry or a proxy client cache entry, you must specify both the IP address and the NBMA address of the node associated with the entry.

In ATM environments, Lucent supports two types of NBMA address formats for establishing SVCs:

ATM End System Address (AESA) formats — AESA formats give service providers using a private ATM network the flexibility to develop an addressing scheme that best suits their network needs; for example, you may find that most CPEs in your network only support a specific AESA address format.

> *AESA Anycast Formats* – AESA Anycast formats give service providers "group address" functionality for each of the AESA address formats. Using the Anycast format, a call is placed to the group address and the network selects one of the members to which the call will be routed. This group address could, for example, represent a group of Internet servers that contain the same information and perform identical functions. It does not matter which of these servers handles the call.

Native E.164 address format — E.164 addresses are phone numbers. This address format is simple and familiar; native E.164 addresses are a convenient choice for service providers using a public ATM network (e.g., RBOCs) that already "own" E.164 address space.

The following sections describe these address formats.

ATM End System Address Formats

NHRP resolves IP addresses to four ATM End System Address (AESA) formats:

Data Country Code (DCC) — For DCC AESA addresses, the initial domain identifier (IDI) is a two-byte data country code field that identifies the country in which this address is registered. These country codes are standardized and defined in ISO reference 3166. *DCC Anycast AESA* provides a group address function for this address type.

International Country Designator (ICD) — For ICD AESA addresses, the IDI field contains the international country designator that uniquely identifies an international organization. The British Standards Organization administers these values. *ICD Anycast AESA* provides a group address function for this address type.

E.164 — For E.164 AESA addresses, the IDI field contains an eight-byte E.164 address. This E.164 address uses the international format and consists of up to fifteen decimal digits. *E.164 Anycast AESA* provides a group address function for this address type.

Custom — Custom AESA addresses enable you to use a customized octet structure and a customized authority and format identifier (AFI).

All AESA address formats consist of 20 octets. Each of these address formats contain the following components:

Initial Domain Part (IDP) — Defines the type of address and the regulatory authority responsible for allocating and assigning the Domain Specific Part. There are two subfields: the AFI and IDI fields.

> *Authority and Format Identifier (AFI)* – The AFI part of the AESA address identifies the authority that allocates the DCC, ICD, or E.164 part of the AESA address, as well as the syntax of the rest of the address. Table 14-1 lists valid AFIs.

Table 14-1. AFI Default Values

Address Type	AFI
DCC	0x39
DCC Anycast	0xBD
ICD	0x47
ICD Anycast	0xC5
E.164	0x45
E.164 Anycast	0xC3

Table 14-1. AFI Default Values (Continued)

Address Type	AFI
Custom	A user-specific code for custom prefixes/addresses. (You must know the appropriate code to enter when defining custom prefixes/addresses.)

Initial Domain Identifier (IDI) – A hex code that identifies the sub-authority that has allocated the address. The format depends on the address types listed in Table 14-2.

Table 14-2. IDI Default Values

Address Type	IDI Description
DCC (including Anycast)	Consists of 2 octets (4 hex digits) that identify the country in which this address is registered. The DCC is generally considered a three-digit quantity with a trailing hex “f” semi-octet. For example, the ANSI IDI of 840 is encoded as 0x840f.
ICD (Anycast)	Consists of 2 octets (4 hex digits) that identify an international organization to which this address is registered. The ICD is generally considered a four-digit quantity. For example, the US GOSIP IDI of “5” is encoded as 0x0005.
E.164 (Anycast)	Consists of 8 octets in BCD format. (1-15 hex digits, plus a trailing Fh; if less than 15 digits are entered, type leading zeros to fill the 8 octets.) Represents an international E.164 address. For example, the E.164 address of 978-555-1212 is encoded as 0x000009785551212f.

Domain Specific Part — Consists of the HO-DSP, EDI, and SEL fields.

High-Order Domain-Specific Part (HO-DSP) – The authority specified in the AFI/IDI octets determines the format of this field. It identifies a segment of address space that is assigned to a particular user or subnetwork. It should be constructed to facilitate routing through interconnected ATM subnetworks. The general format for each address type is listed in Table 14-3.

Table 14-3. HO-DSP Default Values

Address Type	HO-DSP Description
DCC, ICD (including Anycast)	Consists of 10 octets (20 hex digits)
E.164 (Anycast)	Consists of 4 octets (8 hex digits)
Custom	Consists of 12 octets (24 hex digits)

End System Identifier (ESI) – A 6-octet (12 hex digit) field that uniquely identifies the end system within the specified subnetwork. This is typically an IEEE MAC address.

Selector (SEL) – A 1-octet (2 hex digit) field that is not used for ATM routing, but may be used by the end system.

Figure 14-2 shows how the octets are assigned for each AESA address format. Each octet is equivalent to two hex digits.

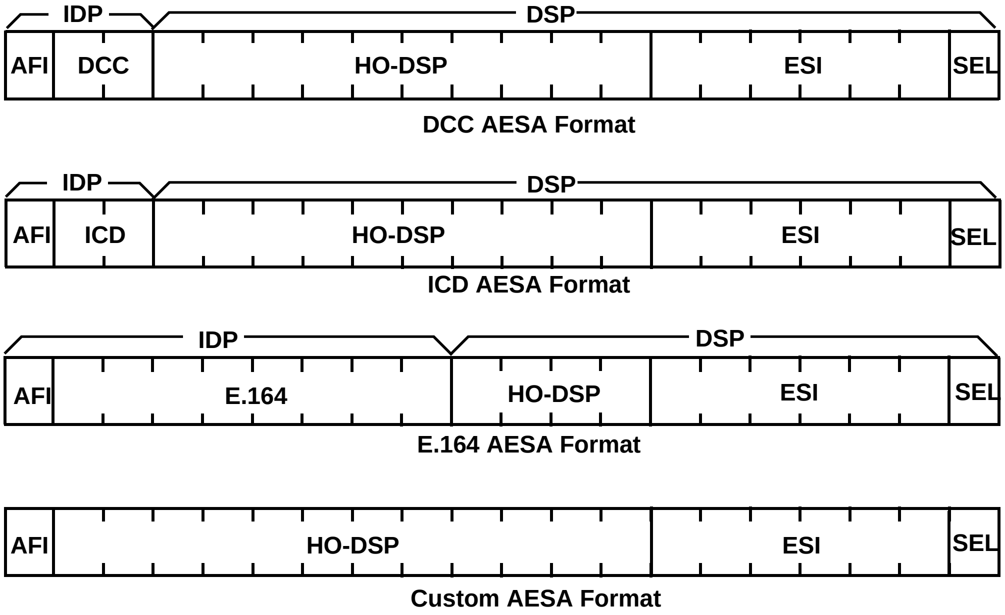

Figure 14-2. AESA Address Formats

Native E.164 Address Format

Native E.164 addresses are the standard Integrated Services Digital Network (ISDN) numbers, including telephone numbers. Native E.164 addresses consist of 1-15 ASCII digits. For example, standard 10-digit United States telephone numbers, such as 508-555-1234, are native E.164 addresses.

Unlike AESA address formats, native E.164 addresses are not broken down into an AFI, HO-DSP, ESI, and SEL portion. When a native E.164 address is translated to E.164 AESA format, the native E.164 address is stored in octets 2-9 of the 20-octet AESA address, while the HO-DSP, ESI, and SEL portions are filled with zeros. Conversely, when an E.164 AESA address is translated to native E.164 address format, the AFI, HO-DSP, ESI, and SEL portions, as well as any leading zeros in the 8-octet AESA E.164 address, are stripped off to produce the native E.164 address.

It is possible to use native E.164 addresses with Network Service Access Point (NSAP) addresses. In this case, the E.164 address is the NBMA address and the NSAP address is the NBMA subaddress. For example, if the switch acts as a gateway between a public network that uses native E.164 addressing and a private network that uses NSAP addresses, the switch is known to the public network by its E.164 address (the NBMA address) and is known to the private network by its NSAP address (NBMA subaddress).

Designing an Address Format Plan

The SVC address formats you select must support the equipment and services your network needs to provide. Keep in mind that some CPEs may not support certain address formats. To avoid address conflicts, apply for globally recognized address space in the ATM formats you need to use.

You use address formats to develop a network numbering plan. Using an AESA address, you can design the IDP portion of an address to target a specific network; then use the HO-DSP portion of the address to identify subnetworks within that network, and use the ESI portion to identify a specific end system.

Regardless of the address format you choose, the network numbering plan should satisfy the following goals:

- Intelligently assign network addresses
- Simplify network topology using a hierarchal organization
- Minimize the size of network routing tables
- Uniquely identify each endpoint
- Provide a high level of network scalability

For more information ATM addressing, see the *NavisCore ATM Configuration Guide*.

About NHRP and IP VPNs

NHRP resources are public resources (that is, they are assigned to the public IP VPN). They cannot be reserved for exclusive use by private IP VPNs. The sources and destinations of customer data that traverse NHRP SVC shortcuts, as well as control traffic, must not be part of a private IP VPN. For example, if an NHRP traffic flow traverses an ATM PVC from the CPE to an ingress IP logical port, the VPI/VCI must not be assigned to a private IP VPN. Rather, it must be assigned to the IP VPN labeled as “public.” See Chapter 16, “Configuring IP Virtual Private Networks” for more information on IP VPNs.

Configuring SVC Node Prefixes

You must configure an SVC node prefix on each egress switch that handles NHRP traffic. When you configure SVC node prefixes:

- Enable Internal Management
- Use an address format that is consistent with the address format you use throughout your network. This is the format you should also use for NHS addresses, NHC addresses, and proxy client addresses. For example, if you use the DCC AESA address format throughout the network, use this format for the SVC node prefix.

See the *NavisCore ATM Configuration Guide* for more information on configuring SVC node prefixes.

Adding and Deleting NHRP Logical Ports

This section describes how to add and delete NHRP logical ports.

Adding an NHRP Logical Port

To add an NHRP logical port:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set all IP LPorts dialog box appears (see Figure 3-2 on page 3-6).
2. Select the switch where the ingress/egress IP logical port or IP Server logical port resides from the Switch Name list at the top of the dialog box. A list of IP logical ports and IP Server logical ports configured on the switch appears in the LPort Name list.
3. Select the ingress/egress IP logical port or IP Server logical port.
4. Choose IP Parameters. The Set IP Parameters dialog box appears (see Figure 14-3).

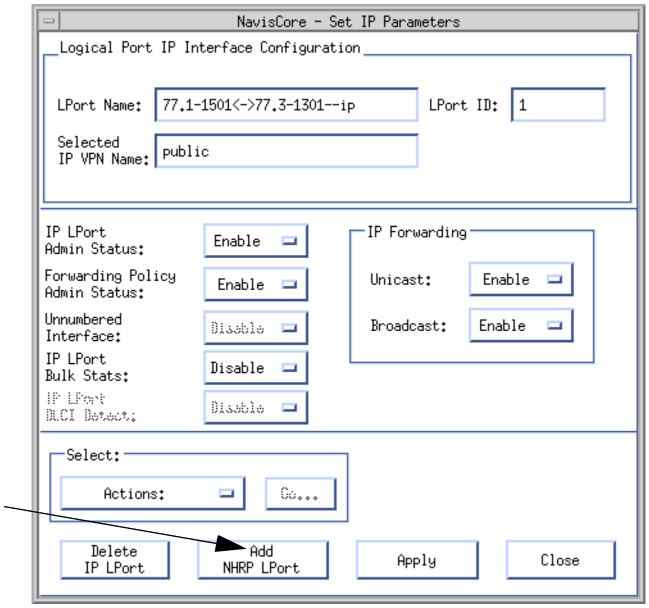

NavisCore - Set IP Parameters

Logical Port IP Interface Configuration

LPort Name: 77.1-1501<->77.3-1301--ip
LPort ID: 1
Selected IP VPN Name: public

IP LPort Admin Status: Enable
Forwarding Policy Admin Status: Enable
Unnumbered Interface: Disable
IP LPort Bulk Stats: Disable
IP LPort DLCI Detect: Disable

IP Forwarding
Unicast: Enable
Broadcast: Enable

Choose Add NHRP LPort

Figure 14-3. Set IP Parameters Dialog Box (With Add NHRP LPort Button)

5. Choose Add NHRP LPort. This action adds the NHRP logical port. The Add NHRP LPort button becomes the Delete NHRP LPort button, allowing you to delete the NHRP logical port whenever it becomes necessary. See "Deleting an NHRP Logical Port" for more information on deleting an NHRP logical port.
6. Choose Close to exit.

Deleting an NHRP Logical Port

Before you delete an NHRP logical port, delete all flow profile associations. Otherwise, NavisCore will not allow you to delete the port. See "Configuring NHRP Logical Port FP/TD Associations" on page 14-64 for more information on associating flow profiles with NHRP logical ports.

To delete an NHRP logical port:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set all IP LPorts dialog box appears (see Figure 3-2 on page 3-6).
2. Select the switch where the NHRP logical port resides from the Switch Name list at the top of the dialog box. A list of IP logical ports and IP Server logical ports configured on the switch appears in the LPort Name list.
3. Select the ingress/egress IP logical port or IP Server logical port associated with the NHRP logical port.
4. Choose IP Parameters. The Set IP Parameters dialog box appears (see Figure 14-3 on page 14-10).
5. Choose Delete NHRP LPort. This action deletes the NHRP logical port. The Delete NHRP LPort button becomes the Add NHRP LPort button, allowing you to add the NHRP logical port whenever it becomes necessary. See "Adding an NHRP Logical Port" for more information on adding an NHRP logical port.
6. Choose Close to exit.

Configuring Node Parameters

NHRP node parameters allow you to specify the number of hops NHRP packets may traverse, tune SVC establishment, and control NHRP logging. These parameters apply to all NHS instances and the proxy client on the selected switch.

You do not have to configure NHRP node parameters. Default values are already configured for you.

To configure NHRP node parameters:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Node Parameters. The Set All NHRP Node Parameters dialog box appears (see Figure 14-4).

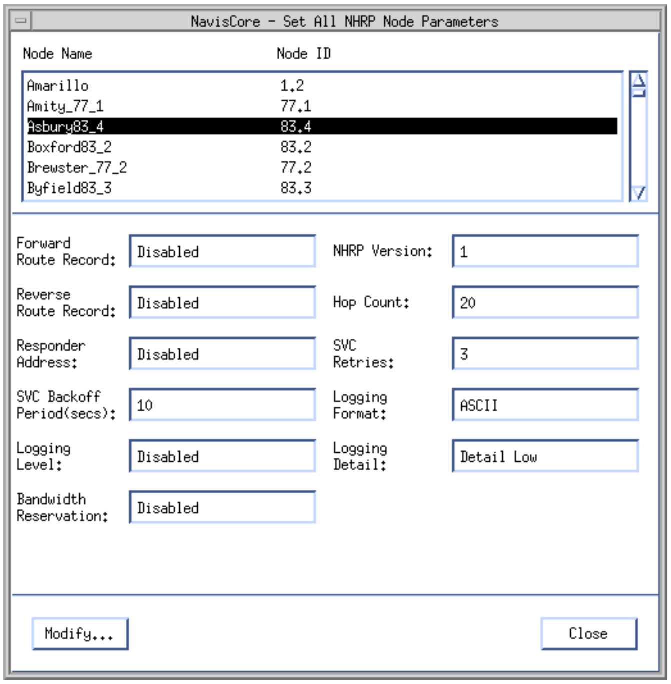

NavisCore - Set All NHRP Node Parameters

Node Name	Node ID
Amarillo	1.2
Amity_77_1	77.1
Asbury83_4	83.4
Boxford83_2	83.2
Brewster_77_2	77.2
Byfield83_3	83.3

Forward Route Record:	Disabled	NHRP Version:	1
Reverse Route Record:	Disabled	Hop Count:	20
Responder Address:	Disabled	SVC Retries:	3
SVC Backoff Period(secs):	10	Logging Format:	ASCII
Logging Level:	Disabled	Logging Detail:	Detail Low
Bandwidth Reservation:	Disabled		

Figure 14-4. Set All NHRP Node Parameters Dialog Box (With Default Values)

The Set All NHRP Node Parameters dialog box allows you to select a switch in the list box at the top of the dialog box, displaying the NHRP node parameters that are currently in effect for that switch at the bottom of the dialog box.

3. Select a switch.
4. Choose Modify to change the NHRP node parameters that are currently in effect for the selected switch. The Set NHRP Node Parameters dialog box appears (see Figure 14-5).

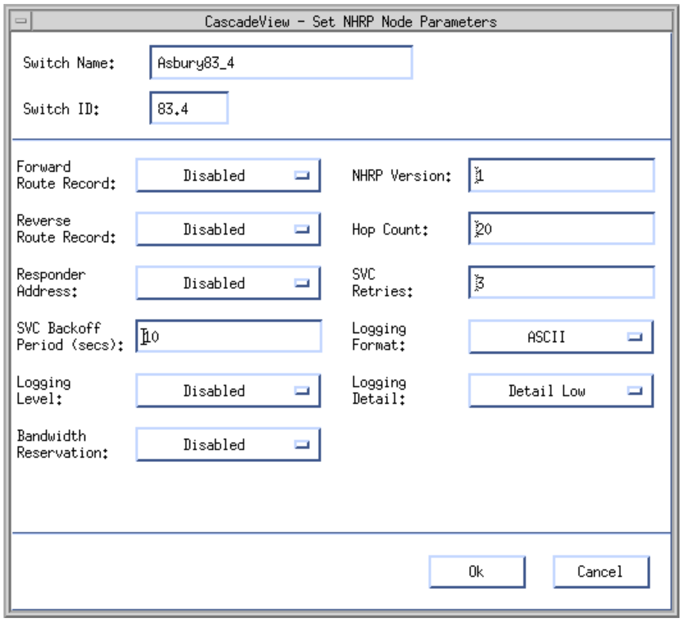

Switch Name: Asbury83_4

Switch ID: 83.4

Forward Route Record: Disabled

NHRP Version: 1

Reverse Route Record: Disabled

Hop Count: 20

Responder Address: Disabled

SVC Retries: 3

SVC Backoff Period (secs): 10

Logging Format: ASCII

Logging Level: Disabled

Logging Detail: Detail Low

Bandwidth Reservation: Disabled

Figure 14-5. Set NHRP Node Parameters Dialog Box

5. Modify the parameters described in Table 14-4.

Table 14-4. NHRP Node Parameters

Field	Action/Description
Switch Name	Displays the name of the selected switch.
Switch ID	Displays the subnetwork number and host ID in the internal IP address of the switch.
Forward Route Record	Choose *Enabled* or *Disabled* (the default) to include or suppress forward route records in NHRP requests and replies. If you enable this parameter, a record of the network- and link-layer addresses of all intermediate NHSs between the source and destination (that is, the forward direction) is included in NHRP requests and replies. The information in route records can help you troubleshoot network problems and detect data link filtering in NBMA networks.
Reverse Route Record	Choose *Enabled* or *Disabled* (the default) to enable or suppress reverse route records in NHRP requests and replies. If you enable this parameter, a record of the network- and link-layer addresses of all intermediate NHSs between the destination and source (that is, the reverse direction) is included in NHRP requests and replies. The information in route records can help you troubleshoot network problems and detect data link filtering in NBMA networks.

Table 14-4. NHRP Node Parameters (Continued)

Field	Action/Description
Responder Address	Choose *Enabled* or *Disabled* (the default) to enable or suppress the responder address option in NHRP request packets. If you enable this option, all NHSs that respond to NHRP requests from this switch will include their respective IP addresses in their NHRP replies.
SVC Backoff Period (msecs)	Specify the period of time, in milliseconds, that the switch waits between retry attempts to establish SVCs (the default is 10). The number of retry attempts that the switch makes is determined by the SVC Retries value. See the description of the SVC Retries field for more information. For example, if you specify 20, the switch will attempt to establish an SVC every 20 milliseconds if previous attempts fail. Either increase this value or increase the SVC Retries value (or increase both) if you find that SVC establishment attempts are timing out.
Logging Level	Choose the logging importance level of NHRP request/reply activity. Keep in mind that when you choose one level, you also choose all the levels above it. For example, if you choose *Warning*, you also choose *Critical* and *Fatal*. The choices are: *Disabled* – (default) Logging is disabled. (Logging levels in order of importance) *Fatal* – No memory available to process NHRP requests and replies. *Critical* – Low amount of memory available to process NHRP requests and replies. Includes the Fatal level. *Warning* – Dropping NHRP requests and replies due to queue overload. Includes the Critical and Fatal levels. *Info-High* – All Registration Requests, Registration Replies, Purge Requests, and Purge Replies are logged. Includes all of the above levels. *Info-Medium* – All Resolution Requests and Resolution Replies are logged. Includes all of the above levels. *Info-Low* – All Error Indication messages are logged. Includes all of the above levels. *Info-Debug* – All Registration Refresh Requests and Registration Refresh Replies are logged. Includes all of the above levels.
Bandwidth Reservation	Choose *Enabled* or *Disabled* (the default) to enable/disable support for backward feedback (CCRMs and BCMs) from ports that use the ATM Flow-Control Processor (ATM FCP) module. These ports reside on CBX 500 switches that are in the path of the SVC shortcut. If you enable this parameter, the reverse traffic descriptor contains a QoS class of Available Bit Rate (ABR) with 34 cells per second of reserved bandwidth. If you disable this parameter, the reverse traffic descriptor is set to 0.

Table 14-4. NHRP Node Parameters (Continued)

Field	Action/Description
NHRP Version	Specify the version of the generic address mapping and management protocol that the switch is supposed to use. The switch transmits this version in the Fixed portion of each NHRP packet. At this time, use the default value (1). No other value is allowed.
Hop Count	Specify the maximum number of NHSs that an NHRP packet may traverse before it is discarded. The default is 20. The default should suffice in most cases. Increase this value only if your network is so large that the selected switch is more than 20 hops (that is, 20 intermediate NHSs) away from other nodes with which it must exchange NHRP packets. Keep in mind that if two Lucent switches that act as NHSs communicate within a Lucent network, and multiple intermediate Lucent switches do not act as NHSs, then the Lucent switches that act as NHSs are only one hop apart.
SVC Retries	Specify the number of retry attempts that the switch makes to establish an SVC after the initial attempt fails. The default is 3. For example, suppose that you set this parameter to 4, and the switch fails to establish an SVC on its first try. The switch will then make up to 4 attempts to establish the SVC before giving up. Either increase this value or increase the SVC Backoff Period value (or increase both) if you find that SVC establishment attempts are timing out.
Logging Format	ASCII is the only supported logging format at this time. You cannot change this field.
Logging Detail	Choose the level of logging detail on the NHRP packets that are processed by the switch. Keep in mind that when you choose one level, you also choose all the levels above it. For example, if you choose *Detail-High*, you also choose *Detail-Medium* and *Detail-Low.* The choices are (in order): • *Detail-Low* (default) • *Detail-Medium* • *Detail-High* • *Detail-Debug* When the network is operating properly, consider setting the level of logging detail to *Detail-Low* or *Detail-Medium*. Setting the logging level to *Detail-High* or *Detail-Debug* can degrade NHRP performance, and should only be used when network problems appear. See Table 14-5 for a description of each of these logging detail levels.

Table 14-5. Logging Detail

Detail Level	Packet Type	Packet Information Logged
Detail-Low	NHRP Registration Request	Client IP address, address family, NBMA address, NBMA subaddress, prefix length, and preference for each client information element (CIE). An NHS maintains a CIE for each client it serves.
	NHRP Registration Reply	Status code, client IP address, address family, NBMA address, NBMA subaddress, prefix length, and preference for each CIE.
	NHRP Resolution Request	Requested destination IP address and the requestor's IP address.
	NHRP Resolution Reply	The status code, next hop IP address, address family, next hop NBMA address, next hop NBMA subaddress, prefix length, and preference (for the first CIE, the one with the highest preference value).
	NHRP Purge Request	Client IP address, address family, NBMA address, NBMA subaddress, and prefix length for each CIE.
	NHRP Purge Reply	Status code, client IP address, address family, NBMA address, NBMA subaddress, and prefix length for each CIE.
	NHRP Error Indication	Error code, source NBMA address, source NBMA subaddress, IP source address, and IP destination address.

Table 14-5. Logging Detail (Continued)

Detail Level	Packet Type	Packet Information Logged
Detail-Medium	NHRP Registration Request	Everything in Detail-Low (Registration Request), plus destination IP address, requestor's IP address, requestor's NBMA address, and requestor's NBMA subaddress.
	NHRP Registration Reply	Everything in Detail-Low (Registration Reply), plus requestor's IP address, requestor's NBMA address, and requestor's NBMA subaddress.
	NHRP Resolution Request	Everything in Detail-Low (Resolution Request), plus address family, requestor's NBMA address, and requestor's NBMA subaddress.
	NHRP Resolution Reply	Everything in Detail-Low (Resolution Reply) for multiple CIEs, plus requestor's IP address, requestor's NBMA address, and requestor's NBMA subaddress.
	NHRP Purge Request	Everything in Detail-Low (Purge Request), plus destination IP address, requestor's IP address, requestor's NBMA address, and requestor's NBMA subaddress.
	NHRP Purge Reply	Everything in Detail-Low (Purge Reply), plus destination IP address, requestor's IP address, requestor's NBMA address, and requestor's NBMA subaddress.
	NHRP Error Indication	Everything in Detail-Low (Error Indication), plus checksum
Detail-High	NHRP Registration Request	All of the above (Registration Request), plus flags, request ID, hold time, and Maximum Transfer Unit (MTU).
	NHRP Registration Reply	All of the above (Registration Reply), plus flags, request ID, hold time, and MTU.
	NHRP Resolution Request	All of the above (Resolution Request), plus flags, request ID, hold time, and MTU.
	NHRP Resolution Reply	All of the above (Resolution Reply), plus flags, request ID, hold time, and MTU.
	NHRP Purge Request	All of the above (Purge Request), plus flags, request ID, and hold time.
	NHRP Purge Reply	All of the above (Purge Reply), plus flags, request ID, and hold time.
	NHRP Error Indication	All of the above (Error Indication).

Table 14-5. Logging Detail (Continued)

Detail Level	Packet Type	Packet Information Logged
Detail-Debug	NHRP Registration Request	All of the above (Registration Request), plus extension information.
	NHRP Registration Reply	All of the above (Registration Reply), plus extension information.
	NHRP Resolution Request	All of the above (Resolution Request), plus extension information.
	NHRP Resolution Reply	All of the above (Resolution Reply), plus extension information.
	NHRP Purge Request	All of the above (Purge Request), plus extension information.
	NHRP Purge Reply	All of the above (Purge Reply), plus extension information.
	NHRP Error Indication	All of the above (Error Indication).

6. Choose OK to enter your changes, or choose Cancel to exit without entering your changes.

Configuring Flow Profiles

You can configure profiles of IP traffic flows called *IP flow profiles*. A flow profile is defined in terms of:

- IP source and destination addresses (with Classless Interdomain Routing [CIDR] masking)
- IP protocol type
- Source and destination TCP/UDP-like protocol ports
- Type of Service (TOS) and TOS mask

After you configure these flow profiles, you associate them with IP logical ports along with traffic descriptors that manage the data flow. You can associate a single profile with multiple IP logical ports on a switch, and/or associate multiple profiles with a single IP logical port. See "Configuring NHRP Logical Port FP/TD Associations" on page 14-64 for more information.

▶ When you associate IP flow profiles with NHRP logical ports, the order in which you associate the IP flow profiles is very important. Switches allocate resources to IP flow profiles based on the order in which they are associated with an NHRP logical port. The IP flow profile that is associated first has the highest resource priority, the IP flow profile associated second has the next highest resource priority, and so on. Thus, make sure you associate IP flow profiles in the following order: from most important to least important.

About Wildcards

When you create a flow profile, you can use wildcards for source and destination IP addresses, and source and destination protocol ports. The wildcard, in effect, indicates that the flow profile applies to “any” or “all” addresses, applications, and so on.

The flow profile controls the flow of traffic in one direction only: from source to destination. To manage the bi-directional flow of traffic between two specific points in the network, you would have to create two or more flow profiles. The source IP address in one flow profile would be the destination IP address in the other profile, and vice versa.

Use of the wildcard IP address gives you a lot of flexibility. For example, you could create a flow profile with a specific source IP address (such as the IP address of a network) and a wildcard destination IP address (0.0.0.0). This profile would then control the flow of traffic from one IP network to any reachable 32-bit destination address.

The use of wildcards also applies to source and destination applications. Through use of the wildcard, you can specify that the flow profile applies to all source and destination applications. Or, you can restrict the flow profile to just specific applications. However, keep in mind that a flow profile is unidirectional, from source to destination. To manage bidirectional traffic flow between applications on two specific endpoints, you must create multiple flow profiles, one or more for each direction.

While the use of wildcards can be beneficial, keep in mind that their use may adversely impact performance under high traffic load conditions. This guideline especially applies to the use of wildcards in the destination IP address. If you specify a destination IP address of 0.0.0.0 (the wildcard destination IP address), keep in mind that, for each traffic flow that matches the criteria in the profile, a separate SVC shortcut will be created for each destination IP address that is detected. For example, if 20 traffic flows match the criteria in the profile, and each flow is destined for a different host (that is, a different IP address), then 20 SVC shortcuts will be created (one for each traffic flow). Under high traffic load conditions, performance may be adversely effected, since an excessive number of SVC shortcuts may be created.

Creating IP Flow Profiles

To create IP flow profiles:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Flow Profiles. The Set All NHRP Node Flow Profiles dialog box appears (see Figure 14-6).

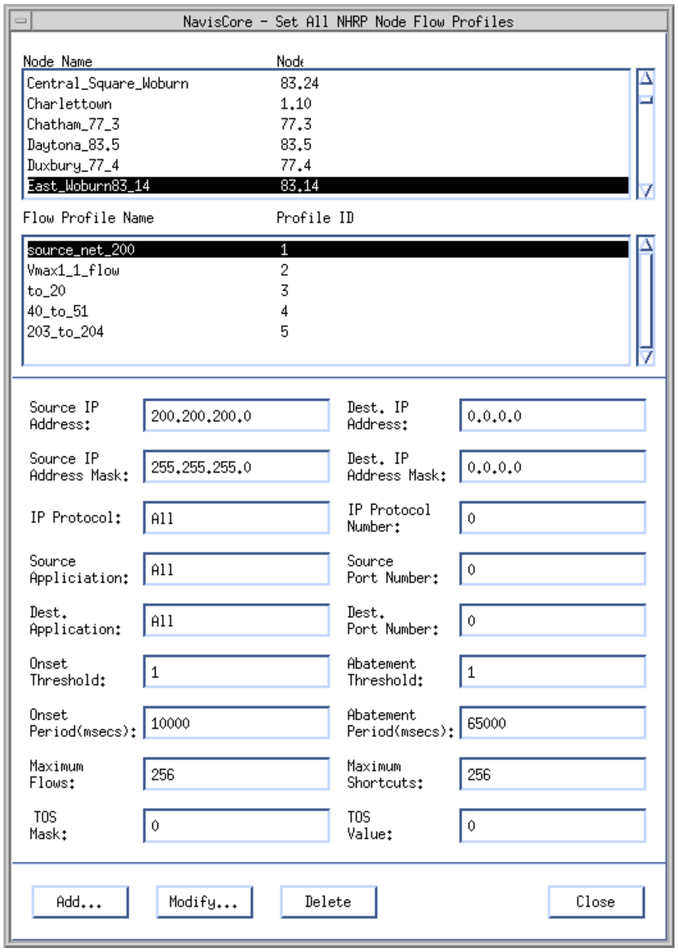

NavisCore - Set All NHRP Node Flow Profiles

Node Name	Node
Central_Square_Woburn	83.24
Charlettown	1.10
Chatham_77_3	77.3
Daytona_83.5	83.5
Duxbury_77_4	77.4
East_Woburn83_14	83.14

Flow Profile Name	Profile ID
source_net_200	1
Vmax1_1_flow	2
to_20	3
40_to_51	4
203_to_204	5

Source IP Address:	200.200.200.0	Dest. IP Address:	0.0.0.0
Source IP Address Mask:	255.255.255.0	Dest. IP Address Mask:	0.0.0.0
IP Protocol:	All	IP Protocol Number:	0
Source Appliciation:	All	Source Port Number:	0
Dest. Application:	All	Dest. Port Number:	0
Onset Threshold:	1	Abatement Threshold:	1
Onset Period(msecs):	10000	Abatement Period(msecs):	65000
Maximum Flows:	256	Maximum Shortcuts:	256
TOS Mask:	0	TOS Value:	0

Figure 14-6. Set All NHRP Node Flow Profiles Dialog Box

The Set All NHRP Node Flow Profiles dialog box allows you to select a switch in the list box at the top of the dialog box. This action displays all the flow profiles configured for that switch in the second list box (the total number of flow profiles per switch cannot exceed 512). You can then select each flow profile, and the parameters for that profile will display in the fields in the bottom half of the dialog box.

3. Select a switch.

4. Choose Add. The Add NHRP Node Flow Profiles dialog box appears (see Figure 14-7).

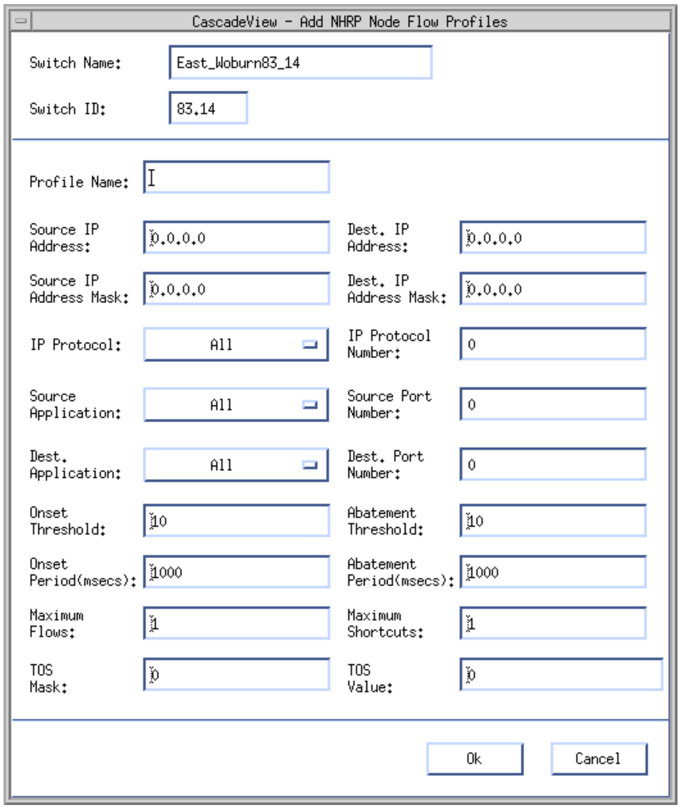

CascadeView - Add NHRP Node Flow Profiles

Switch Name:	East_Woburn83_14		
Switch ID:	83.14		
Profile Name:			
Source IP Address:	0.0.0.0	Dest. IP Address:	0.0.0.0
Source IP Address Mask:	0.0.0.0	Dest. IP Address Mask:	0.0.0.0
IP Protocol:	All	IP Protocol Number:	0
Source Application:	All	Source Port Number:	0
Dest. Application:	All	Dest. Port Number:	0
Onset Threshold:	10	Abatement Threshold:	10
Onset Period(msecs):	1000	Abatement Period(msecs):	1000
Maximum Flows:	1	Maximum Shortcuts:	1
TOS Mask:	0	TOS Value:	0

Figure 14-7. Add NHRP Node Flow Profiles Dialog Box

5. Specify the parameters described in Table 14-6.

Only one of the following parameters may have an ambiguous value: Dest. IP Address, Source Application, or Dest. Application. For Dest. IP Address, an ambiguous value is expressed as 0.0.0.0. For Source Application and Dest. Application, an ambiguous value is expressed by selecting *Ambiguous* from the pull-down menu.

6. Choose OK.

Table 14-6. NHRP IP Flow Profile Parameters

Field	Action/Description
Switch Name	Displays the name of the selected switch.
Switch ID	Displays the subnetwork number and host ID in the internal IP address of the switch.
Profile Name	Specify the name of the profile. The name may not exceed 20 characters (including spaces).
Source IP Address	Specify the IP address of the source of the traffic flow. You may specify any valid IP address in the format a.b.c.d, where a, b, c, and d must be from 0 to 255. The default address (0.0.0.0) acts as a wildcard, which means that the flow profile applies to traffic from all sources. Examples of valid addresses include a class B network address of 152.148.0.0, a class B subnetwork address of 152.148.22.0, and a class B host address of 152.148.22.12.
Source IP Address Mask	Specify the mask for the source IP address. You may specify any valid mask in the format a.b.c.d, where a, b, c, and d must be from 0 to 255. If you use the default source IP address (0.0.0.0), then use the default mask (0.0.0.0). For example, if you specify a source IP address of 131.100.0.0 (a class B network address), you would specify a source IP address mask of 255.255.0.0. In another example, if you specify a source IP address of 131.100.25.0 (a class B address with a subnetwork number of 25), you would specify a mask of 255.255.255.0.

Table 14-6. NHRP IP Flow Profile Parameters (Continued)

Field	Action/Description
IP Protocol	Choose the IP protocol for the flow profile. The default (All) acts as a wildcard. You may choose from the following list of IP protocols: *All* – The flow profile accepts all types of IP protocols. If you leave the default unchanged, “0” displays in the IP Protocol Number field. *TCP* – TCP stands for Transmission Control Protocol. This protocol is connection-oriented, guaranteeing reliable transmission between applications that use it. Examples of applications that use TCP include FTP, Telnet, and HTTP (World Wide Web protocol). If you choose TCP, “6” displays in the IP Protocol Number field. This number is the protocol’s official number as assigned by the Internet Assigned Numbers Authority (IANA). *UDP* – UDP stands for User Datagram Protocol. This protocol is connectionless — it makes a best effort to deliver datagrams between applications. Examples of applications that use UDP include RIP, TFTP, and SNMP. If you choose UDP, “17” displays in the IP Protocol Number field. This number is the protocol’s official number as assigned by the IANA. *ICMP* – ICMP stands for Internet Control Message Protocol. This protocol provides dynamic routing support, such as routing redirects when routes are unavailable. If you choose ICMP, “1” displays in the IP Protocol Number field. This number is the protocol’s official number as assigned by the IANA. *User Specified* – Allows you to specify a protocol not included in this list, such as a proprietary protocol. If you choose *User Specified*, you must enter the IP protocol number assigned to the protocol in the IP Protocol Number field. This number is assigned by the IANA. The IANA can be reached at their Web site (http://www.iana.org).

Table 14-6. NHRP IP Flow Profile Parameters (Continued)

Field	Action/Description
Source Application (*Meaningful only if TCP or UDP is selected for IP Protocol*)	Choose the source application for the flow profile. The list of applications from which you can choose is determined by your choice of IP protocol. *All* – (default) The flow profile applies to all application traffic from source to destination. *BGP* – BGP traffic. *FTP Data* – FTP data traffic. *FTP Control* – FTP control traffic. *Gopher* – Gopher traffic. *IRC* – IRC traffic. *Talk* – Talk traffic. *Telnet* – Telnet traffic. *WWW* – HTTP traffic. *RIP* – RIP traffic. *SNMP* – SNMP traffic. *SNMP Traps* – SNMP traps traffic. *TFTP* – TFTP traffic. *Ambiguous* – The traffic flows are tracked based on discrete source application values (1024 is used for the port value). *User Specified* – Allows you to specify an application that is not in the list (such as a third-party vendor application), and requires you to enter a port number that the application uses in the Source Port Number field. You can obtain this port number from the application developer. The IANA web site (http://www.iana.org) can also help you determine this number, as well as RFC 1700. There is also a listing of many port numbers at the following URL: http://www.isi.edu/in-notes/iana/assignments/port-numbers.
Dest. Application (*Meaningful only if TCP or UDP is selected for IP Protocol*)	Choose the destination application for the flow profile. In most cases, your choices of source and destination application will be the same, since communication between different types of applications is rare. See the description of the Source Application field for information on the choices available to you.
Onset Threshold	Specify the number of IP packets that must be detected within the Onset Period in order to trigger the flow profile. For example, suppose that you set the Onset Threshold to 10 packets and you set the Onset Period to 100 milliseconds. This means that 10 or more packets must be detected within a span of 100 milliseconds in order to trigger a flow onset detection. The default value (10) should suffice in most cases. See "Adjusting Onset and Abatement Parameters" on page 14-30 for more information.

Table 14-6. NHRP IP Flow Profile Parameters (Continued)

Field	Action/Description
Onset Period (msecs)	Specify the span of time, in milliseconds, within which the Onset Threshold is in effect. See the description of the Onset Threshold for more information. The default value (1000) should suffice in most cases. See "Adjusting Onset and Abatement Parameters" on page 14-30 for more information.
Maximum Flows	Specify the maximum number of matching NHRP traffic flows that can be simultaneously detected on a logical port associated with this flow profile. A matching NHRP traffic flow is one that matches the criteria defined in the flow profile. For example, if you use the default value (1), then only one matching NHRP traffic flow can be detected on the associated logical port at any one time. The default value should suffice in most cases. Because each traffic flow uses a shortcut, the Maximum Flows value must be greater than or equal to the Maximum Shortcuts value. Otherwise, you may encounter situations in which flows cannot be supported because insufficient shortcuts are available. See the description of "Maximum Shortcuts" on page 14-29 for more information.
TOS Mask	Specify the Type of Service (TOS) mask, in decimal, for the TOS value (see the description of the TOS field for more information). The mask determines the location of bits required for the TOS value. For example, if you specify a TOS value of 4 (100 in binary), you must specify a compatible TOS mask (such as decimal 4). Valid TOS masks range from 0 to 255. If you specify an invalid TOS mask/value combination, the NMS generates an error. The default TOS mask (0) combined with the default TOS value (0) means that the flow profile will work with any TOS. If you use a TOS mask and TOS value other than 0, make sure you specify values that are compatible with the network equipment (such as CPE) with which the switch exchanges service traffic (such as voice over IP).

Table 14-6. NHRP IP Flow Profile Parameters (Continued)

Field	Action/Description
Dest. IP Address	Specify the IP address of the destination of the traffic flow. You may specify any valid IP address in the format a.b.c.d, where a, b, c, and d must be from 0 to 255. The default address (0.0.0.0) acts as a wildcard, which means that the flow profile will handle traffic to all destinations. If you specify a destination address of 0.0.0.0, the flow profile is triggered whenever the packets in a specific flow to any destination exceeds the Onset Threshold. If you specify a destination IP address of 0.0.0.0, keep in mind that, for each traffic flow that matches the criteria in the profile, a separate SVC shortcut will be created for each destination IP address that is detected. For example, if 20 traffic flows match the criteria in the profile, and each flow is destined for a different host (that is, a different IP address), then 20 SVC shortcuts will be created (one for each traffic flow). Under high traffic load conditions, performance may be adversely effected, since an excessive number of SVC shortcuts may be created. If you specify a network or subnetwork as the destination IP address, a single SVC shortcut will be created to the network or subnetwork. The shortcut will terminate at the last hop to the network or subnetwork. Examples of valid addresses include a class B network address of 152.148.0.0, a class B subnetwork address of 152.148.22.0, and a class B host address of 152.148.22.12.
Dest. IP Address Mask	Specify the mask for the destination IP address. You may specify any valid IP address mask in the format a.b.c.d, where a, b, c, and d must be from 0 to 255. If you use the default destination IP address (0.0.0.0), then use the default destination IP address mask (0.0.0.0). If you specify a destination IP address mask of 0.0.0.0, the flow profile becomes ambiguous (that is, can handle traffic to all destinations). For example, if you specify a destination IP address of 131.100.0.0 (a class B network address), you would specify a destination IP address mask of 255.255.0.0. In another example, if you specify a destination IP address of 131.100.25.0 (a class B subnetwork address), you would specify a destination IP address mask of 255.255.255.0.
IP Protocol Number	Specify the number that identifies the IP protocol only if you chose *User Specified* for the IP Protocol. See the description of the IP Protocol field for more information. If you chose an IP Protocol type other than *User Specified*, NavisCore does not allow you to enter a value in this field, and displays the IP protocol number associated with that choice.

Table 14-6. NHRP IP Flow Profile Parameters (Continued)

Field	Action/Description
Source Port Number	Specify the port number used by the source application only if you chose *User Specified* for the Source Application. You can obtain this port number from the application developer. The IANA web site (http://www.iana.org) can also help you determine this number, as well as RFC 1700. There is also a listing of many port numbers at the following URL: http://www.isi.edu/in-notes/iana/assignments/port-numbers If you chose a Source Application other than *User Specified*, NavisCore does not allow you to enter a value in this field, and displays the port number associated with that choice. See the description of the Source Application field for more information.
Dest. Port Number	Specify the port number used by the destination application only if you chose *User Specified* for the Destination Application. You can obtain this port number from the application developer. The IANA web site (http://www.iana.org) can also help you determine this number, as well as RFC 1700. There is also a listing of many port numbers at the following URL: http://www.isi.edu/in-notes/iana/assignments/port-numbers If you chose a Destination Application other than *User Specified*, NavisCore does not allow you to enter a value in this field, and displays the port number associated with that choice. See the description of the Destination Application field for more information.
Abatement Threshold	Specify the minimum number of packets that must be detected within the Abatement Period in order to keep the flow profile active. For example, suppose that you set the Abatement Threshold to 10 packets and you set the Abatement Period to 100 milliseconds. This means that at least 10 packets must be detected within a span of 100 milliseconds in order to keep the flow profile active. If less than 10 packets are detected within a span of 100 milliseconds, the associated SVC shortcut is taken down. The default value (10) should suffice in most cases. See "Adjusting Onset and Abatement Parameters" on page 14-30 for more information.
Abatement Period (msecs)	Specify the span of time, in milliseconds, within which the Abatement Threshold is in effect. See the description of the Abatement Threshold for more information. The default value (1000) should suffice in most cases. See "Adjusting Onset and Abatement Parameters" on page 14-30 for more information.

Table 14-6. NHRP IP Flow Profile Parameters (Continued)

Field	Action/Description
Maximum Shortcuts	Specify the maximum number of matching NHRP shortcuts that can be simultaneously connected to a logical port associated with this flow profile. A matching NHRP shortcut is one that matches the criteria defined in the flow profile. For example, if you use the default (1), then one shortcut that matches this flow profile can be connected to the associated logical port at any one time. The default should suffice in most cases. Because each traffic flow uses a shortcut, the Maximum Shortcuts value must be less than or equal to the Maximum Flows value. Otherwise, you may encounter situations in which flows cannot be supported because insufficient shortcuts are available. See the description of "Maximum Flows" on page 14-26 for more information.
TOS Value	Specify the Type of Service (TOS) value. This value identifies a traffic flow associated with a particular TOS, such as voice. Values may range from 0 (the default) to 255. The default TOS mask (0) combined with the default TOS value (0) means that the flow profile will work with any TOS. If you specify a non-zero TOS value, you must specify a non-zero TOS mask. See the description of the TOS Mask field for more information. If you use a TOS mask and TOS value other than 0, make sure you specify values that are compatible with the network equipment (such as CPE) with which the switch exchanges service traffic (such as voice over IP).

Adjusting Onset and Abatement Parameters

The onset and abatement parameters — Onset Threshold, Onset Period, Abatement Threshold, and Abatement Period — enable you to tune traffic flow processing. In most cases, the default values for these parameters are sufficient, but you may want to use values other than the defaults if performance is not satisfactory.

When you adjust the parameters, note the following:

- When you increase the Onset Threshold and/or decrease the Onset Period, you increase the risk that the flow profile will not be triggered. This may result in no SVC establishment for some traffic flows, depending on how much you increase the Onset Threshold or decrease the Onset Period.
- When you increase the Onset Period and/or decrease the Onset Threshold, you reduce the risk that the flow profile will not be triggered, thereby increasing the chance of SVC establishment for all traffic flows.
- When you increase the Abatement Threshold and/or decrease the Abatement Period, you increase the risk that some SVCs will be prematurely disconnected (before their associated traffic flows complete).
- When you increase the Abatement Period and/or decrease the Abatement Threshold, you reduce the risk that SVCs will be prematurely disconnected.

Modifying IP Flow Profiles

Before you attempt to modify an IP flow profile, you need to understand the following rules:

- You may modify only IP flow profiles that are not in use. If you attempt to modify an IP flow profile that is in use, the NMS generates an error.
- You must remove all IP flow profile associations with NHRP logical ports before you can modify an IP flow profile. For example, if you have associated an IP flow profile with two NHRP logical ports, you must remove both of those associations before you can modify the IP flow profile. When you finish modifying the IP flow profile, you must reassociate it with both logical ports.

To modify an IP flow profile:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Flow Profiles. The Set All NHRP Node Flow Profiles dialog box appears (see Figure 14-6 on page 14-21).
3. Select a switch. A list of IP flow profiles configured for that switch appears.
4. Select the flow profile you want to modify.
5. Choose Modify. The Modify NHRP Node Flow Profiles dialog box appears. This dialog box is similar to the Add NHRP Node Flow Profiles dialog box (see Figure 14-7 on page 14-22).

6. Modify the fields on the Modify NHRP Node Flow Profiles dialog box as needed. See Table 14-6 on page 14-23 for descriptions of these fields.
7. Choose OK.

Deleting IP Flow Profiles

Before you attempt to delete an IP flow profile, you need to understand the following rules:

- You may delete only IP flow profiles that are not in use. If you attempt to delete an IP flow profile that is in use, the NMS generates an error.
- You must remove all IP flow profile associations with NHRP logical ports before you can delete an IP flow profile. For example, if you have associated an IP flow profile with two NHRP logical ports, you must remove both of those associations before you can delete the IP flow profile.

To delete an IP flow profile:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Flow Profiles. The Set All NHRP Node Flow Profiles dialog box appears (see Figure 14-6).
3. Select a switch. A list of IP flow profiles configured for that switch appears.
4. Select the flow profile you want to delete.
5. Choose Delete.
6. Choose OK when prompted.

Configuring Servers

You can configure one or more NHSs on a switch; however, you have the option to use the default server instead of explicitly configuring NHSs. You then associate each explicitly configured server (or the default server) with an NHRP logical port, one that is used to communicate with a destination. See "Configuring NHRP Logical Port Parameters" on page 14-59 for more information on associating servers with logical ports.

Although NHCs typically register their IP Address/NBMA Address mappings with the NHS automatically, you can also manually create server cache entries for NHCs, if necessary.

The following sections describe how to add, modify, and delete servers and their cache entries.

About the Default Server

Each switch has a default NHS which is automatically associated with all of the trunk logical ports on the switch. The default NHS processes NHRP requests and replies on the trunk logical ports. These requests and replies are received from and sent to other NHSs and the proxy client on switches in the Lucent network. Though intended primarily for use with trunk logical ports, you may associate the default NHS with one or more ingress/egress NHRP logical ports (such as an ATM UNI logical port) that interact with NHCs.

Adding a Server

To add a server:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Servers ⇒ Set Parameters. The Set All NHRP Server Parameters dialog box appears (see Figure 14-8).

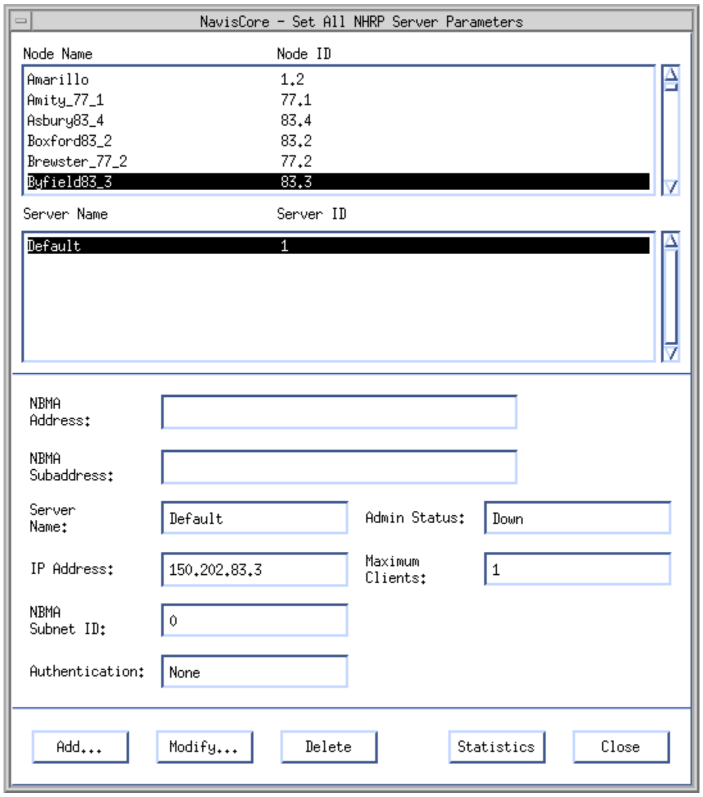

NavisCore - Set All NHRP Server Parameters

Node Name	Node ID
Amarillo	1.2
Amity_77_1	77.1
Asbury83_4	83.4
Boxford83_2	83.2
Brewster_77_2	77.2
Byfield83_3	83.3

Server Name	Server ID
Default	1

NBMA Address:	
NBMA Subaddress:	
Server Name:	Default
Admin Status:	Down
IP Address:	150.202.83.3
Maximum Clients:	1
NBMA Subnet ID:	0
Authentication:	None

Figure 14-8. Set All NHRP Server Parameters Dialog Box

The Set All NHRP Server Parameters dialog box allows you to select a switch in the list box at the top of the dialog box. This action displays all the servers configured for that switch in the second list box. You can then select each server, and the parameters for that server display in the fields in the bottom half of the dialog box. Table 14-7 on page 14-35 describes these fields, except the Current Clients field. This field displays the number of NHCs that are currently registered with the server.

▶ You can display statistics on server activity by choosing the Statistics button. See the *NavisCore Diagnostics Guide* for descriptions of these statistics.

3. Select a switch.

4. Choose Add. The Set NHRP Server dialog box appears (see Figure 14-9).

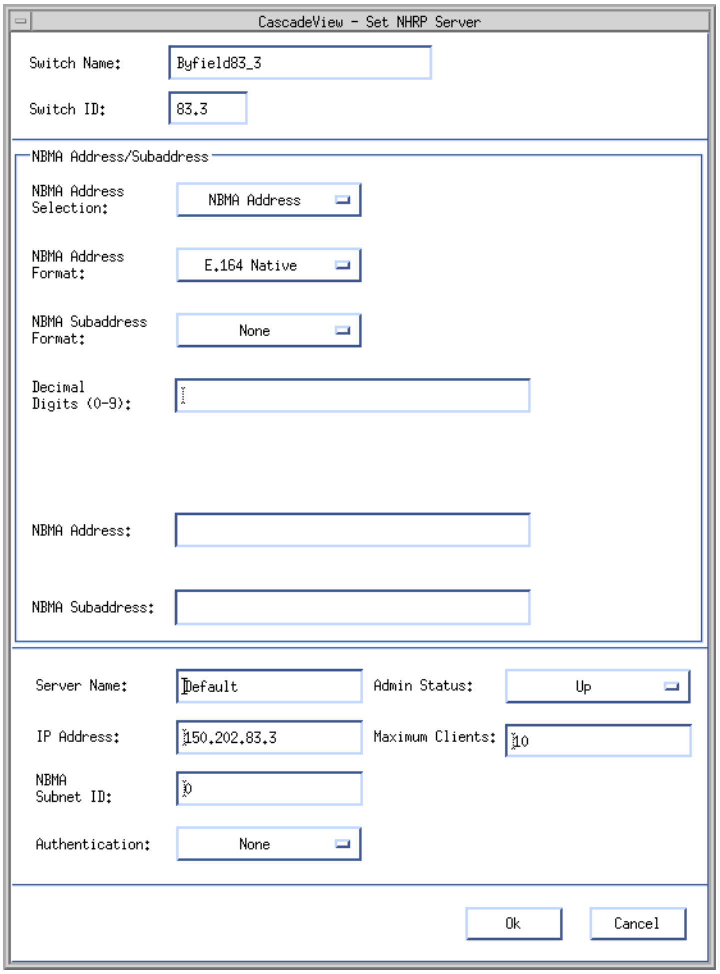

Switch Name: Byfield83_3

Switch ID: 83.3

NBMA Address Selection: NBMA Address

NBMA Address Format: E.164 Native

NBMA Subaddress Format: None

Server Name: Default

Admin Status: Up

IP Address: 150.202.83.3

Maximum Clients: 10

NBMA Subnet ID: 0

Authentication: None

Figure 14-9. Set NHRP Server Dialog Box

5. Specify the parameters described in Table 14-7.

6. Choose OK.

Table 14-7. NHRP Server Parameters

Field	Action/Description
Switch Name	Displays the name of the selected switch.
Switch ID	Displays the subnetwork number and host ID in the internal IP address of the switch.
NBMA Address Selection	Select the type of address you are configuring for the server: *NBMA Address* (default) or *NBMA Subaddress*. You select *NBMA Address* regardless of the address format you use. You select *NBMA Subaddress* only if you use *E.164 with NSAP* addressing. If you use E.164 with NSAP addresses, you specify both an address and a subaddress. Specify the address and subaddress in the following order: 1. In the NBMA Address Selection field, select *NBMA Address*. 2. In the NBMA Address Format field, select *E.164 with NSAP.* 3. In the Decimal Digits field, specify a native E.164 address. 4. Return to the NBMA Address Selection field. Select *NBMA Subaddress*. 5. In the NBMA Address Selection field, select *NBMA Subaddress*. 6. In the NBMA Address Format field, select a valid address format. 7. Specify an AFI if you selected a subaddress format of *Custom AESA*. 8. In the Hex Digits field, specify the hexadecimal subaddress.

Table 14-7. NHRP Server Parameters (Continued)

Field	Action/Description
NBMA Address Format	Select one of the following formats for the server's NBMA address: • *E.164 Native* (default) • *E.164 with NSAP* • *DCC AESA* • *ICD AESA* • *E.164 AESA* • *Custom AESA* See "About NBMA Addressing" on page 14-4 for more information on these formats. Your choice of formats determines how you specify address information in the other fields on the dialog box: ***E.164 Native* and *E.164 with NSAP*** – You must type in a decimal address in the Decimal Digits field (which appears only if you make either of these choices). In addition, if you select *E.164 with NSAP* you must select an NBMA subaddress format and type in a decimal subaddress in the Decimal Digits field. ***DCC AESA, ICD AESA, E.164 AESA,* and *Custom AESA*** – You must type in a hexadecimal address in the Hex Digits field (which, along with the AFI field, appears only if you make any of these selections). With the exception of Customer AESA, the NMS specifies the AFI for you in the AFI field. If you select Custom AESA, you must type in the appropriate AFI in the AFI field yourself.
NBMA Subaddress Format (*E.164 with NSAP only*)	Select one of the following formats for the server's subaddress: • *DCC AESA* • *ICD AESA* • *E.164 AESA* • *Custom AESA* For example, if the switch is connected to a private network that supports DCC AESA addressing, you would select *DCC AESA*. After you select a format, type in a hexadecimal subaddress in the Hex Digits field. In addition, if you select *Custom AESA*, you must also type in an AFI in the AFI field. See "About NBMA Addressing" on page 14-4 for more information on these formats.
AFI (*DCC AESA, ICD AESA, E.164 AESA, and Custom AESA only*)	If you selected *Custom AESA* as the address or subaddress format, type in the two-digit AFI. If you made any other address or subaddress selection, this field is read-only and displays the AFI for the selected format.

Table 14-7. NHRP Server Parameters (Continued)

Field	Action/Description
Hex Digits (0-F) (*DCC AESA, ICD AESA, E.164 AESA, and Custom AESA only*)	Specify the hexadecimal address or subaddress of the server, up to 38 digits long. As you type an address, it displays in the NBMA Address field. As you type a subaddress, it displays in the NBMA Subaddress field.
Decimal Digits (0-9) (*E.164 Native and E.164 with NSAP only*)	Specify the decimal address of the server, up to 15 digits long. As you type an address, it displays in the NBMA Address field.
NBMA Address	Displays the NBMA address as you type it in the Hex Digits or Decimal Digits field. The NBMA Address field is read-only.
NBMA Subaddress	Displays the NBMA subaddress as you type it in the Hex Digits field. The NBMA Subaddress field is read-only.
Server Name	Specify the server's name, which can be up to 20 characters (including spaces). The default server name is "Default."
IP Address	Specify the server's IP address (e.g., 131.100.24.3). The default is the IP address of the switch.
NBMA Subnet ID	Specify the ID of the server's logical NBMA subnetwork. The default is 0.
Authentication	Not supported at this time.
Admin Status	Set the administrative status of the server to *Up* (default) or *Down*.
Maximum Clients	Specify the maximum number of NHCs that are allowed to register with the server. The default is 10.

Modifying a Server

To modify server parameters:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Servers ⇒ Set Parameters. The Set All NHRP Server Parameters dialog box appears (see Figure 14-8 on page 14-33).
3. Select a switch. A list of servers configured for that switch appears.
4. Select the server whose parameters you want to modify.
5. Choose Modify. The Set NHRP Server dialog box appears (see Figure 14-9 on page 14-34).
6. Modify the parameters described in Table 14-7.
7. Choose OK.

Deleting a Server

▶ You cannot delete the default server, and you cannot delete any server if it is associated with a logical port. You must delete all of the server’s logical port associations before you can delete the server.

To delete a server:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Servers ⇒ Set Parameters. The Set All NHRP Server Parameters dialog box appears (see Figure 14-8 on page 14-33).
3. Select a switch. A list of servers configured for that switch appears.
4. Select the server you want to delete.
5. Choose Delete.
6. Choose OK.

Adding Server Cache Entries

A server cache entry maps an IP address of a next hop to its NBMA (i.e., ATM) address. When the server receives an NHRP Resolution Request, the server looks in its cache to determine if one of the entries has an IP address that matches the IP address in the request. If the server finds a match, it returns an NHRP Resolution Reply containing the associated NBMA address. A shortcut can then be established to the destination that is associated with the cache entry.

A server cache entry identifies both the *IP address of the destination* and the *IP address of the next hop* used to reach the destination. In many cases, these addresses are one and the same — they refer to the same network node (that is, the next hop *is* the destination). However, in other cases, these addresses are different.

If the server and the registering client are directly connected by the same logical NBMA subnetwork, then the next hop IP address and destination IP address are the same – that is, the next hop (the registering client) is the destination. Otherwise, the next hop IP address refers to the egress router for the NBMA subnetwork that is closest to the destination, and the destination IP address refers to the destination itself. The NBMA address is mapped to the next hop IP address, not the destination IP address.

Cache entries can be added to the server cache in two ways:

Dynamically — NHCs send their IP/ATM address mappings to the server in NHRP Registration Requests.

Manually — You add the IP/ATM address mappings yourself. You need to do this if you have nodes that do not support NHRP and therefore cannot register their IP/ATM address mappings with the NHS.

To add an entry to the server cache:

1. Select a switch from the network map.

2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Servers ⇒ Set Cache. The Set All NHRP Cache Entries dialog box appears (see Figure 14-10).

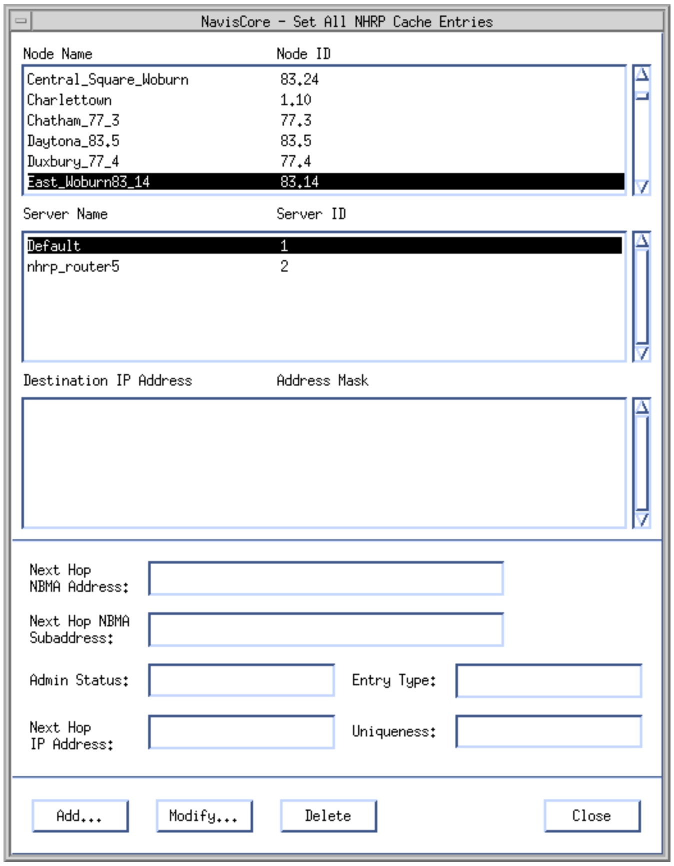

Node Name	Node ID
Central_Square_Woburn	83.24
Charlettown	1.10
Chatham_77_3	77.3
Daytona_83.5	83.5
Duxbury_77_4	77.4
East_Woburn83_14	83.14

Server Name	Server ID
Default	1
nhrp_router5	2

Destination IP Address	Address Mask

Figure 14-10. Set All NHRP Cache Entries Dialog Box (Server)

The Set All NHRP Cache Entries dialog box allows you to select a switch in the list box at the top of the dialog box. This action displays all the servers configured for that switch in the second list box. You can then select each server, and the cache entries for that server display in a list. Table 14-8 on page 14-42 describes these fields.

3. Select a switch. The servers configured on that switch appear.

4. Select the server for which you want to add an NHRP cache entry.

5. Choose Add. The Add Static Cache Entry dialog box appears (see Figure 14-11).

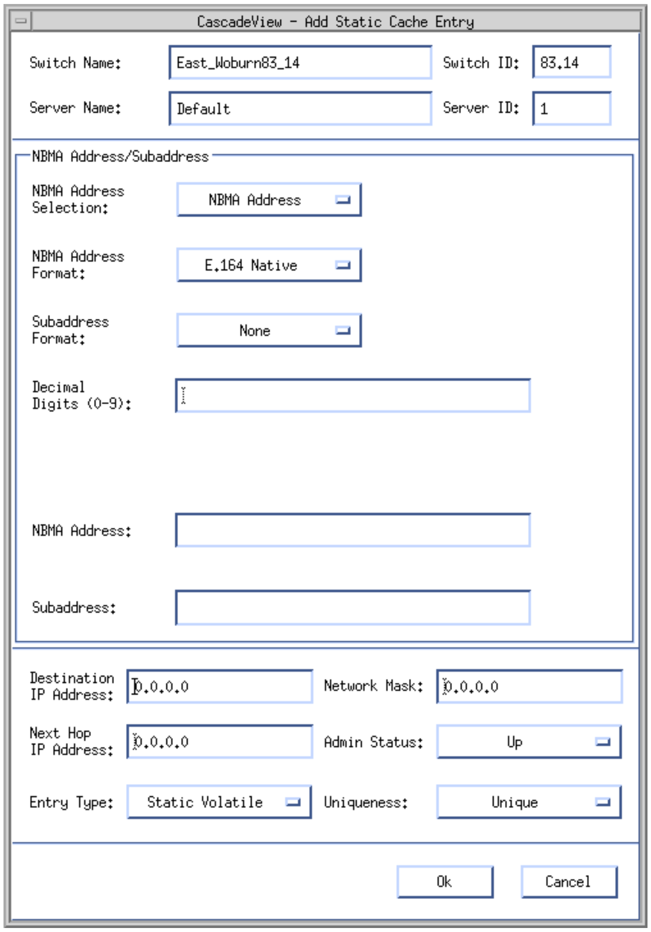

CascadeView - Add Static Cache Entry

Switch Name: East_Woburn83_14
Switch ID: 83.14
Server Name: Default
Server ID: 1

NBMA Address/Subaddress

NBMA Address Selection: NBMA Address
NBMA Address Format: E.164 Native
Subaddress Format: None
Decimal Digits (0-9):
NBMA Address:
Subaddress:

Destination IP Address: 0.0.0.0
Network Mask: 0.0.0.0
Next Hop IP Address: 0.0.0.0
Admin Status: Up
Entry Type: Static Volatile
Uniqueness: Unique

Figure 14-11. Add Static Cache Entry Dialog Box (Server)

6. Specify the parameters described in Table 14-8.

Table 14-8. NHRP Server Cache Entry Parameters

Field	Action/Description
Switch Name	Displays the name of the selected switch.
Switch ID	Displays the subnetwork number and host ID in the internal IP address of the switch.
Server Name	Displays the name of the server.
Server ID	Displays the internally assigned ID of the server.
NBMA Address Selection	Select the type of address you are configuring for the next hop: *NBMA Address* (default) or *NBMA Subaddress*. You select *NBMA Address* regardless of the address format you use. You select *NBMA Subaddress* only if you use *E.164 with NSAP* addressing. If you use E.164 with NSAP addresses, you specify both an address and a subaddress in the following order: 1. In the NBMA Address Selection field, select *NBMA Address*. 2. In the NBMA Address Format field, select *E.164 with NSAP.* 3. In the Decimal Digits field, specify a native E.164 address. 4. Return to the NBMA Address Selection field. Select *NBMA Subaddress*. 5. In the NBMA Address Selection field, select *NBMA Subaddress*. 6. In the NBMA Address Format field, select a valid address format. 7. Specify an AFI if you selected a subaddress format of *Custom AESA*. 8. In the Hex Digits field, specify the hexadecimal subaddress.

Table 14-8. NHRP Server Cache Entry Parameters (Continued)

Field	Action/Description
NBMA Address Format	Select the NBMA address format that the next hop supports: • *E.164 Native* (default) • *E.164 with NSAP* • *DCC AESA* • *ICD AESA* • *E.164 AESA* • *Custom AESA* See "About NBMA Addressing" on page 14-4 for more information on these formats. Your choice of formats determines how you specify address information in the other fields on the dialog box: ***E.164 Native* and *E.164 with NSAP*** – You must type in a decimal address in the Decimal Digits field (which appears only if you make either of these choices). In addition, if you select *E.164 with NSAP* you must select an NBMA subaddress format and type in a decimal subaddress in the Decimal Digits field. ***DCC AESA*, *ICD AESA*, *E.164 AESA*, and *Custom AESA*** – You must type in a hexadecimal address in the Hex Digits field (which, along with the AFI field, appears only if you make any of these selections). With the exception of Customer AESA, the NMS specifies the AFI for you in the AFI field. If you select Custom AESA, you must type in the appropriate AFI in the AFI field yourself.
NBMA Subaddress Format (*E.164 with NSAP only*)	Select the NBMA subaddress format that the next hop supports: • *DCC AESA* • *ICD AESA* • *E.164 AESA* • *Custom AESA* For example, if the next hop is connected to a private network that supports DCC AESA addressing, you would select *DCC AESA*. After you select a format, type in a hexadecimal subaddress in the Hex Digits field. In addition, if you select *Custom AESA*, you must also type in an AFI in the AFI field. See "About NBMA Addressing" on page 14-4 for more information on these formats.
AFI (*DCC AESA, ICD AESA, E.164 AESA, and Custom AESA only*)	If you selected *Custom AESA* as the NBMA address or subaddress format, type in the two-digit AFI. If you made any other address or subaddress selection, this field is read-only and displays the AFI for the selected format.

Table 14-8. NHRP Server Cache Entry Parameters (Continued)

Field	Action/Description
Hex Digits (0-F) (*DCC AESA, ICD AESA, E.164 AESA, and Custom AESA only*)	Specify the hexadecimal address or subaddress of the next hop, up to 38 digits long. As you type an address, it displays in the NBMA Address field. As you type a subaddress, it displays in the NBMA Subaddress field.
Decimal Digits (0-9) (*E.164 Native and E.164 with NSAP only*)	Specify the decimal address of the next hop, up to 15 digits long. As you type an address, it displays in the NBMA Address field.
NBMA Address	Displays the NBMA address of the next hop as you type it in the Hex Digits or Decimal Digits field. The NBMA Address field is read-only.
NBMA Subaddress	Displays the NBMA subaddress of the next hop as you type it in the Hex Digits field. The NBMA Subaddress field is read-only.
Destination IP Address	Specify the destination's IP address (e.g., 131.100.24.3). The default is 0.0.0.0.
Next Hop IP Address	Specify the next hop's IP address (e.g., 131.100.24.3). The default is 0.0.0.0. The next hop IP address should match the destination IP address if the next hop and the destination are the same (that is, the next hop is the destination). This happens when the destination is directly connected to the server by a logical NBMA subnetwork. However, if the destination is not directly connected to the server, it means that an intermediate node must be used to reach the destination, and the next hop IP address must refer to the egress router for the NBMA subnetwork that is closest to the destination.
Entry Type	Select the cache entry type: *Static Volatile* (default) – The entry is volatile and will not be restored after a reset (e.g., a switch reboot). *Static Non Volatile* – The entry is non-volatile and will be restored after a reset (e.g., a switch reboot).
Network Mask	Specify the network mask for the destination IP address (for example, 255.255.255.0).
Admin Status	Set the administrative status of the entry to *Up* (default) or *Down*. If the status is Down, the entry cannot be used to resolve IP/ATM address mappings.

Table 14-8. NHRP Server Cache Entry Parameters (Continued)

Field	Action/Description
Uniqueness	Select the uniqueness value for the entry: *Unique* – (default) Mark this cache entry as unique. It is possible to have multiple NBMA addresses mapped to the same Next Hop IP Address. As a result, you will have multiple cache entries with the same Next Hop IP Address but different NBMA addresses. This allows the NHS to return multiple NBMA addresses to a requesting NHC. In turn, the NHC can establish a virtual circuit to alternate NBMA addresses if one establishment attempt fails. However, the NHC may not want multiple NBMA addresses returned to it. To tell the NHS that it only wants one NBMA address returned, the NHC sets the uniqueness flag in the NHRP Resolution Request. The NHS will then return only the NBMA address from the cache entry that is marked as unique. *Non Unique* – Mark this cache entry as non-unique. This indicates that the NHS will not return the NBMA address in the entry to a requesting NHC if the NHC specifies that it only wants the unique NBMA address.

7. Choose OK.

Modifying a Server Cache Entry

To modify a server cache entry:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Servers ⇒ Set Cache.
3. Select a switch. The list of servers configured on the switch appears.
4. Select the server whose cache entries you want to modify. A list of cache entries appears.
5. Select a cache entry.
6. Choose Modify. The Modify Static Cache Entry dialog box appears.
7. Modify the desired parameters. See Table 14-8 on page 14-42 for descriptions of these parameters.
8. Choose OK.

Deleting a Server Cache Entry

To delete a server cache entry:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Servers ⇒ Set Cache. The Set All NHRP Cache Entries dialog box appears (see Figure 14-10).
3. Select a switch. The list of servers configured on the switch appears.
4. Select the server whose cache entries you want to delete. A list of cache entries appears.
5. Select a cache entry.
6. Choose Delete.
7. Choose OK.

Configuring the Proxy Client

Only one proxy client can exist per switch, and the single proxy client instance is already added for you. However, you still need to:

- Configure the proxy client parameters.
- Enable the proxy client on an NHRP logical port, one that is used to communicate with NHCs and NHSs. See "Configuring NHRP Logical Port Parameters" on page 14-59 for more information on enabling the proxy client on an NHRP logical port.

Although proxy clients automatically cache the IP Address/NBMA Address mappings they receive from the NHS, you can also manually create cache entries for shortcuts to destinations, if necessary.

The following sections describe how to add, modify, and delete proxy clients and their cache entries.

Configuring Proxy Client Parameters

To configure proxy client parameters:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Proxy Client ⇒ Set Parameters. The Set All NHRP Client Parameters dialog box appears (see Figure 14-12).

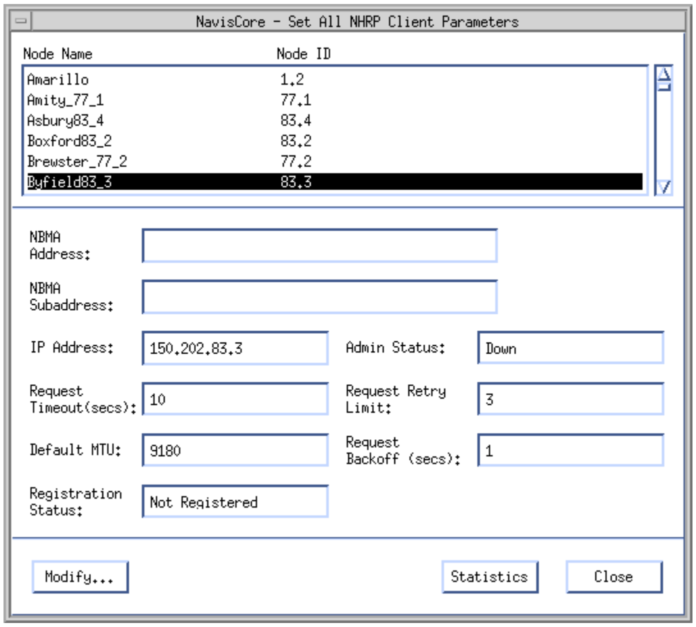

NavisCore - Set All NHRP Client Parameters

Node Name	Node ID
Amarillo	1.2
Amity_77_1	77.1
Asbury83_4	83.4
Boxford83_2	83.2
Brewster_77_2	77.2
Byfield83_3	83.3

NBMA Address:

NBMA Subaddress:

IP Address: 150.202.83.3

Admin Status: Down

Request Timeout(secs): 10

Request Retry Limit: 3

Default MTU: 9180

Request Backoff (secs): 1

Registration Status: Not Registered

Figure 14-12. Set All NHRP Client Parameters Dialog Box

The Set All NHRP Client Parameters dialog box allows you to select a switch in the list box at the top of the dialog box. The proxy client parameters configured for the proxy client instance on the selected switch appear in the bottom half of the dialog box. These parameters are described in Table 14-9 on page 14-49, except for the Registration Status field. This field displays the registration status of the proxy client. The registration status is always “Not Registered,” since the proxy client does not register with an NHS.

> You can display statistics on proxy client activity by choosing the Statistics button. See the *NavisCore Diagnostics Guide* for descriptions of these statistics.

3. Select a switch.
4. Choose Modify. The Set NHRP Client Parameters dialog box appears (see Figure 14-13).

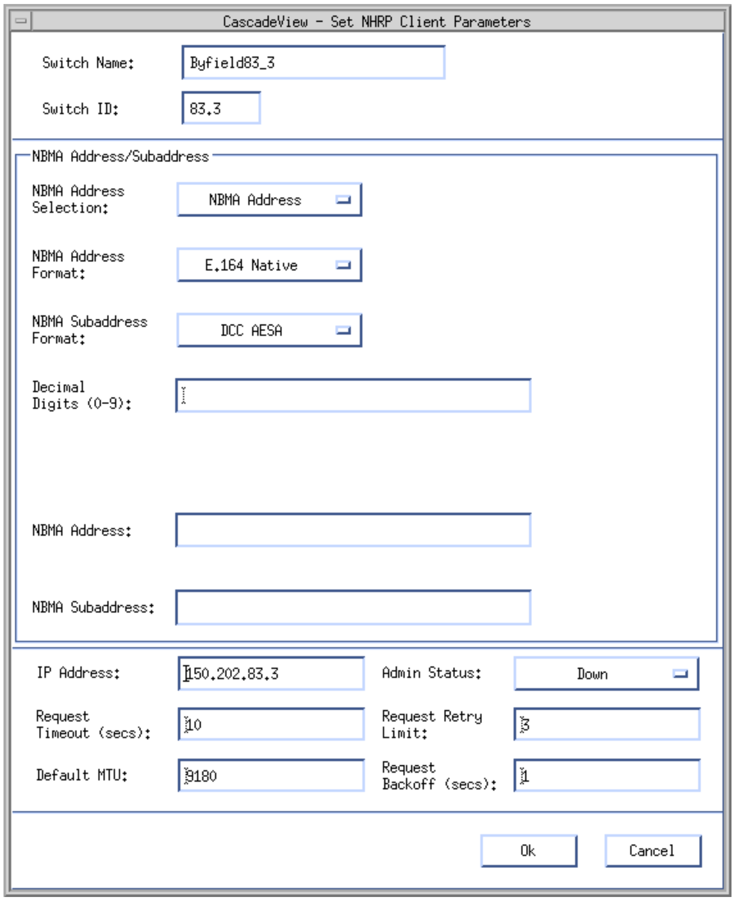

CascadeView - Set NHRP Client Parameters

Switch Name: Byfield83_3

Switch ID: 83.3

NBMA Address/Subaddress

NBMA Address Selection: NBMA Address

NBMA Address Format: E.164 Native

NBMA Subaddress Format: DCC AESA

Decimal Digits (0-9):

NBMA Address:

NBMA Subaddress:

IP Address: 150.202.83.3

Admin Status: Down

Request Timeout (secs): 10

Request Retry Limit: 3

Default MTU: 9180

Request Backoff (secs): 1

Figure 14-13. Set NHRP Client Parameters Dialog Box

5. Specify the parameters described in Table 14-9.
6. Choose OK.

Table 14-9. NHRP Client Parameters

Field	Action/Description
Switch Name	Displays the name of the selected switch.
Switch ID	Displays the subnetwork number and host ID in the internal IP address of the switch.
NBMA Address Selection	Select the type of proxy client NBMA address you are configuring: *NBMA Address* (default) or *NBMA Subaddress*. You select *NBMA Address* regardless of the address format you use. You select *NBMA Subaddress* only if you use *E.164 with NSAP* addressing. If you use E.164 with NSAP addresses, you specify both an address and a subaddress for the proxy client. Specify the address and subaddress in the following order: 1. In the NBMA Address Selection field, select *NBMA Address*. 2. In the NBMA Address Format field, select *E.164 with NSAP.* 3. In the Decimal Digits field, specify a native E.164 address. 4. Return to the NBMA Address Selection field. Select *NBMA Subaddress*. 5. In the NBMA Address Selection field, select *NBMA Subaddress*. 6. In the NBMA Address Format field, select a valid address format. 7. Specify an AFI if you selected a subaddress format of *Custom AESA*. 8. In the Hex Digits field, specify the hexadecimal subaddress.

Table 14-9. NHRP Client Parameters (Continued)

Field	Action/Description
NBMA Address Format	Select one of the following formats for the proxy client's NBMA address: • *E.164 Native* (default) • *E.164 with NSAP* • *DCC AESA* • *ICD AESA* • *E.164 AESA* • *Custom AESA* See "About NBMA Addressing" on page 14-4 for more information on these formats. Your choice of formats determines how you specify address information in the other fields on the dialog box: ***E.164 Native* and *E.164 with NSAP*** – You must type in a decimal address in the Decimal Digits field (which appears only if you make either of these choices). In addition, if you select *E.164 with NSAP* you must select an NBMA subaddress format and type in a decimal subaddress in the Decimal Digits field. ***DCC AESA*, *ICD AESA*, *E.164 AESA*, and *Custom AESA*** – You must type in a hexadecimal address in the Hex Digits field (which, along with the AFI field, appears only if you make any of these selections). With the exception of Customer AESA, the NMS specifies the AFI for you in the AFI field. If you select Custom AESA, you must type in the appropriate AFI in the AFI field yourself.
NBMA Subaddress Format (*E.164 with NSAP only*)	Select one of the following formats for the proxy client's NBMA subaddress: • *DCC AESA* • *ICD AESA* • *E.164 AESA* • *Custom AESA* For example, if the switch is connected to a private network that supports DCC AESA addressing, you would select *DCC AESA*. After you select a format, type in a hexadecimal subaddress in the Hex Digits field. In addition, if you select *Custom AESA*, you must also type in an AFI in the AFI field. See "About NBMA Addressing" on page 14-4 for more information on these formats.
AFI (*DCC AESA, ICD AESA, E.164 AESA, and Custom AESA only*)	If you selected *Custom AESA* as the NBMA address or subaddress format, type in the two-digit AFI. If you made any other address or subaddress selection, this field is read-only and displays the AFI for the selected format.

Table 14-9. NHRP Client Parameters (Continued)

Field	Action/Description
Hex Digits (0-F) (*DCC AESA, ICD AESA, E.164 AESA, and Custom AESA only*)	Specify the hexadecimal address or subaddress of the proxy client, up to 38 digits long. As you type an address, it displays in the NBMA Address field. As you type a subaddress, it displays in the NBMA Subaddress field.
Decimal Digits (0-9) (*E.164 Native and E.164 with NSAP only*)	Specify the decimal address of the proxy client, up to 15 digits long. As you type an address, it displays in the NBMA Address field.
NBMA Address	Displays the NBMA address as you type it in the Hex Digits or Decimal Digits field. The NBMA Address field is read-only.
NBMA Subaddress	Displays the NBMA subaddress as you type it in the Hex Digits field. The NBMA Subaddress field is read-only.
IP Address	Specify the proxy client's IP address (e.g., 131.100.24.3). The default is the IP address of the switch.
Request Timeout (secs)	Specify the number of seconds that the proxy client waits before timing out an NHRP Resolution Request to the NHS. The default is 10. When NHRP Resolution Requests time out, the proxy client re-sends them according to the rules defined by the Request Retry Limit and Request Backoff parameters. See the descriptions of these parameters for more information. The proxy client also re-sends NHRP Resolution Requests according to the rules defined by the Request Retry Limit and Request Backoff parameters if the proxy client receives a negative acknowledgment from the NHS because no IP/ATM address mapping was found.
Default MTU	Specify the default Maximum Transmission Unit (MTU), in bytes, that the client uses to send packets. The default is 9180 bytes. Make sure that the MTU you set is compatible with the rest of the network. If you do not set any default MTU value (that is, you remove the default and leave the field blank), the switch sets the default MTU to the value used by the LIS/LAG.
Admin Status	Set the administrative status of the proxy client to *Up* or *Down* (default).
Request Retry Limit	Specify the number of times that the proxy client will retry NHRP Resolution Requests to the NHS before giving up. Values range from 0 to 65535. The default is 3. A value of 0 specifies that the proxy client will not attempt any retries. A value of 65535 specifies that the client will retry forever. This parameter works in conjunction with the Request Backoff parameter. For example, if you specify a Request Retry Limit of 3 and a Request Backoff of 12 seconds, the proxy client will make three retry attempts, and will wait 12 seconds between each attempt.

Table 14-9. NHRP Client Parameters (Continued)

Field	Action/Description
Request Backoff (secs)	Specify the number of seconds that the proxy client waits before attempting a retry of an NHRP Resolution Request to the NHS. The default is 1. For example, if you use the default, the client will wait 1 second before attempting each NHRP Resolution Request retry. This parameter works in conjunction with the Request Retry Limit parameter. See the description of the Request Retry Limit parameter for more information.

Modifying Proxy Client Parameters

You may want to change proxy client parameters at a later date. To modify proxy client parameters, follow the instructions in the previous section (see "Configuring Proxy Client Parameters" on page 14-47).

Adding Proxy Client Cache Entries

A proxy client cache entry maps an IP address of a next hop to its NBMA (i.e., ATM) address. A shortcut can then be established to the destination that is associated with the cache entry.

A proxy client cache entry identifies both the *IP address of the destination* and the *IP address of the next hop* that is used to reach the destination. In many cases, these addresses are one and the same — they refer to the same network node (that is, the next hop *is* the destination). However, in other cases, these addresses are different.

If the proxy client and the destination are directly connected by the same logical NBMA subnetwork, then the next hop IP address and destination IP address are the same (that is, the next hop is the destination). Otherwise, the next hop IP address refers to the egress router for the NBMA subnetwork that is closest to the destination, and the destination IP address refers to the destination itself. The NBMA address is mapped to the next hop IP address, not the destination IP address.

Cache entries can be added to the server cache in two ways:

Dynamically — The proxy client caches the IP/ATM address resolution mappings received in NHRP Resolution Requests from NHSs. However, these cache entries expire after a certain period of time, or are removed as a result of receiving NHRP Purge Request.

Manually — You add the IP/ATM address mappings yourself. You may want to do this if you do not want your cache entries to expire, or if no NHSs are available to resolve IP/ATM address mappings.

To add an entry to the proxy client cache:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Proxy Client ⇒ Set Cache. The Set All NHRP Cache Entries dialog box appears (see Figure 14-14).

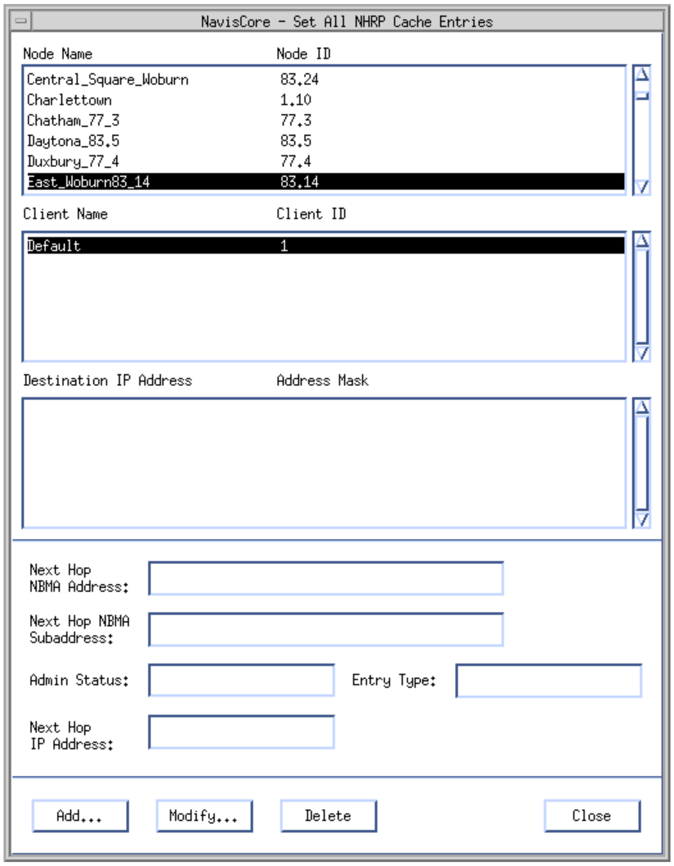

NavisCore - Set All NHRP Cache Entries

Node Name	Node ID
Central_Square_Woburn	83.24
Charlettown	1.10
Chatham_77_3	77.3
Daytona_83.5	83.5
Duxbury_77_4	77.4
East_Woburn83_14	83.14

Client Name	Client ID
Default	1

Destination IP Address	Address Mask

Figure 14-14. Set All NHRP Cache Entries Dialog Box (Proxy Client)

The Set All NHRP Cache Entries dialog box allows you to select a switch in the list box at the top of the dialog box. This action displays the default proxy client ("Default") on that switch in the second list box. You can then select the default proxy client, and the cache entries for the proxy client display in a list. Table 14-10 on page 14-55 describes these fields.

3. Select a switch. The default proxy client ("Default") on that switch appears.
4. Select the default proxy client.

5. Choose Add. The Add Static Cache Entry dialog box appears (see Figure 14-15).

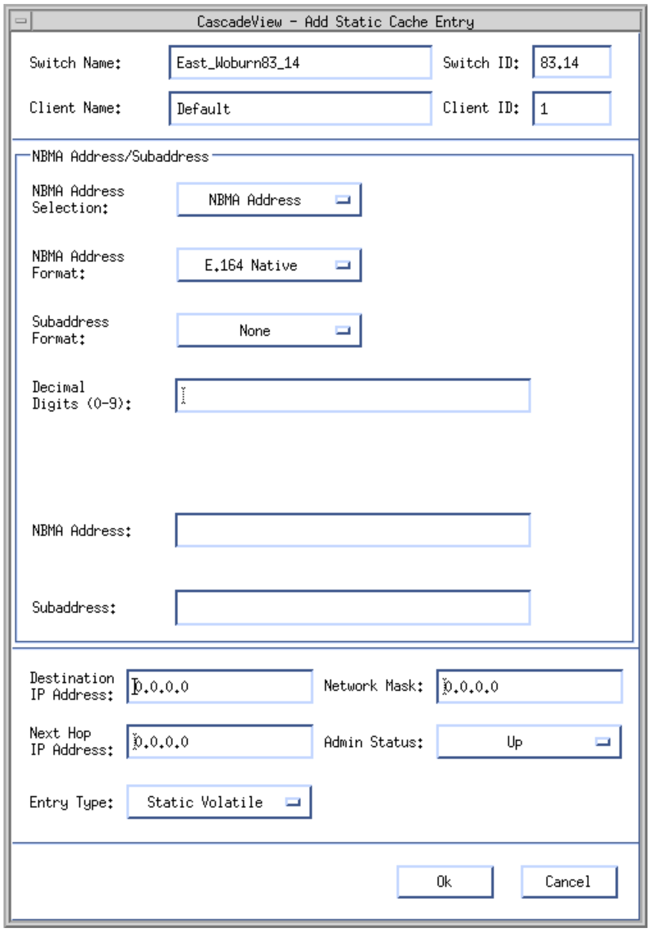

Switch Name: East_Woburn83_14
Switch ID: 83.14
Client Name: Default
Client ID: 1

NBMA Address/Subaddress

NBMA Address Selection: NBMA Address
NBMA Address Format: E.164 Native
Subaddress Format: None
Decimal Digits (0-9):
NBMA Address:
Subaddress:

Destination IP Address: 0.0.0.0
Network Mask: 0.0.0.0
Next Hop IP Address: 0.0.0.0
Admin Status: Up
Entry Type: Static Volatile

Figure 14-15. Add Static Cache Entry Dialog Box (Proxy Client)

6. Specify the parameters described in Table 14-10.

7. Choose OK.

Table 14-10. NHRP Proxy Client Cache Entry Parameters

Field	Action/Description
Switch Name	Displays the name of the selected switch.
Switch ID	Displays the subnetwork number and host ID in the internal IP address of the switch.
Client Name	Displays "Default."
Client ID	Displays the internally assigned ID of the client.
NBMA Address Selection	Select the type of address you are configuring for the next hop: *NBMA Address* (default) or *NBMA Subaddress*. You select *NBMA Address* regardless of the address format you use. You select *NBMA Subaddress* only if you use *E.164 with NSAP* addressing. If you use E.164 with NSAP addresses, you specify both an address and a subaddress in the following order: 1. In the NBMA Address Selection field, select *NBMA Address*. 2. In the NBMA Address Format field, select *E.164 with NSAP.* 3. In the Decimal Digits field, specify a native E.164 address. 4. Return to the NBMA Address Selection field. Select *NBMA Subaddress*. 5. In the NBMA Address Selection field, select *NBMA Subaddress*. 6. In the NBMA Address Format field, select a valid address format. 7. Specify an AFI if you selected a subaddress format of *Custom AESA.* 8. In the Hex Digits field, specify the hexadecimal subaddress.

Table 14-10. NHRP Proxy Client Cache Entry Parameters (Continued)

Field	Action/Description
NBMA Address Format	Select the NBMA address format that the next hop supports: • *E.164 Native* (default) • *E.164 with NSAP* • *DCC AESA* • *ICD AESA* • *E.164 AESA* • *Custom AESA* See "About NBMA Addressing" on page 14-4 for more information on these formats. Your choice of formats determines how you specify address information in the other fields on the dialog box: ***E.164 Native* and *E.164 with NSAP*** – You must type in a decimal address in the Decimal Digits field (which appears only if you make either of these choices). In addition, if you select *E.164 with NSAP* you must select an NBMA subaddress format and type in a decimal subaddress in the Decimal Digits field. ***DCC AESA*, *ICD AESA*, *E.164 AESA*, and *Custom AESA*** – You must type in a hexadecimal address in the Hex Digits field (which, along with the AFI field, appears only if you make any of these selections). With the exception of Customer AESA, the NMS specifies the AFI for you in the AFI field. If you select Custom AESA, you must type in the appropriate AFI in the AFI field yourself.
NBMA Subaddress Format (*E.164 with NSAP only*)	Select the NBMA subaddress format that the next hop supports: • *DCC AESA* • *ICD AESA* • *E.164 AESA* • *Custom AESA* For example, if the next hop is connected to a private network that supports DCC AESA addressing, you would select *DCC AESA*. After you select a format, type in a hexadecimal subaddress in the Hex Digits field. In addition, if you select *Custom AESA*, you must also type in an AFI in the AFI field. See "About NBMA Addressing" on page 14-4 for more information on these formats.
AFI (*DCC AESA, ICD AESA, E.164 AESA, and Custom AESA only*)	If you selected *Custom AESA* as the NBMA address or subaddress format, type in the two-digit AFI. If you made any other address or subaddress selection, this field is read-only and displays the AFI for the selected format.

Table 14-10. NHRP Proxy Client Cache Entry Parameters (Continued)

Field	Action/Description
Hex Digits (0-F) (*DCC AESA, ICD AESA, E.164 AESA, and Custom AESA only*)	Specify the hexadecimal address or subaddress of the next hop, up to 38 digits long. As you type an address, it displays in the NBMA Address field. As you type a subaddress, it displays in the NBMA Subaddress field.
Decimal Digits (0-9) (*E.164 Native and E.164 with NSAP only*)	Specify the decimal address of the next hop, up to 15 digits long. As you type an address, it displays in the NBMA Address field.
NBMA Address	Displays the NBMA address of the next hop as you type it in the Hex Digits or Decimal Digits field. The NBMA Address field is read-only.
NBMA Subaddress	Displays the NBMA subaddress of the next hop as you type it in the Hex Digits field. The NBMA Subaddress field is read-only.
Destination IP Address	Specify the destination's IP address (e.g., 131.100.24.3). The default is 0.0.0.0.
Next Hop IP Address	Specify the next hop's IP address (e.g., 131.100.24.3). The default is 0.0.0.0. The next hop IP address should match the destination IP address if the next hop and the destination are the same (that is, the next hop is the destination). This happens when the destination and the proxy client are directly connected by a logical NBMA subnetwork. However, if the destination and proxy client are not directly connected, an intermediate node must be used to reach the destination, and the next hop IP address must refer to the egress router for the NBMA subnetwork that is closest to the destination.
Entry Type	Select the cache entry type: *Static Volatile* (default) – The entry is volatile and will not be restored after a reset (e.g., a switch reboot). *Static Non Volatile* – The entry is non-volatile and will be restored after a reset (e.g., a switch reboot).
Network Mask	Specify the network mask for the destination IP address (for example, 255.255.255.0).
Admin Status	Set the administrative status of the entry to *Up* (default) or *Down*. If the status is Down, the entry cannot be used to resolve IP/ATM address mappings.

Modifying a Proxy Client Cache Entry

To modify a proxy client cache entry:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Proxy Client ⇒ Set Cache.
3. Select a switch. The name of the proxy client (“Default”) for the switch and associated information appears.
4. Select the default proxy client (“Default”). A list of cache entries appears.
5. Select a cache entry.
6. Choose Modify. The Modify Static Cache Entry dialog box appears.
7. Modify the desired parameters. See Table 14-10 on page 14-55 for descriptions of these parameters.
8. Choose OK.

Deleting a Proxy Client Cache Entry

To delete a proxy client cache entry:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Proxy Client ⇒ Set Cache. The Set All NHRP Cache Entries dialog box appears (see Figure 14-10).
3. Select a switch. The name of the proxy client (“Default”) for the switch and associated information appears.
4. Select the default proxy client (“Default”). A list of cache entries appears.
5. Select a cache entry.
6. Choose Delete.
7. Choose OK.

Configuring NHRP Logical Port Parameters

You must configure NHRP logical ports that will process NHRP requests and responses.

When you configure an NHRP logical port:

- Verify that the NHRP logical port has been added. See "Adding an NHRP Logical Port" on page 14-10 for more information on adding an NHRP logical port.
- Verify the role that the port plays in the network (e.g., ingress and/or egress). If the port is at the ingress/egress of the Lucent network and interfaces with NHCs, associate a server with it. In addition, if you want to provision bandwidth and QoS guarantees for IP traffic flows, enable the proxy client. For trunk logical ports, the default server is already configured and you are not required to do anything.
- Make sure that you have configured NHRP servers and the proxy client. See "Configuring Servers" on page 14-32 and "Configuring the Proxy Client" on page 14-47 for more information.
- Configure other parameters as needed.

To configure logical port parameters:

1. Select a switch from the network map.

2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set LPort Parameters. The Set All NHRP LPort Parameters dialog box appears (see Figure 14-16).

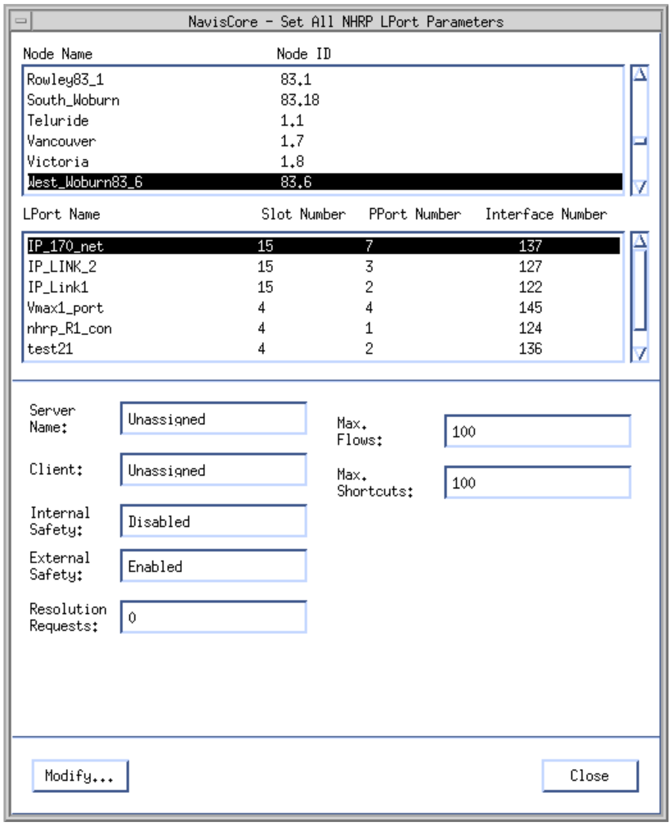

NavisCore - Set All NHRP LPort Parameters

Node Name	Node ID
Rowley83_1	83.1
South_Woburn	83.18
Teluride	1.1
Vancouver	1.7
Victoria	1.8
West_Woburn83_6	83.6

LPort Name	Slot Number	PPort Number	Interface Number
IP_170_net	15	7	137
IP_LINK_2	15	3	127
IP_Link1	15	2	122
Vmax1_port	4	4	145
nhrp_R1_con	4	1	124
test21	4	2	136

Server Name: Unassigned

Client: Unassigned

Internal Safety: Disabled

External Safety: Enabled

Resolution Requests: 0

Max. Flows: 100

Max. Shortcuts: 100

Figure 14-16. Set All NHRP LPort Parameters Dialog Box (Display Mode)

The Set All NHRP LPort Parameters dialog box displays a list of switches in the network. When you select a switch, a list of logical ports on that switch appears. In turn, when you select a logical port, the parameters configured for the logical port appears at the bottom of the dialog box. Table 14-11 provides descriptions of these parameters, except for the Resolution Requests field, which applies only if you associated the proxy client with the logical port. The Resolution Requests field displays the number of NHRP Resolution Requests generated by the logical port.

3. Select the switch that has the NHRP logical port(s) you want to configure. The logical ports on the switch display in the dialog box.
4. Select the NHRP logical port you want to configure.
5. Choose Modify. The Set NHRP LPort Parameters dialog box appears (see Figure 14-17).

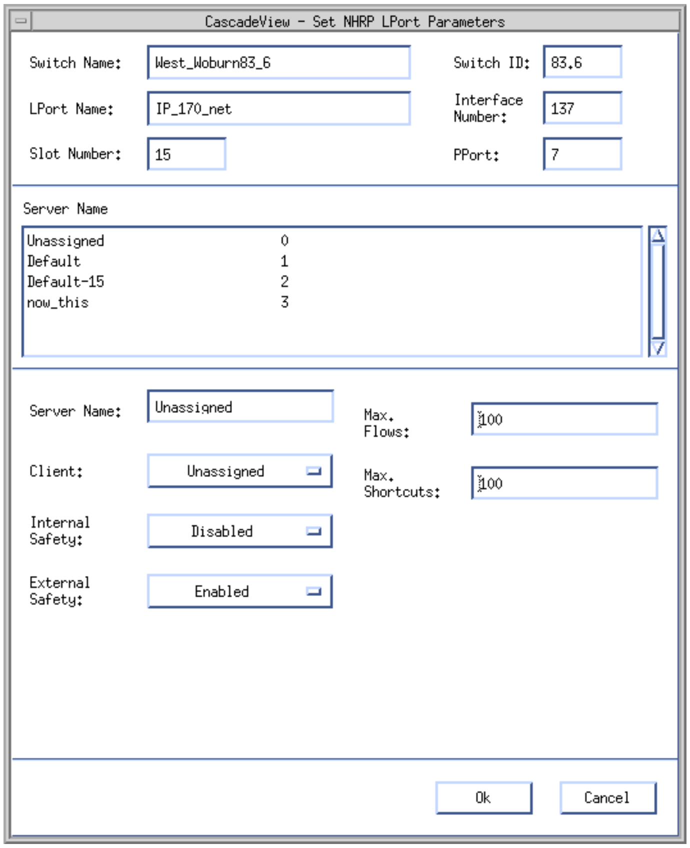

Switch Name: West_Woburn83_6

Switch ID: 83.6

LPort Name: IP_170_net

Interface Number: 137

Slot Number: 15

PPort: 7

Server Name

Unassigned	0
Default	1
Default-15	2
now_this	3

Server Name: Unassigned

Max. Flows: 100

Client: Unassigned

Max. Shortcuts: 100

Internal Safety: Disabled

External Safety: Enabled

Figure 14-17. Set NHRP LPort Parameters Dialog Box

6. Specify the parameters described in Table 14-11.
7. Choose OK.

Table 14-11. NHRP Logical Port Parameters

Field	Action/Description
Switch Name	Displays the name of the selected switch.
LPort Name	Displays the name of the selected NHRP logical port.
Slot Number	Displays the number of the slot in which the I/O module associated with the NHRP logical port is installed.
Switch ID	Displays the subnetwork number and host ID in the internal IP address of the switch.
Interface Number	Displays the internally assigned interface number of the NHRP logical port.
PPort	Displays the physical port associated with the NHRP logical port.

Table 14-11. NHRP Logical Port Parameters (Continued)

Field	Action/Description
Server Name	Specify the NHS associated with this NHRP logical port. To do this, select a server from the list of servers configured for the switch. The name of the server displays in this field. Note that you may select the default NHS for association with non-trunk logical ports (that is, ingress/egress logical ports).
Client	Specify whether the proxy client is enabled on this NHRP logical port. Select *Default* to enable the proxy client. Otherwise, keep the default setting (*Unassigned*). Enabling the proxy client on the NHRP logical port means that the port can generate NHRP Resolution Requests, and thus can detect flow.
Internal Safety	Select the Internal Safety setting: *Enabled* or *Disabled* (the default). If you associate an NHS with this NHRP logical port, and the logical port is an egress port, consider enabling this setting to reduce the risk of creating persistent routing loops. These loops can be created when a proxy client in the Lucent network sends an NHRP Resolution Request and the egress NHS cannot resolve it. By enabling the Internal Safety setting, the egress NHS terminates the request, replies with its own NBMA address, and terminates the shortcut within the Lucent network. See "Preventing Persistent Routing Loops" on page 13-34 for more information on persistent routing loops.
External Safety	Select the External Safety setting: *Enabled* (the default) or *Disabled*. If you associate an NHS with this NHRP logical port, and the logical port is an egress port, consider keeping the default (*Enabled*) to reduce the risk of creating persistent routing loops. These loops can be created when the following scenario takes place: **1.** An NHC outside the Lucent network sends an NHRP Resolution Request. **2.** The destination of the request is off the NBMA, but is not directly connected to the egress NHS. If the External Safety setting is enabled, the egress NHS returns an error in the NHRP Resolution Reply to the external NHC, and no SVC shortcut is established, thereby eliminating the possibility of a persistent routing loop. See "Preventing Persistent Routing Loops" on page 13-34 for more information on persistent routing loops.

Table 14-11. NHRP Logical Port Parameters (Continued)

Field	Action/Description
Max. Flows	Specify the maximum number of matching NHRP traffic flows that can be simultaneously detected on this NHRP logical port. You may not specify a value greater than 512. A matching NHRP traffic flow is one that matches the criteria defined in any flow profile associated with this logical port. See "Configuring Flow Profiles" on page 14-19 and "Configuring NHRP Logical Port FP/TD Associations" on page 14-64 for more information on creating flow profiles and associating them with logical ports. For example, if you use the default value (100), then 100 matching NHRP traffic flows can be detected on this logical port at any one time. There are limits on the maximum number of flows per forwarding engine (that is, the maximum number of flows from all the NHRP logical ports on the forwarding engine that the engine can simultaneously track). See your switch Software Release Notice (SRN) for more information. ***Note**: Each flow profile has its own maximum flows value (see Table 14-6 on page 14-23). Because multiple flow profiles can be associated with an NHRP logical port, the sum of the maximum flows values of all the flow profiles associated with the logical port cannot exceed this Max. Flows value. For example, if you associate two flow profiles with this logical port, and you set the Max. Flows value for the logical port to 50, then the sum of the maximum flows values of the two profiles should not exceed 50.*
Max. Shortcuts	Specify the maximum number of matching NHRP shortcuts that can be simultaneously connected to this logical port. You may not specify a value greater than 512. A matching NHRP shortcut is one that matches the criteria defined in any flow profile associated with this logical port. See "Configuring Flow Profiles" on page 14-19 and "Configuring NHRP Logical Port FP/TD Associations" on page 14-64 for more information on creating flow profiles and associating them with logical ports. For example, if you use the default (100), then 100 shortcuts that match any flow profile associated with this logical port can be connected at any one time. There are limits on the maximum number of shortcuts per forwarding engine (that is, the maximum number of shortcuts from all the logical ports on the forwarding engine that the engine can simultaneously track). See your switch Software Release Notice (SRN) for more information. ***Note**: Each flow profile has its own maximum shortcuts value (see Table 14-6 on page 14-23). Because multiple flow profiles can be associated with an NHRP logical port, the sum of the maximum shortcuts values of all the flow profiles associated with the logical port cannot exceed this Max. Shortcuts value. For example, if you associate two flow profiles with this logical port, and you set the Max. Shortcuts value for the logical port to 50, then the sum of the maximum shortcuts values of the two profiles should not exceed 50.*

Configuring NHRP Logical Port FP/TD Associations

After you create flow profiles, you must associate them with NHRP logical ports to put them into effect.

▶ NavisCore does not allow you to delete an NHRP logical port if it has IP flow profiles associated with it. To delete the NHRP logical port, you must first delete all of the IP flow profile associations. See "Deleting an NHRP Logical Port" on page 14-11 for more information.

When you assign a flow profile to an NHRP logical port, you also specify the traffic descriptors that manage the traffic flow over the SVC that connects the NHRP logical port and the destination. You specify two sets of traffic descriptors:

Primary Traffic Descriptors — The desired traffic descriptors for the traffic flow over the SVC.

Secondary Traffic Descriptors — The alternative or minimum acceptable traffic descriptors associated with the flow profile. When the traffic flow's SVC is set up, both ends negotiate to determine the traffic descriptors to be used for the connection. If not enough resources exist for the network to meet the primary traffic descriptor requirements, then either the alternate traffic descriptors or the minimum acceptable traffic descriptors are used. *Alternate traffic descriptors* define the best possible values that should be used in place of the primary traffic descriptor values. *Minimum acceptable traffic descriptors* define the lowest values that you are willing to accept in place of the primary traffic descriptor values. If the minimum acceptable traffic descriptors are configured, the actual traffic descriptors that are used for the SVC are a negotiated compromise somewhere between the primary and minimum traffic descriptors.

When you associate traffic descriptors with flow profiles, you can specify alternate traffic descriptors or minimum acceptable traffic descriptors — but not both. For more information on alternate and minimum traffic descriptors, see the *ATM User-Network Interface (UNI) Signalling Specification Version 4.0.*

Before You Begin

Before you begin to associate NHRP logical ports, flow profiles, and traffic descriptors, see the *NavisCore ATM Configuration Guide* for traffic descriptor information. Make sure that you use the *NavisCore ATM Configuration Guide* in conjunction with this guide when you associate traffic descriptors with flow profiles.

You should also be aware of the following rules:

- You can associate a single flow profile with multiple NHRP logical ports.
- You can associate multiple flow profiles with a single NHRP logical port, but you cannot associate the same flow profile more than once with the same NHRP logical port.
- If you associate multiple flow profiles with a single NHRP logical port, and these profiles have matching or overlapping parameters, the NHS chooses the most specific profile first, then the second-most specific profile, and so on. For example, the NHS chooses flow profiles with specific source and destination IP addresses before it chooses flow profiles that have 0.0.0.0 configured for the source and destination IP addresses. As another example, if one profile specifies TCP protocol traffic while the other profile specifies the wildcard for protocol traffic, the profile that specifies TCP protocol traffic is chosen to manage TCP traffic flow before the profile that specifies the wildcard.
- Only one of the following parameters may have an ambiguous value in a flow profile: Dest. IP Address, Source Application, or Dest. Application. For Dest. IP Address, an ambiguous value is expressed as 0.0.0.0. For Source Application and Dest. Application, an ambiguous value is expressed by selecting *Ambiguous* from the pull-down menu provided. "Configuring Flow Profiles" on page 14-19 for more information.
- You can associate up to 96 flow profiles with a single NHRP logical port. See your switch software release notice for more information on this NHRP limit as well as other NHRP limits.

Associating a Flow Profile

To associate a flow profile with an NHRP logical port:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set LPort FP/TD Associations. The Set All NHRP LPort Flow Profiles dialog box appears (see Figure 14-18).

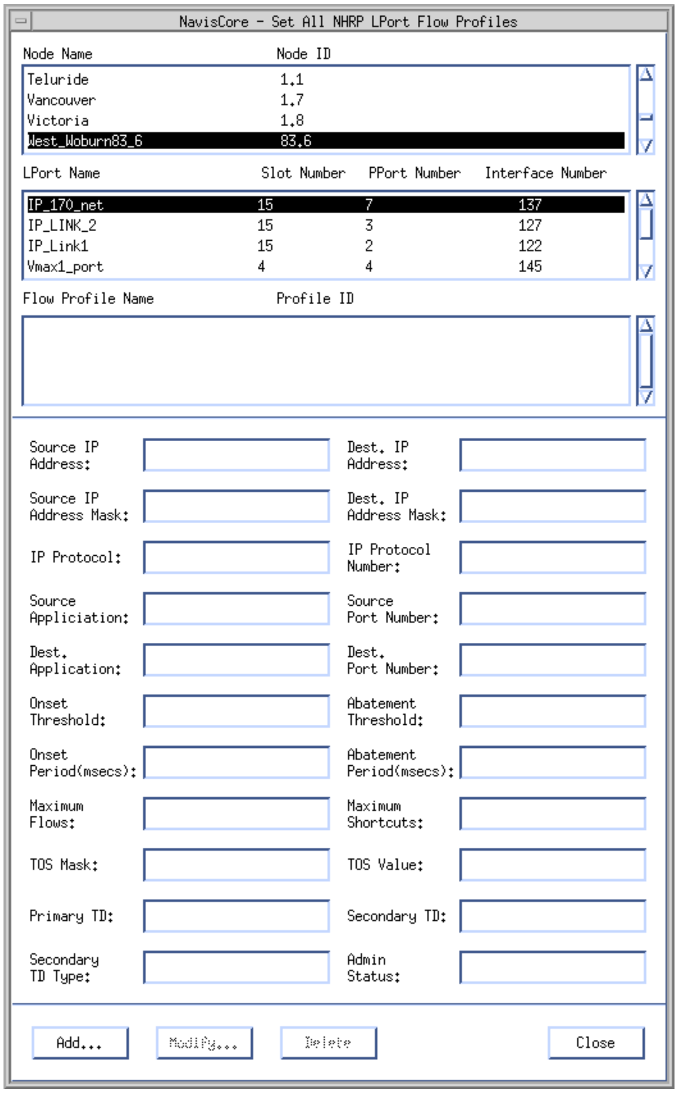

NavisCore - Set All NHRP LPort Flow Profiles

Node Name	Node ID
Telluride	1.1
Vancouver	1.7
Victoria	1.8
West_Woburn83_6	83.6

LPort Name	Slot Number	PPort Number	Interface Number
IP_170_net	15	7	137
IP_LINK_2	15	3	127
IP_Link1	15	2	122
Vmax1_port	4	4	145

Flow Profile Name	Profile ID

Figure 14-18. Set All NHRP LPort Flow Profiles Dialog Box

▶ The order in which you associate IP flow profiles with NHRP logical ports is very important. Switches allocate resources to IP flow profiles based on the order in which they are associated with NHRP logical ports. The IP flow profile that is associated first has the highest resource priority, the IP flow profile associated second has the next highest resource priority, and so on. Associate IP flow profiles in the following order: from most important to least important.

This dialog box displays a list of switches in the top list box. You can select a switch to display all of its NHRP logical ports, and then select an NHRP logical port to display its associated flow profiles. You can then select a flow profile, and its parameters appear in the fields in the bottom half of the dialog box. These fields are described in Table 14-6 on page 14-23, except for the following fields:

Primary TD — The ID of the primary traffic descriptors associated with the flow profile.

Secondary TD Type — The type of secondary traffic descriptors: *Alternate* or *Minimum*.

Secondary TD — The ID of the secondary traffic descriptors associated with the flow profile.

Admin Status — The administrative status of the flow profile association: *Up* or *Down*. Setting the administrative status to *Down* disables the flow profile for use on the logical port, but does not delete the association.

3. Select the switch that has the logical port with which you want to associate the flow profile. A list of the logical ports on the switch appears.
4. Select the logical port.

5. Choose Add. The Set NHRP LPort Flow Profile/TD Associations dialog box appears (see Figure 14-19).

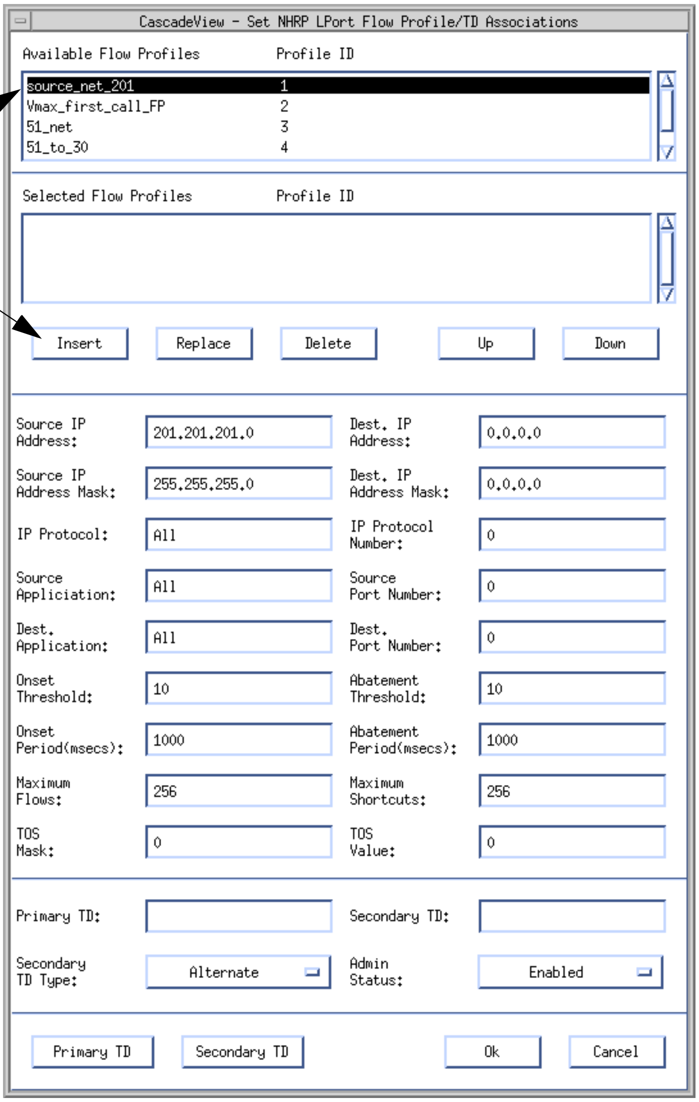

To associate the first profile:

1. Select the profile.

2. Choose Insert.

CascadeView - Set NHRP LPort Flow Profile/TD Associations

Available Flow Profiles	Profile ID
source_net_201	1
Vmax_first_call_FP	2
51_net	3
51_to_30	4

Selected Flow Profiles	Profile ID

Source IP Address:	201.201.201.0	Dest. IP Address:	0.0.0.0
Source IP Address Mask:	255.255.255.0	Dest. IP Address Mask:	0.0.0.0
IP Protocol:	All	IP Protocol Number:	0
Source Appliciation:	All	Source Port Number:	0
Dest. Application:	All	Dest. Port Number:	0
Onset Threshold:	10	Abatement Threshold:	10
Onset Period(msecs):	1000	Abatement Period(msecs):	1000
Maximum Flows:	256	Maximum Shortcuts:	256
TOS Mask:	0	TOS Value:	0
Primary TD:		Secondary TD:	
Secondary TD Type:	Alternate	Admin Status:	Enabled

Figure 14-19. Set NHRP LPort Flow Profile/TD Associations Dialog Box (No Profiles Added)

The Set NHRP LPort Flow Profile/TD Associations dialog box displays two flow profile lists: a list of available profiles and a list of selected profiles (that is, profiles that have been associated with the logical port).

6. Select a profile from the list of available profiles. The parameters for the selected profile display in the fields on the dialog box. These fields are described in Table 14-6 on page 14-23.

7. Skip this step if you are adding the first profile association. Select an insertion point in the list. To do this, select the profile association at the point in the list where you want the new profile association to appear. The new profile association appears after the selected position.

8. Choose Insert. NavisCore inserts the selected profile into the list of selected profiles (see Figure 14-20).

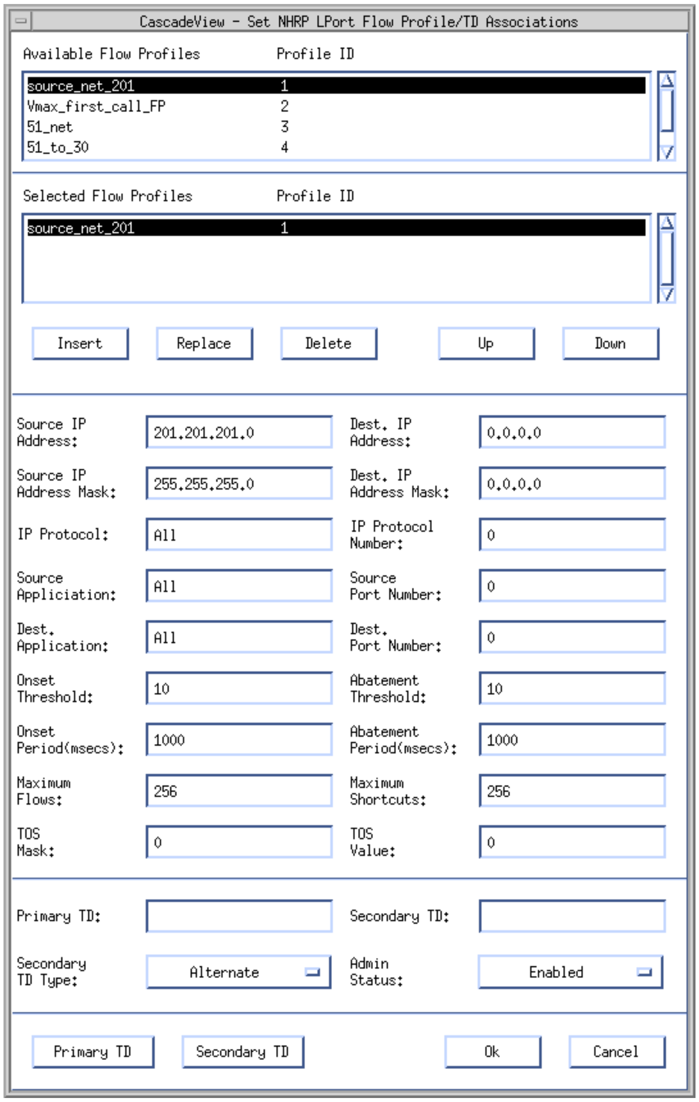

CascadeView - Set NHRP LPort Flow Profile/TD Associations

Available Flow Profiles	Profile ID
source_net_201	1
Vmax_first_call_FP	2
51_net	3
51_to_30	4

Selected Flow Profiles	Profile ID
source_net_201	1

Source IP Address:	201.201.201.0	Dest. IP Address:	0.0.0.0
Source IP Address Mask:	255.255.255.0	Dest. IP Address Mask:	0.0.0.0
IP Protocol:	All	IP Protocol Number:	0
Source Appliciation:	All	Source Port Number:	0
Dest. Application:	All	Dest. Port Number:	0
Onset Threshold:	10	Abatement Threshold:	10
Onset Period(msecs):	1000	Abatement Period(msecs):	1000
Maximum Flows:	256	Maximum Shortcuts:	256
TOS Mask:	0	TOS Value:	0
Primary TD:		Secondary TD:	
Secondary TD Type:	Alternate	Admin Status:	Enabled

Figure 14-20. Set NHRP Flow Profile/TD Associations Dialog Box (One Profile)

At this point, you are ready to configure traffic descriptors for the profile.

9. Select the associated flow profile from the list of selected profiles.

10. Choose Primary TD. The Set All ATM Traffic Descriptors dialog box appears (see Figure 14-21). See the *NavisCore ATM Configuration Guide* for a description of the fields on this dialog box.

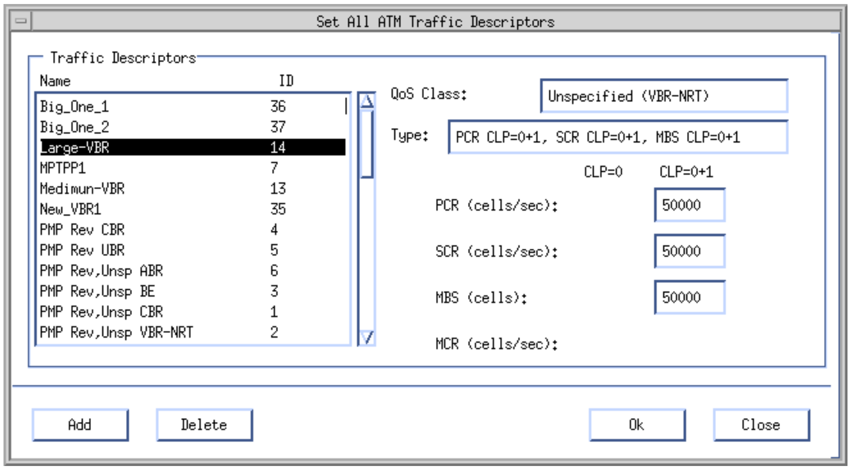

Name	ID
Big_One_1	36
Big_One_2	37
Large-VBR	14
MPTPP1	7
Medimun-VBR	13
New_VBR1	35
PMP Rev CBR	4
PMP Rev UBR	5
PMP Rev,Unsp ABR	6
PMP Rev,Unsp BE	3
PMP Rev,Unsp CBR	1
PMP Rev,Unsp VBR-NRT	2

QoS Class: Unspecified (VBR-NRT)

Type: PCR CLP=0+1, SCR CLP=0+1, MBS CLP=0+1

	CLP=0	CLP=0+1
PCR (cells/sec):		50000
SCR (cells/sec):		50000
MBS (cells):		50000
MCR (cells/sec):		

Figure 14-21. Set All ATM Traffic Descriptors Dialog Box

The Set All ATM Traffic Descriptors dialog box displays all the traffic descriptors defined on the switch.

11. Perform one of the following actions:

    a. Select a traffic descriptor that has already been defined from the list provided. Then, choose OK. When the Set NHRP LPort Flow Profile/TD Associations dialog box appears, proceed to step 16 to add a secondary traffic descriptor.

    b. Choose add to create a new traffic descriptor. The Add Traffic Descriptor dialog box appears (see Figure 14-22).

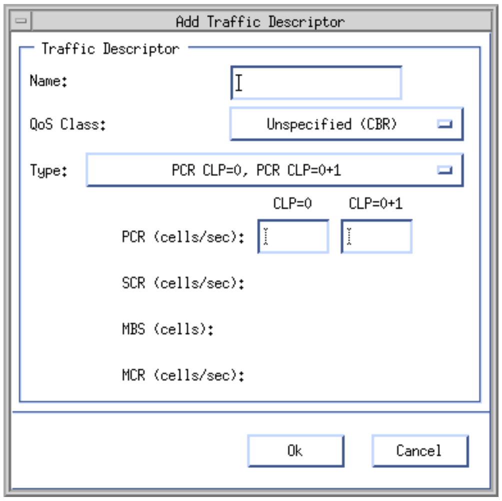

Name:

QoS Class: Unspecified (CBR)

Type: PCR CLP=0, PCR CLP=0+1

	CLP=0	CLP=0+1
PCR (cells/sec):		
SCR (cells/sec):		
MBS (cells):		
MCR (cells/sec):		

Figure 14-22. Add Traffic Descriptor Dialog Box

12. Specify the information required to create a traffic descriptor in the fields on the dialog box. See the *NavisCore ATM Configuration Guide* for details.

13. Choose OK. The Set All ATM Traffic Descriptors dialog box appears. The name of the newly created traffic descriptor appears in the list of traffic descriptors.

14. Select the newly created traffic descriptor.

15. Choose OK. The Set NHRP LPort Flow Profile/TD Associations dialog box appears.

16. Verify that the flow profile association for which you are configuring traffic descriptors is still selected in the list of selected profiles.

17. Choose Secondary TD from the Set NHRP LPort Flow Profile/TD Associations dialog box. The Set All ATM Traffic Descriptors dialog box appears (see Figure 14-21). See the *NavisCore ATM Configuration Guide* for a description of the fields on this dialog box.

18. Perform one of the following actions:

    a. Select a traffic descriptor that has already been defined from the list provided. Then, choose OK. When the Set NHRP LPort Flow Profile/TD Associations dialog box appears, proceed to step 23 to specify whether the traffic descriptor is the alternative or minimum acceptable traffic descriptor.

    b. Choose add to create a new traffic descriptor. The Add Traffic Descriptor dialog box appears (see Figure 14-22).

19. Specify the information required to create a traffic descriptor in the fields on the dialog box. See the *NavisCore ATM Configuration Guide* for details.

20. Choose OK. The Set All ATM Traffic Descriptors dialog box appears. The name of the newly created traffic descriptor appears in the list of traffic descriptors.

21. Select the newly created traffic descriptor.

22. Choose OK. The Set NHRP LPort Flow Profile/TD Associations dialog box appears.

23. Verify that the flow profile association for which you are configuring traffic descriptors is still selected in the list of selected profiles.

24. Select the Secondary TD Type from the Set NHRP LPort Flow Profile/TD Associations dialog box: *Alternate* (the default) or *Minimum*.

25. Choose OK. The Show NHRP LPort Flow Profile/TD dialog box appears, displaying the flow profile association you just created.

Modifying a Flow Profile Association

You can modify a flow profile association's parameters, such as its primary and secondary traffic descriptors and its administrative status. To modify a flow profile association:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set LPort FP/TD Associations. The Set All NHRP LPort Flow Profiles dialog box appears (see Figure 14-18).
3. Select the switch that has the NHRP logical port whose flow profile associations you want to modify. A list of the logical ports on the switch appears.
4. Select the NHRP logical port associated with the flow profile association you want to modify. A list of flow profiles associated with the NHRP logical port appears.
5. Select the flow profile you want to modify.
6. Choose Modify. The Set NHRP LPort Flow Profile/TD Associations dialog box appears (see Figure 14-19).
7. Change the primary traffic descriptors, secondary traffic descriptors, secondary traffic descriptors type, or administrative status as described in the previous section, "Associating a Flow Profile" on page 14-66.
8. Choose OK from the Set NHRP LPort Flow Profile/TD Associations dialog box when you finish changing parameters.

Replacing a Flow Profile Association

You can replace one flow profile association with another flow profile association. To replace a flow profile association:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set LPort FP/TD Associations. The Set All NHRP LPort Flow Profiles dialog box appears (see Figure 14-18).
3. Select the switch that has the NHRP logical port with which you want to associate the flow profile. A list of the logical ports on the switch appears.
4. Select the NHRP logical port.
5. Choose Add. The Set NHRP LPort Flow Profile/TD Associations dialog box appears (see Figure 14-19).

   The Set NHRP LPort Flow Profile/TD Associations dialog box displays two flow profile lists: a list of available profiles and a list of selected profiles (that is, profiles that have been associated with the logical port).

6. In the list of available profiles, select the profile with which you want to replace the associated profile.
7. In the list of selected profiles, select the flow profile association that you want to replace.
8. Choose Replace. The two selected flow profiles move from one list to the other — the flow profile that you selected in the list of available profiles moves to the list of selected profiles, and the flow profile you selected in the list of selected profiles moves to the list of available profiles.
9. Change the primary traffic descriptors, secondary traffic descriptors, secondary traffic descriptors type, or administrative status as described in the previous section, "Associating a Flow Profile" on page 14-66.
10. Choose OK from the Set NHRP LPort Flow Profile/TD Associations dialog box when you are finished.

Deleting a Flow Profile Association

You can delete flow profile associations in two ways:

- Delete all flow profile associations for a selected NHRP logical port. You perform this task from the Set All NHRP LPort Flow Profiles dialog box (see Figure 14-18).
- Delete a single flow profile association for a selected NHRP logical port. You perform this task from the Set NHRP LPort Flow Profile/TD Associations dialog box (see Figure 14-20).

Deleting All Flow Profile Associations

To delete all flow profile associations for a selected NHRP logical port:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set LPort FP/TD Associations. The Set All NHRP LPort Flow Profiles dialog box appears (see Figure 14-18).
3. Select the switch that has the NHRP logical port from which you want to delete the flow profile association. A list of the logical ports on the switch appears.
4. Select the NHRP logical port. A list of flow profiles associated with the logical port appears.
5. Choose Delete.

Deleting a Single Flow Profile Association

To delete a single flow profile association:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set LPort FP/TD Associations. The Set All NHRP LPort Flow Profiles dialog box appears (see Figure 14-18).
3. Select the switch that has the NHRP logical port from which you want to delete the flow profile association. A list of the logical ports on the switch appears.
4. Select the NHRP logical port.
5. Choose Modify. The Set NHRP Flow Profile/TD Associations dialog box appears (see Figure 14-20).
6. Select the flow profile association you want to delete from the list of selected profiles.
7. Choose Delete.

Configuring Log Parameters

Log parameters allow you to specify:

- The network workstation and directory where NHRP logging information is stored
- When the NHRP logs are flushed from the switch to the workstation

> The switch “flushes” log information to a workstation over the network using the Trivial File Transfer protocol (TFTP). The workstation must run the TFTP server in order to receive log information from the switch.

To configure NHRP log parameters:

1. Select a switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set NHRP ⇒ Set Log Parameters. The Set All NHRP Log Parameters dialog box appears (see Figure 14-23).

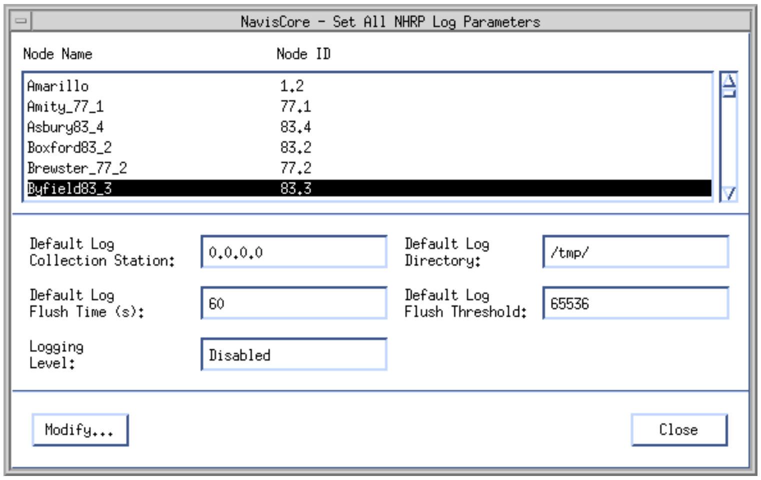

NavisCore - Set All NHRP Log Parameters

Node Name	Node ID
Amarillo	1.2
Amity_77_1	77.1
Asbury83_4	83.4
Boxford83_2	83.2
Brewster_77_2	77.2
Byfield83_3	83.3

Default Log Collection Station: 0.0.0.0

Default Log Directory: /tmp/

Default Log Flush Time (s): 60

Default Log Flush Threshold: 65536

Logging Level: Disabled

Figure 14-23. Set All NHRP Log Parameters Dialog Box

3. Select the switch you want to configure and choose Modify. The Set NHRP Log Parameters dialog box appears (see Figure 14-24).

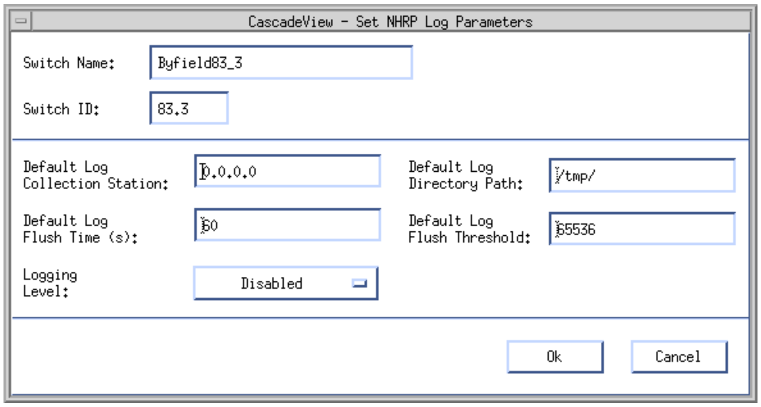

CascadeView - Set NHRP Log Parameters

Switch Name: Byfield83_3

Switch ID: 83.3

Default Log Collection Station: 0.0.0.0

Default Log Directory Path: /tmp/

Default Log Flush Time (s): 60

Default Log Flush Threshold: 65536

Logging Level: Disabled

Figure 14-24. Set NHRP Log Parameters Dialog Box

4. Specify the parameters in the fields described in Table 14-12.

Table 14-12. NHRP Log Parameters

Field	Action/Description
Switch Name	Displays the name of the selected switch.
Switch ID	Displays the subnetwork number and host ID in the internal IP address of the switch.
Default Log Collection Station	Specify the IP address (e.g., 131.100.45.12) of the NMS that collects the NHRP logging information from the switch. Keep in mind that NMS performance could be adversely effected if you use a single NMS to collect logs from many switches. Try to distribute log collection among multiple NMSes, if possible.
Default Log Flush Time (s)	Specify the number of seconds, from 0 to 65535, that logs are stored on the switch before they are flushed to the NMS. The default is 60 seconds. A value of 0 means that logs are never flushed due to timeout.

Table 14-12. NHRP Log Parameters (Continued)

Field	Action/Description
Logging Level	Select the logging importance level of NHRP request/reply activity on the switch. Keep in mind that, when you choose one level, you also choose all the levels above it. For example, if you choose *Warning*, you also choose *Critical* and *Fatal*. The choices are: *Disabled* – (default) Logging is disabled. (Logging levels in order of importance) *Fatal* – No memory available to process NHRP requests and replies. *Critical* – Low amount of memory available to process NHRP requests and replies. Includes the Fatal level. *Warning* – Dropping NHRP requests and replies due to queue overload. Includes the Critical and Fatal levels. *Info-High* – All Registration Requests, Registration Replies, Purge Requests, and Purge Replies are logged. Includes all of the above levels. *Info-Medium* – All Resolution Requests and Resolution Replies are logged. Includes all of the above levels. *Info-Low* – All Error Indication messages are logged. Includes all of the above levels. *Info-Debug* – All Registration Refresh Requests and Registration Refresh Replies are logged. Includes all of the above levels.
Default Log Directory Path	Specify the pathname of the directory where logs are stored. The pathname must end with a slash (for example, /tmp/). The default is /tmp/. Make sure that the TFTP server can write to this directory (use the chmod 777 command to give the TFTP server access, if needed). The format of NHRP log file names is similar to the format of bulk statistics log file names. NHRP log file names begin with “nhrp,” followed by the IP address of the switch and a timestamp that identifies the date and time. The file contains a record of events, one entry for each event. Events are recorded on a second-by-second basis.
Default Log Flush Threshold	Specify the number of bytes of RAM that logs can occupy on the switch before they are flushed to the workstation. The value may range from 1 to 262144 (256 KB). The default is 65536 (64 KB). For example, if you specify 80000, the switch flushes the logs to the NMS when the logs occupy 80000 bytes of RAM.

5. Choose OK.

15

Configuring IP Multicast Routing

This chapter provides:

- An overview of IP multicast routing
- An overview of IP multicast routing protocols
- An overview of Lucent’s implementation of IP multicast routing
- Instructions on how to configure IP multicast routing in a Lucent network

Overview of IP Multicast Routing

IP multicast routing provides dynamic, real-time communications between network hosts. It is commonly used to deliver multimedia traffic (such as video) over local- and wide-area IP networks.

Before the emergence of IP multicast routing, IP networks typically supported two types of host-to-host communications:

Unicast Communication — One host communicates with another host. For example, one host transfers files to another host using the File Transfer Protocol (FTP). Unicast communication makes efficient use of network resources, but is limited to communication between two hosts.

Broadcast Communication — One host simultaneously communicates with all the other hosts on its LAN. Broadcast communication enables one host to communicate with many hosts simultaneously, but can potentially be an excessive consumer of network resources because it is indiscriminate — that is, a broadcast cannot be selectively targeted to a group of hosts. As a result, broadcast communication cannot scale to networks that have a large number of hosts.

IP multicast routing combines the bandwidth efficiency of unicast communication with the ability of broadcast communication to reach multiple hosts at once. Using IP multicast routing, a host sends data to a selected group of hosts called a *multicast group*. Each multicast group is identified by a special Class D IP address called a *group address*, and each member of the group shares this same address. Class D IP addresses begin with the bits 1110 and may range from 224.0.0.0 to 239.255.255.255, except for the addresses reserved by the Internet Assigned Numbers Authority (IANA). The IANA web site (http://www.iana.org) provides more information on reserved addresses.

Class D addresses differ somewhat from Class A, B, and C addresses. Unlike Class A, B, and C addresses, Class D addresses are not divided into network, subnetwork, and host parts. A Class D address is a simple, non-hierarchical address that identifies a group of multicast hosts.

Group members do not have to share the same physical link (e.g., an Ethernet LAN). Members of the same group may be dispersed on various LANs and WANs across the network. As the multicast datastream traverses the network, it is replicated by routers only on interfaces that are used to reach the multicast destinations, making efficient use of network bandwidth.

Figure 15-1 shows a host sending a multicast data stream to a multicast group.

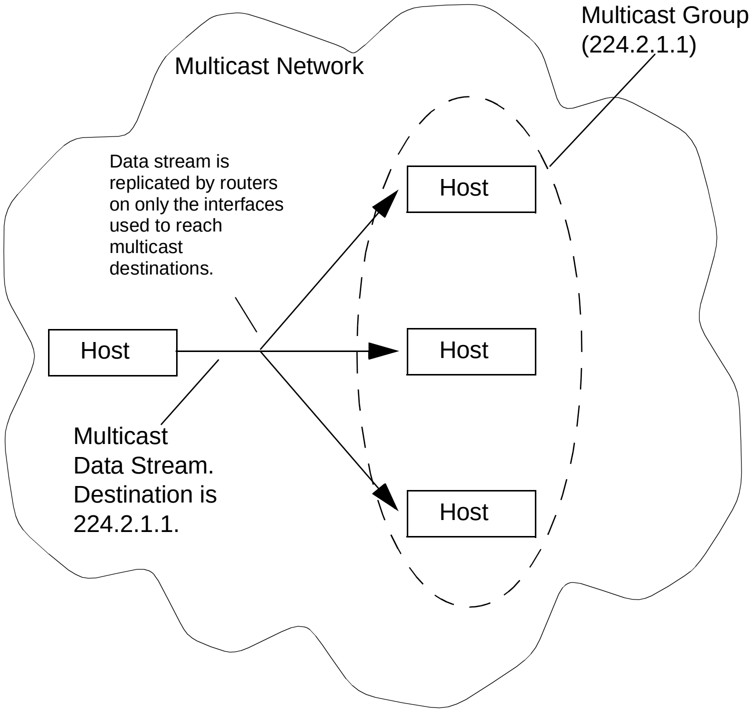

Figure 15-1. Multicast Transmission

A host joins a multicast group through use of the Internet Group Management Protocol (IGMP). This protocol allows hosts to register as members of a particular multicast group. Once hosts register as multicast group members, routers in the network track them dynamically. Lucent currently supports IGMP.

IP multicast traffic is transmitted from the source to the destinations via a spanning tree that connects all the hosts in the group. Different IP multicast routing protocols use different techniques to construct these trees. All multicast traffic is distributed through this tree once it is constructed.

The next section describes the techniques and protocols for constructing spanning trees.

IP Multicast Routing Protocols

IP multicast routing protocols use one of two methods to distribute routing information, depending on how the multicast group members are connected in the network:

Dense-mode — Dense-mode routing protocols are based on the assumptions that the multicast group members are densely distributed throughout the network (that is, many of the subnets contain at least one group member) and that bandwidth is always available. Dense-mode multicast routing protocols rely on a broadcast-like technique called *flooding* to propagate routing information to all routers in the network. Dense-mode routing protocols include Distance Vector Multicast Routing Protocol (DVMRP), Multicast Open Shortest Path First (MOSPF), and Protocol-Independent Multicast — Dense Mode (PIM-DM).

Sparse-mode — Sparse-mode routing protocols assume that the multicast group members are sparsely distributed throughout the network and bandwidth is not necessarily available (for example, users may be connected via slow-speed links). The term “sparse-mode” does not imply that the group has a few members — the term only means that the members are widely dispersed. In this case, using flooding to propagate routing information would unnecessarily waste network bandwidth and, as a result, could cause serious performance problems. Sparse-mode multicast routing protocols must rely on selective techniques to set up and maintain multicast trees. Sparse-mode routing protocols include Core-Based Trees (CBT) and Protocol-Independent Multicast — Sparse Mode (PIM-SM).

The dense-mode multicast routing protocols are the more common of the two types of multicast routing protocols currently in use. At this time, Lucent supports the dense-mode multicast routing protocols DVMRP and MOSPF. Lucent also supports a technique called *tunneling* that can be used to route IP Multicast traffic over portions of an inter-network that do not have multicast capabilities.

Distance Vector Multicast Routing Protocol

DVMRP was the first protocol developed to support IP multicast routing. DVMRP is described in RFC 1075 and in an Internet Draft update. It has been widely used on the Internet Multicast Backbone (MBONE), which is the multicast equivalent of the traditional unicast Internet.

DVMRP routers forward multicast datagrams over *virtual interfaces* (VIFs), which correspond to broadcast links and point-to-point links. An example of a broadcast link is an Ethernet LAN; an example of a point-to-point link is PPP. A router that supports DVMRP can use a special VIF, called a DVMRP tunnel, to send DVMRP messages to other DVMRP routers through non-DVMRP routers.

▶ Routers cannot forward multicast datagrams directly over Non-Broadcast Multiple Access (NBMA) networks, such as ATM and Frame Relay networks. Instead, DVMRP treats these networks as a DVMRP tunnel.

DVMRP routers collectively build a different distribution tree for each source of multicast traffic and its destination host group. Each distribution tree is a minimum spanning tree from the multicast source at the root to all the multicast destinations at the leaves. The distribution tree provides the shortest path between the source and each multicast destination in the group, based on the DVMRP metric. This metric (configurable through the NMS) is typically the number of hops in the path.

The routers collectively construct a tree on demand when a source begins to transmit datagrams to a multicast group. The datagrams follow this tree from the source to all group members. The routers along the paths to the group members replicate the datagrams only at branches that lead to the group members. The routers “prune” branches that do not lead to group members. The tree construction has two phases: a broadcast phase and a pruning phase.

Broadcast Phase

During the broadcast phase, DVMRP assumes initially that every host on the network is part of the multicast group. The designated router on the subnet where the multicast traffic’s source host resides transmits multicast datagrams to all adjacent routers. Each adjacent router then selectively forwards the datagrams to other routers until all multicast group members receive the datagrams.

Upon receiving a multicast datagram, a router checks its DVMRP unicast routing table (DVMRP has its own unicast routing protocol) to determine the interface associated with the shortest path back to the source. DVMRP routers periodically exchange information to keep these routing tables up-to-date, insuring that all routers have a consistent network view.

If the multicast datagram arrived on the interface associated with the shortest path back to the source, the router identifies the following in its multicast routing tables:

- The destination multicast group
- Interfaces on which messages addressed to that group should be forwarded

The router then forwards the multicast datagram to all adjacent routers except for the router from which the datagram is received. This routing technique — *Reverse Path Forwarding* — insures that:

- No routing loops are in the spanning tree
- The tree includes the shortest paths from the source to all destinations

Using information in its DVMRP unicast routing tables, a DVMRP router can selectively forward datagrams to only those adjacent routers that can further construct the multicast spanning tree (that is, those adjacent routers that should be used to reach the members of the destination multicast group, except for the router from which the datagram was received).

The relationship between routers along a multicast spanning tree is called *downstream dependency.* As datagrams flow from a multicast source to a multicast destination, each router forwards datagrams "downstream" to one or more adjacent routers. When an "upstream" router forwards datagrams to a "downstream" router, the "downstream" router is a *downstream dependent* of the upstream router. In turn, the "downstream" router may have its own downstream dependents to which it forwards datagrams, and so on.

As part of the spanning tree construction process, downstream dependents send *poison reverse route reports* to upstream adjacent routers. A *poison reverse route report* notifies an upstream router that it has a downstream dependent reachable over the interface on which the upstream router received the report. The upstream router will then forward multicast traffic over that interface.

For example, Router1 receives a poison reverse route report from Router2 on InterfaceA. Router1 now knows that it has a downstream dependent (Router2) reachable over InterfaceA. Router1 will forward multicast traffic over InterfaceA.

A router never forwards multicast traffic over an interface if the interface meets both of the following criteria:

- No poison reverse route report is received on the interface
- No multicast group members are directly connected by the interface (for example, no multicast group members are directly connected via an Ethernet LAN)

Pruning Phase

During the pruning phase, DVMRP eliminates branches of the tree that do not lead to any multicast group members. To determine group membership, DVMRP relies on IGMP to maintain group membership information in the routers. When a router determines that no hosts beyond it are members of the multicast group, it sends a prune message to its upstream router. Routers must update source and destination group mapping information in their tables so that they know which branches have been pruned from the tree. This process continues until all of the useless branches are eliminated from the tree.

Once constructed, the multicast spanning tree is used to transmit multicast messages from the source to the destination multicast members. Each router in the path forwards messages over only the interfaces that can be used to reach group members. Since new members may join the group at any time, and since these new members may depend on one of the pruned branches to receive the multicast transmission, DVMRP reconstructs the spanning tree from time to time.

▶ In two cases, routers can forward data over pruned branches:

– If a member joins a multicast group, and that member resides on a branch that was previously pruned, a downstream neighbor sends a graft message upstream. This message tells the routers that the previously pruned branch is part of the multicast tree again.

– Pruned branches have timers associated with them. When these timers expire, routers can forward data to the branch until they receive a new prune message.

Figure 15-2 illustrates construction of a DVMRP multicast spanning tree.

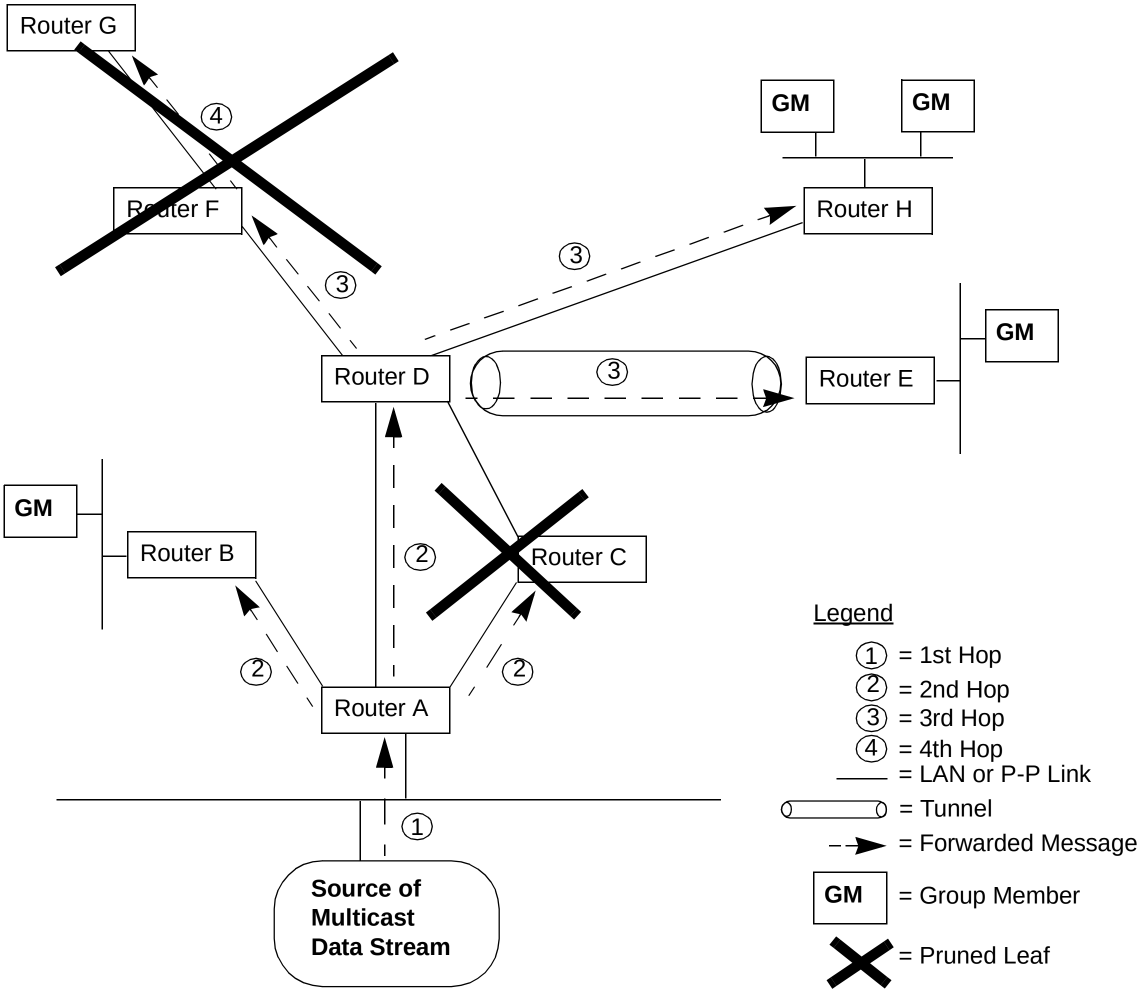

Figure 15-2. Constructing a DVMRP Multicast Spanning Tree

In Figure 15-2, the spanning tree is constructed as follows:

1. The multicast datagram originates from the source and reaches Router A (first hop). Router A looks in its DVMRP unicast routing table and determines that it must replicate the datagram and forward it to downstream dependents. Router A forwards the datagram on all of its interfaces except for the one on which the datagram was received.

2. The message reaches Router B, Router C, and Router D (second hop). Router B delivers the datagram to the group member on its LAN. Router D forwards the message (third hop) to Router H and Router E, which in turn deliver the message to the group members on their respective LANs.

3. Notice that Router C and Router D do not exchange messages in the third hop. This exchange does not occur because Router C and Router D know, based on their DVMRP unicast routing tables, that the interface they use to communicate with each other is associated with a path to the source of the DVMRP message, but is not the shortest path to the source. In terms of reaching the source network, Router D is not a downstream dependent of Router C.

   When multiple paths to a multicast source exist, a router will send a poison reverse route report upstream only on the interface associated with the path with the lowest cost (that is, the best path). For example, Router D never sends a poison reverse route report to Router C, but does send a poison reverse route report to Router A. Since Router C never receives a poison reverse route report on the interface it uses to reach Router D, it will not know that a downstream dependent is reachable on that interface. Thus, Router C will not forward multicast traffic to Router D.

4. The message reaches Router G. Because Router G is a leaf router with no group members on its subnet, it sends a prune message back to Router F. Router F then sends a prune message to Router D. Note that Router C also sends a prune message to Router A.

Figure 15-3 shows the final, constructed spanning tree.

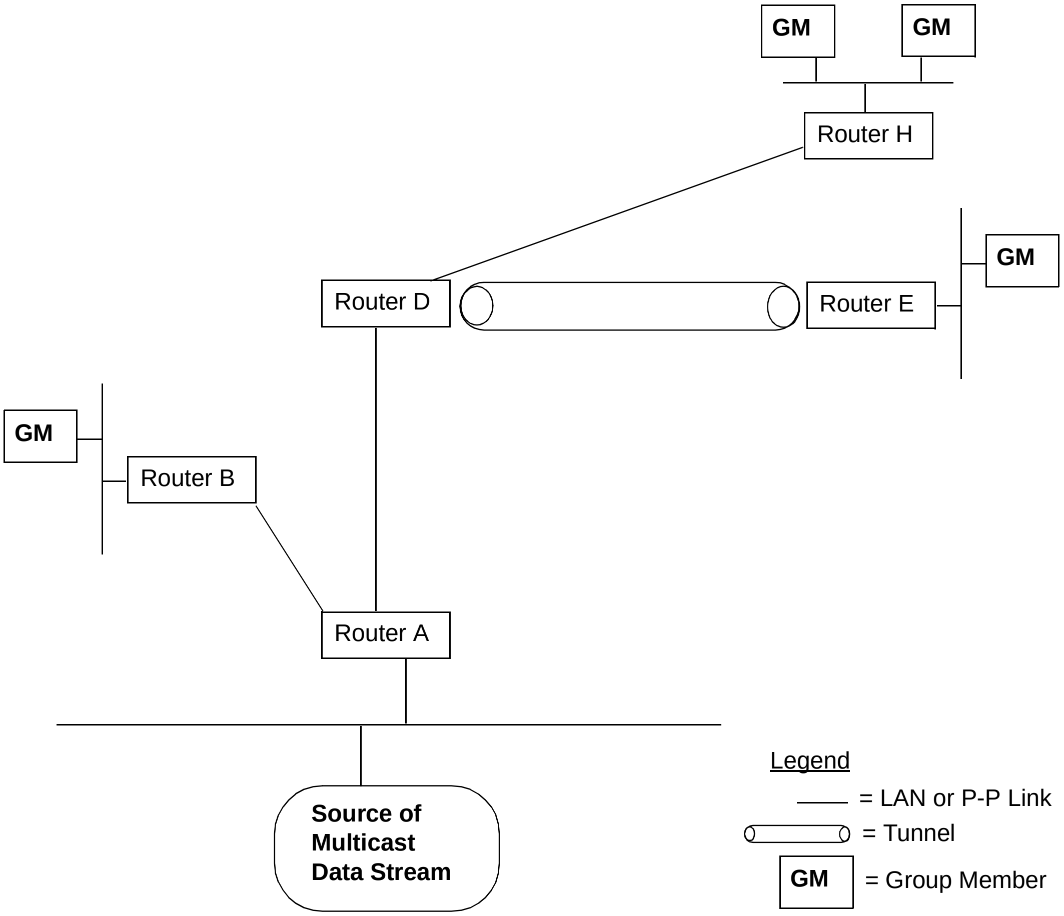

Figure 15-3. Constructed DVMRP Multicast Spanning Tree

Multicast Open Shortest Path First

Just as DVMRP includes its own unicast routing protocol, MOSPF relies on OSPF as its unicast routing protocol. Before proceeding further with a description of MOSPF, it is helpful to quickly review OSPF.

OSPF is a unicast routing protocol that routes messages along least-cost paths where cost is expressed as a link-state metric. Along with the number of hops in a path, other factors that can influence the cost of a path include:

Load-balancing Information — In an attempt to balance traffic on the network, OSPF might assign a lower cost to a link with very low traffic than the cost it might assign to a heavily used link

Application's Desired Quality of Service — If an application requires low latency, a path that includes, for example, a satellite link should be assigned a high cost.

OSPF can partition a network into multiple routing domains. For example, you could partition a network into multiple routing domains where each domain is controlled by a single organization.

MOSPF (defined in RFC 1584) is designed for use in a single routing domain. In a network that uses OSPF and MOSPF, each router maintains a topological view of the entire network that is regularly updated. Routers use this link-state information from this view to construct multicast distribution trees.

Each MOSPF router uses IGMP to periodically collect information about multicast group membership. Along with the link-state information, the group membership information is flooded to all other routers in the routing domain. Routers update their internal link-state information based on information that they receive from adjacent routers. Because each router knows the topology of the entire network, each router can then calculate a least-cost spanning tree on its own, with the multicast source as the root and the group members as leaves. This tree is the path for routing multicast traffic from the source to each of the group members.

MOSPF uses the Dijkstra algorithm to compute a shortest-path tree. A separate calculation is required for each source and its associated destination group. For efficiency, a router only makes this calculation when it receives the first datagram in a stream. Once the tree is constructed, the router stores the information for later use in routing datagrams from that stream.

Figure 15-4 shows how an MOSPF tree is constructed.

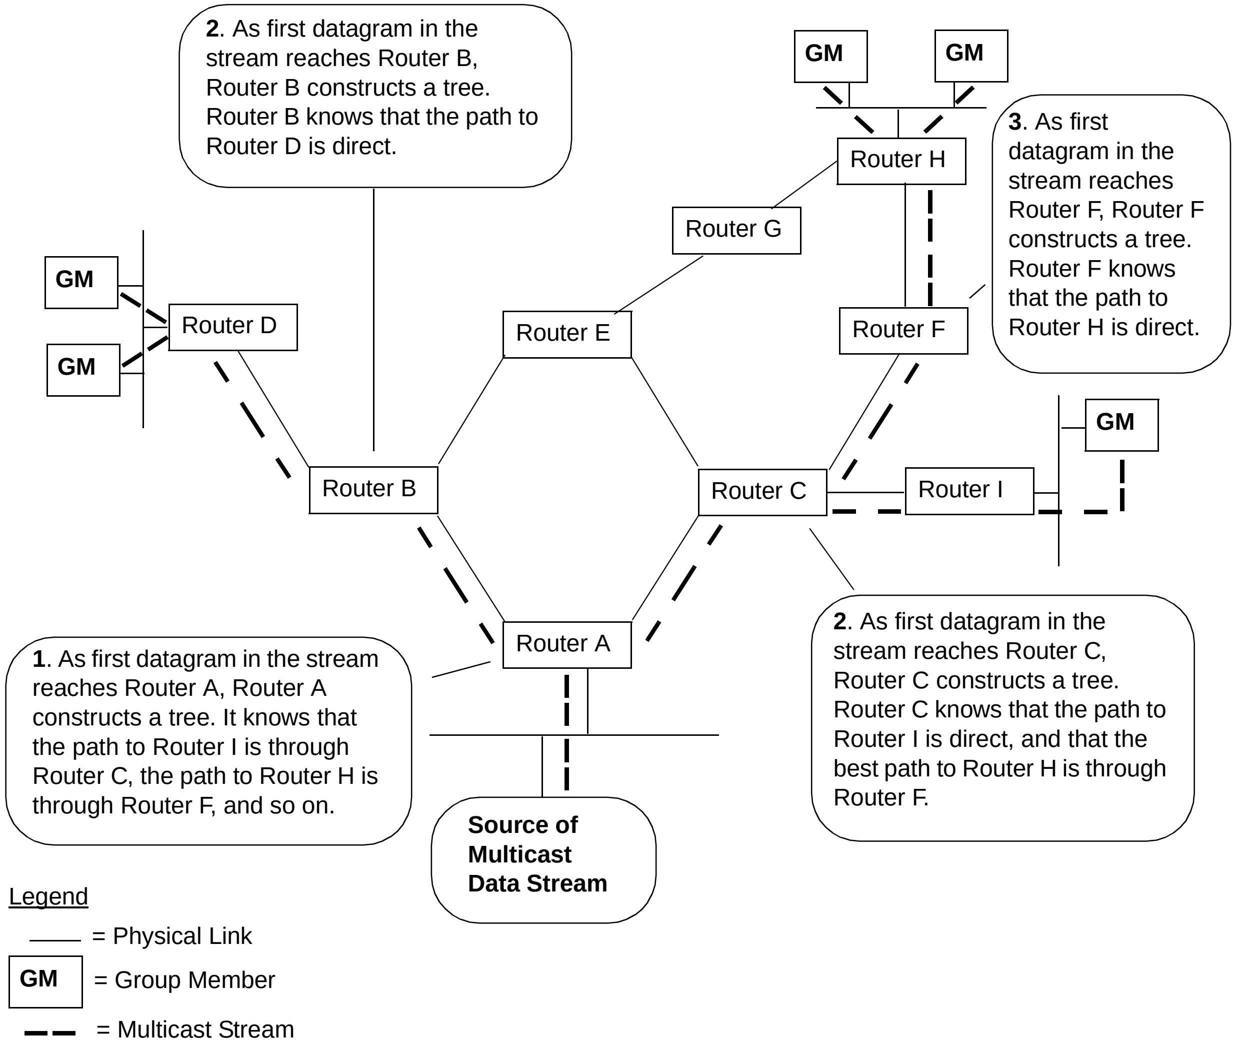
2. As first datagram in the stream reaches Router B, Router B constructs a tree. Router B knows that the path to Router D is direct.
GM
GM
GM
GM
GM
Router H
Router G
Router D
Router E
Router F
3. As first datagram in the stream reaches Router F, Router F constructs a tree. Router F knows that the path to Router H is direct.
Router B
Router C
Router I
Router A
1. As first datagram in the stream reaches Router A, Router A constructs a tree. It knows that the path to Router I is through Router C, the path to Router H is through Router F, and so on.
2. As first datagram in the stream reaches Router C, Router C constructs a tree. Router C knows that the path to Router I is direct, and that the path to Router H is through Router F.
Source of Multicast Data Stream
Legend
= Physical Link
GM = Group Member
= Multicast Stream

Figure 15-4. Constructing an MOSPF Tree

The multicast transmission initiates the scenario illustrated in Figure 15-4. When it receives a message, each router calculates exactly the same distribution tree as its predecessors and uses the tree to forward the message.

Note that Router C chooses the least cost path to reach Router H (that is, routes the multicast stream through Router F instead of Router E). Keep in mind that MOSPF can use OSPF to determine the best path to take to a destination.

Tunneling

Multicast tunneling involves encapsulating multicast packets in unicast packets (IP datagrams). Once encapsulated, multicast packets can be routed through unicast networks, such as the Internet, that do not support multicast routing. Multicast tunneling is also used to route multicast traffic over ATM and Frame Relay circuits. Tunneling can play an important part during a transition from unicast networking to multicast networking.

Figure 15-5 illustrates the use of tunneling.

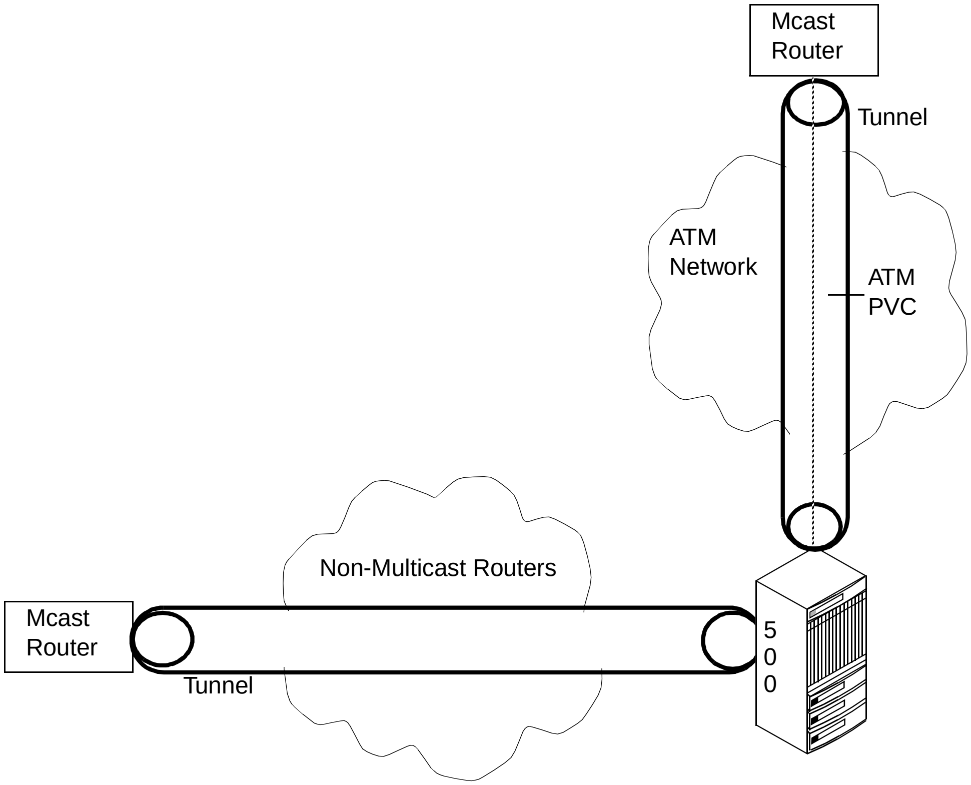

Figure 15-5. Tunneling

Lucent Implementation of IP Multicast Routing

Lucent supports IGMP, DVMRP, and MOSPF in its implementation of IP multicast routing. The rest of this section describes how Lucent implements these protocols.

IGMP Implementation

Lucent switches support Version 1 and Version 2 of IGMP on Ethernet logical ports. Version 1 is defined in RFC 1112, and Version 2 is defined in RFC 2236.

Lucent switches use IGMP to learn the existence of multicast group members on the directly attached Ethernet LAN. Once the switch acquires this knowledge, DVMRP and MOSPF on the Lucent switch can use it to construct multicast spanning trees.

Lucent switches may act as *queriers* or *non-queriers*. As the name implies, a querier issues IGMP queries on the physical network. A non-querier does not issue IGMP queries.

On each physical network (e.g., an Ethernet LAN), there is only one querier. When a multicast router starts up, it is a querier by default, but it relinquishes its querier status and becomes a non-querier when it receives a query from a router with a lower IP address. (The multicast router with the lowest IP address on the LAN will always be the querier). If a router does not receive a query message from another router within a certain period of time, it assumes that it is the only router on the LAN and remains a querier.

► In Lucent's implementation of IGMP Version 1, the switch is always a querier. It never relinquishes its querier status.

When acting as a querier, to learn the existence of multicast group members, the Lucent switch broadcasts membership query request messages on the attached Ethernet LAN. There are two types of membership query requests:

General Query — Used to learn which multicast groups have members on the LAN.

Group-Specific Query — Used to learn if a specific multicast group has members on the LAN.

Hosts on the LAN that want to join the group respond with membership report messages. Hosts that respond with membership reports may receive multicast datagrams destined for the group.

For Version 2 only, when hosts initially join a group, they may send unsolicited membership report messages. After sending the initial membership report message in response to the membership query request, hosts may send the Lucent switch an unsolicited membership report message as a guard against the loss or corruption of the initial membership report message.

When a host wants to end its membership in a group (and it is the last multicast group member on its LAN), the host notifies the Lucent switch (and all other routers in the network) by sending it a leave group message. The message is addressed to a special multicast group called the all-routers group, which uses the reserved Class D IP address of 224.0.0.2. The leave group message is supported in Version 2 only.

Figure 15-6 shows Lucent switches on two different Ethernet LANs exchanging IGMP messages with hosts. A host on one LAN is joining a multicast group; a host on the other LAN is leaving a multicast group. The figure assumes that the Lucent switches and hosts all run Version 2 of IGMP.

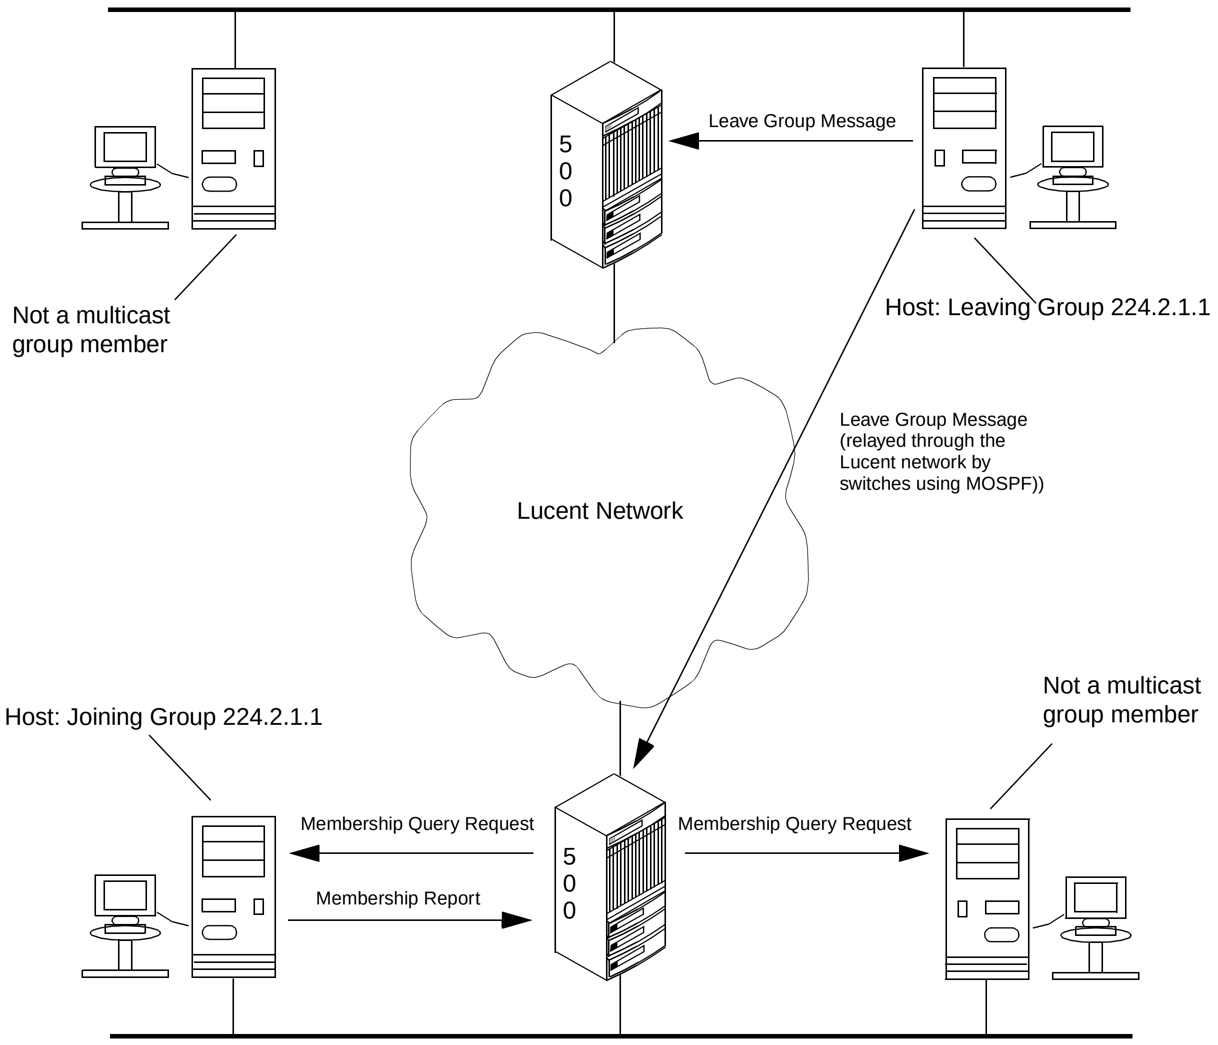
Leave Group Message
5
0
0
Not a multicast group member
Host: Leaving Group 224.2.1.1
Leave Group Message (relayed through the Lucent network by switches using MOSPF))
Lucent Network
Host: Joining Group 224.2.1.1
Not a multicast group member
Membership Query Request
Membership Report
5
0
0
Membership Query Request

Figure 15-6. IGMP Message Exchange

DVMRP and MOSPF Implementation

Lucent's DVMRP and MOSPF implementations interoperate to support IP multicast routing. On Lucent switches, DVMRP and MOSPF share a common *multicast forwarding cache*, which contains all of the routing information used by DVMRP and MOSPF.

Outside the Lucent network, Lucent switches can use DVMRP and MOSPF to exchange multicast traffic with third-party equipment such as routers.

Within the Lucent network, Lucent switches use MOSPF to route all multicast traffic (both MOSPF and DVMRP traffic) over connections called *multicast label switched paths* (multicast LSPs). Switches create these connections automatically — your only management task is to verify that multicast LSPs are enabled on the switches that require them. See Chapter 12, "Configuring Label Switched Paths" for details.

To DVMRP on ingress and egress switches, the Lucent network appears as a LAN. DVMRP routes multicast traffic across the Lucent network as if it were routing the traffic over an Ethernet network.

You may encounter situations where DVMRP and MOSPF may have to interoperate outside the Lucent network. For example, a Lucent switch may be connected to both a network that supports DVMRP and a network that supports MOSPF. In this case, DVMRP exports the source networks (networks that are sources of multicast traffic) it knows about to MOSPF, and MOSPF then advertises these networks throughout the MOSPF routing domain.

By the same token, MOSPF exports the source networks it knows about to DVMRP, and DVMRP in turn advertises these networks to other DVMRP routers. If DVMRP learns more than half of its source networks (or more) on its own, DVMRP instructs MOSPF to advertise a default route. However, if DVMRP learns half of its source networks (or more) through MOSPF, DVMRP advertises all of the individual routes to MOSPF.

To export routes from DVMRP to MOSPF, you do not have to perform any configuration tasks. However, to export routes from MOSPF to DVMRP, you have to perform certain configuration tasks. See "Planning Your MOSPF Configuration" on page 15-24 for more information.

Figure 15-7 shows a mixed network environment in which Lucent switches and third-party equipment use both DVMRP and MOSPF to route multicast traffic. On Switch 3, notice how DVMRP and MOSPF pass routing information back and forth between the external MOSPF network and the DVMRP network.

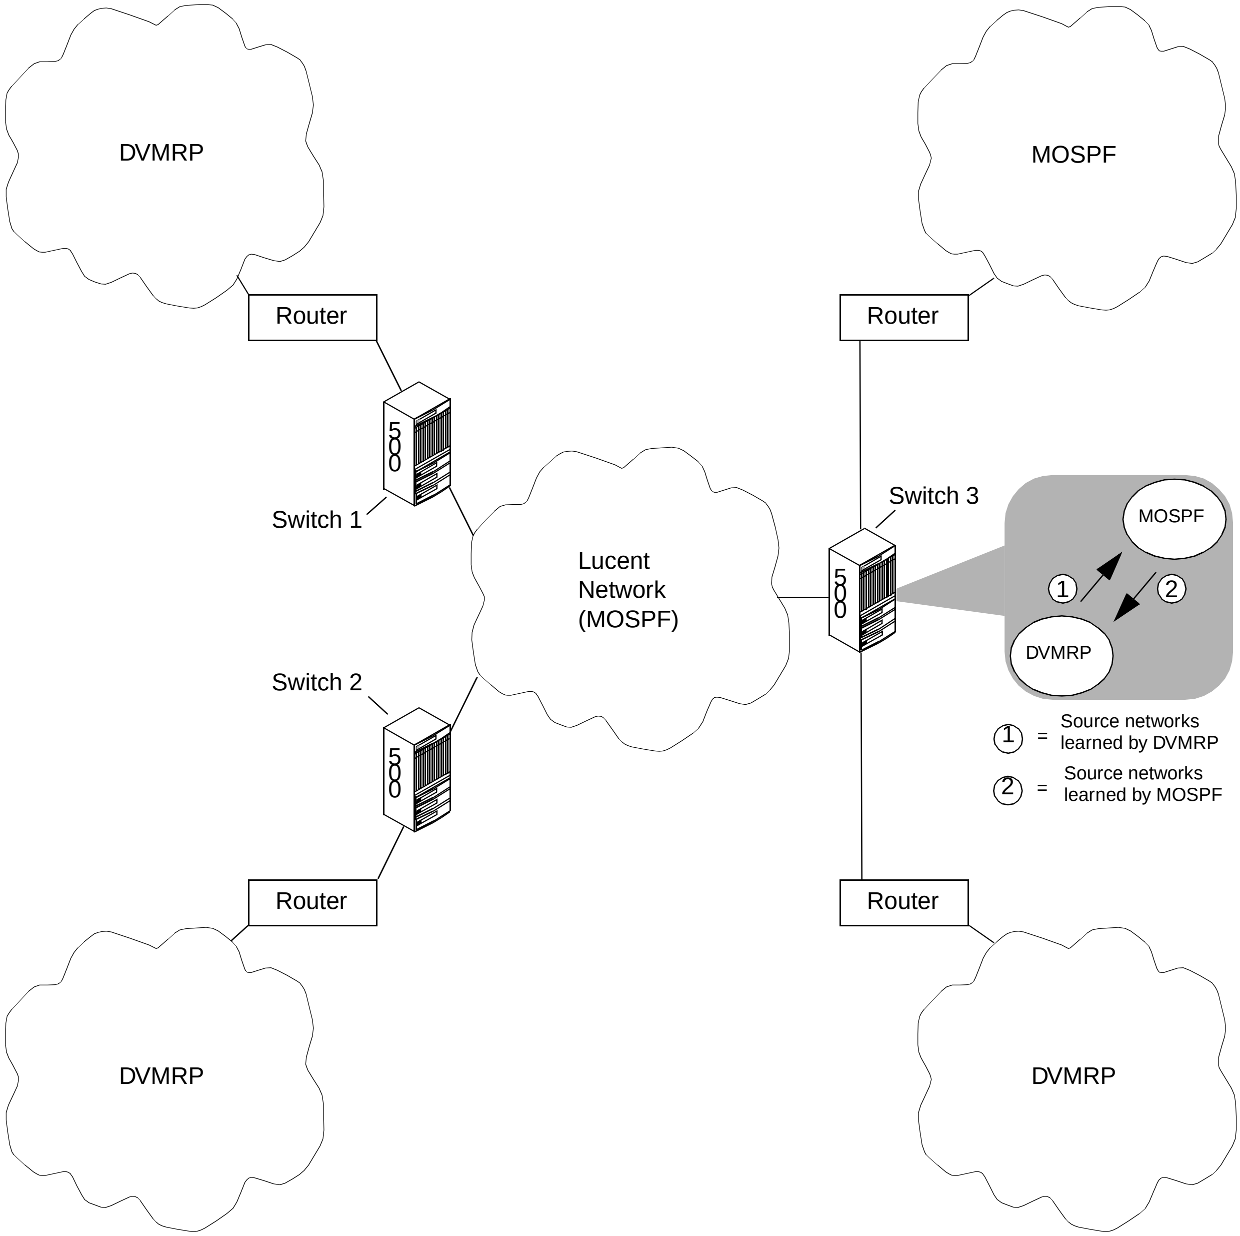

Figure 15-7. Mixed DVMRP/MOSPF Network Environment

Planning IP Multicast Routing

Before you configure IP multicast routing, take some time to plan your configuration. The rest of this section helps you to plan for configuring IGMP, DVMRP, and MOSPF.

Since MOSPF and DVMRP rely on IGMP, consider configuring IGMP first (if necessary). Then, configure DVMRP and MOSPF.

Keep in mind that you may encounter a situation where your switch requires DVMRP and MOSPF capabilities, but does not require IGMP. For example, your switch may not be connected to any Ethernet LANs that have multicast group members, but it still may have to forward multicast traffic using DVMRP or MOSPF.

Verifying Basic IP Connectivity

Before you configure IP multicast routing, verify that the switch is capable of basic IP broadcast and unicast communications on all relevant interfaces. All IP logical ports, IP server logical ports, and/or Ethernet logical ports should be operational.

Planning Your IGMP Configuration

To plan your IGMP configuration, identify Ethernet logical ports that are associated with Ethernet LANs where multicast group members reside. You will have to configure IGMP on these logical ports.

Figure 15-8 illustrates IGMP configuration requirements for Ethernet logical ports. In this figure, a Lucent switch has two Ethernet logical ports. One requires IGMP to be configured; the other one does not.

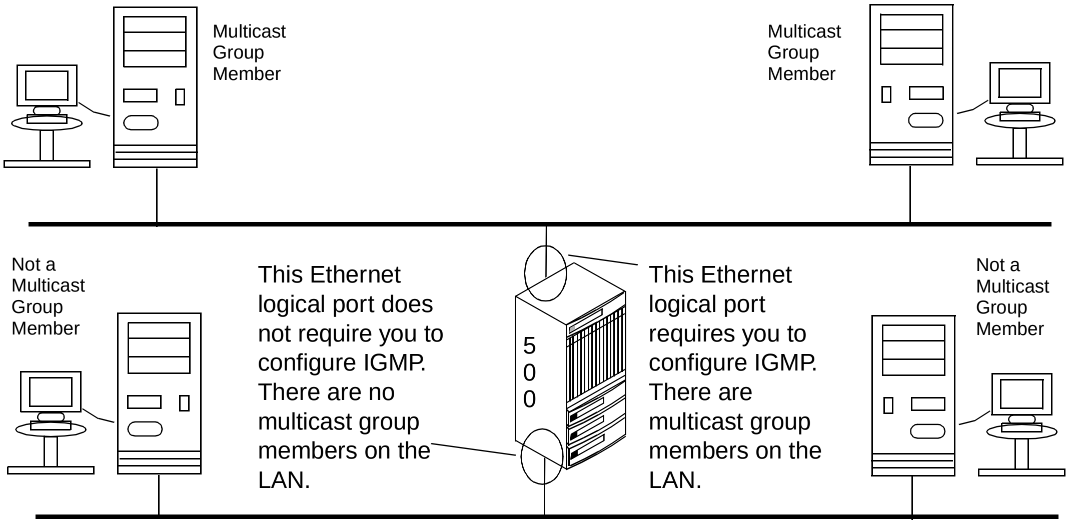
Multicast Group Member
Multicast Group Member
Not a Multicast Group Member
This Ethernet logical port does not require you to configure IGMP. There are no multicast group members on the LAN.
500
This Ethernet logical port requires you to configure IGMP. There are multicast group members on the LAN.
Not a Multicast Group Member

Figure 15-8. IGMP Configuration Requirements for Ethernet Logical Ports

Planning Your DVMRP Configuration

To plan your DVMRP configuration, identify physical interfaces and tunnels, and identify scoped boundaries. The rest of this section describes these tasks.

Identifying Physical Interfaces and Tunnels

Identify all the physical interfaces and tunnels with which you will need to associate VIFs. A physical interface is an interface to an Ethernet LAN or a Point-to-Point Protocol (PPP) line. A tunnel is either a path through a non-multicast router or a Frame Relay or ATM circuit. These physical interfaces and tunnels use DVMRP as the multicast protocol.

► Because DVMRP depends on the services of a unicast distance-vector routing protocol, add a RIP interface to any IP interface that uses DVMRP as the multicast routing protocol. For example, if you configure DVMRP for an IP interface to an Ethernet LAN, you would configure a RIP interface for that IP interface. As another example, if an IP interface acts as a tunnel endpoint, you would configure a RIP interface for that IP interface. See "Configuring RIP at the Logical Port" on page 7-1 for more information on configuring a RIP interface.

You do not associate VIFs with physical interfaces that use MOSPF as the only multicast protocol, unless you tunnel through the MOSPF network to DVMRP nodes.

Figure 15-9 shows some interfaces that should be associated with VIFs and a non-VIF that is associated with an interface to an MOSPF-only network.

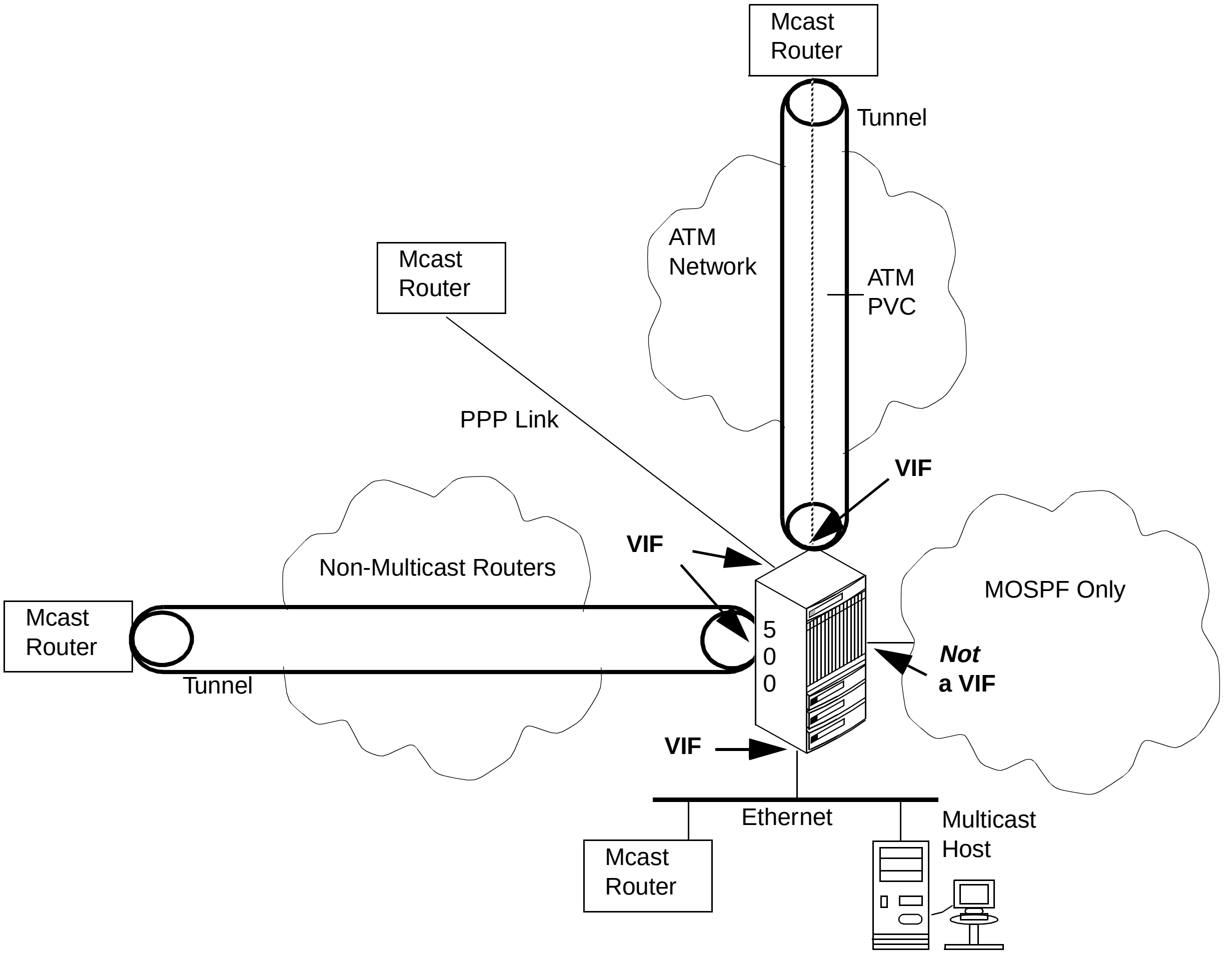

Figure 15-9. Sample VIFs and a Non-VIF

When you configure a VIF that is associated with an Ethernet or PPP physical interface, you must specify the IP address assigned to the IP logical port configured for that physical interface and the IP address of the subnetwork, as well as other parameters. For example, when you configure a VIF associated with the Ethernet logical port in Figure 15-10, you would specify 131.100.28.3 for the local (IP) address and 131.100.28.0 for the subnetwork address. See "Configuring DVMRP VIFs for Fast Ethernet and PPP Logical Ports" on page 15-32 for more information.

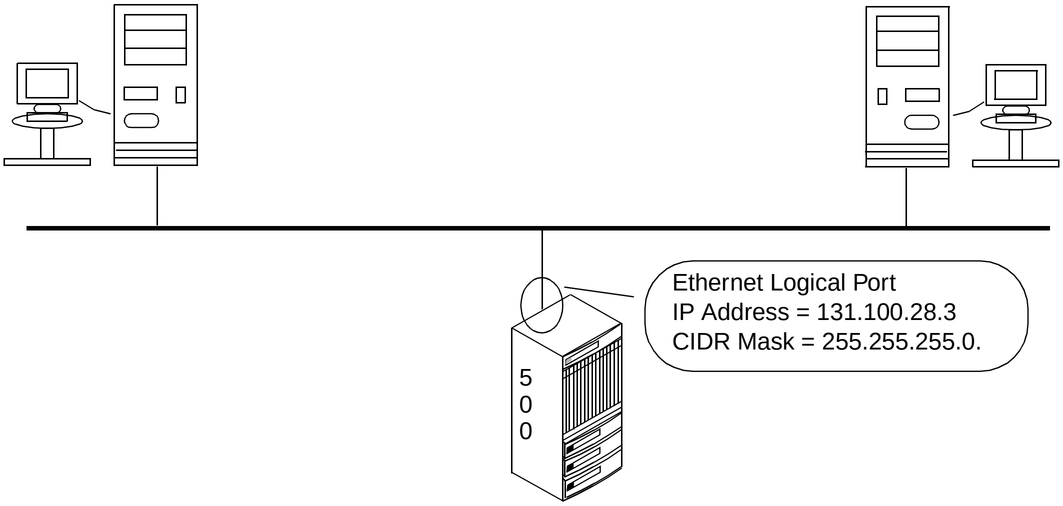

Figure 15-10. Sample IP Address and CIDR Mask for Ethernet Logical Ports

When you configure a VIF that is associated with a tunnel, you must specify the IP address assigned to the IP logical port or IP server logical port that the switch uses to interface to the tunnel, and the IP address of the node (e.g., a router) at the other end of the tunnel. For example, when you configure VIFs associated with the two tunnels in Figure 15-11, you would specify the local and remote IP addresses listed in Table 15-1 (as well as other parameters).

> You cannot configure an IP logical port associated with a physical interface (for example, an Ethernet or PPP interface) as a VIF and have that same interface act as a tunnel endpoint.

> Lucent switches route DVMRP traffic to each other without requiring you to configure any VIFs, as long as the switches are connected by trunks. However, if the switches are not connected by trunks, you must configure a DVMRP tunnel between the switch endpoints.

See "Configuring VIFs for Tunnels" on page 15-38 for more information on configuring VIFs for tunnels. See Chapter 3, "Configuring IP Logical Ports and IP Servers" for more information on IP logical ports and IP server logical ports.

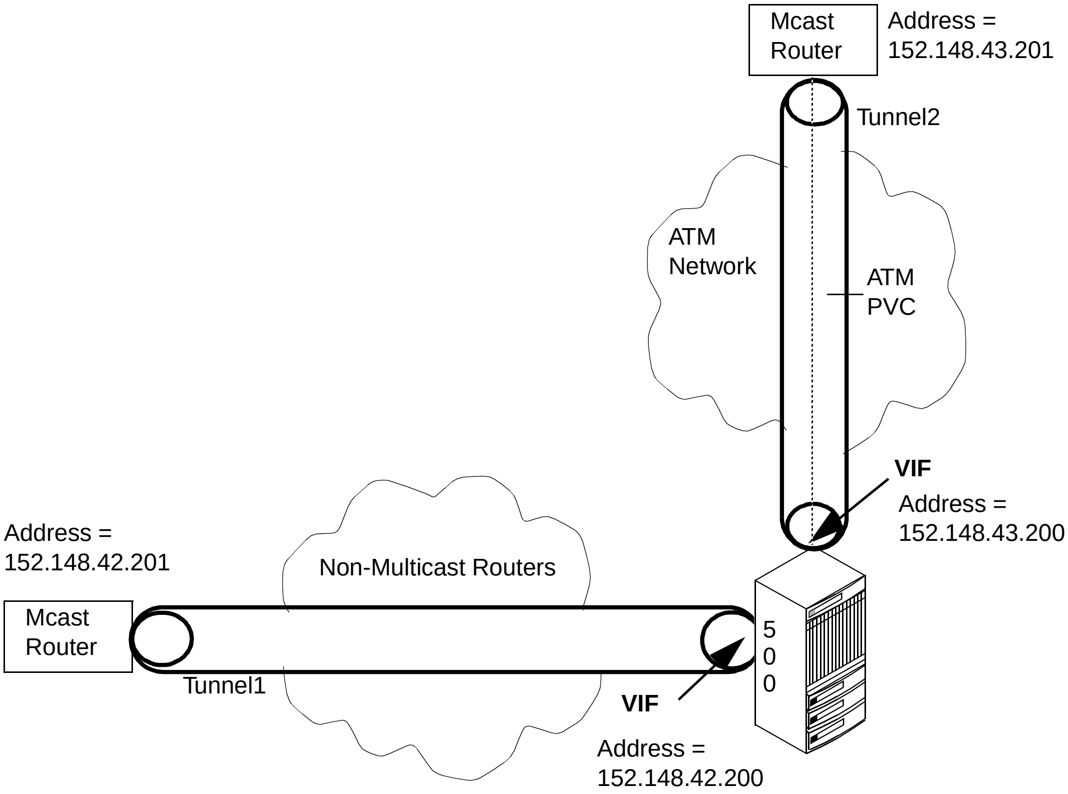

Figure 15-11. Sample IP Addresses for VIFs Associated with Tunnels

Table 15-1. Sample Local and Remote IP Addresses

VIF	Local IP Address	Remote IP Address
Tunnel1	152.148.42.200	152.148.42.201
Tunnel2	152.148.43.200	152.148.43.201

Identifying Scoped Boundaries

Multicast addresses in the range 239.0.0.0 to 239.255.255.255 are administratively scoped IP addresses. This range of addresses is reserved for use within private networks (such as a corporate enterprise network), but should not be used to send multicast traffic through public networks.

You can configure a specific VIF so that the switch does not forward private multicast traffic on it in either direction. To accomplish this task, NavisCore allows you to configure a range of administratively scoped IP addresses on a VIF-by-VIF basis. When you specify this address range, the VIF will not forward any multicast datagrams with a destination address that falls in the range. NavisCore allows you to specify multiple address ranges for a VIF.

Address ranges are expressed by specifying a base address and a mask. The base address specifies the lowest address in the range. The mask, combined with the base address, specifies the highest address in the range. For example, the following base address/mask combination tells the switch not to forward (on a given VIF) all multicast datagrams with destination addresses in the range 239.0.0.0 to 239.255.255.255:

Base Address — 239.0.0.0
Mask — 255.0.0.0

Figure 15-12 illustrates the above address/mask combination. The figure shows two private networks connected by a public network. For the VIF associated with the Ethernet logical port on Switch 2, the administrator has configured the address/mask combination shown above. This means that the switch will not forward any multicast datagrams with destination addresses in the 239.0.0.0-to-239.255.255.255 range in either direction (that is, datagrams originating from the private network to which Switch 2 is connected, or from Switch 1's private network).

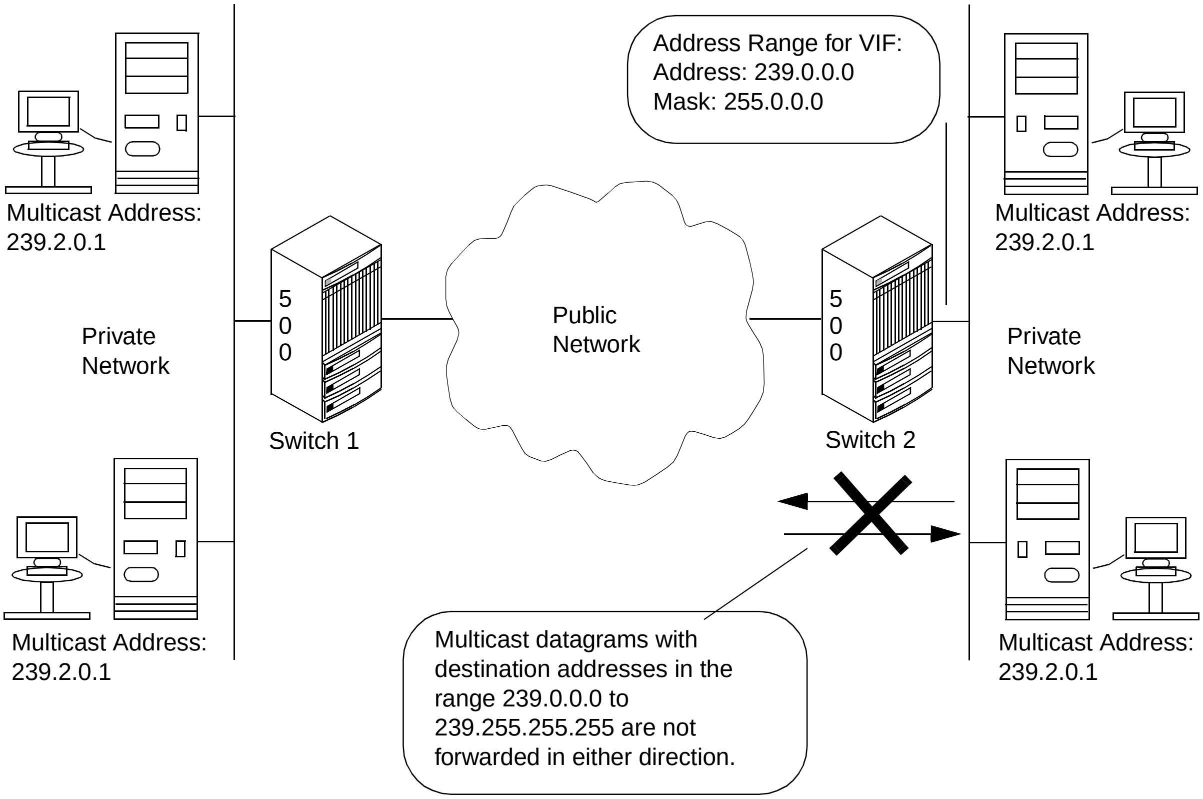

Figure 15-12. Sample Scoped Boundary

Planning Your MOSPF Configuration

To plan your MOSPF configuration, perform the following tasks:

- Verify IP OSPF router IDs
- Identify OSPF interfaces that must have multicast traffic forwarding configured
- Verify that multicast LSPs are enabled on switches
- Plan for DVMRP interoperability

The rest of this section describes these tasks.

Verify IP OSPF Router IDs

MOSPF is tightly integrated with the IP Navigator instance of OSPF. You must configure an IP OSPF router ID on either:

- Switches connected to trunks that belong to IP OSPF areas
- Switches that have at least one OSPF interface configured on an IP logical port

Before you perform any other configuration tasks, verify that you have met the IP OSPF router ID requirement. See "Planning Router IDs" on page 9-22 and "Configuring IP OSPF Router IDs" on page 9-42 for more information on configuring an IP OSPF router ID.

Identifying OSPF Interfaces

Make sure that OSPF interfaces are created on the IP logical ports that will handle MOSPF traffic. MOSPF is enabled on OSPF interfaces by default. To forward multicast traffic, make sure that the Multicast Forwarding parameter for the OSPF interface is set to the default value "Multicast" (in rare situations, you may want to set it to "Unicast"). Figure 15-13 shows IP logical ports with OSPF interfaces that handle MOSPF traffic.

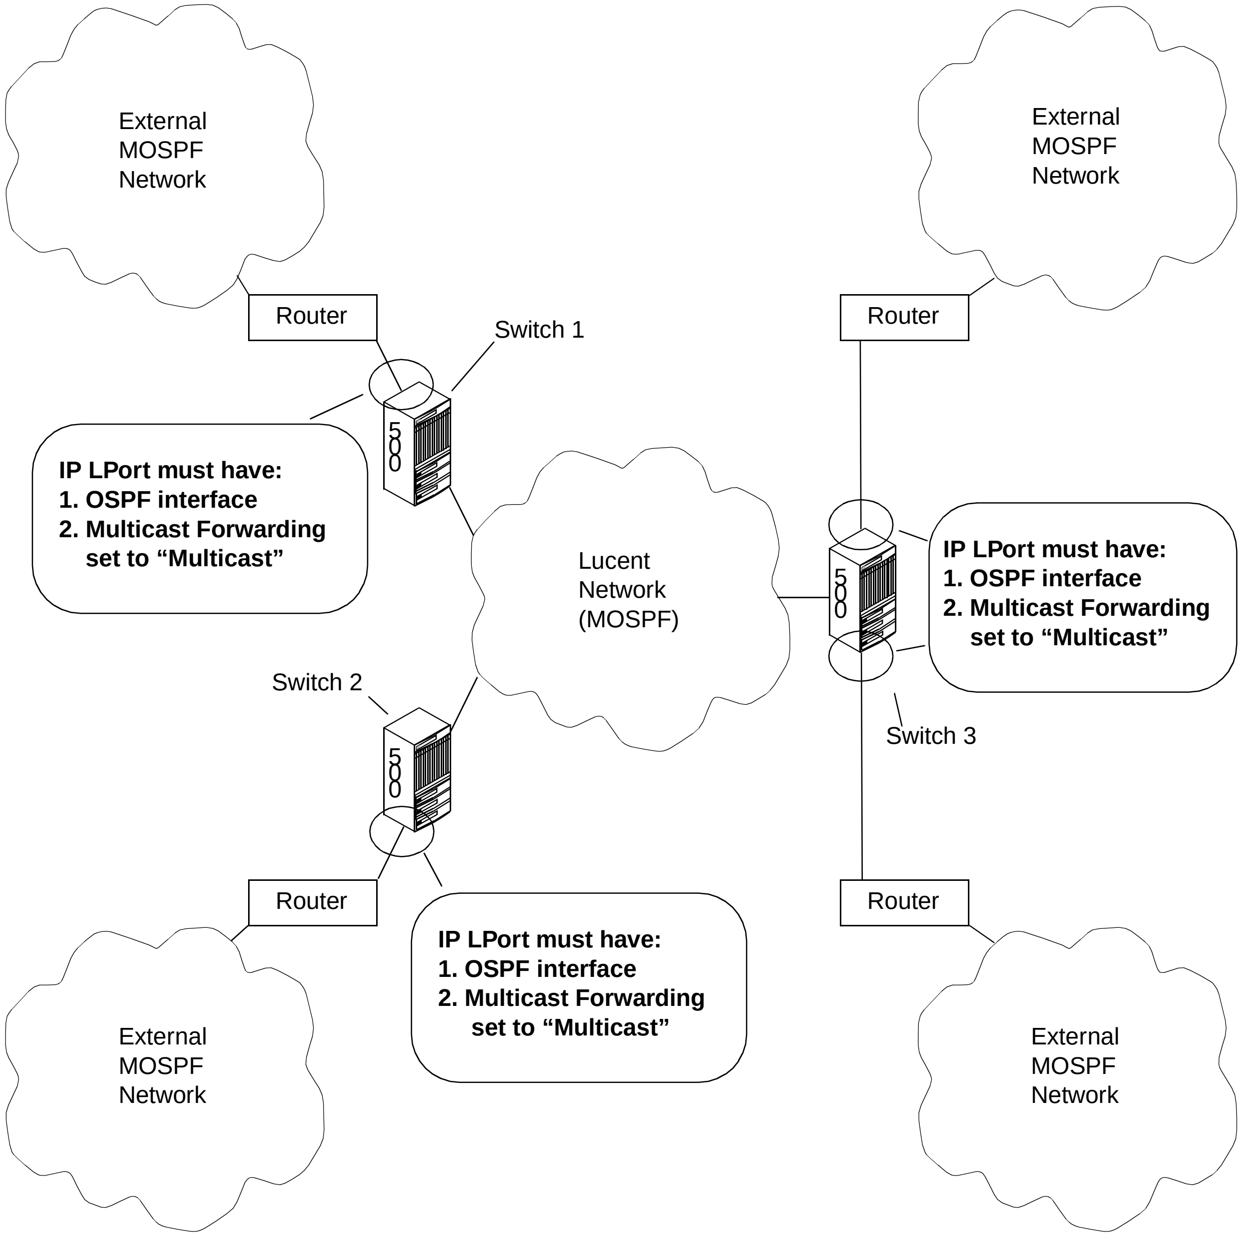

Figure 15-13. IP Logical Ports That Require OSPF Interfaces

For more information, see “Configuring Multicast Traffic Forwarding” on page 15-43.

Verifying That Multicast LSPs are Enabled

Make sure that multicast LSPs are enabled on all switches in the Lucent network that will forward multicast traffic. See "Enabling Multicast LSPs" on page 15-44 for more information.

Planning DVMRP Interoperability

To plan your network so that MOSPF and DVMRP can interoperate, perform the following tasks:

- Verify that you have configured one IP loopback address on each switch where DVMRP traffic accesses the Lucent network
- Identify route maps to export MOSPF routes to DVMRP

Verifying IP Loopback Addresses

In order for DVMRP traffic to pass through a Lucent network, you must configure IP loopback addresses on the ingress/egress switches (that is, the "edge" switches). On each of these switches, you configure one IP loopback address.

When you configure an IP loopback address, you enter the following:

IP Address — When you configure an IP loopback address to support DVMRP-MOSPF interoperability, always specify the IP OSPF router ID of the switch as the IP address. See "Planning Router IDs" on page 9-22 for more information on the IP OSPF router ID.

IP OSPF Area ID — When you specify an IP OSPF area ID, specify the IP OSPF area ID of any OSPF area of which the switch is a member (for example, the IP OSPF area ID of a trunk connected to the switch).

For example, in Figure 15-4, there are four "edge" switches that require an IP loopback address. All of these switches require you to configure an IP loopback address.

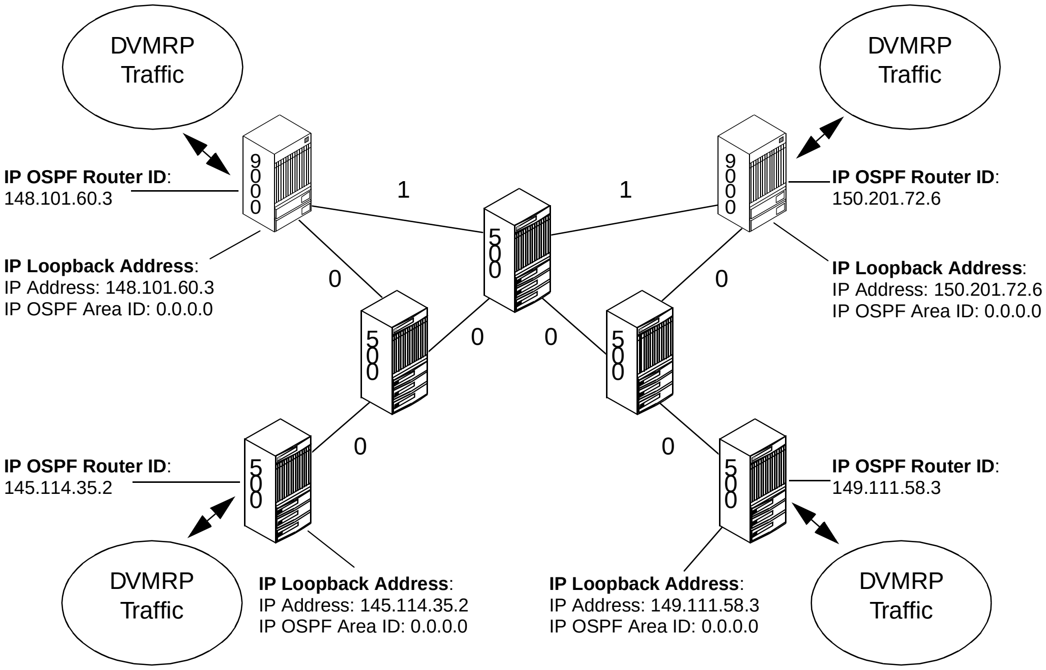
DVMRP Traffic
DVMRP Traffic
IP OSPF Router ID: 148.101.60.3
9000
1
1
500
9000
IP OSPF Router ID: 150.201.72.6
IP Loopback Address: IP Address: 148.101.60.3 IP OSPF Area ID: 0.0.0.0
0
0
IP Loopback Address: IP Address: 150.201.72.6 IP OSPF Area ID: 0.0.0.0
500
0
0
500
IP OSPF Router ID: 145.114.35.2
500
0
0
500
IP OSPF Router ID: 149.111.58.3
DVMRP Traffic
IP Loopback Address: IP Address: 145.114.35.2 IP OSPF Area ID: 0.0.0.0
IP Loopback Address: IP Address: 149.111.58.3 IP OSPF Area ID: 0.0.0.0
DVMRP Traffic

Figure 15-14. IP Loopback Address Requirements

See “Configuring IP Loopback Addresses” on page 8-31 for information on configuring an IP loopback address.

Identifying Route Maps to Export Routes to DVMRP

If DVMRP on a switch needs to know routes learned by MOSPF, you can use route maps to export routes from MOSPF to DVMRP. See “Exporting MOSPF Routes to DVMRP” on page 15-44 for more information.

Configuring IGMP

You configure IGMP on Fast Ethernet logical ports associated with LANs where multicast group members reside. To configure IGMP:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set all IP LPorts dialog box appears (see Figure 3-2 on page 3-6).
2. Select the switch where the Fast Ethernet logical port resides from the Switch Name list at the top of the dialog box. A list of IP logical ports configured on the switch appears in the LPort Name list.
3. Select the Fast Ethernet logical port.
4. Choose IP Parameters. The Set IP Parameters dialog box appears.

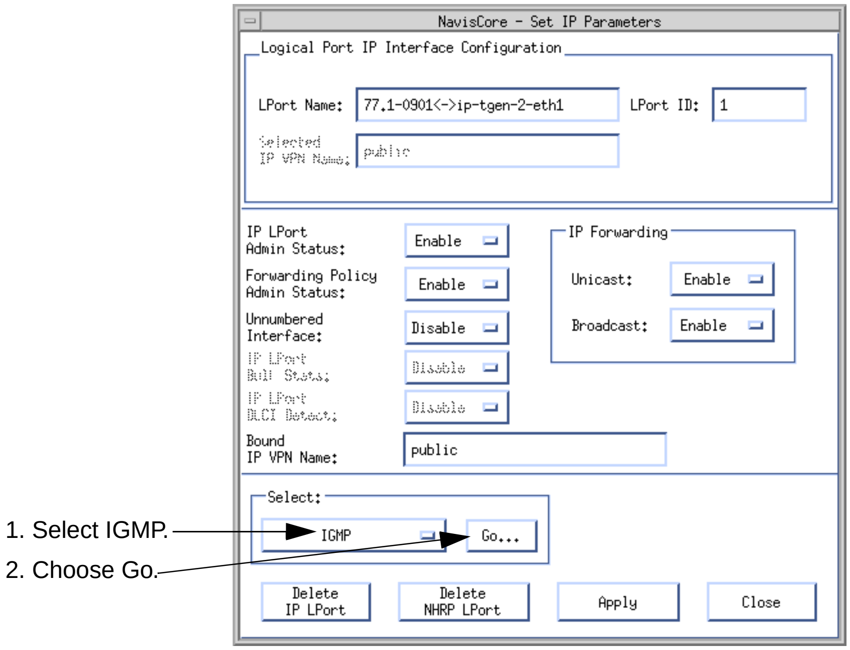

NavisCore - Set IP Parameters

Logical Port IP Interface Configuration

LPort Name: 77.1-0901<->ip-tgen-2-eth1 LPort ID: 1

Selected IP VPN Name: public

IP LPort Admin Status: Enable

Forwarding Policy Admin Status: Enable

Unnumbered Interface: Disable

IP LPort Bull Stats: Disable

IP LPort DLCI Detect: Disable

Bound IP VPN Name: public

IP Forwarding

Unicast: Enable

Broadcast: Enable

Select: IGMP

Figure 15-15. Set IP Parameters Dialog Box (With IGMP Selected)

5. Select IGMP.
6. Choose Go. The Set IGMP dialog box appears.

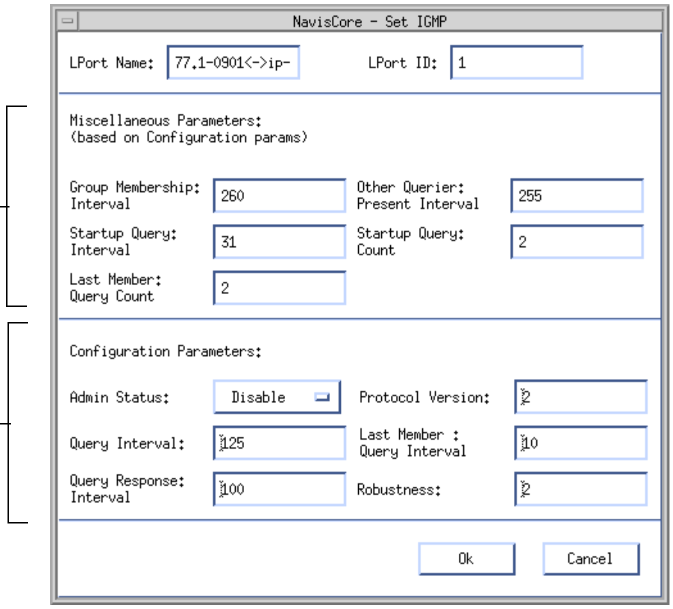

NavisCore - Set IGMP

LPort Name: 77.1-0901<->ip-

LPort ID: 1

Miscellaneous Parameters:
(based on Configuration params)

Group Membership: Interval: 260

Other Querier: Present Interval: 255

Startup Query: Interval: 31

Startup Query: Count: 2

Last Member: Query Count: 2

Configuration Parameters:

Admin Status: Disable

Protocol Version: 2

Query Interval: 125

Last Member : Query Interval: 10

Query Response: Interval: 100

Robustness: 2

Do not enter parameters here. These parameters are calculated based on parameters you enter in the bottom of the dialog box.

Enter parameters here.

Figure 15-16. Set IGMP Dialog Box

7. Enter the IGMP configuration parameters in the fields in the bottom half of the dialog box. You cannot enter IGMP configuration parameters in the fields in the top half of the dialog box, as they are calculated based on the configuration parameters you enter in the bottom half of the dialog box. Table 15-2 describes both the configuration parameters you enter and the configuration parameters that are calculated automatically.

Table 15-2. IGMP Configuration Parameters

Parameter	Description
Parameters You Enter	
Admin Status	Specify the administrative status of IGMP on the selected logical port. Choose Enable to enable IGMP on the port. Choose Disable (the default) to disable IGMP on the port. Enable IGMP on the port only if multicast group members are on the Ethernet LAN connected to the port.
Query Interval	Specify the number of seconds between general queries transmitted on the LAN associated with the selected logical port. To determine whether any multicast group members are on the LAN, the switch periodically transmits general queries to solicit multicast group membership information. Specify a value from 2 to 999 seconds (the default is 125). Keep in mind that, as you enter larger values, IGMP general queries are sent less often.

Table 15-2. IGMP Configuration Parameters (Continued)

Parameter	Description
Query Response Interval	Specify the maximum response time, in tenths of a second, that IGMP inserts into general query datagrams. When a multicast group member receives the general query datagram, it must respond within the amount of time specified by the maximum response time value in the datagram. Specify a value from 1 to 100 tenths of a second (the default is 100). The number of seconds you specify for the query response interval must be less than the number of seconds you specify for the query interval. For example, if you use the default (100 tenths of a second or 10 seconds), you must specify a value of at least 11 for the query interval. Keep in mind that, as you enter larger values, traffic on the LAN becomes less bursty because responses become spread out over a larger time interval.
Protocol Version	Specify the IGMP version (1 or 2) that is in use on the LAN (the default is 2).
Last Member Query Interval	Specify the maximum response time, in tenths of a second, that IGMP inserts into group specific query datagrams. When the switch receives a leave group message from a multicast host, and the host belongs to a multicast group with members reachable via the interface on which the leave group message is received, the switch sends group specific queries to the members of that multicast group. When a member of the multicast group receives the group specific query datagram, it must respond with a membership report within the amount of time specified by the maximum response time value in the datagram. If no member responds within the maximum response time, the routers assume that no members of the group are on the LAN. Specify a value from 1 to 100 tenths of a second (the default is 10). Keep in mind that, as you enter smaller values, you reduce the amount of time the switch has to detect the loss of the last member of a group on the LAN.
Robustness	Specify a value that allows you to tune for expected packet loss on the LAN. If you expect a significant amount of packet loss (for example, you have a high collision rate), increase this value. Specify a value from 2 to 10 (the default is 2).
Parameters Calculated Automatically	
Group Membership Interval	Specifies the amount of time (in seconds) that must elapse before the switch determines that no more members of a group are on the attached LAN. This timer is refreshed whenever the switch receives a membership report. This timer is calculated as follows: *(Robustness * Query Interval) + (Query Response Interval/10)*

Table 15-2. IGMP Configuration Parameters (Continued)

Parameter	Description
Startup Query Interval	Specifies the number of seconds between transmissions of general queries when the switch first starts up and assumes that it is a querier. This timer is calculated as follows: *Query Interval * .25*
Last Member Query Count	Specifies the number of group specific queries that the switch sends before it assumes that no members of a multicast group are currently on the LAN. The last member query count is equal to the robustness value.
Other Querier Present Interval	Specifies the number of seconds that the switch waits to receive a query message from another multicast router after the switch starts up. The switch assumes the role of querier if this timer expires before the switch receives a query message. This timer is calculated as follows: *(Robustness * Query Interval)* *+ (0.5 * (Query Response Interval/10))*
Startup Query Count	Specifies the number of queries that the switch sends upon startup. The startup query count is equal to the robustness variable.

8. Choose OK.

Configuring DVMRP VIFs

The way in which you configure DVMRP VIFs differs, depending on whether you are configuring DVMRP VIFs for:

- Fast Ethernet and PPP logical ports
- Tunnels

The rest of this section describes how to configure DVMRP VIFs for Fast Ethernet logical ports and PPP links, and for tunnels.

Configuring DVMRP VIFs for Fast Ethernet and PPP Logical Ports

To configure DVMRP VIFs for Fast Ethernet and PPP logical ports:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set all IP LPorts dialog box appears (see Figure 3-2 on page 3-6).
2. Select the switch where the Fast Ethernet logical port or PPP logical port resides from the Switch Name list at the top of the dialog box. A list of IP logical ports configured on the switch appears in the LPort Name list.
3. Select the Fast Ethernet or PPP logical port.
4. Choose IP Parameters. The Set IP Parameters dialog box appears.

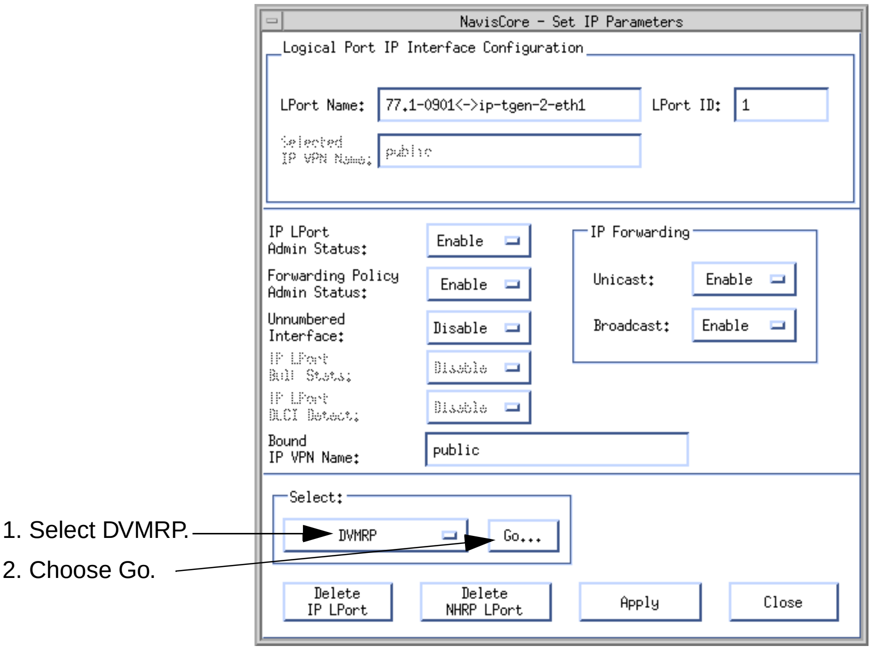
NavisCore - Set IP Parameters

Logical Port IP Interface Configuration

LPort Name: 77.1-0901<->ip-tgen-2-eth1
LPort ID: 1
Selected IP VPN Name: public

IP LPort Admin Status: Enable
Forwarding Policy Admin Status: Enable
Unnumbered Interface: Disable
IP LPort Bulk Stats: Disable
IP LPort DLCI Detect: Disable
Bound IP VPN Name: public

IP Forwarding
Unicast: Enable
Broadcast: Enable

Select: DVMRP

1. Select DVMRP.
2. Choose Go.

Figure 15-17. Set IP Parameters Dialog Box (With DVMRP Selected)

5. Select DVMRP.
6. Choose Go. The Set DVMRP Interface dialog box appears.

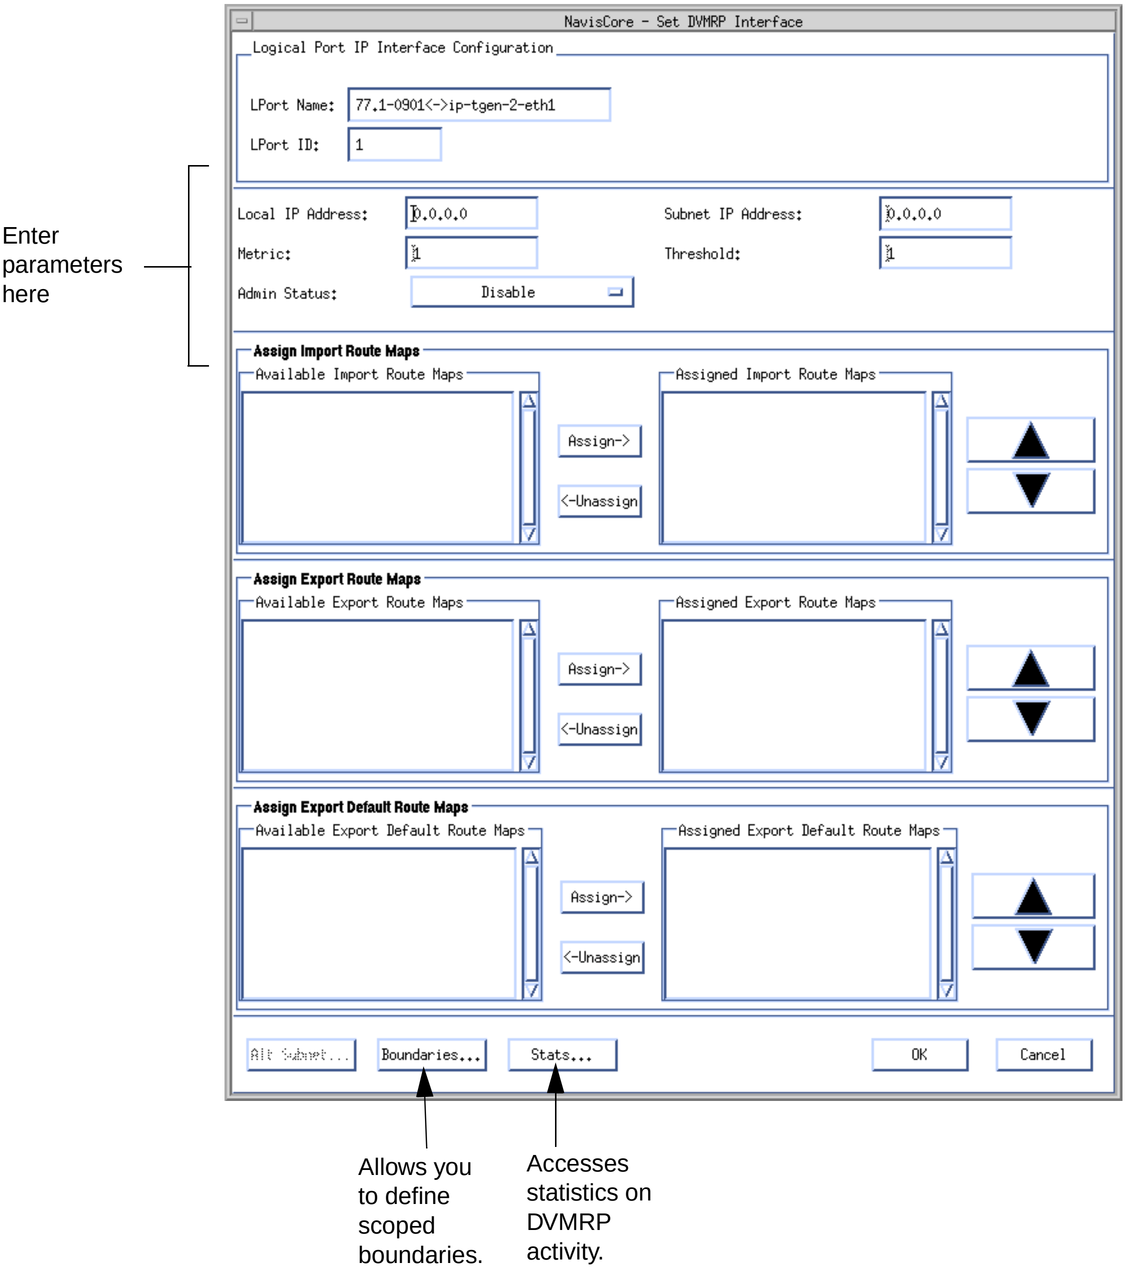

NavisCore - Set DVMRP Interface

Logical Port IP Interface Configuration

LPort Name: 77.1-0901<->ip-tgen-2-eth1

LPort ID: 1

Local IP Address: 0.0.0.0

Subnet IP Address: 0.0.0.0

Metric: 1

Threshold: 1

Admin Status: Disable

Assign Import Route Maps
Available Import Route Maps
Assigned Import Route Maps
Assign->
<-Unassign
Assign Export Route Maps
Available Export Route Maps
Assigned Export Route Maps
Assign->
<-Unassign
Assign Export Default Route Maps
Available Export Default Route Maps
Assigned Export Default Route Maps
Assign->
<-Unassign
Alt Subnet...
Boundaries...
Stats...
OK
Cancel

Enter parameters here

Allows you to define scoped boundaries.

Accesses statistics on DVMRP activity.

Figure 15-18. Set DVMRP Interface Dialog Box

Table 15-3 describes the buttons on the Set DVMRP Interface dialog box.

Table 15-3. Set DVMRP Interface Dialog Box Buttons

Button	Function
Alt Subnet	Not supported.
Boundaries	Allows you to define scoped boundaries. See "Configuring Scoped Boundaries for Ethernet and PPP Ports" on page 15-37 for more information.
Stats	Displays statistics on DVMRP activity. See the *NavisCore Diagnostics Guide* for more information.
OK	Saves your DVMRP configuration changes.
Cancel	Exits the dialog box.

7. Enter the DVMRP interface configuration parameters in the fields on the dialog box. Table 15-4 describes these parameters.

Table 15-4. DVMRP Interface Configuration Parameters

Parameter	Description
Local IP Address	Specify the IP address assigned to the associated Fast Ethernet IP logical port interface or PPP IP logical port interface (e.g., 152.148.12.107).
Subnet IP Address	The subnetwork address assigned to the Fast Ethernet LAN or PPP IP logical port interface (e.g., 152.148.12.0). Do not specify this address. NavisCore calculates it for you.
Metric	Specify the metric for the interface. Although you may specify a number from 1 to 32, use the default (1). Metrics are hop counts that assign costs to routes. As the metric for a route increases, so does its cost. If you have multiple routes to the same destination, the one with the lowest metric is preferred. As routes are propagated through the network, a route's metric continues to increment each time it passes through an interface, such as an interface to an Ethernet LAN. For example, if you set the metric of the interface you are configuring to 1, then each route that passes through the interface will have its metric incremented by one. For Ethernet LANs, the default metric (1) is sufficient in most cases.

Table 15-4. DVMRP Interface Configuration Parameters (Continued)

Parameter	Description
Threshold	Specify the Time-to-Live (TTL) threshold. This threshold allows you to control the scope of multicast transmission. The switch will only forward a multicast datagram across an interface if the value of the TTL field in the IP header is greater than the TTL threshold assigned to the interface. For example, if you set the TTL threshold to 5, the switch will only forward multicast datagrams that have a TTL field greater than 5 across the interface. The default is 1, which allows most multicast datagrams to be forwarded. Valid values range from 1 to 255. To prevent all multicast datagrams from being forwarded across the interface, set the TTL threshold to a high value, such as 255.
Admin Status	Specify the administrative status of DVMRP on the selected Ethernet logical port. Choose Enable to enable DVMRP on the port. Choose Disable (the default) to disable DVMRP on the port. Enable DVMRP on the port only if it must forward multicast traffic.
Assign Import Route Maps	
Available Import Route Maps	The import route maps that are available for assignment to this DVMRP interface. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, “Configuring Route Policies.” To display the parameters for any listed route map, double-click on the map.
Assigned Import Route Maps	The import route maps that are assigned to this DVMRP interface. All incoming routes on this interface are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.

Table 15-4. DVMRP Interface Configuration Parameters (Continued)

Parameter	Description
Assign Export Route Map	
Available Export Route Maps	The export route maps that are available for assignment to this DVMRP interface. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, “Configuring Route Policies.” To display the parameters for any listed route map, double-click on the map.
Assigned Export Route Maps	The export route maps that are assigned to this DVMRP interface. All outgoing routes on this interface are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.
Assign Export Default Route Maps	
Available Export Default Route Maps	The export default route maps that are available for assignment to this DVMRP interface. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. For more information about creating route maps, see Chapter 11, “Configuring Route Policies.” To display the parameters for any listed route map, double-click on the map.
Assigned Export Default Route Maps	The export default route maps that are assigned to this DVMRP interface. All outgoing routes on this interface are filtered using the assigned route maps in the listed sequence. Use the Assign button to move a route map from the Available to the Assigned list. Use the Unassign button to move a route map from the Assigned to the Available list. Use the up and down arrows to change the sequence of the route maps in the Assigned list. IP Navigator executes the route maps in the sequence that they are ordered in this list. Route maps *should be ordered from most specific to least specific.* To display the parameters for any listed route map, double-click on the map.

8. When you finish configuring the DVMRP VIF, you may choose OK to save your changes and exit, or configure scoped boundaries. The next section describes how to configure scoped boundaries.

Configuring Scoped Boundaries for Ethernet and PPP Ports

You can configure one or more scoped boundaries for each VIF associated with an Ethernet logical port or a PPP logical port. To configure a scoped boundary:

1. From the Set DVMRP Interface dialog box (see Figure 15-18 on page 15-33), choose Boundaries. The IP Multicast Scoped Boundary Table dialog box appears (see Figure 15-19).

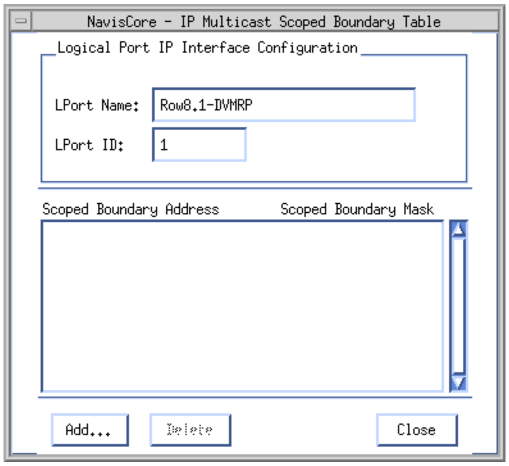
NavisCore - IP Multicast Scoped Boundary Table

Logical Port IP Interface Configuration

LPort Name: Row8.1-DVMRP

LPort ID: 1

Figure 15-19. IP Multicast Scoped Boundary Table Dialog Box (Ethernet/PPP)

The IP Multicast Scoped Boundary Table dialog box lists all of the defined scoped boundaries.

2. Choose Add. The Add IP Multicast Scoped Boundary Address dialog box appears (see Figure 15-20).

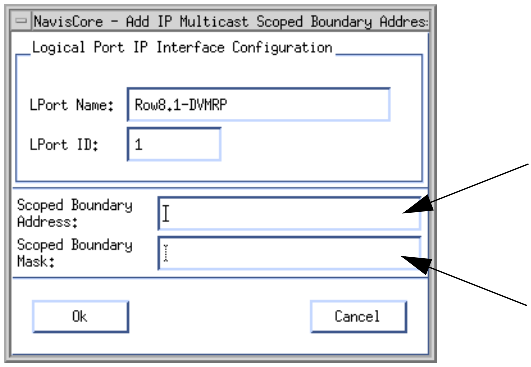
NavisCore - Add IP Multicast Scoped Boundary Addres

Logical Port IP Interface Configuration

LPort Name: Row8.1-DVMRP

LPort ID: 1

Enter the base address of the boundary here.
For example: 239.0.0.0

Enter the mask for the boundary here.
For example: 255.0.0.0

Figure 15-20. Add IP Multicast Scoped Boundary Address Dialog Box (Ethernet/PPP)

3. Enter the base address of the boundary in the Scoped Boundary Address field.

4. In the Scoped Boundary Mask field, enter the mask that, combined with the base address, defines the full range of destination addresses within the scoped boundary. See "Identifying Scoped Boundaries" on page 15-22 if you need more information.
5. Choose OK. The IP Multicast Scoped Boundary Table dialog box appears, displaying the address and mask of the newly added scoped boundary.

To delete a scoped boundary:

1. From the IP Multicast Scoped Boundary Table dialog box (see Figure 15-19 on page 15-37), select the scoped boundary that you want to delete from the list of scoped boundaries.
2. Choose Delete.
3. Choose OK when prompted.

Configuring VIFs for Tunnels

To configure DVMRP VIFs for tunnels:

1. Select the appropriate switch from the network map.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set DVMRP Tunnels. The DVMRP Virtual Interface Parameters dialog box appears (see Figure 15-21).

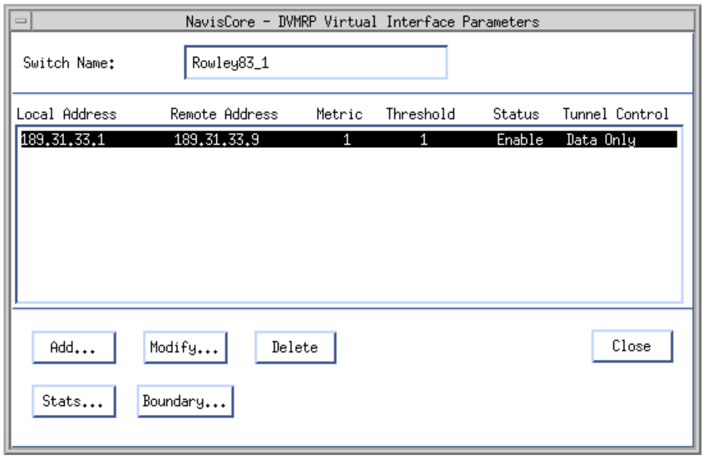

Switch Name: Rowley83_1

Local Address	Remote Address	Metric	Threshold	Status	Tunnel Control
189.31.33.1	189.31.33.9	1	1	Enable	Data Only

Figure 15-21. DVMRP Virtual Interface Parameters Dialog Box

The DVMRP Virtual Interface Parameters dialog box lists all of the VIFs you have configured for tunnels links. The buttons on the dialog box are described in Table 15-5. The fields on the dialog box are described in Table 15-6.

Table 15-5. DVMRP Virtual Interface Parameters Dialog Box Buttons

Button	Function
Add	Allows you to add a tunnel VIF.
Modify	Allows you to modify a tunnel VIF.
Delete	Allows you to delete a tunnel VIF.
Stats	Displays statistics on DVMRP activity. See the *NavisCore Diagnostics Guide* for more information.
Boundary	Allows you to define scoped boundaries for the tunnel VIF. See "Configuring Scoped Boundaries for Tunnels" on page 15-41 for more information.
Close	Exits the dialog box.

3. Choose Add. The Add DVMRP Tunnel dialog box appears.

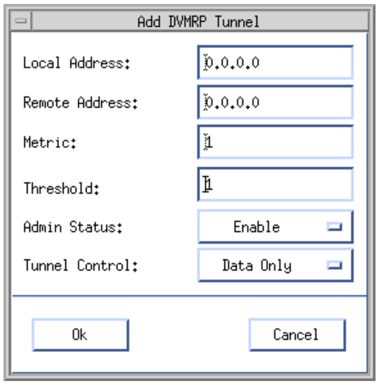

Local Address: 0.0.0.0
Remote Address: 0.0.0.0
Metric: 1
Threshold: 1
Admin Status: Enable
Tunnel Control: Data Only

Figure 15-22. Add DVMRP Tunnel Dialog Box

4. Complete the fields described in Table 15-6.

Table 15-6. DVMRP Interface Configuration Parameters

Parameter	Description
Local Address	Specify the IP address assigned to the IP interface that the switch uses to transmit multicast datagrams across the tunnel. See Chapter 3, "Configuring IP Logical Ports and IP Servers" for more information on IP interfaces.
Remote Address	Specify the IP address of the node (e.g., a multicast router) at the other end of the tunnel.
Metric	Specify the metric for the interface. You may specify a number from 1 to 32. The default (1) should suffice in most cases unless you expect significant network congestion on the interface. Metrics are hop counts that assign costs to routes. As the metric for a route increases, so does its cost. If you have multiple routes to the same destination, the one with the lowest metric is preferred. As routes are propagated through the network, a route's metric continues to increment each time it passes through an interface, such as an interface to a tunnel. For example, if you set the metric of the interface you are configuring to 1, then each route that passes through the interface will have its metric incremented by one.
Threshold	Specify the Time-to-Live (TTL) threshold. This threshold allows you to control the scope of multicast transmission. The switch will only forward a multicast datagram across an interface if the value of the TTL field in the IP header is greater than the TTL threshold assigned to the interface. For example, if you set the TTL threshold to 5, the switch will only forward multicast datagrams that have a TTL field greater than 5 across the interface. The default is 1, which allows most multicast datagrams to be forwarded. Valid values range from 1 to 255. To prevent all multicast datagrams from being forwarded across the interface, set the TTL threshold to a high value, such as 255.
Admin Status	Specify the administrative status of DVMRP on the tunnel. Choose Enable (the default) to enable DVMRP on the tunnel. Choose Disable to disable DVMRP.
Tunnel Control	Specify how you want DVMRP messages (e.g., prune messages) and multicast packets to be sent through the tunnel: *Data Only* – (default) DVMRP messages are sent as unicast packets (that is, no IP-in-IP encapsulation) through the tunnel. All other packets use IP-in-IP encapsulation. *All Packets* – All packets (including DVMRP) sent through the tunnel use IP-in-IP encapsulation.

5. Choose OK when you finish entering parameters. The DVMRP Virtual Interface Parameters dialog box appears (see Figure 15-21 on page 15-38), displaying the new VIF you just added.

You can also modify VIF parameters and delete VIFs. To modify a VIF:

1. From the DVMRP Virtual Interface Parameters dialog box (see Figure 15-21 on page 15-38), select the VIF you want to modify.
2. Choose Modify. The Modify DVMRP Tunnel dialog box appears. This dialog box is similar in appearance to the Add DVMRP Tunnel dialog box (see Figure 15-22 on page 15-39).
3. Modify the appropriate parameters. See Table 15-6 on page 15-40 for descriptions of the parameters.
4. Choose OK when you are finished.

To delete a VIF:

1. From the DVMRP Virtual Interface Parameters dialog box (see Figure 15-21 on page 15-38), select the VIF you want to delete.
2. Choose Delete.
3. Choose OK when prompted.

Configuring Scoped Boundaries for Tunnels

Optionally, you can configure scoped boundaries for the tunnel VIFs you add. To configure a scoped boundary for a tunnel:

1. From the DVMRP Virtual Interface Parameters dialog box (see Figure 15-21 on page 15-38), select the tunnel VIF for which you want to define a scoped boundary.
2. Choose Boundary. The IP Mcast Scoped Boundary Table dialog box appears (see Figure 15-23).

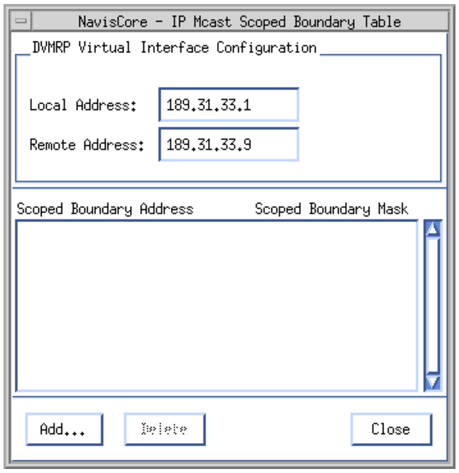

NavisCore - IP Mcast Scoped Boundary Table

DVMRP Virtual Interface Configuration

Local Address: 189.31.33.1

Remote Address: 189.31.33.9

Figure 15-23. IP Mcast Scoped Boundary Table Dialog Box (Tunnels)

3. Choose Add. The Add IP Mcast Boundary Interface Dialog Box appears (see Figure 15-24).

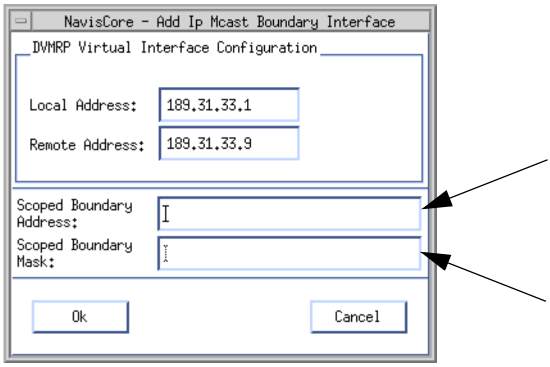

NavisCore - Add Ip Mcast Boundary Interface

DVMRP Virtual Interface Configuration

Local Address: 189.31.33.1

Remote Address: 189.31.33.9

Scoped Boundary Address:

Scoped Boundary Mask:

Enter the base address of the boundary here.
For example: 239.0.0.0

Enter the mask for the boundary here.
For example: 255.0.0.0

Figure 15-24. Add IP Mcast Boundary Interface Dialog Box (Tunnels)

4. Enter the base address of the boundary in the Scoped Boundary Address field.
5. In the Scoped Boundary Mask field, enter the mask that, combined with the base address, defines the full range of destination addresses within the scoped boundary. See "Identifying Scoped Boundaries" on page 15-22 if you need more information.
6. Choose OK. The IP Mcast Scoped Boundary Table dialog box appears, displaying the address and mask of the newly added scoped boundary.

To delete a scoped boundary:

1. From the IP Mcast Scoped Boundary Table dialog box (see Figure 15-23 on page 15-41), select the scoped boundary that you want to delete from the list of scoped boundaries.
2. Choose Delete.
3. Choose OK when prompted.

Configuring MOSPF

When you configure MOSPF, you configure multicast traffic forwarding (at the data link level) on the necessary OSPF interfaces, export MOSPF routes to DVMRP, and verify the multicast LSPs are enabled. The rest of this section tells you how to perform these tasks.

Configuring Multicast Traffic Forwarding

MOSPF is automatically enabled when you enable OSPF (that is, you have configured an OSPF interface). However, you have to specify how the switch forwards multicast traffic on the data link level using the OSPF interface, as follows:

1. Access the Set IP Interface Addresses dialog box for the IP logical port. See "Setting the IP Interface Address" on page 3-14 for more information on accessing this dialog box.
2. Do one of the following:
   - If you are adding OSPF parameters to the IP logical port, choose Add OSPF. The Add OSPF Interface dialog box appears (see Figure 15-25). See "Configuring IP OSPF on the IP Logical Port" on page 9-30 for more information on this dialog box.

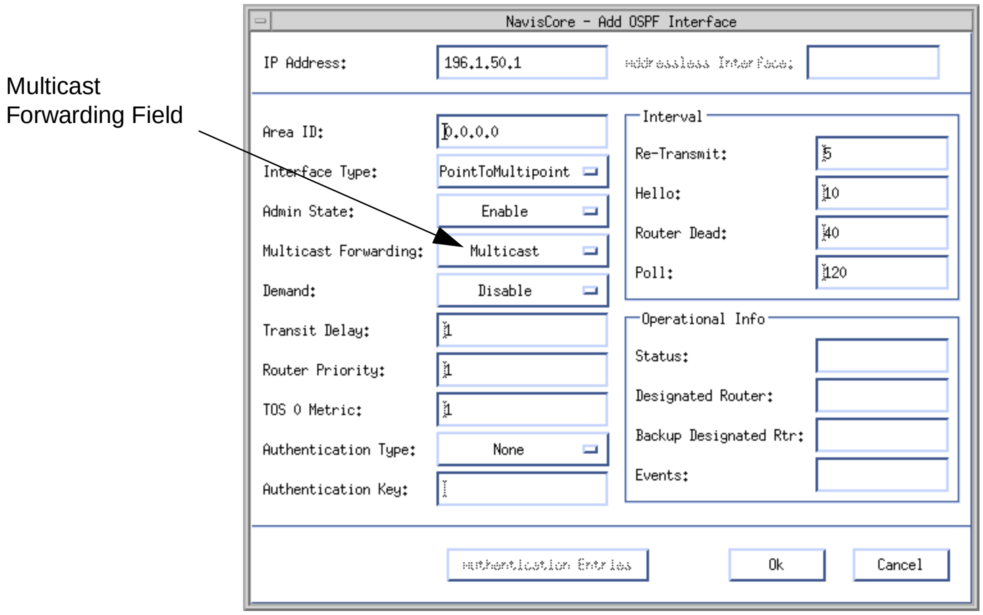

Multicast Forwarding Field

NavisCore - Add OSPF Interface

IP Address: 196.1.50.1
Addressless Interface:

Area ID: 0.0.0.0
Interface Type: PointToMultipoint
Admin State: Enable
Multicast Forwarding: Multicast
Demand: Disable
Transit Delay: 1
Router Priority: 1
TOS 0 Metric: 1
Authentication Type: None
Authentication Key:

Interval
Re-Transmit: 5
Hello: 10
Router Dead: 40
Poll: 120

Operational Info
Status:
Designated Router:
Backup Designated Rtr:
Events:

Figure 15-25. Add OSPF Interface Dialog Box (With Multicast Selected)

   - If you are modifying the existing OSPF parameters for the logical port, choose Modify OSPF. The Modify OSPF Interface dialog box appears. This dialog box is almost identical in appearance to the Add OSPF Interface dialog box.

3. In the Multicast Forwarding field, specify one of the following:

Multicast — (default) The OSPF interface forwards multicast traffic to a multicast data link address. Do not change the default unless you want to block multicast traffic or (in rare circumstances) forward multicast traffic to a unicast data link address.

Unicast — The OSPF interface forwards multicast traffic to a unicast data link address (Ethernet MAC address, Frame Relay DLCI, ATM VPI/VCI, etc.).

Blocked — The OSPF interface does not forward multicast traffic (but MOSPF continues to run).

4. When you finish setting parameters, choose OK.

Exporting MOSPF Routes to DVMRP

To export routes from MOSPF to DVMRP, you set up a route map using MOSPF as the source (that is, the “From” choice) and DVMRP as the destination (that is, the “To” choice). See Chapter 11, “Configuring Route Policies” for more information.

Enabling Multicast LSPs

MOSPF uses multicast LSPs to forward traffic through a Lucent network. You need to verify that multicast LSPs are enabled on all the switches in the network (they are enabled by default). Multicast LSPs are enabled on the Set IP Parameters dialog box for the switch. See Chapter 12, “Configuring Label Switched Paths” for more information on multicast LSPs.

16

Configuring IP Virtual Private Networks

An IP Virtual Private Network (VPN) is a collection of IP network resources that a public carrier or service provider reserves for private use. Typical applications of IP VPNs include:

Corporate Intranets — A network within an enterprise, which may consist of many interlinked local area networks and may also use leased-lines in the wide-area network. A corporate intranet may or may not include connections through one or more gateways to the outside Internet. The main purpose of an intranet is, in most cases, to enable employees to share company information and computing resources. An intranet can also be used to facilitate work groups and teleconferencing.

Corporate Extranets — Secure commerce-enabled networks that electronically link distributed organizations or individuals over the Internet in a public, semi-public, or private forum. These emerging networks establish *virtual firewalls* to extend the benefits of a company's Intranet and enable collaborative business applications across multiple organizations. Secure, collaborative, and interactive workspaces serve to connect companies and their customers, suppliers, and other stakeholders, and produce efficiencies in such representative business models as electronic commerce collaborative publishing, and supply-chain management.

The rest of this chapter explains how to plan for and configure IP VPNs.

Understanding an IP VPN

In a traditional IP enterprise network, all resources are owned and controlled by a single organization. To users within the organization, the network appears as a separate routing domain, such as the one shown in Figure 16-1.

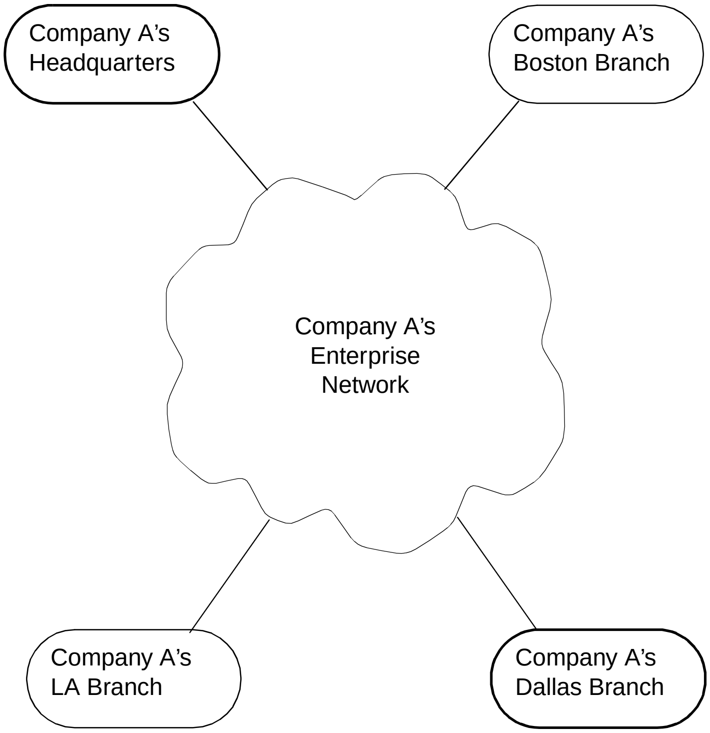
Company A's Headquarters
Company A's Boston Branch
Company A's Enterprise Network
Company A's LA Branch
Company A's Dallas Branch

Figure 16-1. Traditional IP Enterprise Network

When a public carrier or service provider reserves resources for IP VPNs, each IP VPN has its own view — that is, users of the IP VPN only see the resources reserved for them. Although multiple VPNs may share the same physical topology, each VPN appears to its users as if it were a separate routing domain.

To a network manager, an IP VPN also appears as a separate network. NavisCore allows the network manager to select a specific VPN to be managed. This allows the network manager to see only those resources, such as routes and route maps, configured for the VPN.

Figure 16-2 shows two customers — Customer A and Customer B — sharing a common public physical topology to link their respective headquarters and branch offices through IP VPNs. The customer's equipment accesses the network at switches at the edge of the Lucent network.

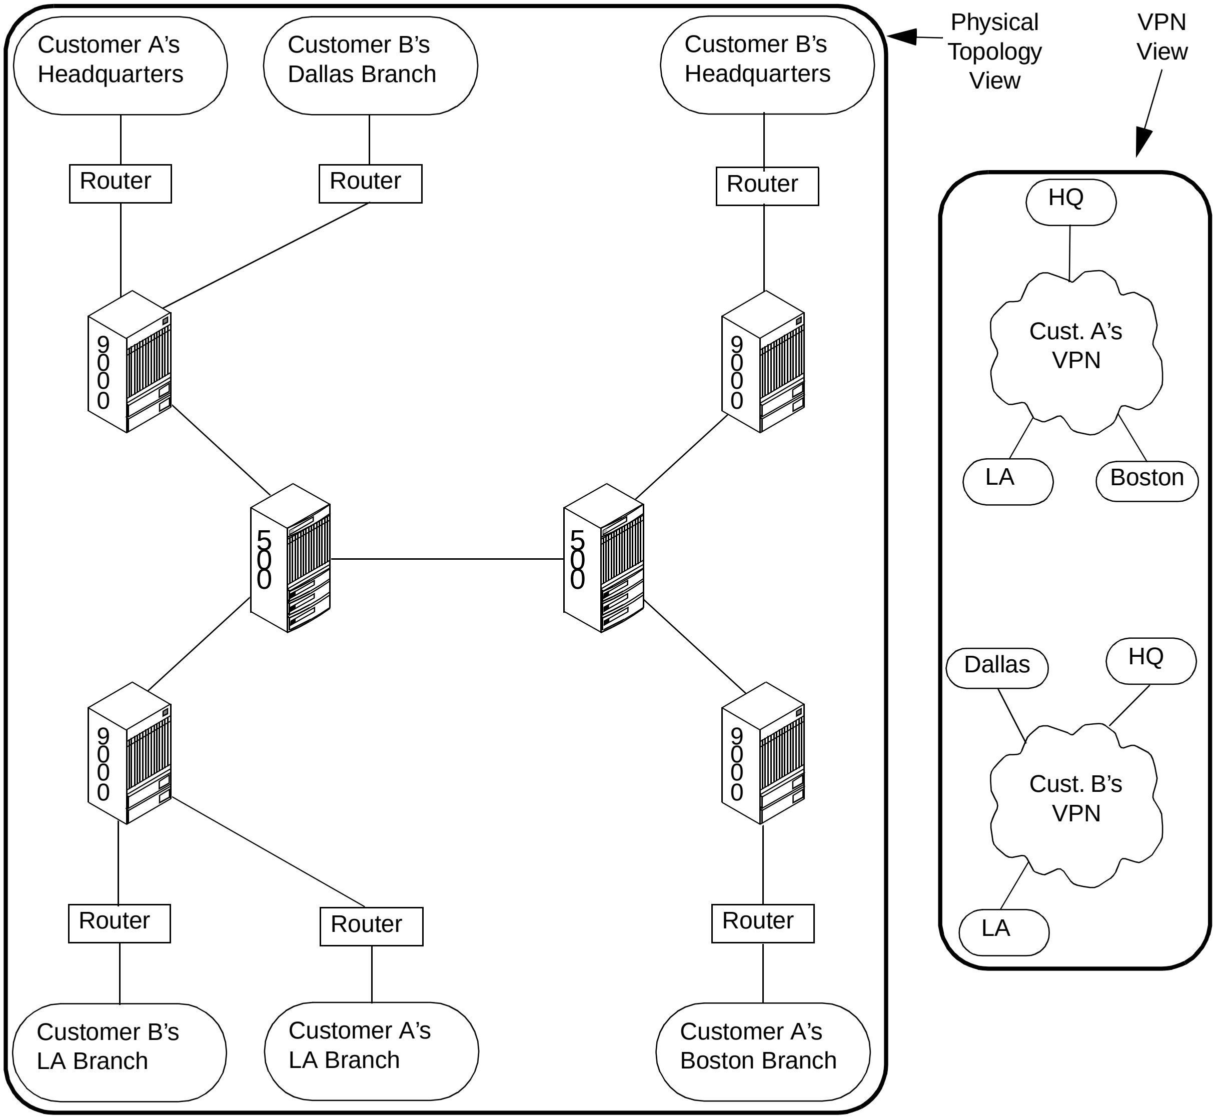

Figure 16-2. Multiple IP VPNs

IP Navigator supports IP VPNs through the implementation of *virtual routers* on Lucent switches. Unlike a physical router, a virtual router exists as a logical entity in IP Navigator switch software. However, like a physical router, a virtual router performs routing functions, such as IP packet forwarding.

An IP VPN has its own virtual router on each switch that it uses to access the network (that is, switches at the edge of the Lucent network). Like a physical router, a virtual router has its own set of IP resources, such as its own private routing table, ARP cache, and route maps. Virtual routers reside on edge switches only. For example, the edge switches in Figure 16-2 (the B-STDX 9000 switches) would maintain virtual routers, but the switches inside the Lucent network (the CBX 500 switches) would not.

An edge switch that provides three different IP VPNs with network access would maintain three virtual routers, one for each IP VPN. Each virtual router has its own routing tables, ARP cache, and so on.

To understand how IP VPNs work, you need to understand how IP VPN traffic accesses the Lucent network and how IP VPN traffic is routed through the Lucent network. The rest of this section explains how Lucent switches handle IP VPN traffic.

How IP VPN Traffic Accesses the Lucent Network

On a Lucent switch, IP logical ports provide IP VPNs with network access, as follows:

- For IP logical ports that use Frame Relay as the data link protocol, each IP VPN that uses the logical port has one or more DLCIs. You associate each DLCI with a particular IP VPN.
- For IP server logical ports and IP logical ports that use ATM as the data link protocol, each IP VPN that uses the port has one or more VPI/VCIs. For IP server logical ports, keep in mind that you must create an IP server PVC for each VPI/VCI you want to assign to an IP VPN. You associate each VPI/VCI with a particular IP VPN.
- For IP logical ports that use Ethernet or PPP as the data link protocol, a single IP VPN owns the port. This means that the Ethernet or PPP interface can be assigned to one (and only one) IP VPN.

When you assign a Frame Relay DLCI or an ATM VPI/VCI on a switch to an IP VPN for the first time, you create a virtual router on that switch for that IP VPN. (For Ethernet and PPP interfaces, you assign the entire IP logical port to an IP VPN, and thereby create the virtual router.) The virtual router then performs IP routing functions exclusively for the IP VPN on that switch. Figure 16-3 illustrates virtual routers.

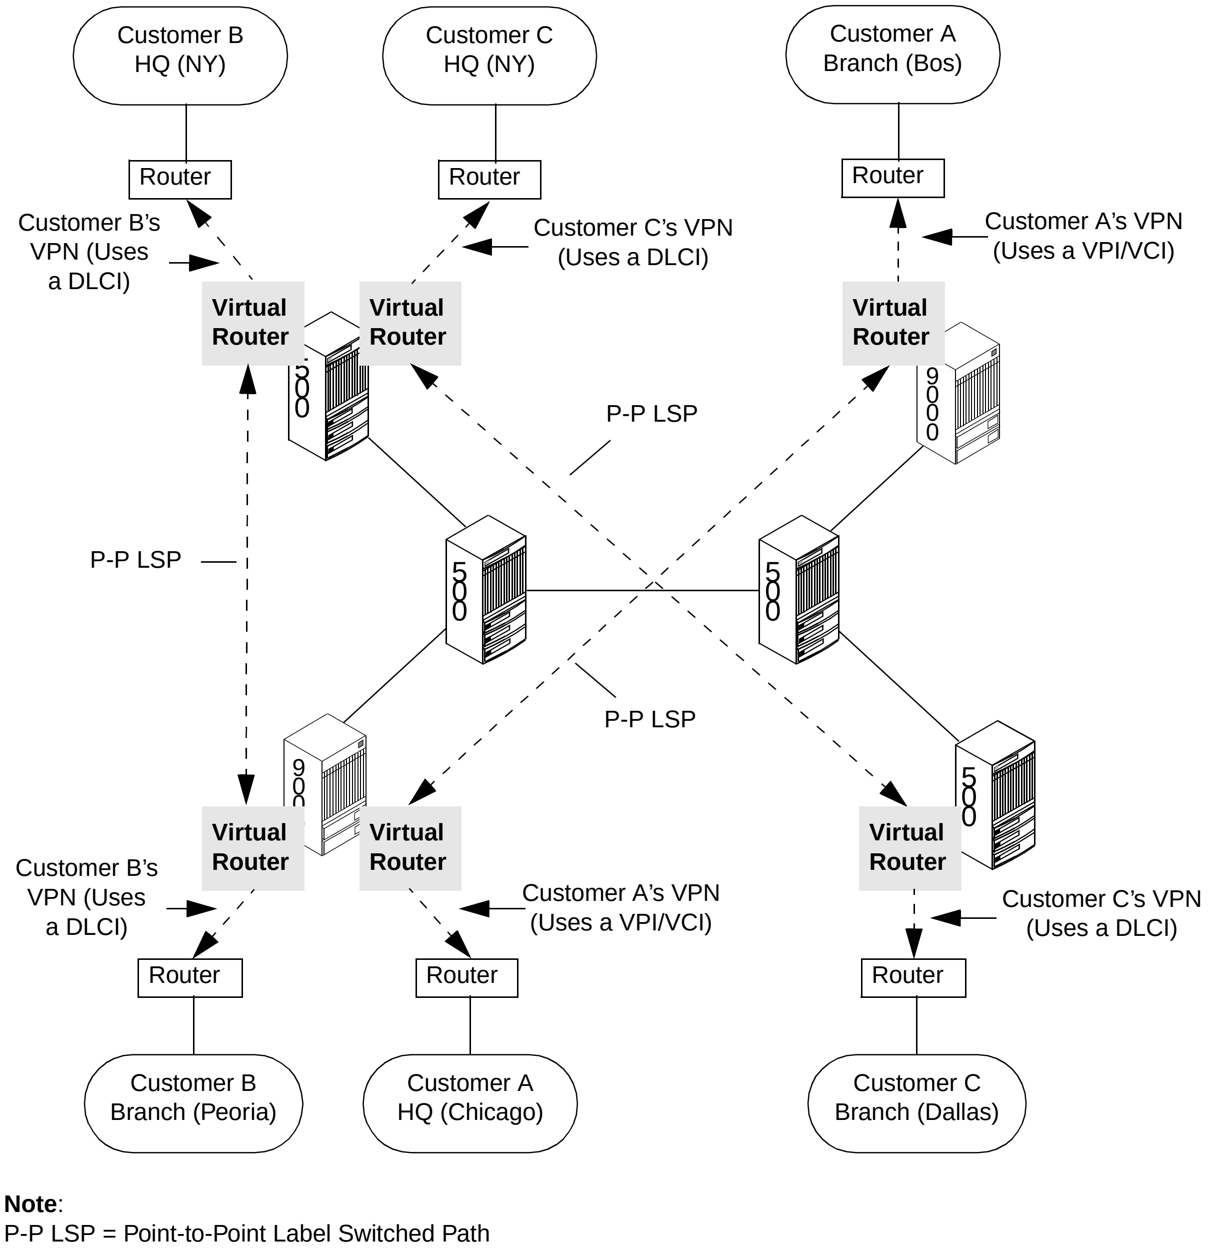

Figure 16-3. Virtual Routers

As Figure 16-3 illustrates, IP VPN users access the IP VPN via a virtual router at the edge of the Lucent network. The virtual router uses an Ethernet interface, a PPP interface, a DLCI, or a VPI/VCI on an ingress IP logical port or IP server logical port. The Ethernet interface, PPP interface, DLCI, or VPI/VCI is bound to an ingress IP interface.

In Figure 16-3, also notice that IP VPNs use point-to-point label switched paths (LSPs) to transport traffic through the Lucent network between virtual routers. A point-to-point LSP allows IP Navigator to directly forward IP packets between two Lucent switches. All IP VPN traffic between the virtual routers on the two switches travels through the point-to-point LSP. See "How IP VPN Traffic is Routed Through the Lucent Network" for more information on point-to-point LSPs.

How IP VPN Traffic is Routed Through the Lucent Network

While ingress IP interfaces allow IP VPN traffic to access the Lucent network, *VPN cloud IP interfaces* allow IP VPN traffic to pass through the Lucent network. For a specific IP VPN, these interfaces are created automatically on the CP or SP when the IP VPN's virtual router is added to the switch (that is, when you assign a DLCI, VPI/VCI, Ethernet logical port, or PPP logical port to an IP VPN). The interfaces become active when you make the switch a trunk endpoint.

VPN cloud IP interfaces are configured just like other IP interfaces. Like ingress IP interfaces, you create them manually and assign each one an IP address, an MTU size, and so on.

► You cannot define an IP VPN cloud IP interface on a switch until you configure an IP loopback address for the public IP VPN on that switch. See "Identify IP OSPF Router ID and Loopback Address Requirements" on page 16-21 for more information.

IP VPN cloud interfaces are associated with a special logical port called the *IP VPN cloud logical port*. When you create a trunk logical port on a switch, NavisCore automatically creates a VPN cloud logical port for each VPN that has a virtual router on the switch. The virtual routers use the IP interfaces associated with the VPN cloud logical ports to route IP VPN traffic through the Lucent network.

Figure 16-4 shows Customer A's VPN from Figure 16-3 on page 16-5. Customer A's virtual routers on the edge switches use IP VPN cloud IP interfaces to route data. Notice the distinction between the IP VPN cloud IP interfaces, which are maintained on the CP/SP, and the ingress IP interfaces to which the VPI/VCIs are bound, which are maintained on the switch card.

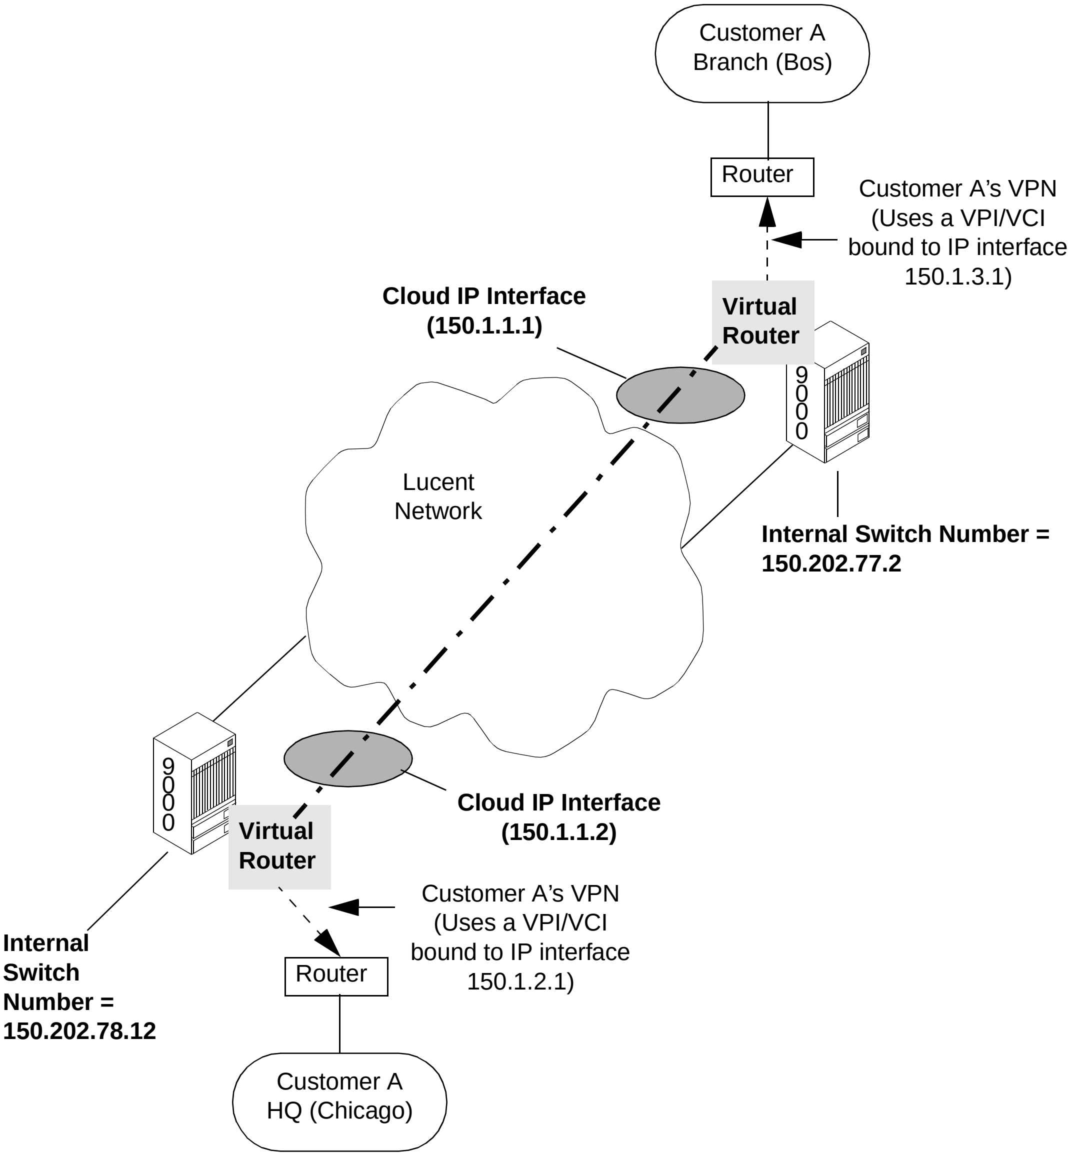

Figure 16-4. IP VPN Cloud IP Interfaces

All virtual routers in the same IP VPN use ARP to map cloud interface IP addresses to the internal switch address (the IP address assigned to a switch when it is installed). In effect, the internal switch address acts as a MAC address. ARP entries for cloud IP interfaces can be created dynamically or statically. See Chapter 6, "Configuring Static ARP Entries" for more information on creating static ARP entries for cloud IP interfaces.

Within the Lucent network, IP VPN virtual routers use the Open Shortest Path First (OSPF) routing protocol, the Routing Information Protocol (RIP), and static routes to route IP VPN traffic.

Each VPN is automatically assigned a Class D multicast address in the format 239.192.*x.y*, where *x.y* is the ID that is automatically assigned to the IP VPN. Each virtual router belonging to the VPN automatically joins the multicast group for that VPN (239.192.*x.y*). All the virtual routers in the same IP VPN share the same Class D multicast address. Figure 16-5 illustrates how multicast addresses are assigned to virtual routers from the same VPN.

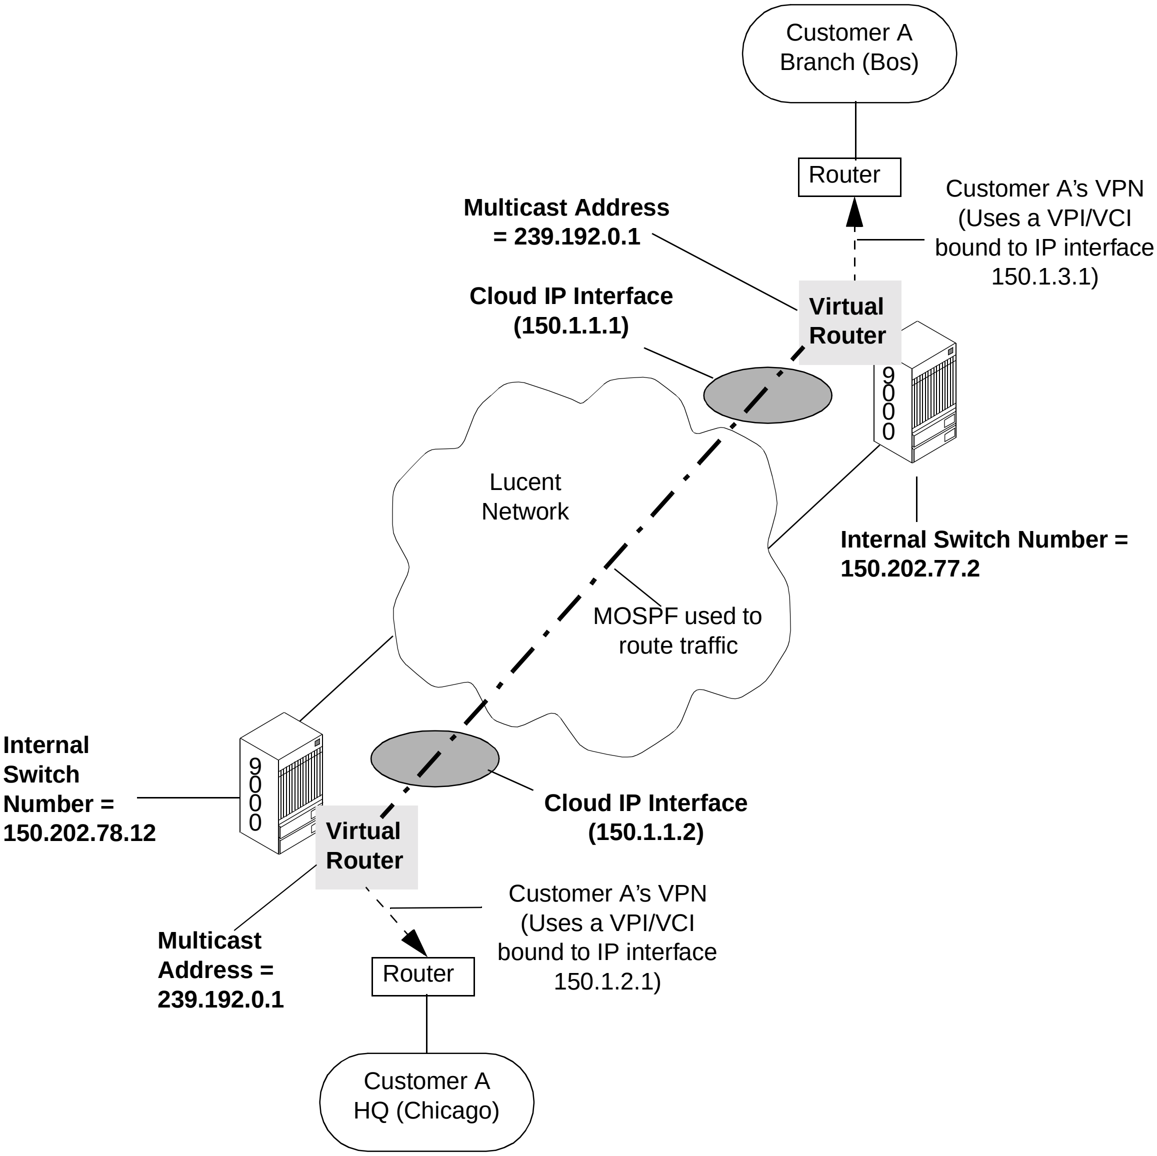

Figure 16-5. Multicast Addresses Assigned to Virtual Routers

Routing and Forwarding

Within the Lucent network, IP packets belonging to an IP VPN can be forwarded between virtual routers in the following ways:

- Using point-to-point LSPs configured explicitly for that VPN. Multiple point-to-point LSPs can connect two switches. However, for a single VPN, you can configure one (and only one) point-to-point LSP between two switches. For example, suppose that you have three VPNs that require point-to-point LSPs between the same two switches. You can configure one point-to-point LSP for each VPN, but you cannot configure two point-to-point LSPs for any of the VPNs or configure one point-to-point LSP to support two VPNs.
- Using point-to-point LSPs configured for public use (available to all VPNs and public traffic).
- Using the default, best-effort Multipoint-to-Point Tunnel (MPT) LSP, which is always available. This MPT LSP is overridden by point-to-point LSPs configured explicitly for public use.
- Using hop-by-hop forwarding.

IP Navigator prioritizes IP VPN traffic to be forwarded between virtual routers in a Lucent network in the following order:

1. A point-to-point LSP configured explicitly for the VPN.
2. A public, configured point-to-point LSP between the switches where the virtual routers reside.
3. The public, default MPT LSP between the switches where the virtual routers reside.
4. Hop-by-hop forwarding.

► In some cases, MPT LSPs may be selected over public point-to-point LSPs if the MPT LSPs provide more bandwidth and are of a lower cost than the point-to-point LSPs.

Figure 16-6 shows VPN traffic being transported between two edge switches via point-to-point LSPs. Notice the following details:

- VPN1 and VPN2 have point-to-point LSPs configured explicitly for each of them.
- A configured public point-to-point LSP is available for use by either VPN1 or VPN2.
- The default public MPT LSP is available for use by either VPN1 or VPN2.

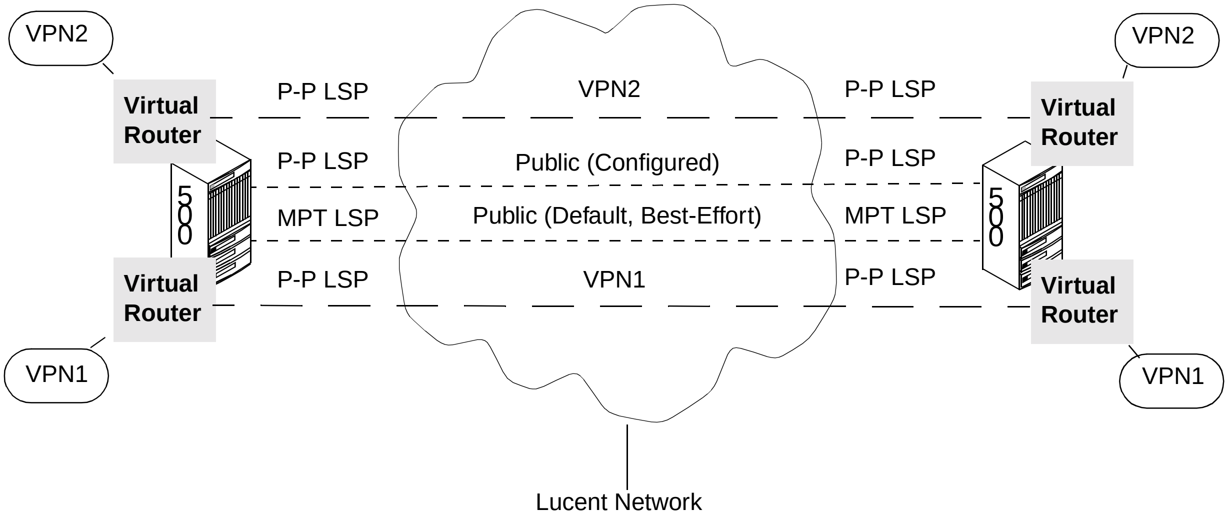

Figure 16-6. LSPs Transporting IP VPN Traffic

Typically, point-to-point LSPs connect switches at the network edge. Traffic is transported between the switches according to the QoS requirements configured for the point-to-point LSP. See Chapter 12, “Configuring Label Switched Paths” for more information on point-to-point LSPs and MPT LSPs.

How IP VPNs Differ From VNN VPNs

In addition to IP VPNs, Lucent provides support for Virtual Network Navigator (VNN) VPNs. Whereas VNN VPNs allow you to reserve OSI Layer 2 resources (for example, trunks and circuits) for a customer, IP VPNs allow you to reserve OSI Layer 3 resources (for example, IP interfaces and IP routing tables) for a customer.

It is possible to create a VNN VPN and an IP VPN for the same customer. The VNN VPN would reserve OSI Layer 2 resources for the customer, and the IP VPN would reserve OSI Layer 3 resources for the customer.

VNN VPNs are managed through the Lucent Parameters ⇒ Set All Virtual Private Networks selection on the NavisCore Administer menu. See the following guides for more information on VNN VPNs:

- *NavisCore Frame Relay Configuration Guide*
- *NavisCore ATM Configuration Guide*

About the Public IP VPN

The public IP VPN is a special IP VPN that reserves network resources for public IP traffic — IP traffic that does not originate from a private IP VPN customer. You assign IP network resources to the public IP VPN in the same way that you assign IP network resources to private IP VPNs. The IP VPN ID of the public IP VPN is always zero (0).

Planning an IP VPN

Planning an IP VPN is very similar to planning a traditional IP network, and managing a logical router is very similar to managing a physical router. Like a traditional IP network, an IP VPN acts as a separate routing domain, complete with its own routing tables and route maps. When you log in to an IP VPN, you can manage it as if it were a distinct physical network.

Divide IP VPN Network Management Tasks

IP VPN management responsibilities are divided among two groups of network managers:

- Network managers from the service provider (such as an ISP)
- Private network administrators (PNAs) from the IP VPN customer (such as a corporation)

Table 16-1 describes the division of IP VPN management responsibilities.

Table 16-1. IP VPN Management Responsibilities

Group	Task
Service Provider	Manage physical network infrastructure (for example, switches, DS0s, DS1s, DS3s, and trunks).
	Manage Frame Relay, ATM, PPP, and Ethernet logical ports (and associated IP logical ports).
	Manage Frame Relay and ATM circuits (that is, DLCIs and VPI/VCIs).
	Manage route maps, IP OSPF router IDs, and IP loopback addresses.
	Manage IP VPN definitions (that is, add, modify, and delete VPN definitions).
Customer	Manage ingress IP interfaces and IP VPN cloud IP interfaces, including binding DLCIs and VPI/VCIs to ingress IP interfaces.
	Manage point-to-point LSPs within the VPN.
	Manage routing tables, OSPF, route maps, and other IP network resources within the VPN.

Make sure that customers have the necessary access to NavisCore to manage their IP VPNs. Customers in a VPN should also be able to telnet into switches belonging to that VPN to perform management tasks.

Determine a Unique Name and Password

Determine a unique name that identifies the IP VPN. It is a good idea to base the IP VPN name on the name of the organization(s) it serves (for example, XYZCompanyVPN).

Also, determine a password that customers must supply when they access the switch console. When a customer supplies this password, any console commands that the customer issues apply to the customer's IP VPN only.

Identify Ingress Ports, DLCIs, and VPI/VCIs

Identify ingress IP logical ports and IP server logical ports at the edge of the Lucent network that will act as IP VPN access points for users. Also, identify the DLCIs and VPIs/VCIs that you will assign to each customer on each ingress IP logical port or IP server logical port. Observe the following rules:

- Only a single IP VPN can be assigned to an IP logical port that uses either Ethernet or PPP as the data link protocol.
- For each IP logical port that uses Frame Relay as the data link protocol, you can assign multiple DLCIs to a single VPN, or you can distribute the DLCIs among multiple VPNs. This rule applies to both the B-STDX 8000/9000 and the CBX 500 switches.
- For each IP logical port or IP server logical port that uses ATM as the data link protocol, you can assign multiple VPI/VCIs to a single IP VPN, or you can distribute the VPI/VCIs among multiple IP VPNs.

> ▶ For IP server logical ports, you must create an IP server PVC for each VPI/VCI that you assign to an IP VPN. For example, suppose you have an IP server logical port, and you want to assign one VPI/VCI to one IP VPN and another VPI/VCI to another IP VPN. You would create two IP server PVCs, one for each IP VPN.

- The sum of DLCIs and VPIs/VCIs you assign to a single virtual router (i.e., IP VPN) on a switch cannot exceed the maximum protocol connection ID limit you configure for the IP VPN. The default maximum protocol connection ID limit is 9, but this can be changed by the service provider.

Figure 16-7 illustrates the relationship between an IP logical port that uses Ethernet as the data link protocol and an IP VPN.

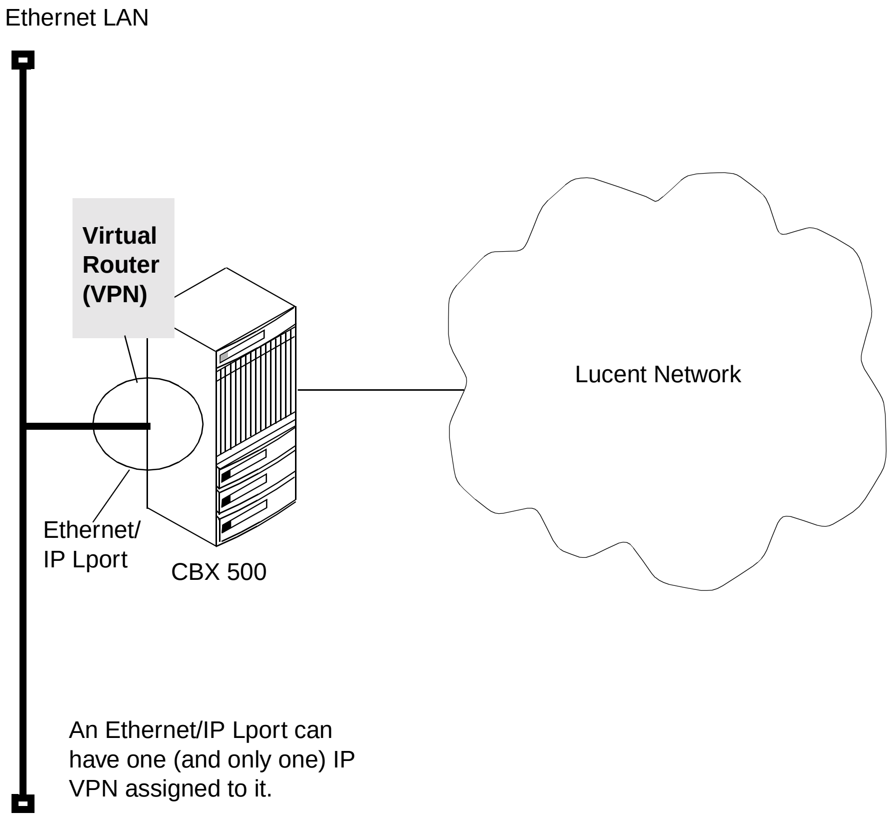

Figure 16-7. Ethernet/VPN Relationship

Figure 16-8 illustrates the relationship between an IP logical port that uses PPP as the data link protocol and an IP VPN.

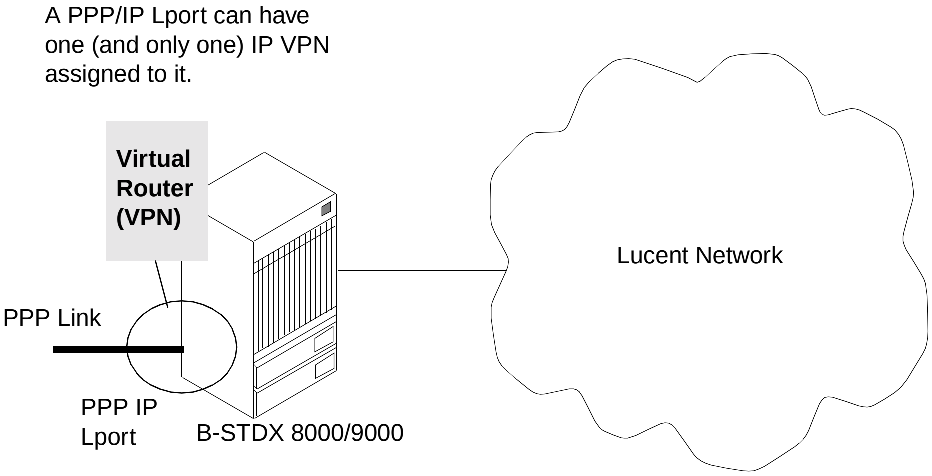

Figure 16-8. PPP/VPN Relationship

Figure 16-9 illustrates the relationship between IP VPNs and DLCIs.

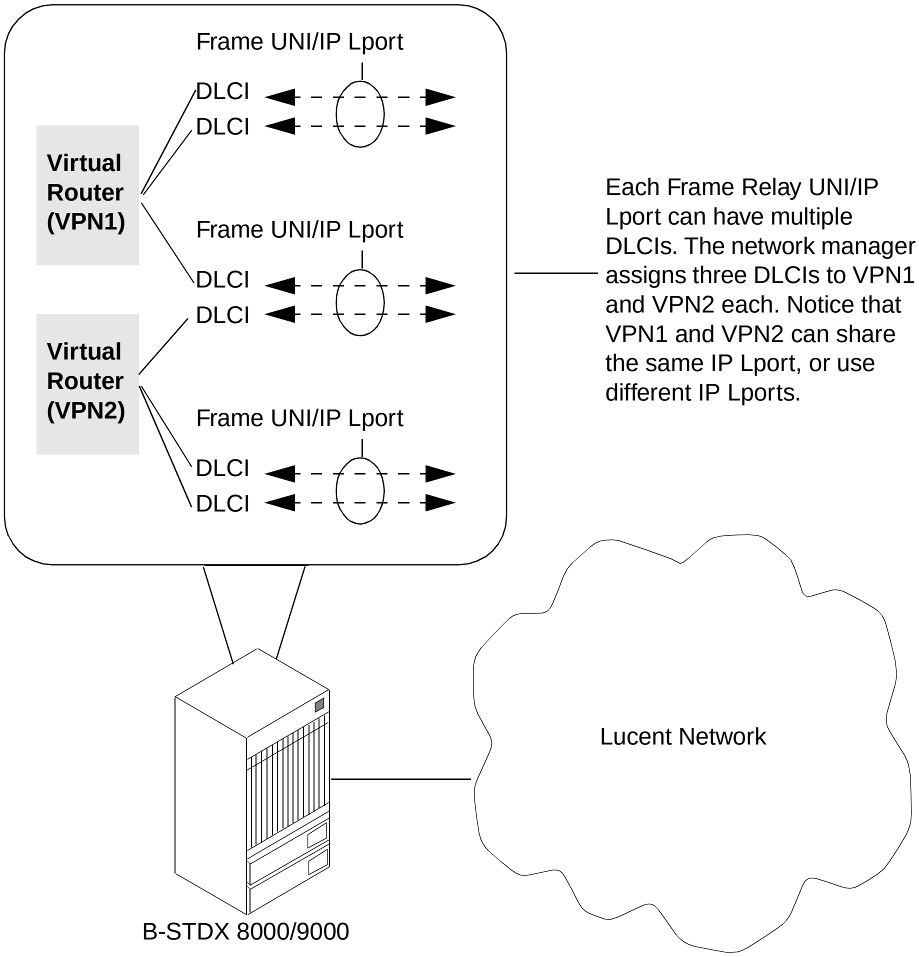

Each Frame Relay UNI/IP Lport can have multiple DLCIs. The network manager assigns three DLCIs to VPN1 and VPN2 each. Notice that VPN1 and VPN2 can share the same IP Lport, or use different IP Lports.

Figure 16-9. IP VPN/DLCI Relationship

Figure 16-10 illustrates the relationship between IP VPNs and VPI/VCIs on a B-STDX 8000/9000 switch.

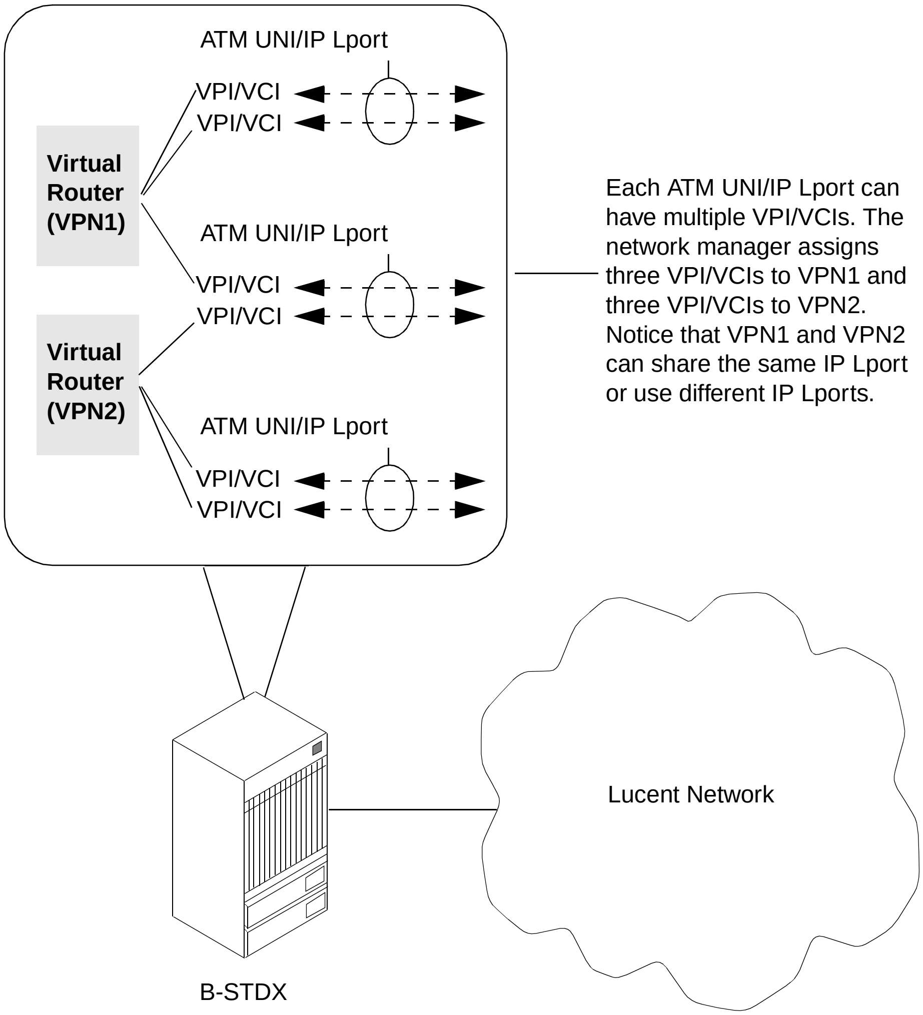

Each ATM UNI/IP Lport can have multiple VPI/VCIs. The network manager assigns three VPI/VCIs to VPN1 and three VPI/VCIs to VPN2. Notice that VPN1 and VPN2 can share the same IP Lport or use different IP Lports.

Figure 16-10. IP VPN/VPI/VCI Relationship (ATM/IP Ports on the B-STDX)

Figure 16-11 illustrates the relationship between IP VPNs and VPI/VCIs for IP server logical ports on the CBX 500 switch.

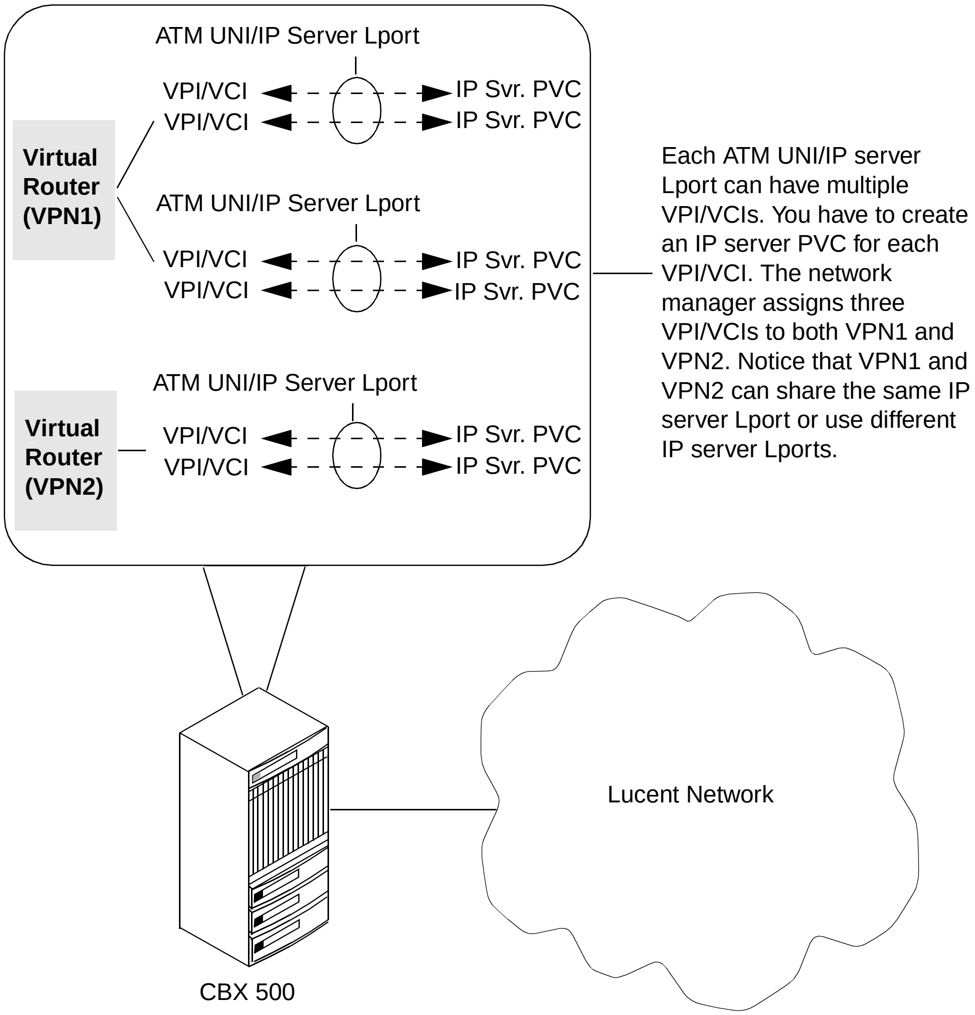

Each ATM UNI/IP server Lport can have multiple VPI/VCIs. You have to create an IP server PVC for each VPI/VCI. The network manager assigns three VPI/VCIs to both VPN1 and VPN2. Notice that VPN1 and VPN2 can share the same IP server Lport or use different IP server Lports.

Figure 16-11. VPN/VPI/VCI Relationship (IP Server Logical Ports)

Identify Ingress IP Interfaces

On a given edge switch, each IP VPN has its own ingress IP interfaces. Each IP interface is defined by an IP address, CIDR mask, and so on. You manually create these IP interfaces, assign them to IP VPNs through NavisCore, and then bind these interfaces to DLCIs and VPI/VCIs. By binding DLCIs and VPI/VCIs to an IP VPN's IP interfaces, you effectively assign the DLCIs and VPI/VCIs to the IP VPN. No other IP VPN can own these resources. See "Configuring Ingress IP Interfaces for IP VPNs" on page 16-40 for details. Figure 16-12 shows examples of ingress IP interface bindings.

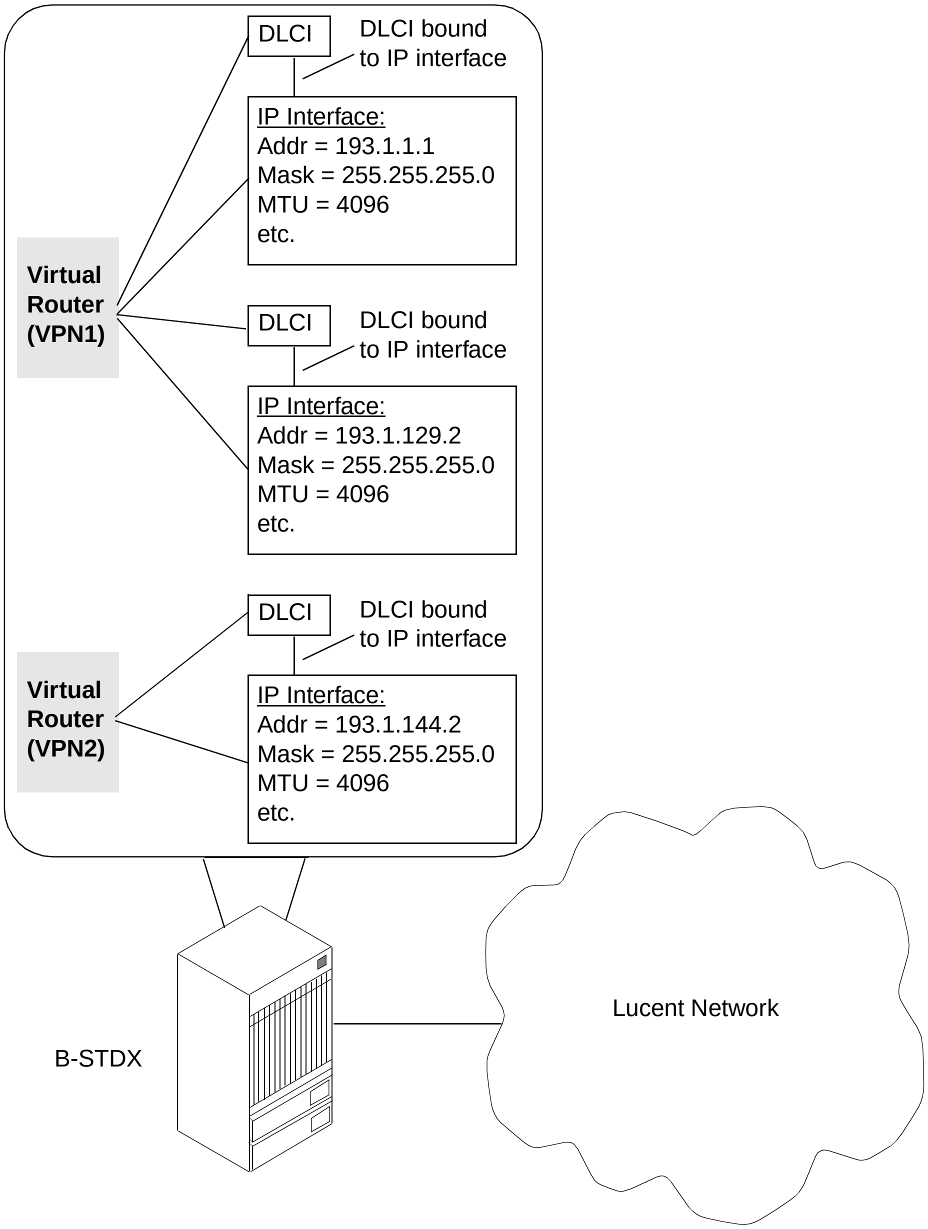

Figure 16-12. DLCIs Bound to Ingress IP Interfaces

Identify IP VPN Cloud IP Interfaces

You need to configure one VPN cloud IP interface per IP VPN per switch. For example, if the switch supports three IP VPNs (that is, supports three different virtual routers), you need to configure at least three VPN cloud IP interfaces on the switch, one per IP VPN.

► Configuring more than one IP VPN cloud IP interface per IP VPN per switch is not recommended.

► You cannot define an IP VPN cloud IP interface on a switch until you configure an IP loopback address for the public IP VPN on that switch. See "Identify IP OSPF Router ID and Loopback Address Requirements" on page 16-21 for more information.

When you configure VPN cloud IP interfaces for a specific IP VPN, make sure that the IP addresses of all the interfaces have the same subnetwork number, just as if you were configuring IP interfaces on an Ethernet LAN. For a specific IP VPN, the virtual routers on all of the edge switches communicate with each other through their VPN cloud IP interfaces as if they share a single Ethernet LAN.

In Figure 16-13, two switches support two IP VPN virtual routers each, and one switch supports one IP VPN virtual router. On each switch that supports two IP VPN virtual routers, you would have to configure a VPN cloud IP interface for each IP VPN. On the third switch, you would have to configure only one VPN cloud IP interface.

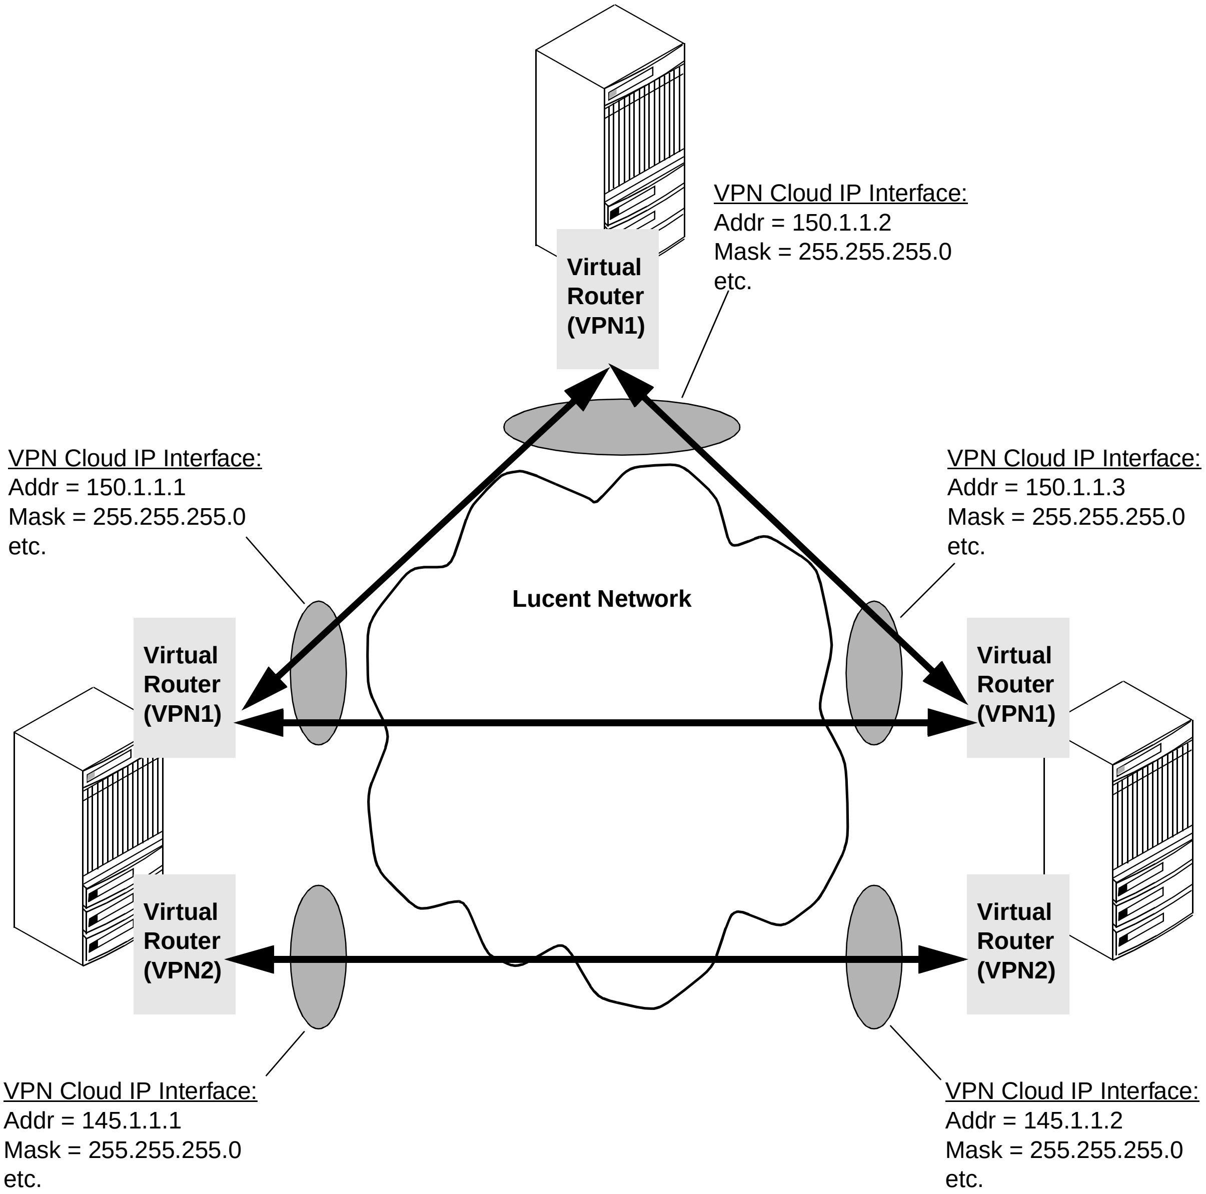

Figure 16-13. Identifying VPN Cloud IP Interfaces

In Figure 16-13, notice that for each IP VPN, all of the VPN cloud IP interfaces have the same subnetwork number. All of the VPN cloud IP interfaces in VPN1 are in 150.1.1.0 and all of the VPN cloud IP interfaces in VPN2 are in 145.1.1.0.

Identify Point-to-Point LSPs

Identify the point-to-point LSPs that will connect the virtual routers on the edge switches in the Lucent network. Keep in mind the following rules:

- Using point-to-point LSPs configured explicitly for that VPN. Multiple point-to-point LSPs can connect two switches. However, for a single VPN, you can configure one (and only one) point-to-point LSP between two switches. For example, suppose that you have three VPNs that require point-to-point LSPs between the same two switches. You can configure one point-to-point LSP for each VPN, but you cannot configure two point-to-point LSPs for any of the VPNs or configure one point-to-point LSP to support two VPNs.
- You can configure point-to-point LSPs for public use (available to all IP VPNs and public traffic).
- A default, best-effort, MPT LSP is always available for public use. This MPT LSP is overridden by point-to-point LSPs explicitly configured for public use.

See Chapter 12, "Configuring Label Switched Paths" for more information on configuring point-to-point LSPs and MPT LSPs.

Verify MOSPF Support

Multicast Open Shortest Path First (MOSPF) is the multicast routing protocol used by virtual routers to discover other group members and to exchange routing information. If your VPNs will use MOSPF, verify that the switches in the Lucent network that carry IP VPN traffic support MOSPF. Note that you do not have to configure MOSPF on any switches that carry IP VPN traffic. See "How IP VPN Traffic is Routed Through the Lucent Network" on page 16-6 for more information on how MOSPF is used to exchange routing information.

Verify RIP/OSPF Support

If your VPN will use the Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) routing protocols to route IP VPN traffic through the network, verify that the switches in the Lucent network that carry IP VPN traffic support RIP or OSPF. See "How IP VPN Traffic is Routed Through the Lucent Network" on page 16-6 for more information.

Identify IP OSPF Router ID and Loopback Address Requirements

In order to support IP VPN traffic, edge switches (that is, switches that provide ingress IP interfaces) require that you configure:

- An IP OSPF router ID configured for the public IP VPN (that is, IP VPN 0). Note that, in some cases, the switch configures an IP OSPF router ID automatically. See "Planning Router IDs" on page 9-22 for more information.

- One IP loopback address configured for the public IP VPN (that is, IP VPN 0). This address is required for IP multicasting. You cannot configure an IP VPN cloud IP interface on a switch until you configure an IP loopback address for the public IP VPN on that switch.

Switches within the Lucent network that are not edge switches, but are connected to intermediate trunks that carry VPN traffic, require an IP OSPF router ID only.

Figure 16-14 illustrates IP OSPF router ID and IP loopback address requirements.

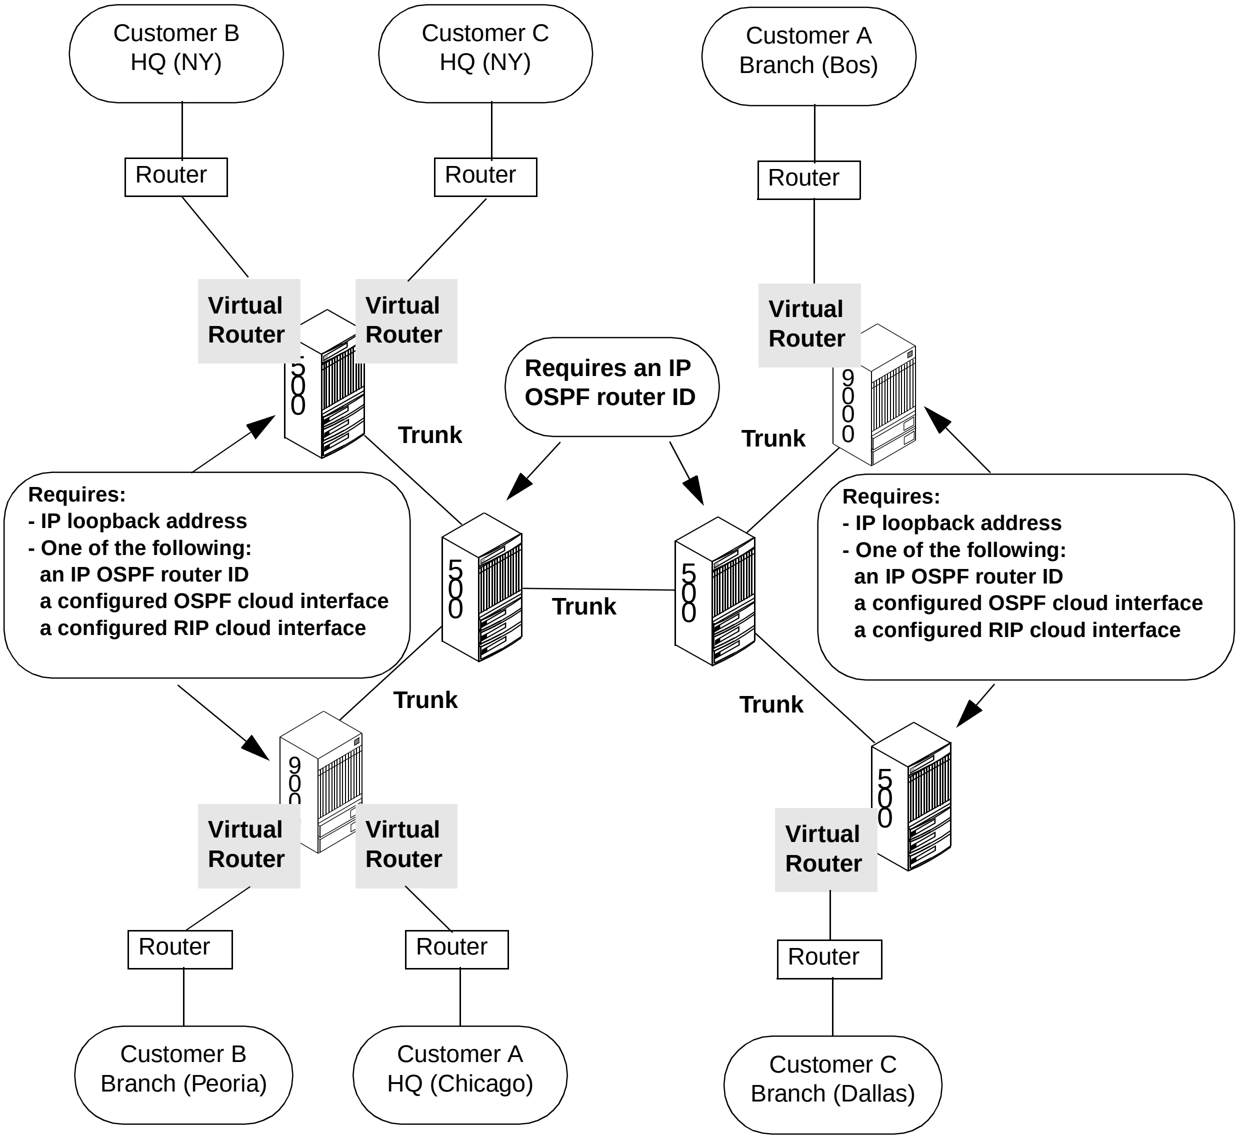

Figure 16-14. IP OSPF Router ID and IP Loopback Address Requirements

For information on the IP OSPF router ID, see “Planning Router IDs” on page 9-22. For information on configuring IP loopback addresses, see “Configuring IP Loopback Addresses” on page 8-31.

► When configuring IP OSPF router IDs and IP loopback addresses to support IP VPN traffic, make sure that you are in the context of the public IP VPN (that is, VPN 0). You do this by selecting the IP VPN called “Public.” See “Selecting the IP VPN” on page 16-33 for more information.

Determine IP VPN Resource Limits

You can set upper limits on the specific IP resources that can be configured for the VPN on switches that contain the VPN virtual routers. For each VPN, these limits are enforced on a per switch basis. All switches participating in a VPN should use the same limits. These limits include:

- Maximum protocol connection IDs per switch. A protocol connection ID is either a DLCI (for Frame Relay connections) or a VPI/VCI (for ATM connections).
- Maximum RIP interfaces per switch.
- Maximum static ARP entries per switch.
- Maximum routes per switch.
- Maximum network access lists per switch.
- Maximum route maps per switch.
- Maximum route maps per RIP interface.
- Maximum OSPF virtual links per switch.
- Maximum IP interfaces per switch.
- Maximum OSPF interfaces per switch.
- Maximum static routes per switch.
- Maximum network filters per switch.
- Maximum network filters per access list per switch.
- Maximum network access lists per route map per switch.
- Maximum OSPF neighbors per switch.
- Maximum OSPF area aggregates per switch.

Figure 16-15 shows two switches at the edge of the Lucent network. Each switch contains a virtual router for a single IP VPN (Customer A’s VPN). When the network manager defines the IP VPN, the manager configures a single set of resource limits that are enforced on each switch.

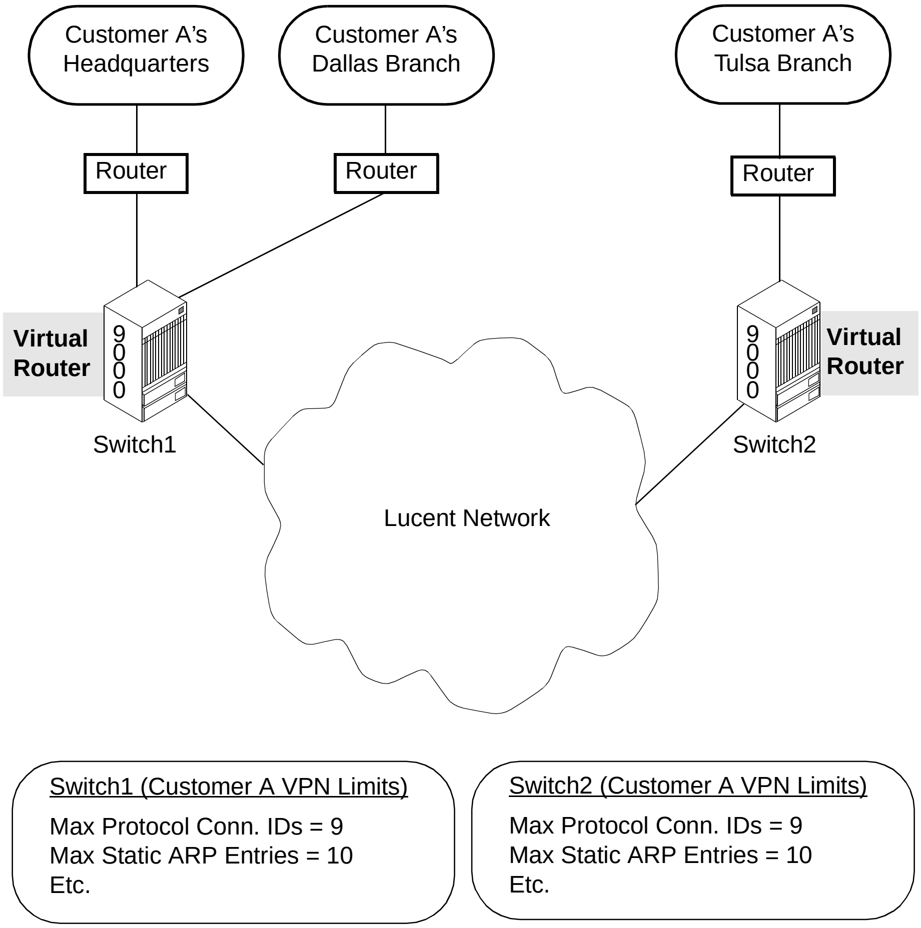
Customer A's Headquarters
Customer A's Dallas Branch
Customer A's Tulsa Branch
Router
Router
Router
Virtual Router
9000
Switch1
9000
Virtual Router
Switch2
Lucent Network
Switch1 (Customer A VPN Limits)
Max Protocol Conn. IDs = 9
Max Static ARP Entries = 10
Etc.
Switch2 (Customer A VPN Limits)
Max Protocol Conn. IDs = 9
Max Static ARP Entries = 10
Etc.

Figure 16-15. IP VPN Limits Per Switch

Identify IP VPN Memory Requirements

Verify that edge switches where VPN logical ports reside have sufficient memory. The software component that implements VPN functionality requires approximately 20 KB to 80 KB of memory on the processor module and each IOM/IOP. Also, each VPN has its own set of dedicated resources such as routing tables, and these resources consume memory. See the Software Release Notice for your switch for more information.

VPN Network Design Recommendations

When you are designing a VPN network to span multiple VNN trunk areas, the following VPN designs are suggested to ensure quality of data processing.

- Meshed point-to-point LSPs – Provision a fully-meshed network of point-to-point LSP circuits between VPN member switches.

- Match virtual router gateways and VNN Area Border Routers – Configure the VNN ABRs as virtual router gateways, which requires that the VNN ABRs have two Cloud interfaces (one for each subnet to which the VNN ABR is a member).

Meshed Point-to-Point LSP Solution

Figure 16-16 illustrates a network where switches A, B, D and E are members of a virtual private network (VPN1). Switches B and D are Area Border Routers. For data to be forwarded from CPE 1 to CPE 2, the switches need a full mesh of point-to-point LSPs configured between the VPN member switches. This point-to-point LSP mesh ensures that the data from CPE 1 to CPE 2 will be forwarded from end to end by point-to-point LSPs.

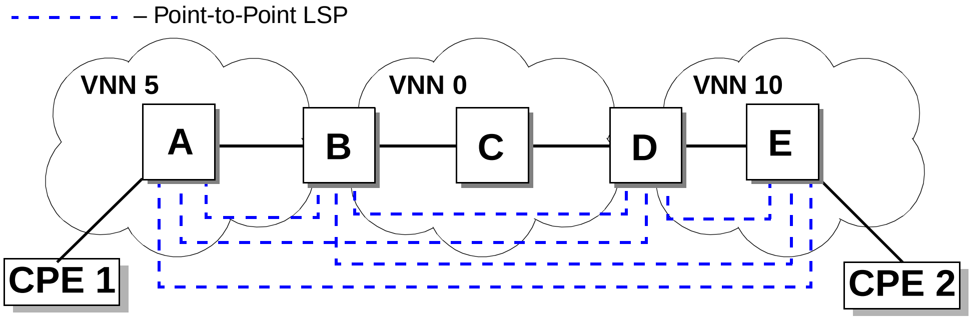

Figure 16-16. Meshed Point-to-Point LSP Solution

Virtual Router Gateway Solution

Figure 16-17 illustrates a network consisting of multiple VNN areas Each area has virtual private network members that have VPN cloud interfaces in a particular subnet (subnet 1, 2, 3, or 4). Within each area, a single VPN member has an additional Cloud interface in another subnet (Subnet 0). This node acts as the Virtual Router Gateway (VRGW) for all the VPN members in a VNN area.

Each of the VPN members in subnet 1, 2, 3, and 4 has a point-to-point LSP to the Virtual Router Gateway in the VPN. All the VRGWs are connected in a full mesh across the VNN areas. To send data from CPE 1 to CPE 2, CP1 would forward the data to a switch in Subnet 1. That switch would then forward the data over a point-to-point LSP within Subnet 1 to VRGW1, the virtual gateway router for Subnet 1. VRGW1 would forward the data on the point-to-point LSP circuit to VRGW4, and from there to through the transit switch to CPE 2.

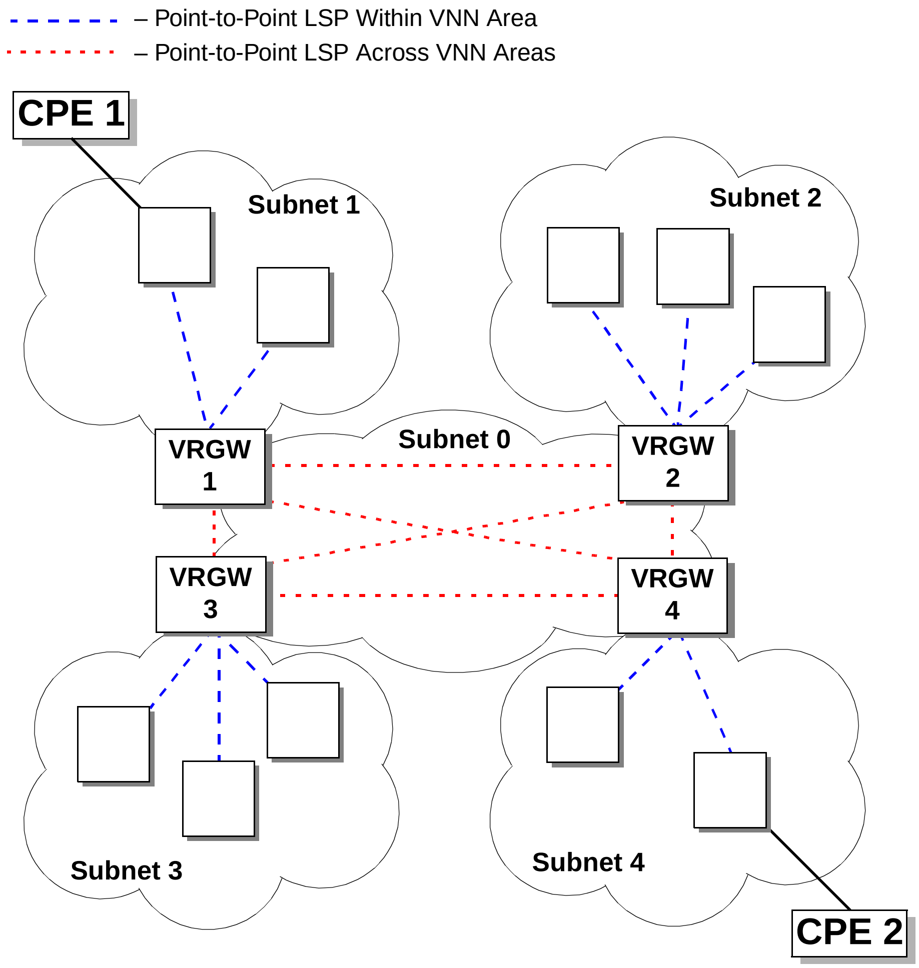

Figure 16-17. Virtual Router Gateway Solution

P VPN Configuration Flowchart

Figure 16-18 summarizes the process of configuring IP VPNs.

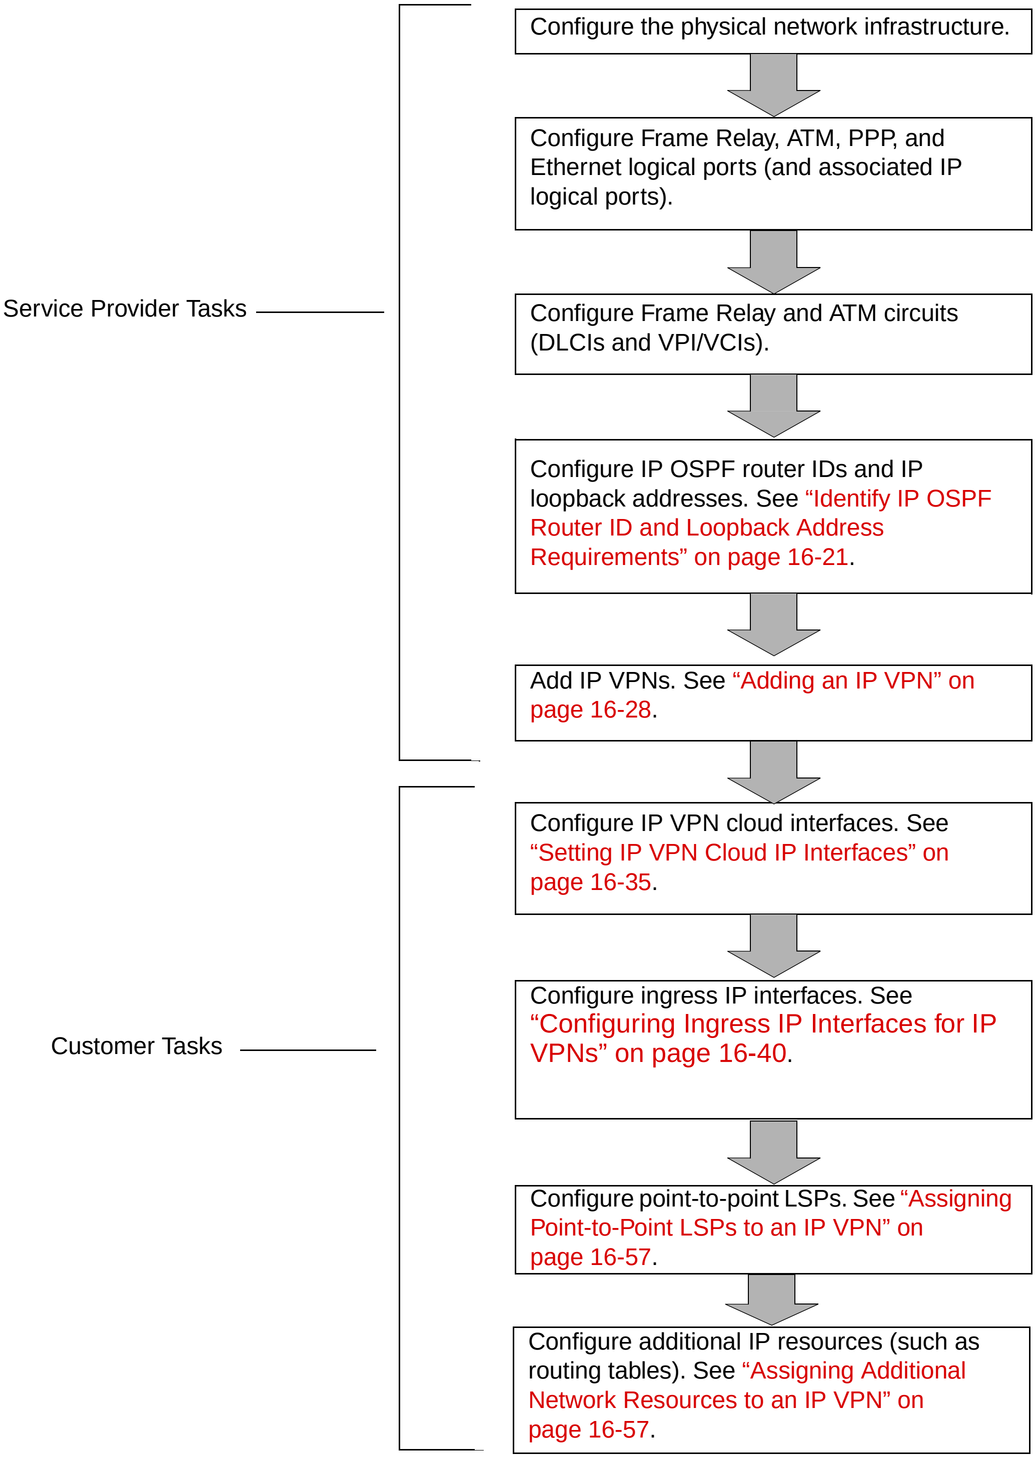

Configure the physical network infrastructure.

Configure Frame Relay, ATM, PPP, and Ethernet logical ports (and associated IP logical ports).

Configure Frame Relay and ATM circuits (DLCIs and VPI/VCIs).

Configure IP OSPF router IDs and IP loopback addresses. See "Identify IP OSPF Router ID and Loopback Address Requirements" on page 16-21.

Add IP VPNs. See "Adding an IP VPN" on page 16-28.

Configure IP VPN cloud interfaces. See "Setting IP VPN Cloud IP Interfaces" on page 16-35.

Configure ingress IP interfaces. See "Configuring Ingress IP Interfaces for IP VPNs" on page 16-40.

Configure point-to-point LSPs. See "Assigning Point-to-Point LSPs to an IP VPN" on page 16-57.

Configure additional IP resources (such as routing tables). See "Assigning Additional Network Resources to an IP VPN" on page 16-57.

Figure 16-18. IP VPN Configuration Flowchart

Adding an IP VPN

To add an IP VPN:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP VPNs. The Set All IP Virtual Private Networks dialog box appears (see Figure 16-19).

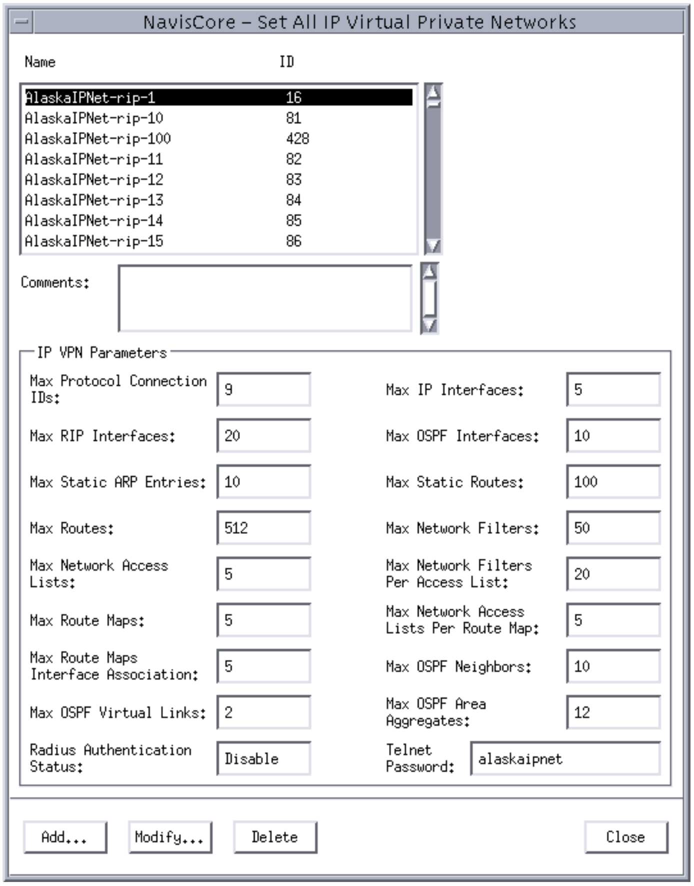

Name	ID
AlaskaIPNet-rip-1	16
AlaskaIPNet-rip-10	81
AlaskaIPNet-rip-100	428
AlaskaIPNet-rip-11	82
AlaskaIPNet-rip-12	83
AlaskaIPNet-rip-13	84
AlaskaIPNet-rip-14	85
AlaskaIPNet-rip-15	86

Comments:

IP VPN Parameters

Max Protocol Connection IDs:	9	Max IP Interfaces:	5
Max RIP Interfaces:	20	Max OSPF Interfaces:	10
Max Static ARP Entries:	10	Max Static Routes:	100
Max Routes:	512	Max Network Filters:	50
Max Network Access Lists:	5	Max Network Filters Per Access List:	20
Max Route Maps:	5	Max Network Access Lists Per Route Map:	5
Max Route Maps Interface Association:	5	Max OSPF Neighbors:	10
Max OSPF Virtual Links:	2	Max OSPF Area Aggregates:	12
Radius Authentication Status:	Disable	Telnet Password:	alaskaipnet

Figure 16-19. Set All IP Virtual Private Networks Dialog Box

The Set All IP Virtual Private Networks dialog box lists each VPN and its associated IP parameters. See Table 16-2 for descriptions of the buttons on the dialog box. See Table 16-3 for descriptions of the parameters.

Table 16-2. Set All IP Virtual Private Networks Dialog Box: Buttons

Button	Function
Add	Adds a VPN.
Modify	Modifies a VPN. See "Modifying an IP VPN" on page 16-31 for details.
Delete	Deletes a VPN. See "Deleting an IP VPN" on page 16-31 for details.

2. Choose Add. The Add IP Virtual Private Network dialog box appears (see Figure 16-20).

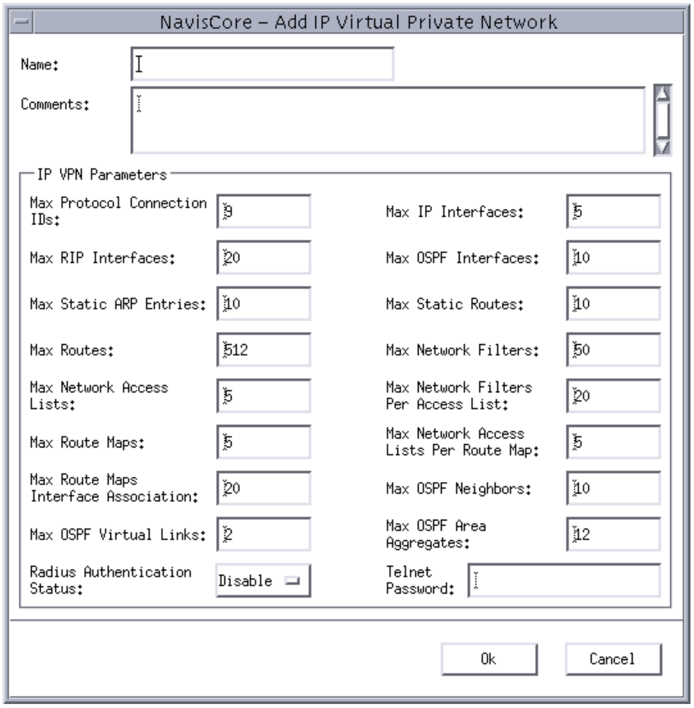

NavisCore – Add IP Virtual Private Network

Name:

Comments:

IP VPN Parameters

Max Protocol Connection IDs: 9
Max IP Interfaces: 5
Max RIP Interfaces: 20
Max OSPF Interfaces: 10
Max Static ARP Entries: 10
Max Static Routes: 10
Max Routes: 512
Max Network Filters: 50
Max Network Access Lists: 5
Max Network Filters Per Access List: 20
Max Route Maps: 5
Max Network Access Lists Per Route Map: 5
Max Route Maps Interface Association: 20
Max OSPF Neighbors: 10
Max OSPF Virtual Links: 2
Max OSPF Area Aggregates: 12
Radius Authentication Status: Disable
Telnet Password:

Figure 16-20. Add IP Virtual Private Network Dialog Box

3. Complete the fields described in Table 16-3.

Table 16-3. Add IP Virtual Private Network Dialog Box: Fields

Field	Description
Name	Specify the IP VPN name. The VPN name must be at least one character long.
Comments	Specify comments related to the VPN (for example, “Customer A’s VPN”).
IP VPN Parameters — Sets limits on IP Navigator resources for each IP VPN per switch.	
Max Protocol Connection IDs	Specify the maximum number of concurrent protocol connection IDs per switch. A protocol connection ID is either a DLCI (for Frame Relay connections) or a VPI/VCI (for ATM connections). The default is 9. Valid values range from 0 to 2147483647. For example, if you specify 9 (the default), the switch can support a total of 9 DLCIs and VPI/VCIs for the IP VPN.
Max RIP Interfaces	Specify the maximum number of RIP interfaces per switch. The default is 20. Valid values range from 0 to 2147483647.
Max Static ARP entries	Specify the maximum number of static ARP entries per switch. The default is 10. Valid values range from 0 to 2147483647.

Table 16-3. Add IP Virtual Private Network Dialog Box: Fields (Continued)

Field	Description
Max Routes	Specify the maximum number of routes that can be stored in the IP VPN's private routing table per switch. The default is 512. Valid values range from 0 to 2147483647.
Max Network Access Lists	Specify the maximum number of network access lists per switch. The default is 5. Valid values range from 0 to 2147483647.
Max Route Maps	Specify the maximum number of route maps per switch. The default is 5. Valid values range from 0 to 2147483647.
Max Route Maps Interface Association	Specify the maximum number of route maps per RIP interface. The default is 5. Valid values range from 0 to 2147483647.
Max OSPF Virtual Links	Specify the maximum number of IP OSPF virtual links per switch. The default is 2. Valid values range from 0 to 2147483647.
Radius Authentication Status	Specify whether Radius Authentication is used to authenticate IP VPN logins. Specify *Enable* if you want Radius Authentication to authenticate IP VPN logins. Otherwise, specify *Disable* (the default). *Note: When you use RADIUS authentication for a private VPN (that is, a VPN with an ID greater than 0), you must append the VPN ID to users' login names in the RADIUS database. For example, if you want user edison to log in to VPN 999, you must enter his username as edison-999 in the RADIUS database. When edison logs in, he would enter "edison" at the login prompt, and the switch would append the VPN ID ("-999") before forwarding the login to the RADIUS server for authentication.*
Max IP Interfaces	Specify the maximum number of IP interfaces per switch. The default is 5. Valid values range from 0 to 2147483647.
Max OSPF Interfaces	Specify the maximum number of OSPF interfaces per switch. You can create OSPF interfaces for the IP VPN IP cloud logical port or the ingress IP interfaces. The default is 10. Valid values range from 0 to 2147483647.
Max Static Routes	Specify the maximum number of static routes (from 0 to 2147483647) that can be configured in the IP VPN's private routing table per switch. The default is 10.
Max Network Filters	Specify the maximum number of network filters per switch. The default is 50. Valid values range from 0 to 2147483647.
Max Network Filters Per Access List	Specify the maximum number of network filters that can be associated with an access list per switch. The default is 20. Valid values range from 0 to 2147483647. For example, if you specify 20 (the default), no more than 20 network filters can be associated with any access list on a switch.

Table 16-3. Add IP Virtual Private Network Dialog Box: Fields (Continued)

Field	Description
Max Network Access Lists Per Route Map	Specify the maximum number of network access lists that can be associated with a route map per switch. The default is 5. Valid values range from 0 to 2147483647. For example, if you specify 5 (the default), no more than 5 network access lists can be associated with a route map on a switch.
Max OSPF Neighbors	Specify the maximum number of IP OSPF neighbors for the IP VPN per switch. The default is 10. Valid values range from 0 to 2147483647.
Max OSPF Area Aggregates	Specify the maximum number of IP OSPF area aggregates for the IP VPN per switch. The default is 12. Valid values range from 0 to 2147483647.
Telnet Password	Specify the password you will use to access the switch console. Once you access the switch console with this password, all console commands you issue retrieve information on your VPN's resources only. The password can be from 6 to 20 characters long. You must specify a password. Otherwise, the NMS will not allow you to add the IP VPN.

4. Choose OK.

Modifying an IP VPN

To modify an IP VPN:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP VPNs. The Set All IP Virtual Private Networks dialog box appears (see Figure 16-19 on page 16-28).

2. Select the IP VPN you want to modify.

3. Choose Modify. The Modify IP Virtual Private Networks dialog box appears. This dialog box is similar to the Add IP Virtual Private Networks dialog box (see Figure 16-20 on page 16-29).

4. Modify the fields described in Table 16-3 on page 16-29.

5. Choose OK.

Deleting an IP VPN

To delete an IP VPN:

1. Delete all resources associated with the IP VPN, such as interfaces, DLCI/VPI bindings, static routes and so on. Otherwise, you will not be able to delete the IP VPN.

2. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP VPNs. The Set All IP Virtual Private Networks dialog box appears (see Figure 16-19 on page 16-28).
3. Select the IP VPN you want to delete.
4. Choose Delete.
5. Choose OK when prompted.

Selecting the IP VPN

After you add an IP VPN, NavisCore allows you to enter the context of that VPN. Once in the context of the IP VPN, all network management operations you perform are for that IP VPN only.

> While you are in the context of a private IP VPN, you cannot manage physical ports, cards, and trunks. To manage these resources, you must be in the context of the public IP VPN.

To enter the context of an IP VPN, you must select it. NavisCore provides you with two ways to select an IP VPN:

- Using an Administer menu selection.
- Choosing the Select IP VPN button, which appears on many IP Navigator dialog boxes. This button allows you to enter an IP VPN context without having to exit the dialog box you are currently using.

Selecting the IP VPN from the Administer Menu

To select an IP VPN from the Administer menu:

1. From the Administer menu, select Lucent VNN/IP VPN ⇒ Select IP VPN. The Select IP VPN dialog box appears (see Figure 16-21).

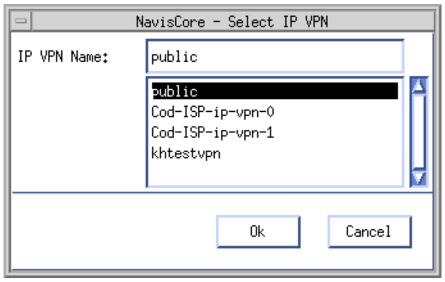

NavisCore - Select IP VPN

IP VPN Name: public

public
Cod-ISP-ip-vpn-0
Cod-ISP-ip-vpn-1
khtestvpn

Figure 16-21. Select IP VPN Dialog Box

2. Select an IP VPN.
3. Choose OK.

Network operations you perform are now in the context of the selected IP VPN.

Selecting the IP VPN Using the Select IP VPN Button

Many NavisCore IP Navigator dialog boxes display a Select IP VPN button. For example, the Set All IP LPorts dialog box is one of the dialog boxes that displays this button (see Figure 16-22).

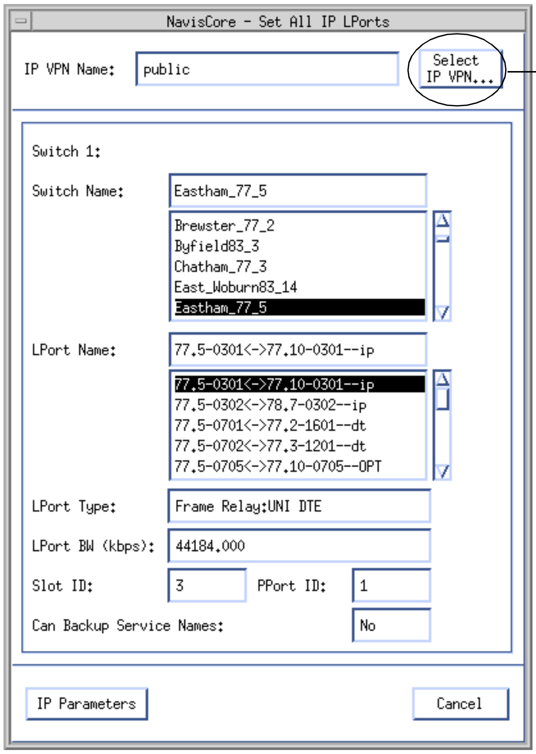

NavisCore - Set All IP LPorts

IP VPN Name: public

Switch 1:

Switch Name: Eastham_77_5

Brewster_77_2
Byfield83_3
Chatham_77_3
East_Woburn83_14
Eastham_77_5

LPort Name: 77.5-0301<->77.10-0301--ip

77.5-0301<->77.10-0301--ip
77.5-0302<->78.7-0302--ip
77.5-0701<->77.2-1601--dt
77.5-0702<->77.3-1201--dt
77.5-0705<->77.10-0705--OPT

LPort Type: Frame Relay:UNI DTE

LPort BW (kbps): 44184.000

Slot ID: 3 PPort ID: 1

Can Backup Service Names: No

Select IP VPN button

Figure 16-22. Set All IP LPorts Dialog Box (With Select IP VPN Button)

To select an IP VPN using the Select IP VPN button:

1. Choose Select IP VPN at any of the dialog boxes displaying this button. The Select IP VPN dialog box appears (see Figure 16-21 on page 16-33).
2. Select an IP VPN.
3. Choose OK. This action returns you to the dialog box from where you chose the Select IP VPN button.

Network operations you perform are now within the context of the selected IP VPN.

Setting IP VPN Cloud IP Interfaces

Before you define IP VPN cloud IP interfaces, make sure that you have defined the IP VPN. See "Adding an IP VPN" on page 16-28 for more information.

▶ You cannot define an IP VPN cloud IP interface on a switch until you configure an IP loopback address for the public IP VPN on that switch. See "Identify IP OSPF Router ID and Loopback Address Requirements" on page 16-21 for more information.

For each IP VPN, you define a VPN cloud IP interface on each switch where the IP VPN's virtual router resides. It is recommended that you create no more than one VPN cloud IP interface per IP VPN per switch.

To add VPN cloud IP interfaces:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears (see Figure 16-22 on page 16-34).
2. Choose Select IP VPN. The Select IP VPN dialog box appears (see Figure 16-21 on page 16-33).
3. Select the appropriate IP VPN.
4. Choose OK. The Set All IP LPorts dialog box appears. However, this time the dialog box lists the VPN Cloud IP logical port, and the name of the VPN you selected appears in the IP VPN name field (see Figure 16-23).

▶ If resources associated with any IP logical ports have been assigned to the IP VPN, those IP logical ports are also displayed. For example, if a DLCI has been assigned to an IP VPN, the IP logical port associated with the DLCI is displayed.

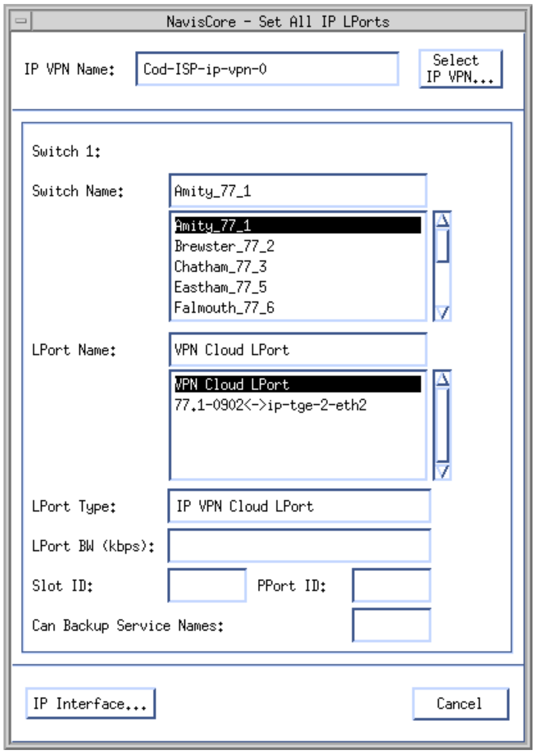

NavisCore - Set All IP LPorts

IP VPN Name: Cod-ISP-ip-vpn-0

Switch 1:

Switch Name: Amity_77_1

Amity_77_1
Brewster_77_2
Chatham_77_3
Eastham_77_5
Falmouth_77_6

LPort Name: VPN Cloud LPort

VPN Cloud LPort
77.1-0902<->ip-tge-2-eth2

LPort Type: IP VPN Cloud LPort

LPort BW (kbps):

Slot ID:

PPort ID:

Can Backup Service Names:

Figure 16-23. Set All IP LPorts Dialog Box (with VPN Cloud LPort)

5. Select the edge switch where you want to configure the VPN cloud IP interfaces.
6. The VPN Cloud LPort for that switch appears.
7. Select the VPN Cloud LPort.
8. Choose IP Interface. The Set Cloud IP Interface Addresses dialog box appears (see Figure 16-24).

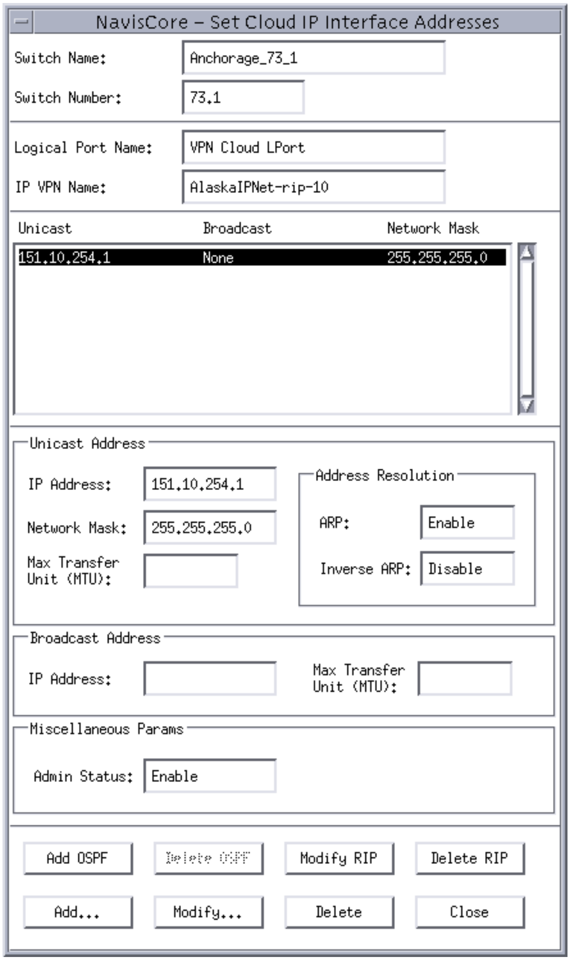

Switch Name: Anchorage_73_1

Switch Number: 73.1

Logical Port Name: VPN Cloud LPort

IP VPN Name: AlaskaIPNet-rip-10

Unicast	Broadcast	Network Mask
151.10.254.1	None	255.255.255.0

IP Address: 151.10.254.1

Network Mask: 255.255.255.0

Max Transfer Unit (MTU):

ARP: Enable

Inverse ARP: Disable

IP Address:

Max Transfer Unit (MTU):

Admin Status: Enable

Figure 16-24. Set Cloud IP Interface Addresses Dialog Box

The Set Cloud IP Interface Addresses dialog box functions the same way as the Set IP Interface Addresses dialog box — only within the context of the selected VPN. From the Set Cloud IP Interface Addresses dialog box, you can add and modify IP interfaces for the selected VPN, create an OSPF interface, and so on. See “Setting the IP Interface Address” on page 3-14 for more information on the functions of the Set IP Interface Addresses dialog box.

9. Choose Add. The Set IP Interface Address dialog box appears (see Figure 16-25).

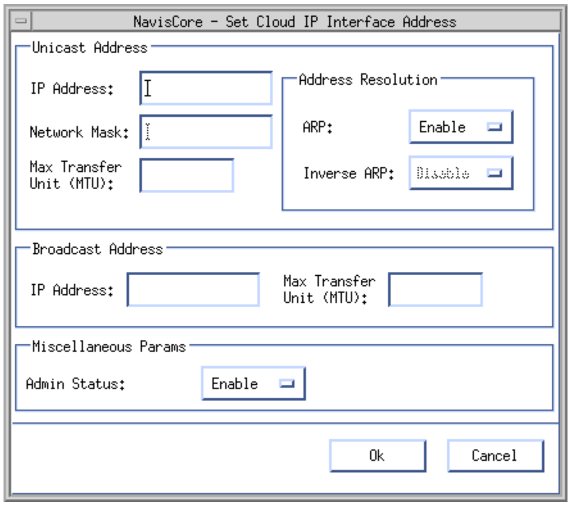

IP Address:
Network Mask:
Max Transfer Unit (MTU):
ARP: Enable
Inverse ARP: Disable
IP Address:
Max Transfer Unit (MTU):
Admin Status: Enable

Figure 16-25. Set Cloud IP Interface Address Dialog Box

10. Complete the fields described in Table 16-4.

Table 16-4. IP VPN Cloud IP Interface Address Fields

Field	Action/Description
Unicast Address	
IP Address	The IP address for this interface. Keep in mind that, once you configure an IP address, the IP address cannot be changed. To make any modifications to the IP address, you must delete the interface and re-add it with the modified IP address. Make sure that all the VPN cloud IP interfaces in the same IP VPN are in the same subnetwork. For example, suppose you have to configure two VPN cloud IP interfaces that are in the same IP VPN. You could assign 180.170.10.1 to one interface and 180.170.10.2 to the other interface. The common subnetwork number is 180.170.10.0.
Network Mask	The mask used to determine the subnet of this IP interface. Once this value is set, you cannot use the Modify Interface Address function to modify the network mask value. In order to change the network mask, you must delete the IP interface and then add a new one using the correct network mask.
Max Transfer Unit (MTU)	Not supported for IP VPN cloud interfaces.
Address Resolution	
ARP	Select one of the following options: *Enable* – (default) Enables the Address Resolution Protocol (ARP). Make sure that you keep ARP enabled. Keep in mind that all of the IP VPN cloud IP interfaces that are in the same IP VPN communicate with each other as if they were all on a common Ethernet LAN. *Disable* – Disables the ARP.

Table 16-4. IP VPN Cloud IP Interface Address Fields (Continued)

Field	Action/Description
Inverse ARP	Not supported for IP VPN cloud interfaces.
Broadcast Address	
IP Address	Not supported for IP VPN cloud interfaces.
Max Transfer Unit (MTU)	Not supported for IP VPN cloud interfaces.
Miscellaneous Params	
Admin Status	Select one of the following options: *Enable* – (default) Enables the IP interface. *Disable* – Disables the IP interface.

11. Choose OK. The Set Cloud IP Interface Addresses dialog box appears (see Figure 16-24 on page 16-37), displaying the newly created IP interface.
12. Choose Add OSPF to add an OSPF interface. The Add OSPF Interface dialog box appears.
13. Choose OK to accept the defaults. The Set Cloud IP Interface Addresses dialog box appears (see Figure 16-24 on page 16-37).

▶ When you add an OSPF interface on a VPN cloud logical port, it is recommended that you *do not* modify the defaults — especially the Area ID.

You can now manage the interface like other IP interfaces. If RIP is used as the routing protocol over the VPN cloud, then it can be configured instead of OSPF on the cloud interfaces. For example, you can configure RIP for the cloud interface just as you would for a logical port. For more information, see "Configuring RIP at the Logical Port" on page 7-1

Configuring Ingress IP Interfaces for IP VPNs

Before you configure ingress IP interfaces for an IP VPN, make sure that you have:

- Added the IP VPN. See "Adding an IP VPN" on page 16-28 for details.
- Created the Frame Relay, ATM, Ethernet, and PPP logical ports with which the IP logical ports and IP server logical ports will be associated. See the *NavisCore Frame Relay Configuration Guide*, *NavisCore ATM Configuration Guide*, and Chapter 2, "Configuring Ethernet Logical Ports" for more information.
- Created Frame Relay and ATM PVCs (that is, DLCIs and VPI/VCIs) that will be used by the IP VPNs to connect into the Lucent network. See the *NavisCore Frame Relay Configuration Guide* and the *NavisCore ATM Configuration Guide* for details.
- Created the IP logical port or IP server logical port with which the IP interface is associated. See Chapter 3, "Configuring IP Logical Ports and IP Servers" for more information on creating IP logical ports and IP server logical ports.

The way in which you configure ingress IP interfaces for IP VPNs depends on the type of data link protocol the interface supports (Ethernet, PPP, Frame Relay, or ATM). Table 16-5 describes the rules for assigning ingress IP interfaces to IP VPNs.

Table 16-5. Ingress IP Interface Assignment Rules

Interface Type	You assign...	See...
IP logical port that supports Ethernet or PPP.	An IP logical port (and its configured IP interface) to a single IP VPN.	"Assigning an Ethernet or PPP IP Logical Port to an IP VPN" on page 16-41.
IP logical port that supports Frame Relay.	All of the port's DLCIs to a single IP VPN or distribute the DLCIs to multiple IP VPNs.	"Assigning IP Logical Port DLCIs or VPI/VCIs to IP VPNs" on page 16-42.
IP logical port or IP server logical port that supports ATM[1].	All of the port's VPI/VCIs to a single IP VPN or distribute the VPI/VCIs to multiple IP VPNs.	"Assigning IP Logical Port DLCIs or VPI/VCIs to IP VPNs" on page 16-42.

[1] In addition, for an IP server logical port, you must create an IP server PVC for each VPI/VCI you want to assign to an IP VPN. See "IP Server PVCs on the CBX 500" on page 3-33 for more information on creating IP Server PVCs.

After you assign ports, DLCIs, and VPI/VCIs to an IP VPN, connect the IP VPN's virtual routers at the edge of the Lucent network with point-to-point LSPs, and configure additional IP resources (such as route maps and static routes) for the IP VPN. See "Assigning Point-to-Point LSPs to an IP VPN" on page 16-57 and "Assigning Additional Network Resources to an IP VPN" on page 16-57.

Assigning an Ethernet or PPP IP Logical Port to an IP VPN

Before you assign an IP logical port that uses Ethernet or PPP for the data link protocol to an IP VPN, make sure that you have added and configured the IP logical port (including the IP interface) as described in "Configuring IP Logical Ports" on page 3-5.

> ► You may assign an IP logical port that uses Ethernet or PPP as the data link protocol to only one IP VPN.

You create the assignment from the Set IP Parameters dialog box. To assign an Ethernet or PPP port to an IP VPN:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set all IP LPorts dialog box appears (see Figure 16-22 on page 16-34).
2. Select the switch where the Ethernet or PPP logical port resides from the Switch Name list at the top of the dialog box. A list of IP logical ports configured on the switch appears in the LPort Name list.
3. Select the Ethernet or PPP logical port.
4. Choose IP Parameters. The Set IP Parameters dialog box appears (see Figure 16-26).

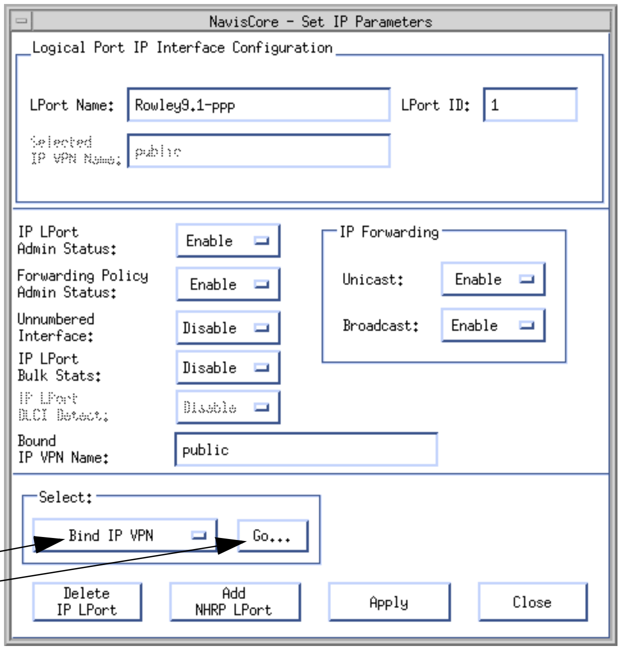

NavisCore - Set IP Parameters

Logical Port IP Interface Configuration

LPort Name: Rowley9.1-ppp
LPort ID: 1
Selected IP VPN Name: public

IP LPort Admin Status: Enable
Forwarding Policy Admin Status: Enable
Unnumbered Interface: Disable
IP LPort Bulk Stats: Disable
IP LPort DLCI Detect: Disable
Bound IP VPN Name: public

IP Forwarding
Unicast: Enable
Broadcast: Enable

Select: Bind IP VPN
Go...

Delete IP LPort | Add NHRP LPort | Apply | Close

1. Select Bind IP VPN.

2. Choose Go.

Figure 16-26. Set IP Parameters Dialog Box (PPP Port)

5. Select Bind IP VPN.

6. Choose Go. The Select IP VPN dialog box appears (see Figure 16-21 on page 16-33).
7. Select the IP VPN (the default is “public,” which means that the Ethernet or PPP logical port can be used by public network traffic). Note that only one VPN can use a PPP or Ethernet logical port.
8. Choose OK.

> If the IP logical port was previously bound to another private IP VPN, an error message appears at this time, notifying you that you must delete all IP interfaces configured for the IP logical port. When an IP logical port is bound to a private IP VPN, you cannot bind the port to another private IP VPN until you delete the IP interfaces associated with the IP logical port.

Assigning IP Logical Port DLCIs or VPI/VCIs to IP VPNs

Before you assign a Frame Relay DLCI or ATM VPI/VCI (and associated parameters) to an IP VPN, make sure that you have added and configured the associated IP logical port as described in “Configuring IP Logical Ports” on page 3-5.

> The DLCI information in this section applies to IP logical ports that use Frame Relay as the data link protocol on both B-STDX switches and CBX 500 switches. The VPI/VCI information in this section applies to IP logical ports that use ATM as the data link protocol on B-STDX switches only.

For each IP logical port that supports Frame Relay, you can assign multiple DLCIs to a single VPN, or distribute the DLCIs among multiple VPNs.

For each IP logical port that supports ATM, you can assign multiple VPI/VCIs to a single VPN, or distribute the VPI/VCIs among multiple VPNs.

The procedure for assigning DLCIs and VPI/VCIs is divided into two parts:

- Assign a DLCI or a VPI/VCI to an IP VPN.
- Bind a DLCI or a VPI/VCI to one of the IP VPN’s IP interfaces, which you can add during the binding process.

To assign a DLCI or VPI/VCI to a VPN:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set all IP LPorts dialog box appears (see Figure 16-22 on page 16-34).
2. Select the switch where the IP logical port that supports Frame Relay or ATM resides from the Switch Name list at the top of the dialog box. A list of IP logical ports configured on the switch appears in the LPort Name list.
3. Select the IP logical port that supports Frame Relay or ATM.

4. Choose IP Parameters. The Set IP Parameters dialog box appears (see Figure 16-27 and Figure 16-28).

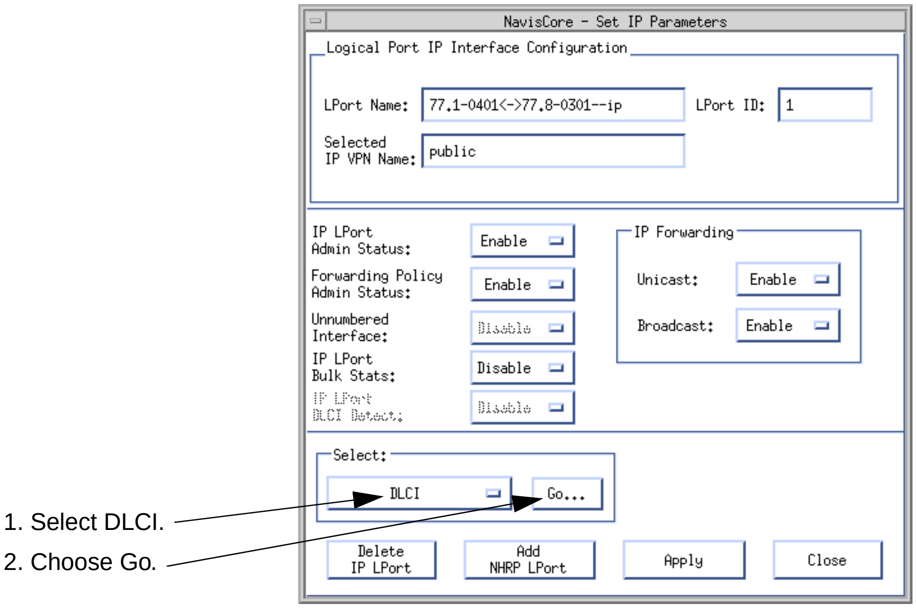

NavisCore - Set IP Parameters

Logical Port IP Interface Configuration

LPort Name: 77.1-0401<->77.8-0301--ip

LPort ID: 1

Selected IP VPN Name: public

IP LPort Admin Status: Enable

Forwarding Policy Admin Status: Enable

Unnumbered Interface: Disable

IP LPort Bulk Stats: Disable

IP LPort DLCI Detect: Disable

IP Forwarding
Unicast: Enable
Broadcast: Enable

Select: DLCI

1. Select DLCI.
2. Choose Go.

Figure 16-27. Set IP Parameters Dialog Box (DLCI)

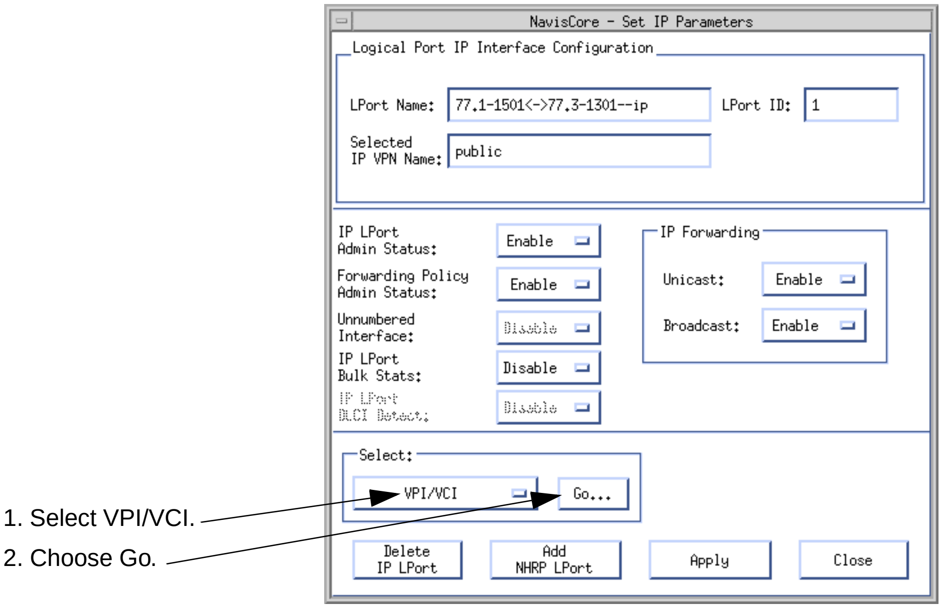

NavisCore - Set IP Parameters

Logical Port IP Interface Configuration

LPort Name: 77.1-1501<->77.3-1301--ip

LPort ID: 1

Selected IP VPN Name: public

IP LPort Admin Status: Enable

Forwarding Policy Admin Status: Enable

Unnumbered Interface: Disable

IP LPort Bulk Stats: Disable

IP LPort DLCI Detect: Disable

IP Forwarding
Unicast: Enable
Broadcast: Enable

Select: VPI/VCI

1. Select VPI/VCI.
2. Choose Go.

Figure 16-28. Set IP Parameters Dialog Box (VPI/VCI)

5. Select DLCI or VPI/VCI.

6. Choose Go. The Set All IP Interface Data Link IDs dialog box appears (see Figure 16-29 and Figure 16-30).

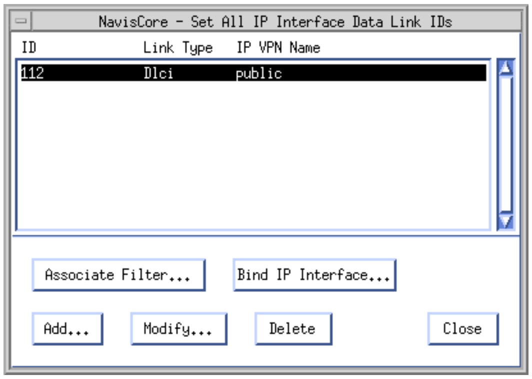

NavisCore - Set All IP Interface Data Link IDs

ID	Link Type	IP VPN Name
112	Dlci	public

Figure 16-29. Set All IP Interface Data Link IDs Dialog Box (DLCI)

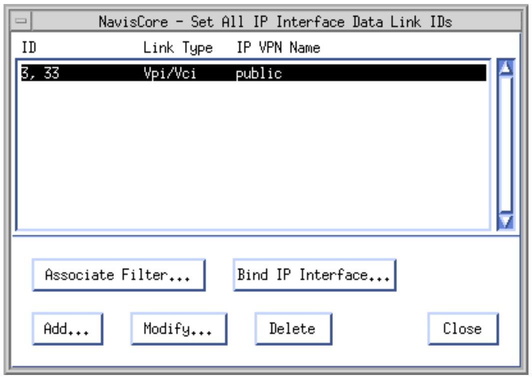

NavisCore - Set All IP Interface Data Link IDs

ID	Link Type	IP VPN Name
3, 33	Vpi/Vci	public

Figure 16-30. Set All IP Interface Data Link IDs Dialog Box (VPI/VCI)

The Set All IP Interface Data Link IDs dialog box displays a list of DLCIs or VPI/VCIs and the IP VPNs to which you have assigned them. Table 16-6 describes the buttons on the dialog box.

Table 16-6. Set All IP Interface Data Link IDs Dialog Box Buttons

Button	Function
Associate Filter	Associates an IP packet filter with the Frame Relay circuit identified by the DLCI or the ATM circuit identified by the VPI/VCI. See Chapter 4, “Configuring IP Packet Filters” for more information on IP packet filters and associating them with circuits.
Bind IP Interface	Binds an IP interface to the DLCI or VPI/VCI.
Add	Assigns a DLCI or VPI/VCI to a VPN.
Modify	Modifies a DLCI/VPN or VPI/VCI/VPN assignment.
Delete	Deletes a DLCI/VPN or VPI/VCI/VPN assignment.

If you do not want to make any new DLCI or VPI/VCI assignments, proceed to "Binding an IP Interface to a DLCI or a VPI/VCI" on page 16-53. Otherwise, proceed to the next step.

7. Choose Add. The Add Protocol Connection ID dialog box appears.

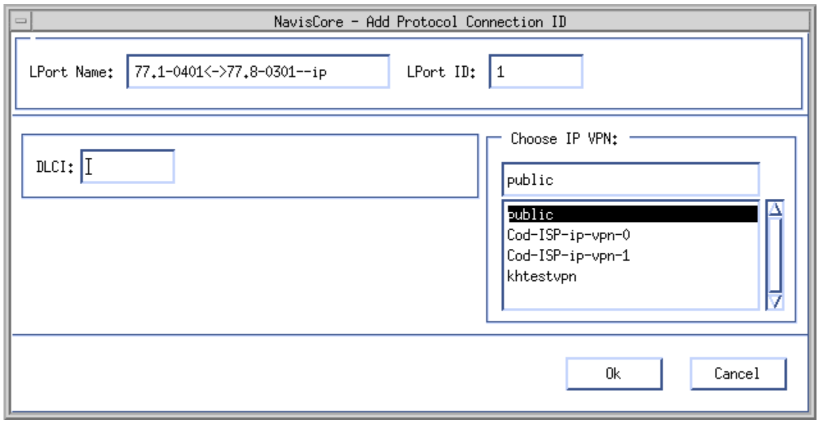

NavisCore - Add Protocol Connection ID

LPort Name: 77.1-0401<->77.8-0301--ip

LPort ID: 1

DLCI:

Choose IP VPN: public

Figure 16-31. Add Protocol Connection ID Dialog Box (DLCI)

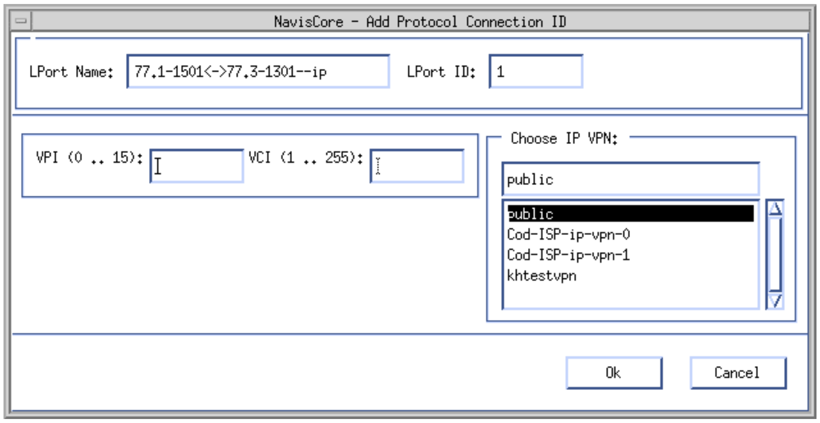

NavisCore - Add Protocol Connection ID

LPort Name: 77.1-1501<->77.3-1301--ip

LPort ID: 1

VPI (0 .. 15):

VCI (1 .. 255):

Choose IP VPN: public

Figure 16-32. Add Protocol Connection ID Dialog Box (VPI/VCI)

8. Enter a valid DLCI in the DLCI field, or enter a valid VPI and VCI in the VPI and VCI fields. See the *NavisCore Frame Relay Configuration Guide* for more information on DLCIs. See the *NavisCore ATM Configuration Guide* for more information on VPI/VCIs.

9. Select an IP VPN.

10. Choose OK. The Set All IP Interface Data Link IDs dialog box appears (see Figure 16-29 on page 16-44), displaying the new DLCI/VPN or VPI/VCI/VPN assignment.

11. Bind an IP interface to the DLCI or VPI/VCI. See "Binding an IP Interface to a DLCI or a VPI/VCI" on page 16-53 for more information.

Assigning IP Server Logical Port VPI/VCIs to IP VPNs

Before you assign an IP server logical port's VPI/VCI (and associated parameters) to an IP VPN, verify the following:

- You have added and configured the IP server logical port as described in Chapter 3, "Configuring IP Logical Ports and IP Servers."
- You have configured an IP server PVC on the IP server logical port for each VPI/VCI you want to assign to an IP VPN.

▶ IP server logical ports apply to CBX 500 switches only.

The procedure for creating an IP VPN interface is divided into two parts:

- Assign a VPI/VCI to an IP VPN
- Bind the VPI/VCIs to an IP interface

You can assign a VPI/VCI to an IP VPN in two ways:

- By selecting Lucent IP Parameters ⇒ Set IP Servers ⇒ Set IP Server LPorts from the Administer menu
- By selecting Lucent IP Parameters ⇒ Set IP Servers ⇒ Set IP Server PVCs from the Administer menu

Assigning VPI/VCIs through the Set IP Server LPorts Selection

For IP server logical ports, you can assign VPI/VCIs to IP VPNs through the Set IP Server LPorts selection, as follows:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set IP Servers ⇒ Set IP Server LPorts. The Show IP Servers dialog box appears (see Figure 16-33).

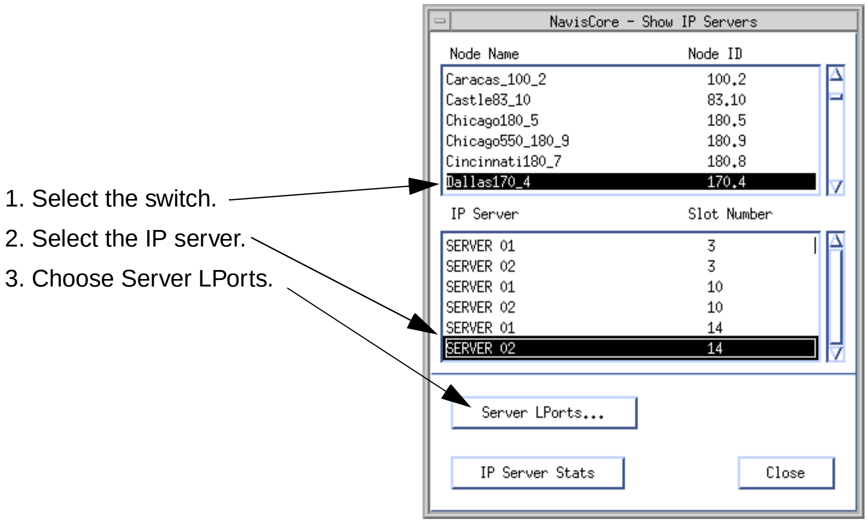

NavisCore - Show IP Servers

Node Name	Node ID
Caracas_100_2	100.2
Castle83_10	83.10
Chicago180_5	180.5
Chicago550_180_9	180.9
Cincinnati180_7	180.8
Dallas170_4	170.4

IP Server	Slot Number
SERVER 01	3
SERVER 02	3
SERVER 01	10
SERVER 02	10
SERVER 01	14
SERVER 02	14

1. Select the switch.
2. Select the IP server.
3. Choose Server LPorts.

Figure 16-33. Show IP Servers Dialog Box

2. Select the switch where the IP server logical port resides. A list of IP servers on the switch appears.
3. Select the IP server that contains the IP server logical port.
4. Choose Server LPorts. The Set All Logical Ports in IP Server PPort dialog box appears (see Figure 16-34).

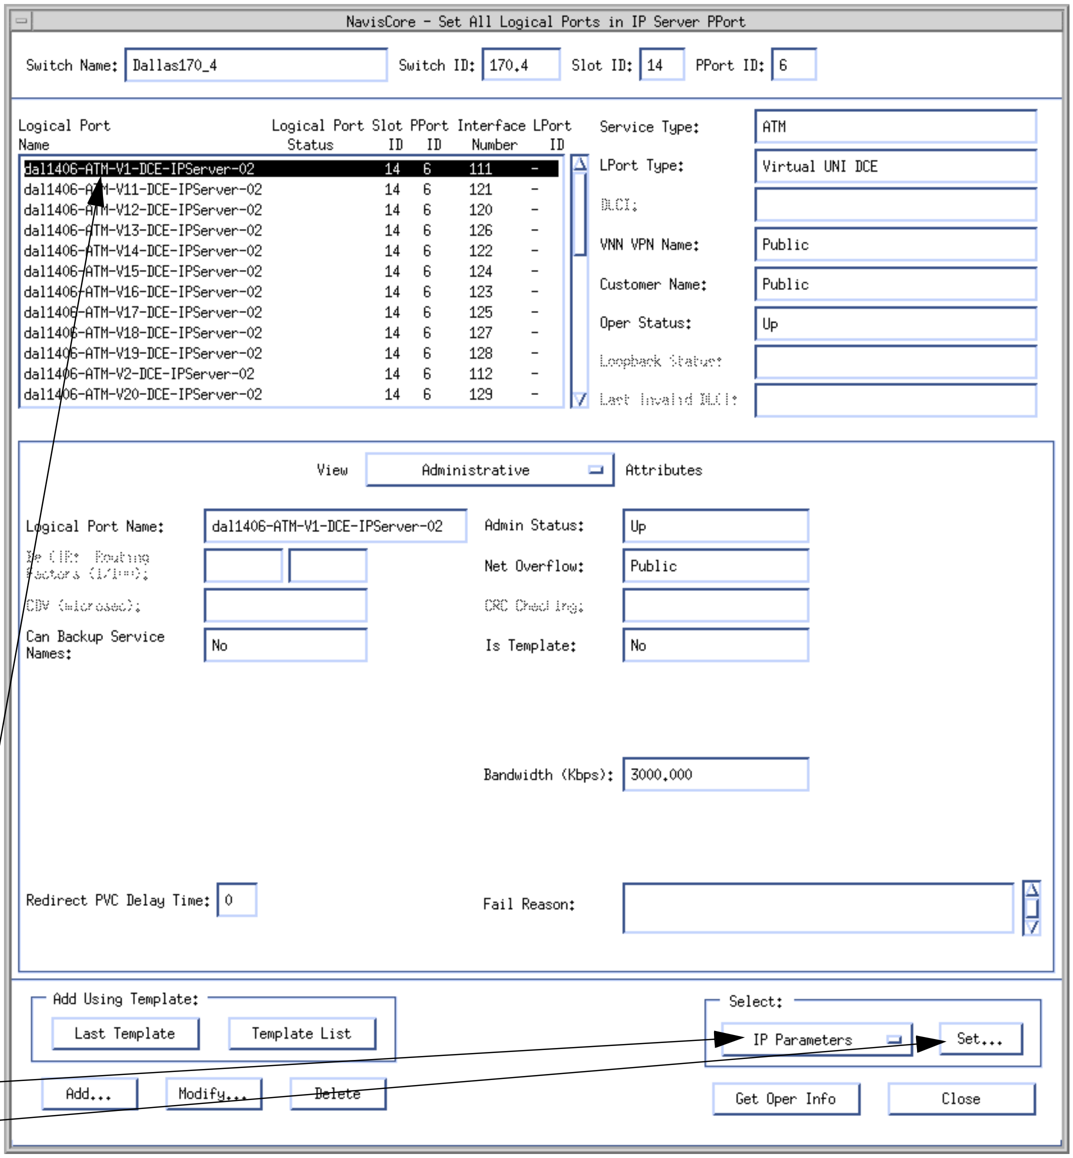

NavisCore - Set All Logical Ports in IP Server PPort

Switch Name: Dallas170_4 Switch ID: 170.4 Slot ID: 14 PPort ID: 6

Logical Port Name	Logical Port Status	Slot ID	PPort ID	Interface Number	LPort ID
dal1406-ATM-V1-DCE-IPServer-02		14	6	111	-
dal1406-ATM-V11-DCE-IPServer-02		14	6	121	-
dal1406-ATM-V12-DCE-IPServer-02		14	6	120	-
dal1406-ATM-V13-DCE-IPServer-02		14	6	126	-
dal1406-ATM-V14-DCE-IPServer-02		14	6	122	-
dal1406-ATM-V15-DCE-IPServer-02		14	6	124	-
dal1406-ATM-V16-DCE-IPServer-02		14	6	123	-
dal1406-ATM-V17-DCE-IPServer-02		14	6	125	-
dal1406-ATM-V18-DCE-IPServer-02		14	6	127	-
dal1406-ATM-V19-DCE-IPServer-02		14	6	128	-
dal1406-ATM-V2-DCE-IPServer-02		14	6	112	-
dal1406-ATM-V20-DCE-IPServer-02		14	6	129	-

Service Type: ATM
LPort Type: Virtual UNI DCE
DLCI:
VNN VPN Name: Public
Customer Name: Public
Oper Status: Up
Loopback Status:
Last Invalid DLCI:

View Administrative Attributes

Logical Port Name: dal1406-ATM-V1-DCE-IPServer-02
Eq CIR: Routing Factors (1/1000):
CDV (microsec):
Can Backup Service Names: No
Admin Status: Up
Net Overflow: Public
CRC Checking:
Is Template: No
Bandwidth (Kbps): 3000.000
Redirect PVC Delay Time: 0
Fail Reason:

1. Select an IP Server logical port.
2. Select IP Parameters.
3. Choose Set.

Figure 16-34. Set All Logical Ports in IP Server PPort Dialog Box

5. Select the IP server logical port whose VPI/VCI will be associated with the IP VPN.
6. Select IP Parameters.
7. Choose Set. The Set IP Parameters dialog box appears (see Figure 16-35).

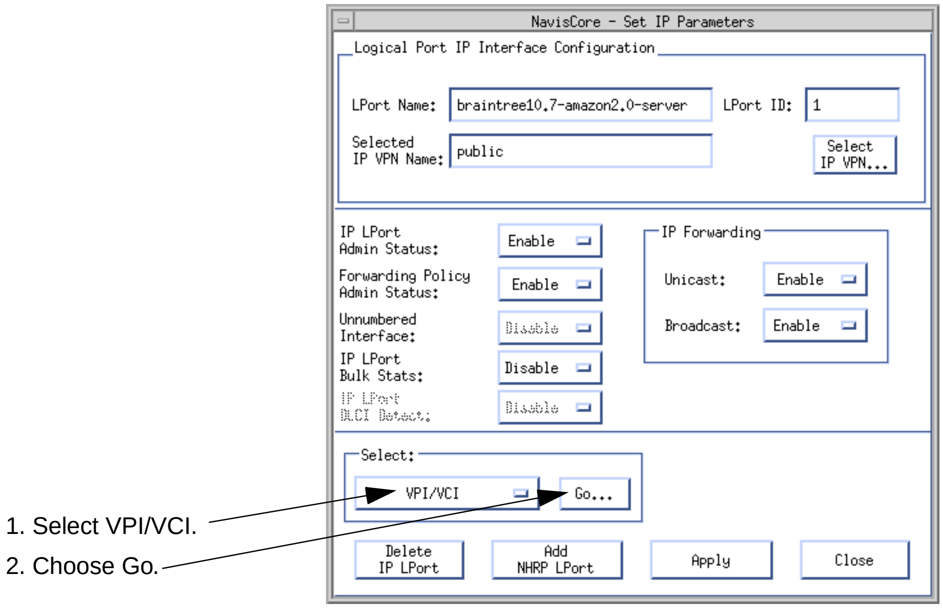

LPort Name: braintree10.7-amazon2.0-server
LPort ID: 1
Selected IP VPN Name: public
IP LPort Admin Status: Enable
Forwarding Policy Admin Status: Enable
Unnumbered Interface: Disable
IP LPort Bulk Stats: Disable
IP LPort DLCI Detect: Disable
Unicast: Enable
Broadcast: Enable
Select: VPI/VCI

1. Select VPI/VCI.
2. Choose Go.

Figure 16-35. Set IP Parameters Dialog Box (IP Server VPI/VCI)

8. Select VPI/VCI.
9. Choose Go. The Set All IP Interface Data Link IDs dialog box appears (see Figure 16-36).

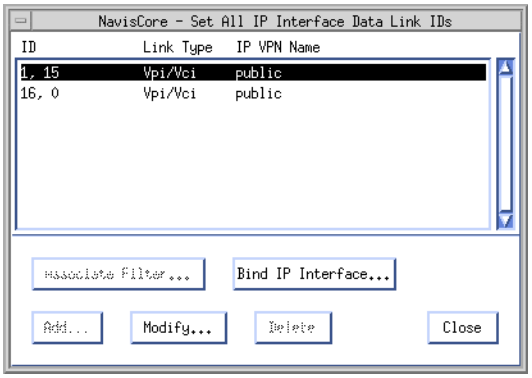

ID	Link Type	IP VPN Name
1, 15	Vpi/Vci	public
16, 0	Vpi/Vci	public

Figure 16-36. Set All IP Interface Data Link IDs Dialog Box (IP Server VPI/VCI)

The Set All IP Interface Data Link IDs dialog box displays the VPI/VCIs for the IP server logical port and the IP VPNs to which they are assigned. Each VPI/VCI is associated with an IP server PVC. Table 16-7 describes the buttons on the dialog box.

Table 16-7. Set All IP Interface Data Link IDs Dialog Box (For VPI/VCI) Buttons

Button	Function
Bind IP Interface	Binds an IP interface to the VPI/VCI.
Modify	Modifies the VPI/VCI assignment. Choose Modify to assign the VPI/VCI to a VPN other than "Public."

If you do not want to change the VPI/VCI assignments (VPI/VCIs are assigned to the public IP VPN by default), proceed to "Binding an IP Interface to a DLCI or a VPI/VCI" on page 16-53. Otherwise, proceed to the next step.

10. Choose Modify. The Modify Protocol Connection ID dialog box appears (see Figure 16-37).

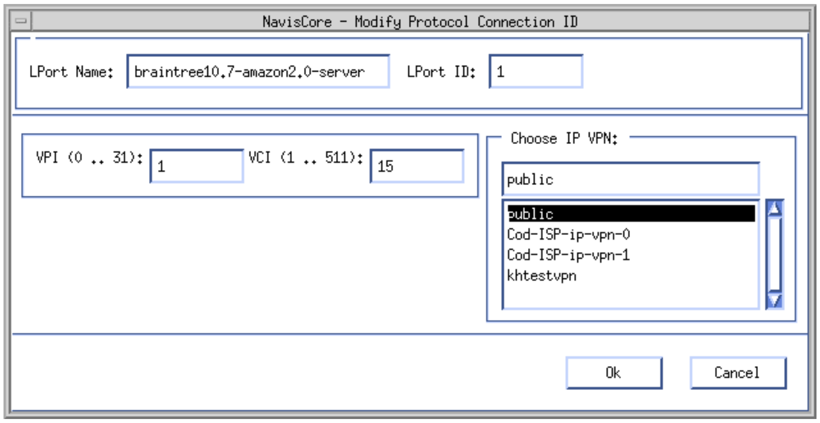

LPort Name: braintree10.7-amazon2.0-server

LPort ID: 1

VPI (0 .. 31): 1

VCI (1 .. 511): 15

Choose IP VPN: public

public
Cod-ISP-ip-vpn-0
Cod-ISP-ip-vpn-1
khtestvpn

Figure 16-37. Modify Protocol Connection ID Dialog Box (IP Server VPI/VCI)

11. Select an IP VPN. Note that you cannot modify the VPI and VCI values on this dialog box. You must modify these values by modifying the associated IP server PVC.

12. Choose OK. The Set All IP Interface Data Link IDs dialog box appears (see Figure 16-36 on page 16-49), displaying the new VPI/VCI/VPN assignment.

13. Bind an IP interface to the VPI/VCI. See "Binding an IP Interface to a DLCI or a VPI/VCI" on page 16-53 for more information.

Assigning VPI/VCIs through the Set IP Server PVCs Selection

For IP server logical ports, you can assign VPI/VCIs to IP VPNs through the Set IP Server PVCs selection, as follows:

1. On the network map, select a switch where an IP server PVC endpoint is configured.
2. From the Administer menu, select Lucent IP Parameters ⇒ Set IP Servers ⇒ Set IP Server PVCs. The Set All IP Server PVCs dialog box appears (see Figure 16-38).

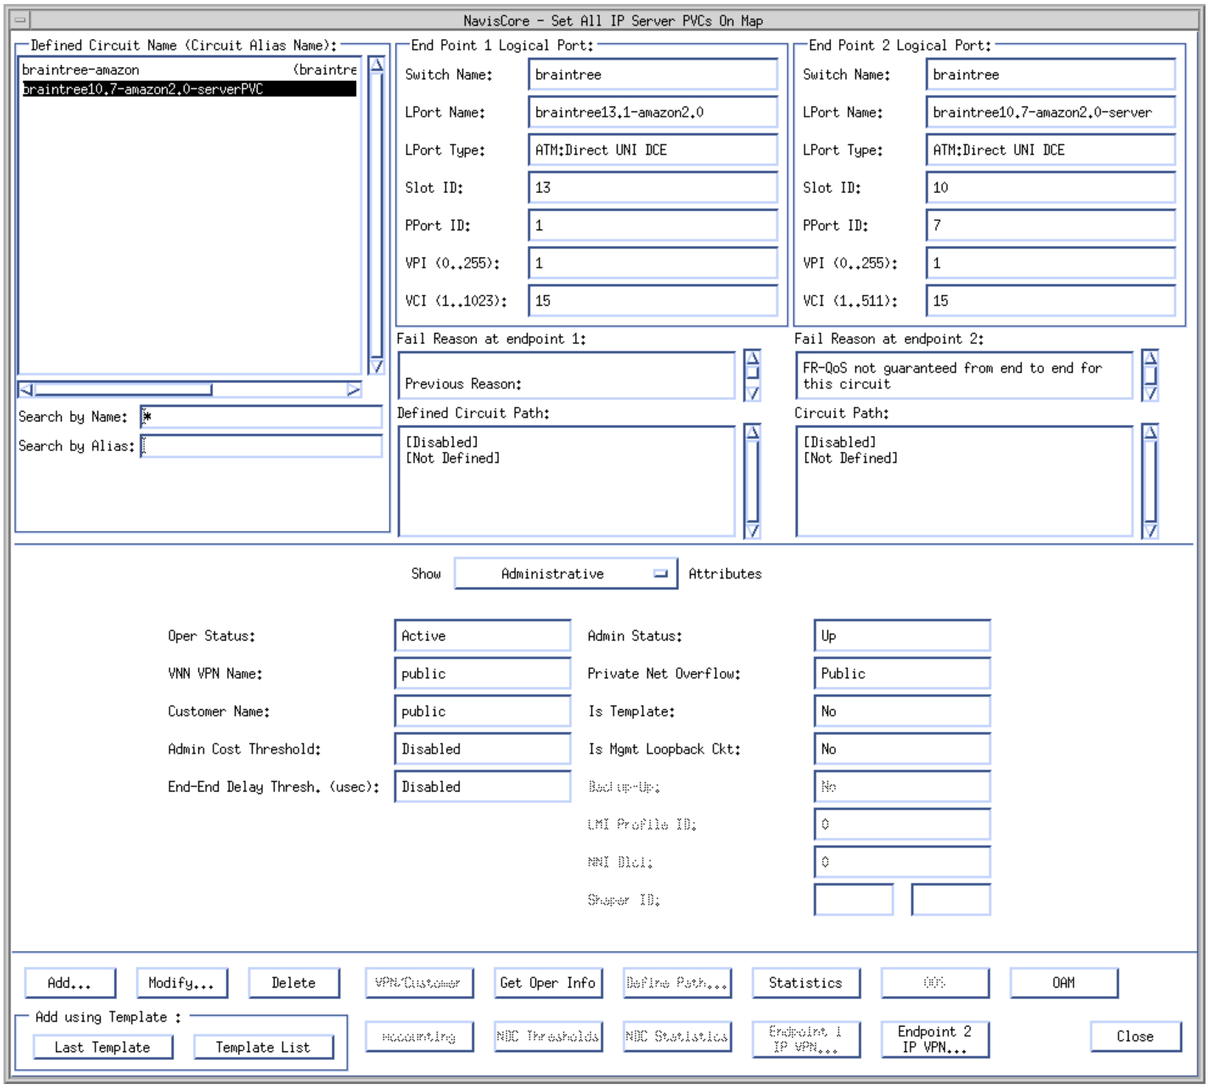
NavisCore - Set All IP Server PVCs On Map

Defined Circuit Name (Circuit Alias Name):
braintree-amazon (braintre
braintree10.7-amazon2.0-serverPVC

Search by Name: *
Search by Alias:

End Point 1 Logical Port:		End Point 2 Logical Port:	
Switch Name:	braintree	Switch Name:	braintree
LPort Name:	braintree13.1-amazon2.0	LPort Name:	braintree10.7-amazon2.0-server
LPort Type:	ATM:Direct UNI DCE	LPort Type:	ATM:Direct UNI DCE
Slot ID:	13	Slot ID:	10
PPort ID:	1	PPort ID:	7
VPI (0..255):	1	VPI (0..255):	1
VCI (1..1023):	15	VCI (1..511):	15

Fail Reason at endpoint 1:
Previous Reason:

Fail Reason at endpoint 2:
FR-QoS not guaranteed from end to end for this circuit

Defined Circuit Path:
[Disabled]
[Not Defined]

Circuit Path:
[Disabled]
[Not Defined]

Show Administrative Attributes

Oper Status:	Active	Admin Status:	Up
VNN VPN Name:	public	Private Net Overflow:	Public
Customer Name:	public	Is Template:	No
Admin Cost Threshold:	Disabled	Is Mgmt Loopback Ckt:	No
End-End Delay Thresh. (usec):	Disabled	Backup-Up:	No
		LMI Profile ID:	0
		NNI Dlci:	0
		Shaper ID:	

Figure 16-38. Set All IP Server PVCs Dialog Box

▶ If the circuit names do not appear immediately in the Defined Circuit Name list box, insert the cursor in the blank Search by Name field and press Return. This action will display a list of configured circuits.

3. Select the IP server PVC. Use the Search by Name and Search by Alias fields to enter wildcard characters:

- Use an * to match any number of characters
- Use a ? to match a single character
- Type * to match the * character
- Type \? to match the ? character
- Type \\ to match the \ character

See "Creating an IP Server PVC" on page 3-33 for a description of all the fields and buttons on the Set All IP Server PVCs dialog box.

4. Choose either Endpoint 1 IP VPN or Endpoint 2 IP VPN. The Set All IP Interface Data Link IDs dialog box appears (see Figure 16-36 on page 16-49). Note that if one of the endpoints is not an IP server logical port (for example, an ATM UNI logical port), that associated endpoint button is grayed out.

5. Choose Modify. The Modify Protocol Connection ID dialog box appears (see Figure 16-37 on page 16-50).

6. Select an IP VPN. Note that you cannot modify the VPI and VCI values on this dialog box. You must modify these values by modifying the associated IP server PVC.

7. Choose OK. The Set All IP Interface Data Link IDs dialog box appears (see Figure 16-36 on page 16-49), displaying the new VPI/VCI/VPN assignment.

8. Bind an IP interface to the VPI/VCI. See "Binding an IP Interface to a DLCI or a VPI/VCI" on page 16-53 for more information.

Binding an IP Interface to a DLCI or a VPI/VCI

For each IP VPN, you create its own IP interfaces on a switch. You then must bind one of these interfaces to each DLCI or VPI/VCI you assign to an IP VPN. You bind and create IP interfaces during the same process.

To bind an IP interface to a DLCI or VPI/VCI (and create it):

1. At the Set All IP Interface Data Link IDs dialog box (see Figure 16-29 on page 16-44, Figure 16-30 on page 16-44, or Figure 16-36 on page 16-49), select the DLCI or VPI/VCI to which you want to bind an IP interface.
2. Choose Bind IP Interface. The Bind IP Interface Address to Protocol ID dialog box appears (see Figure 16-39).

> Although the dialog box in Figure 16-39 is for a DLCI, the dialog box for a VPI/VCI is almost identical. The only difference is that the dialog box for a VPI/VCI displays the VPI and VCI instead of the DLCI.

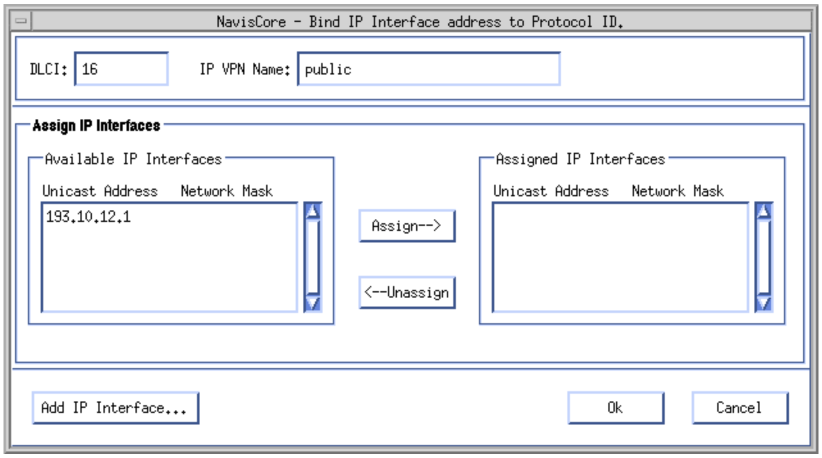

NavisCore - Bind IP Interface address to Protocol ID.

DLCI: 16

IP VPN Name: public

Available IP Interfaces: 193.10.12.1

Figure 16-39. Bind IP Interface Address to Protocol ID Dialog Box (DLCI)

3. Choose Add IP Interface to add an IP interface to which the DLCI or VPI/VCI will be bound. The Set IP Interface Address dialog box appears (see Figure 16-40).

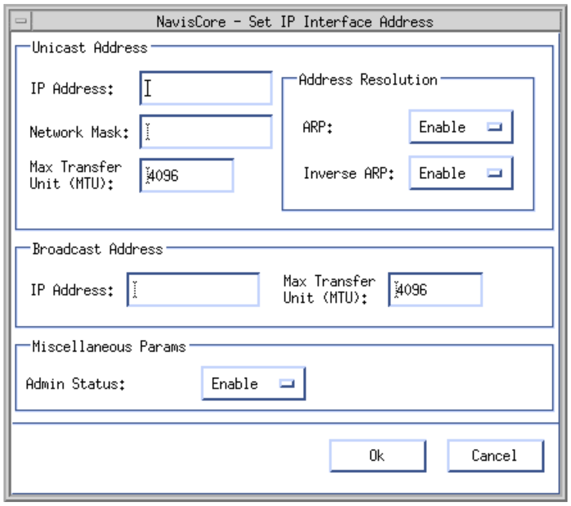

NavisCore - Set IP Interface Address

Unicast Address

IP Address:

Network Mask:

Max Transfer Unit (MTU): 4096

Address Resolution

ARP: Enable

Inverse ARP: Enable

Broadcast Address

IP Address:

Max Transfer Unit (MTU): 4096

Miscellaneous Params

Admin Status: Enable

Figure 16-40. Set IP Interface Address Dialog Box

4. Complete the fields in Table 16-8.

Table 16-8. Ingress IP Interface Address Fields

Field	Action/Description
Unicast Address	
IP Address	The IP address for this interface. Keep in mind that, once you configure an IP address, the IP address cannot be changed. To make any modifications to the IP address, you must delete the interface and re-add it with the modified IP address.
Network Mask	The mask used to determine the subnet of this IP interface. Once this value is set, you cannot use the Modify Interface Address function to modify the network mask value. In order to change the network mask, you must delete the IP interface and then add a new one using the correct network mask.
Max Transfer Unit (MTU)	The maximum size of a packet that can be sent through the physical port. The default value for this field varies depending on the logical port type as follows: **LPort Type** **Default** ATM 9180 Frame Relay 4096 Ethernet 1500
Address Resolution	
ARP (*Frame Relay and Ethernet only*)	Select one of the following options: *Enable* – (default) Enables the Address Resolution Protocol (ARP). *Disable* – Disables the ARP.

Table 16-8. Ingress IP Interface Address Fields (Continued)

Field	Action/Description
Inverse ARP (*Frame Relay and ATM only*)	Select one of the following options: *Enable* – (default) Enables the Inverse Address Resolution Protocol (InARP). *Disable* – Disables the InARP.
Broadcast Address	
IP Address	The address used by this interface for subnet broadcasting.
Max Transfer Unit (MTU)	The maximum size of a packet that can be sent through the physical port. The default value for this field varies depending on the logical port type as follows: **LPort Type** **Default** ATM 9180 Frame Relay 4096 Ethernet 1500
Miscellaneous Params	
Admin Status	Select one of the following options: *Enable* – (default) Enables the IP interface. *Disable* – Disables the IP interface.

5. Choose OK. The Bind IP Interface Address to Protocol ID dialog box (see Figure 16-39 on page 16-53) appears. By default, the IP interface is bound to the DLCI or VPI/VCI you selected in step 1 on page 16-53 (that is, it appears in the Assigned IP Interfaces list). If you want, you can select the IP interface and choose Unassign to make the IP interface available to other DLCIs or VPI/VCIs.

6. Choose OK.

▶ At any time, you can delete the IP interface binding by selecting the IP interface from the list of assigned IP interfaces and choosing Unassign. Choose OK after you have deleted the binding.

Managing IP VPN Ingress IP Interfaces

To manage ingress IP interfaces for an IP VPN once you have configured them:

1. From the Administer menu, select Lucent IP Parameters ⇒ Set All IP LPorts. The Set All IP LPorts dialog box appears (see Figure 16-22 on page 16-34).
2. Choose Select IP VPN. The Select IP VPN dialog box appears (see Figure 16-21 on page 16-33).
3. Select the IP VPN to which the ingress IP interface was assigned.
4. Choose OK. The Set All IP LPorts dialog box appears (see Figure 16-22 on page 16-34).
5. Select the switch where the ingress IP interface resides. A list of IP logical ports appears.
6. Select the IP logical port associated with the ingress IP interface.
7. Choose IP Parameters. The Set IP Parameters dialog box appears.
8. Select Actions ⇒ IP Interface.
9. Choose Go. The Set IP Interface Addresses dialog box appears, displaying all of the ingress IP interfaces assigned to the selected IP VPN on the IP logical port.

From the Set IP Interface Addresses dialog box, you can manage all of the IP VPN's ingress IP interfaces associated with the selected IP logical port. For example, you can add additional ingress IP interfaces, modify ingress IP interfaces, and delete ingress IP interfaces.

You can also add, modify, and delete an IP OSPF interface. See Chapter 9, "Configuring IP OSPF and VNN OSPF" for more information on managing IP OSPF interfaces.

Assigning Point-to-Point LSPs to an IP VPN

You can configure point-to-point LSPs to connect the VPN’s virtual routers at the edge of the Lucent network. See “Understanding an IP VPN” on page 16-2, “Identify Point-to-Point LSPs” on page 16-21, and Chapter 12, “Configuring Label Switched Paths” for more information.

Assigning Additional Network Resources to an IP VPN

You can assign additional IP network resources to an IP VPN. These resources include the same kinds of resources you can manage for traditional IP networks, such as static routes, static ARP entries, route maps, and so on.

To assign additional IP network resources for an IP VPN, you must first select the VPN (see “Selecting the IP VPN” on page 16-33). By selecting a VPN, management tasks you perform will be in the context of a single VPN. For example, if you add a static route, that route is added for the selected VPN only.

Once you select a VPN, the NavisCore dialog boxes you access display the selected VPN’s name. Figure 16-41, Figure 16-42, and Figure 16-43 provide sample dialog boxes that display the name of a selected VPN. Notice that all of these dialog boxes display the Select IP VPN button, which allows you to access an IP VPN from the current dialog box you are viewing.

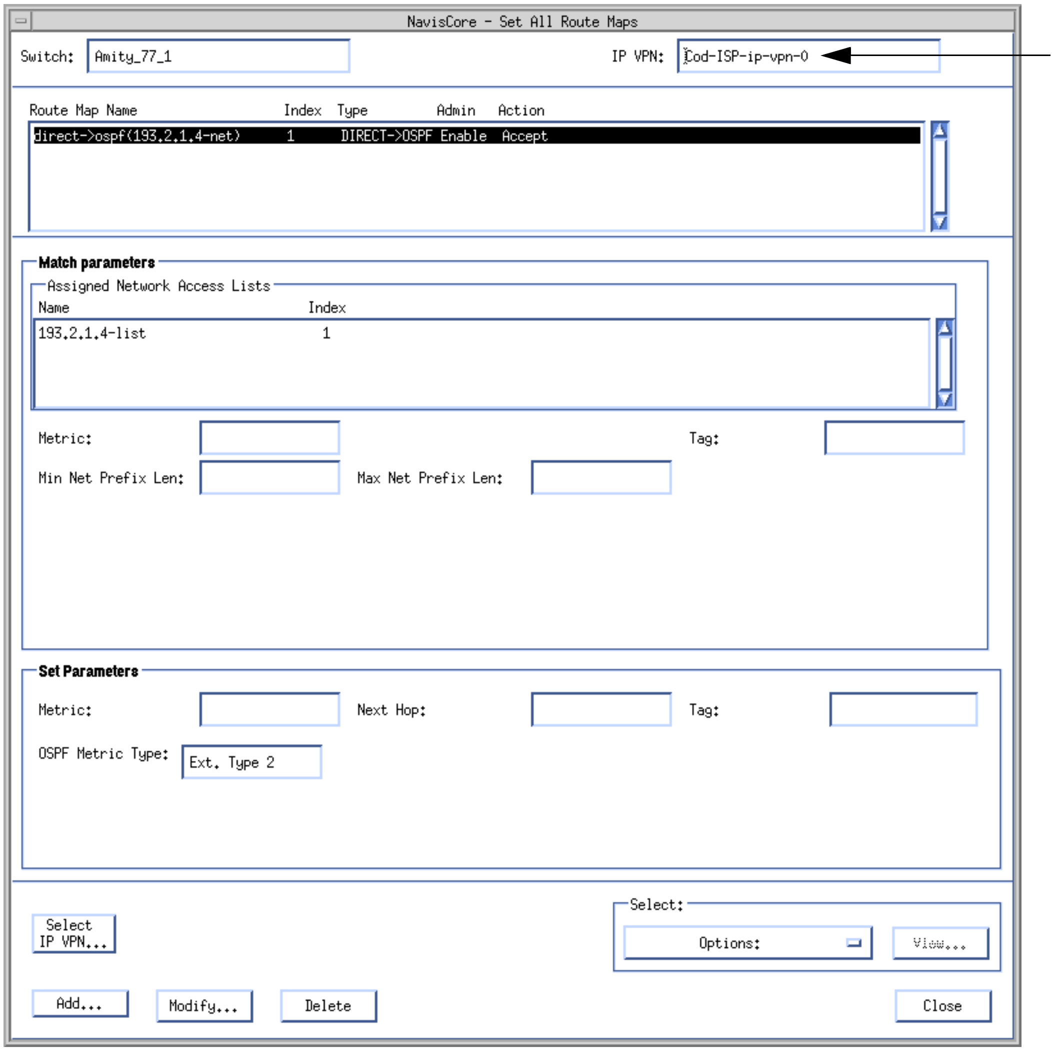

NavisCore - Set All Route Maps

Switch: Amity_77_1

IP VPN: Cod-ISP-ip-vpn-0

Route Map Name	Index	Type	Admin	Action
direct->ospf(193.2.1.4-net)	1	DIRECT->OSPF	Enable	Accept

Match parameters

Assigned Network Access Lists

Name	Index
193.2.1.4-list	1

Metric:

Tag:

Min Net Prefix Len:

Max Net Prefix Len:

Set Parameters

Metric:

Next Hop:

Tag:

OSPF Metric Type: Ext. Type 2

Figure 16-41. Sample Set All Route Maps Dialog Box

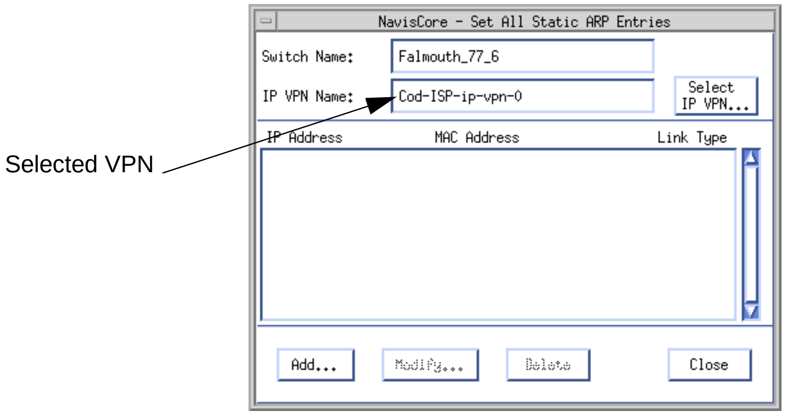

NavisCore - Set All Static ARP Entries

Switch Name: Falmouth_77_6

IP VPN Name: Cod-ISP-ip-vpn-0

Figure 16-42. Sample Set All Static ARP Entries Dialog Box

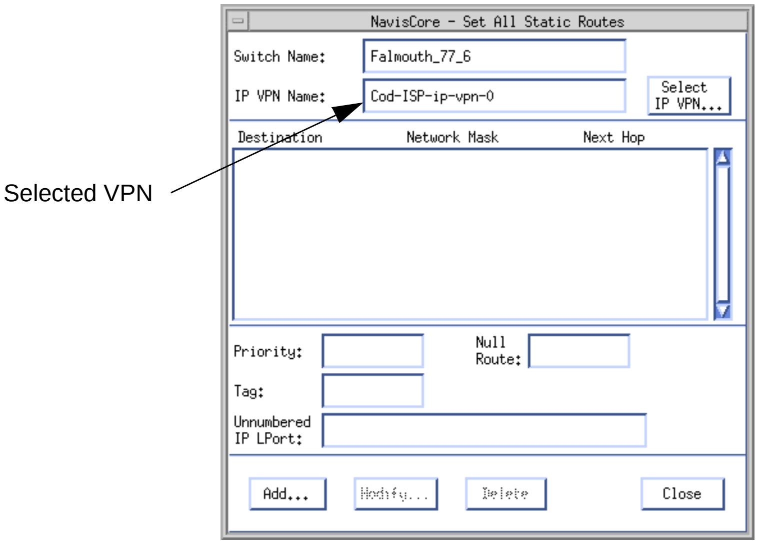

NavisCore - Set All Static Routes

Switch Name: Falmouth_77_6

IP VPN Name: Cod-ISP-ip-vpn-0

Figure 16-43. Sample Set All Static Routes Dialog Box

Table 16-9 provides more information on the IP network resources you can manage in the context of a private IP VPN.

Table 16-9. Private IP VPN Resource Management and References

Network Resource	See...
IP logical ports and associated resources (except for RIP)	Chapter 3, “Configuring IP Logical Ports and IP Servers”
IP OSPF interfaces	“Configuring IP OSPF on the IP Logical Port” on page 9-30
IP OSPF neighbors	“Configuring IP OSPF Neighbors” on page 9-34
IP OSPF area aggregates	“Configuring IP OSPF Area Aggregates” on page 9-36
IP OSPF virtual links	“Configuring IP OSPF Virtual Links” on page 9-38
IP OSPF route maps	“Configuring IP OSPF Route Maps” on page 9-41
IP OSPF router IDs	“Configuring IP OSPF Router IDs” on page 9-42
Network filters	“Adding a Network Filter” on page 11-13
Network access lists	“Adding a Network Access List” on page 11-15
RIP interfaces	“Configuring RIP at the Logical Port” on page 7-1
Route maps	“Adding Route Maps” on page 11-18
Static routes	“Configuring a Static Route” on page 10-2
Static ARP entries	“Defining a Static ARP Entry” on page 6-2
Cloud VPN IP interfaces	“Setting IP VPN Cloud IP Interfaces” on page 16-35
IP server logical ports	“Configuring IP Server Logical Ports” on page 3-28
IP server PVCs	“IP Server PVCs on the CBX 500” on page 3-33
Point-to-point LSPs	“Configuring Point-to-Point LSP Connections” on page 12-13
Forwarding policies	“Defining a Forwarding Policy” on page 5-28

Network resources not mentioned in Table 16-9 (such as NHRP resources and packet filters) cannot be reserved for private IP VPNs. They are public network resources only.

Telneting Into an IP VPN

You can telnet into a switch console within the context of an IP VPN. You can then manage resources assigned to the IP VPN on that switch using console commands. This feature allows customers to telnet from CPE into switches.

To telnet into a switch console within the context of an IP VPN:

1. Issue the telnet command and supply an IP address assigned to the IP VPN on the switch. The IP address can be assigned to an ingress IP interface or to an IP VPN cloud interface.
2. When prompted for a password, specify the IP VPN's telnet password. This password was defined when the IP VPN was added. See "Adding an IP VPN" on page 16-28 for more information on adding an IP VPN.

▶

Do not use the HP OpenView telnet feature (accessed through Misc ⇒ Terminal Connect) to telnet to a switch within an IP VPN context. Instead, use a telnet program that is not part of HP OpenView.

If Radius authentication is enabled for the IP VPN, a Radius server must authenticate the telnet login. Otherwise, the login will fail.

It is not possible to telnet from one IP VPN to another; you can only telnet within an IP VPN.

Avoid starting a telnet session while you are in a telnet session. This uses additional memory.

For example, suppose that an IP VPN has an ingress IP interface (193.2.1.5) and an IP VPN cloud interface (132.125.1.1) on a switch. You could specify either IP address when you issue the telnet command. You decide to use the IP address assigned to the ingress IP interface:

```
telnet 193.2.1.5
```

The telnet password defined for the IP VPN is "boston." When prompted for the password, you would specify "boston."

You can also telnet to a switch within the context of the public IP VPN. Once logged in, you can issue the **`login ipvpn`** command to enter the context of a private IP VPN. See the next section for more information.

Logging in to an IP VPN at the Switch Console

You can telnet to the switch in the context of the public IP VPN and then log in to a private IP VPN at the switch console. If you do this, any console commands you issue act on the IP VPN only.

When you telnet to the switch, you can login at the switch console with the standard console password. Then, issue the **`login ipvpn`** command as follows:

```
login ipvpn vpn_id
```

You specify the ID of the IP VPN for *vpn_id*. For example, the following command logs you into IP VPN 2:

```
login ipvpn 2
```

After you issue the command, you are prompted for a password. Enter the telnet password that was defined for the IP VPN when it was added. See "Adding an IP VPN" on page 16-28 for more information on adding an IP VPN.

To exit an IP VPN, type **`exit`** and press Return at the console prompt.

To determine the IDs of IP VPNs that have a virtual router on the switch, type the following command:

```
show ip vrouter
```

The console output displays the IDs of all the IP VPNs in the "VPN" column.

Through NavisCore, you can also view IP VPN IDs by selecting Lucent IP Parameters ⇒ Set All IP VPNs from the NavisCore Administer menu.

A

PRAM Upload

This appendix describes the uses for the Upload PRAM feature and supported IP objects. For complete details about Upload PRAM, see the *NavisCore NMS Getting Started Guide*.

Using the Upload PRAM Command

Occasionally, the switch configuration file for a specific I/O module and the configuration stored in the NMS database do not match. This situation can occur when you upgrade your switch software, use a network management product to manage the switch, or use the MIB to change a switch configuration.

▶ If you remove an I/O Module from one switch and install this module in a second switch, you get a PRAM conflict. This happens because the module contains an unknown configuration. Do not use PRAM upload to clear this condition. Instead, use the Erase PRAM function to clear PRAM on this module; then reconfigure the module. Refer to the *NavisCore Getting Started Guide* for more information.

To resolve PRAM conflicts, use the Upload PRAM function to view the switch configuration file stored in PRAM. This enables you to compare the configuration file in the switch (PRAM) to the configuration file in the NMS database.

Using Upload PRAM After Configuring IP Objects

▶ Upload PRAM currently does not support MPT Point-to-Point connections. All other IP objects are supported. See the *NavisCore Console Command Reference Guide* for more information.

Before you use the Upload PRAM function, review the following points:

- If you configure IP parameters using the console commands instead of using NavisCore, you need to upload the switch configuration to the NMS database. See the *NavisCore Getting Started Guide* for detailed instructions about how to upload a switch configuration file.
- You can use Upload PRAM to add objects from switch PRAM to the NMS database, as long as the objects being added do not conflict with existing objects in the database; for example, if the NMS database already contains a switch with a switch that you are adding, there would be a conflict.
- Due to the interdependency of objects with other objects in the database, *be careful* when you use Upload PRAM to delete objects from the database. In general, do not create a situation where there are dangling objects (i.e., an object without a parent) in the switch before applying Upload PRAM.

  For example, deleting a logical port without first deleting all associated individual addresses or address screens, creates dangling objects and causes a problem during the Upload PRAM process.

B

Troubleshooting IP Navigator Problems

This appendix describes how to troubleshoot and understand IP Navigator problems on Lucent switches.

Identifying IP Navigator Problems

The first thing to remember when troubleshooting IP Navigator problems on a Lucent switch is that the problem may not be directly related to IP. The problem may be related instead to one of the following causes:

- Hardware failure, such as a bad cable.
- Data link protocols that encapsulate IP packets.

IP packets that pass through Lucent switches are encapsulated in one of the following protocols:

Frame Relay — Frame headers encapsulate IP packets that are forwarded over Frame Relay circuits.

ATM — ATM cell headers encapsulate IP packets that are forwarded over ATM circuits.

Ethernet — Ethernet headers encapsulate IP packets that are forwarded over Ethernet LANs.

PPP — PPP headers encapsulate IP packets that are forwarded over PPP interfaces.

When IP Navigator problems occur, you should first determine if Frame Relay, ATM, or Ethernet problems also exist. Check for traps, changes in network map object colors, LEDs changing from green to red, and user complaints that might indicate the presence of Frame Relay, ATM, or Ethernet problems. If problems in any of these areas coincide with IP problems, one (or more) of these areas may be the source of the problem. For more information on troubleshooting Frame Relay problems and ATM problems, refer to the *Naviscore Troubleshooting Guide* (Product Code: 80104).

Isolating IP Navigator Problems

If you determine that the problem is related to IP Navigator, then the problem is probably the result of configuration errors, such as a misconfigured IP address, ARP entry, routing table entry, or route maps.

You can isolate IP Navigator problems by performing the following steps:

1. Determine which nodes are unreachable (for example, switches, routers, hosts).
2. Determine the physical paths to the unreachable nodes.
3. Check the connections and configurations of each node and interface along the path (for example, routing tables, ARP entries, IP forwarding statistics).
4. Consult other administrative personnel. For example, you may determine that the source of the problem lies with customer premise equipment (CPE) at a customer site, and you need the assistance of administrative personnel at that site to correct the problem.

Verifying IP Navigator Problems

Use the NMS and console commands to help you verify IP Navigator problems. In particular, use the IP attributes and statistics described in the *NavisCore Diagnostics Guide* (Product Code: 80105) to verify the presence of problems, as well as traditional IP troubleshooting tools such as `ping` and `traceroute.` Additionally, use the `show ip forward statistics` command, described in "IP Forwarding Console Commands" on page B-66.

IP Navigator Limitations

The IP Navigator limits for this release are described in the "IP Navigator Limits" section in the Software Release Notice (SRN) for your switch software. Unless otherwise noted, the public and private IP VPN limits apply to both the B-STDX 8000/9000 and the CBX 500.

TCP/IP Programs

TCP/IP uses the following two programs in troubleshooting IP problems:

`ping` — An industry standard TCP/IP program that tests the connectivity with a remote host by sending Internet Control Message Protocol (ICMP) echo commands and then waiting for the corresponding replies. If `ping` receives at least one echo, it can deduce that the remote host is operational. `ping` is described in RFC2151.

`traceroute` — An industry standard TCP/IP program that lets you see the route that IP datagrams follow from one host to another. `traceroute` is described in RFC2151.

▶ You can retrieve RFC2151 from ***http://www.ietf.org/rfc/rfc2151.txt*** on the internet.

ping Program

The ICMP protocol is handled by the active CP/SP card. ICMP echo response packets may be lost when the CP/SP is processing higher priority packets such as BGP updates.

IP Source Address Selection in a public VPN

Lucent switches use the following process to select the IP source address used for transmitting an ICMP echo request or reply packet in a public VPN:

- If the outgoing interface is an IP logical port (lport), the switch uses the highest numbered IP address on that lport.
- If the logical port does not have an IP address (for example, a trunk or unnumbered PPP interface) and the destination address is not known via VNN, the switch uses the highest numbered IP loopback address on the switch.
- If the logical port does not have an IP address (for example, a trunk or unnumbered PPP interface) and the destination address is known via VNN, the switch uses the internal switch address.
- If a loopback address is not configured, the switch uses the highest numbered IP interface on the switch.
- If there are no configured loopback addresses or IP addresses, then the switch will not send the packet.

IP Source Address Selection in a private VPN

Lucent switches use the following process to select the IP source address used for transmitting an ICMP echo request or reply packet in a private VPN:

- If the outgoing interface is an IP lport, the switch uses the highest numbered IP address on that lport.
- If the logical port does not have an IP address, the switch uses the highest numbered IP interface on the switch, including the cloud interface.

`ping` Extended Options

To run the `ping` program with extended options, you must log in as a privileged user. The following options are available when `ping` is run with extended options:

Destination address – Specify the destination IP.

Number of datagrams — Specify the number of packets to transmit. The default is 5 and there is no range.

Data Length — Specify the packet data length. The default is 100 and the range is 1 to 65507.

Vary data length — Specify the default value by entering `yes,` or enter `no` to specify the starting data length (range is 0 to 100) and data length increment (range is 1 to 65507).

Timeout in seconds — Specify the timeout in seconds. The default is 2 and the range is 1 to 120.

Source address — Specify the source address. The local IP address should be an address on the Lucent switch.

Set DF Bit — Specify the Set DF Bit. The default is `no.` All fragmented packets will be discarded. Enter `yes` to set the fragment bit to forward fragmented packets.

Data pattern — Select the default or enter a unique data pattern.

The `ping` program may timeout if any of the following conditions occur:

- The user specified an incorrect IP address.
- The destination IP interface is down.
- There is a problem with the cabling.

Syntax:

```
ping { ip_address }
```

The following is an example of `ping` program output:

```
Rowley83_1# ping 150.202.83.3

Sending 5 ICMP echo requests to 150.202.83.3
ICMP: echo reply rcvd, src: 150.202.83.3, dst: 150.202.83.1, icmp_seq=0, time: 7.ms,
len: 108 bytes
ICMP: echo reply rcvd, src: 150.202.83.3, dst: 150.202.83.1, icmp_seq=1, time: 7.ms,
len: 108 bytes
ICMP: echo reply rcvd, src: 150.202.83.3, dst: 150.202.83.1, icmp_seq=2, time: 7.ms,
len: 108 bytes
----150.202.83.3 PING Statistics----
5 packets transmitted, 5 packets received, 0% loss
round-trip (ms) min/avg/max = 7/7/8


Rowley83_1## ping

Destination address: 150.202.83.3
Number of datagrams [5]:
Data length [100]:
Vary data length? [no]:
Timeout in seconds [2]:
Source address: 150.202.83.1
Set DF bit? [no]:
Data pattern [0x1234]:

Sending 5 ICMP echo requests to 150.202.83.3
/ Data length: 100 octets / Timeout: 2 secs / DF bit: not set /
Source: 150.202.83.1 / Data pattern: 0x1234 /
ICMP: echo reply rcvd, src: 150.202.83.3, dst: 150.202.83.1, icmp_seq=0, time: 7.ms,
len: 108 bytes
ICMP: echo reply rcvd, src: 150.202.83.3, dst: 150.202.83.1, icmp_seq=1, time: 20.ms,
len: 108 bytes
ICMP: echo reply rcvd, src: 150.202.83.3, dst: 150.202.83.1, icmp_seq=2, time: 8.ms,
len: 108 bytes
----150.202.83.3 PING Statistics----
5 packets transmitted, 5 packets received, 0% loss
round-trip (ms) min/avg/max = 7/10/20
```

traceroute Program

The `traceroute` program is handled by the CP/SP card. By default, `traceroute` provides the hop-by-hop path that an IP datagram takes to get to an IP destination.

Usage of the *mpt* keyword shows the ingress switch and the egress switch in the circuit data path that the MPT takes to get to an IP destination. A transit node will not respond to a traceroute request because the LSP leaf node at the ingress of the cloud will forward any packet with a TTL > 1 over an LSP to the root node. A transit switch will respond to a traceroute request if MPTs are disabled and an IP interface or loopback is configured.

`traceroute` traces an IP route until one of the following conditions occur:

- The IP datagram reaches its final destination.
- The connection break point is identified.

Syntax:

```
traceroute [mpt ip-address | ip-address]
```

► The `traceroute` command functions the same for either a public VPN 0 or a private VPN.

traceroute IP Address Selection

Lucent switches use the following process to select the IP address used for replying to a traceroute request:

- If the outgoing interface is an IP logical port (lport), the switch uses the highest numbered IP address on that lport.
- If the logical port does not have an IP address (for example, a trunk or unnumbered PPP interface) and the destination address is not known via VNN, the switch uses the highest numbered IP loopback address on the switch.
- If the logical port does not have an IP address (for example, a trunk or unnumbered PPP interface) and the destination address is known via VNN, the switch uses the internal switch address.
- If a loopback address is not configured, the switch uses the highest numbered IP interface on the switch.
- If there are no configured loopback addresses or IP addresses, then the switch will not send the packet.

IP Navigator Diagnostic Utilities

This section describes how to use the IP trace utility and the LSP trace utility and describes how to:

- Create a trace profile.
- Start a trace.
- Use the IP Trace commands with IP/VPNs.

Sample trace outputs, IP Trace command syntax, ctr command syntax, and tr command syntax are also included in this section.

IP Trace Utility

The IP trace utility is a packet trace tool for troubleshooting IP-related network problems. This utility has a command line interface only; you cannot use it from the NMS.

To set up, start, and stop the IP packet trace, you use the `set ip trace profile` command. You can view the packet trace setup through the `show ip trace profile` command.

▶ You must be in debug mode to issue the `set ip trace profile` command. To enter debug mode, issue the `enable debug` command from the switch console and specify the required password when prompted.

The IP trace utility screen output displays:

- IP packet header information.
- Data, called the *payload*, that follows the IP header. The payload is in hexadecimal format.

See “Sample Trace Output” on page B-14 for an example of trace output.

You can perform a trace of incoming and/or outgoing IP packets on any of the following:

- IP logical ports
- Ethernet management port on the B-STDX 8000/9000 CP, CBX 500 SP or GX 500 NP.
- Control port, which acts as an interface between the CP/SP/NP and the IOMs/IOPs on the switch

Trace filters allow you to trace IP packets selectively. These filters include the fields in the IP packet header (such as source IP address and destination port).

Creating a Trace Profile

A trace profile consists of three parts:

- Trace filters
- Payload length
- Trace counter

► Trace profiles are stored in memory only. They are never written to a permanent storage area. When you reset a card, trace profile information is lost.

Trace Filters

A trace profile includes one or more trace filters, which specify values that the IP packet must match to be traced. For example, only IP packets that meet both of the following criteria are traced:

- IP packets with a destination IP address of 131.100.110.1
- IP packets that carry HTTP data

If you do not specify a filter value, the default (“any”) is assumed. For example, if you do not specify a source IP address, any source IP address is accepted.

You can specify the following trace filters as part of a trace profile:

Source IP address with optional mask — Specify a valid source IP address. A mask can be assigned although it is optional.

Destination IP address with optional mask — Specify a valid destination IP address. A mask can be assigned although it is optional.

Source port — Specify the source port. You can specify a name or a decimal number. The command supports the following names: TELNET, FTP, FTPDATA, HTTP, SNMP, SNMPTRP, TFTP, and BGP. Port numbers are described in RFC1700.

Destination port — Specify the destination port. You can specify a name or a decimal number. The command supports the following names: TELNET, FTP, FTPDATA, HTTP, SNMP, SNMPTRP, TFTP, and BGP. Port numbers are described in RFC1700.

Protocol — Specify the protocol type. You can specify a name or a decimal number. The command supports the following names: ICMP, IGMP, IPIP, TCP, UDP, EGP, OSPF, and RAW. Protocol numbers are described in RFC1700.

Type of Service (TOS) with optional mask — Specify a hexadecimal number that identifies the TOS (with optional mask).

Transmission direction (send, receive, or any) — Specify the direction of transmission (send, receive, or any).

▶ You can retrieve RFC1700 from ***http://www.ietf.org/rfc/rfc1700.txt*** on the Internet.

You specify each filter on a separate command line. For example, suppose that you want to filter IP packets on the IP logical port identified by interface number 20, and you want to specify three filters: a source IP address of 150.140.10.5, a source port of HTTP, and you are only interested in filtering received packets. You would specify the following commands to accomplish this:

```
set ip trace profile port 20 source_ip 150.140.10.5
set ip trace profile port 20 source_port HTTP
set ip trace profile port 20 direction receive
```

See "IP Trace Command Syntax" on page B-15 for more information about the syntax that you use to specify filters.

Payload Length

The payload length is the number of bytes of data following the IP header to be traced. You can specify a payload length in the range of 0-100 for the control port and the Ethernet management port, and 0-32 for IOMs/IOPs. The default value is 0 (zero).

▶

> The payload starts immediately after the IP header and includes the TCP/UDP header, if present. Any IP header options are not included.
>
> For all IOPB cards and IOP3 cards on B-STDX 8000/9000 switches, a full 32 bytes of payload cannot be traced. For these IOPs, no more than 12 payload bytes can be traced in the *send* direction. All 32 bytes are traced in the *receive* direction.

The `data_len` keyword on the `set ip trace profile` command line allows you to specify the payload length. For example, the following command specifies a payload length of 32 bytes:

```
set ip trace profile port 20 data_len 32
```

See "IP Trace Command Syntax" on page B-15 for more information about command syntax.

Trace Counter

After you start a trace the trace counter increments each time an IP packet that matches the trace filter is detected. If you want to reset the counter, use the `count` option. For example, the following command resets the counter to zero:

```
set ip trace profile port 20 count 0
```

See "IP Trace Command Syntax" on page B-15 for more information about command syntax.

When you view a profile using the `show ip trace profile` command, the count appears in the format value/value (for example, 5/0). The first value displays the number of IP packets that were displayed on the console screen. The second value displays the number of IP packets that could not be displayed on the console screen for performance reasons. For example, a count of 12/2 means that 12 packets were displayed, but two packets were not.

Deleting Trace Profiles

To delete a trace profile, use the `deleted` keyword. For example, the following command deletes the trace profile associated with an IP logical port identified by interface number 10:

```
set ip trace profile port 10 deleted
```

The following command deletes the trace profile associated with all the ports on the switch:

```
set ip trace profile port all deleted
```

See "IP Trace Command Syntax" on page B-15 for more information about command syntax.

Starting a Trace

► You should create trace profiles before you start tracing. If you start tracing, and no trace filters exist, all of the filters are set to "any" by default and all IP packets are output to the console screen.

To start tracing, you must:

- Enable logging on the switch by issuing the following command from the switch console:

  ```
  set logging on
  ```

- Start tracing by using the `set ip trace profile` command. You can start tracing on a single port, all the ports associated with an IOM/IOP, or all ports on the switch.

On a Single Port

To start tracing on a single port, issue the following command:

```
set ip trace profile port port_id on
```

The *`port_id`* variable identifies the port and can be set to any of the following values:

Interface # — An interface number that identifies an IP logical port. To determine the interface numbers of IP logical ports on the switch, use the `show ip lport` command. The interface numbers appear in the IpLport column of the command output.

ethernet — Specifies the Ethernet management port on the CP/SP/NP.

control — Specifies the control port on the CP/SP/NP.

When you start tracing, the trace output appears on the console screen. See "Sample Trace Output" on page B-14 for more information. Some examples are shown below:

Examples:

To start tracing on the IP logical port identified by interface number 20.

```
set ip trace profile port 20 on
```

To start tracing on the Ethernet management port.

```
set ip trace profile port ethernet on
```

To start tracing on the control port.

```
set ip trace profile port control on
```

To stop tracing on a port, use the `off` keyword. See "IP Trace Command Syntax" on page B-15 for more information about command syntax.

On All Ports Associated With an IOM/IOP

To start tracing on all ports associated with an IOM/IOP, issue the following command:

```
set ip trace profile slot_id on
```

The *`slot_id`* variable identifies the slot number of the IOM/IOP.

When you start tracing, the trace output appears on the console screen. See "Sample Trace Output" on page B-14 for more information.

Example:

```
set ip trace profile slot 4 on
```

This command starts tracing for all of the IP logical ports associated with the IOM/IOP in slot 4.

To stop tracing on the IOM/IOP, use the off switch. See "IP Trace Command Syntax" on page B-15 for more information about command syntax.

On All Ports on the Switch

To start tracing on all ports on the switch, issue the following command:

```
set ip trace profile slot all on
```

When you start tracing, the trace output appears on the console. See "Sample Trace Output" on page B-14 for more information.

To stop tracing on the entire switch, use the off switch.

See "IP Trace Command Syntax" on page B-15 for more information about command syntax.

Using the IP Trace Commands With IP VPNs

This section describes how to use the `set ip trace profile` and `show ip trace profile` commands with IP VPNs.

set ip trace profile

You can issue the `set ip trace profile` command only if you are logged into the public IP VPN (IP VPN 0) at the switch console. See "Configuring IP Virtual Networks" in Chapter 16 for information about a public IP VPN and for information about logging in to an IP VPN.

Although you can issue the `set ip trace profile` command while logged in to the public IP VPN only, you can still set up and start traces for other IP VPNs. To do this, specify an IP VPN ID at the end of the command line. For example, the following command starts a trace on the IP logical port identified by interface number 3 for IP VPN 2:

```
set ip trace profile port 3 on vpn 2
```

If you do not specify an IP VPN ID, the public IP VPN is assumed, by default.

To view the IP VPN trace packets, you must log in to the IP VPN at the switch console after you start the trace. To log in to an IP VPN at the switch console, you must provide the Telnet password assigned to that IP VPN when it was added. See the *NavisCore IP Navigator Configuration Guide* for a description of the Telnet password.

When you start a trace, only the contents of the IP packets associated with the IP VPN (IP VPN 2 in the above example) are displayed. When you configure trace profiles, only the profiles for the specified IP VPN are affected.

> Ethernet management ports on CPs and SPs do not handle IP VPN traffic. To start a trace on these ports, do not specify any IP VPN IDs on the `set ip trace profile` command line.

`show ip trace profile`

Unlike the `set ip trace profile` command, you can issue the `show ip trace profile` command while logged in to any IP VPN at the switch console. When you issue the `show ip trace profile` command, it applies to the *current IP VPN* unless you specify the ID of another IP VPN at the end of the command line.

Sample Trace Output

The following sample illustrates trace output for an IP packet:

```
xx.xx ETH ip SND  ckt=N/A  prot=UDP tos=0x01 ttl=64 len=93
        s=154.72.1.5 d=152.148.30.129 sp=SNMP dp=34750 dlen=73
        0x00 A1 87 BE 00 49 41 01 30 3F 02 01 00 04 07 63
        61 73 63 61


xx.xx CTL ip RCV  ckt=N/A  prot=ICMP tos=0x00 ttl=225 len=84
        s=154.72.1.2 d=154.72.1.5  type=ECHO-REQUEST
        0x08 00 EF C6 4D 47 00 00 35 78 0B 29 00 06 8F 47 08 09
        0A 0B
```

The trace output provides the following information:

- Header information that differs depending on the protocol
- TCP and UDP data length and port
- IP packet type and port
- DLCI or VCI/VPI circuit identifier

IP Trace Command Syntax

The syntax of the `set ip trace profile` command is as follows:

```
set ip trace profile port [interface # | ethernet | control]
option

set ip trace profile port [interface # ]option vpn vpn_id

set ip trace profile slot [slot # | all ] option

set ip trace profile slot [slot # | all] option vpn vpn_id

options:
   [on | off | deleted]
   source_ip [ipaddr | ipaddr ipmask | any]
   dest_ip [ipaddr | ipaddr ipmask | any]
   source_port [port # | port name| any]
   dest_port [port # | port name | any]
   protocol [protocol # | protocol name | any]
   tos [tos hex value | tos_hex_value hex_mask ] | any]
   direction [send | receive | any]
   data_len payload_length
   count value
```

The syntax of the `show ip trace profile` command is as follows:

```
show ip trace profile port [interface # | ethernet | control ]

show ip trace profile port [interface # ] vpn vpn_id

show ip trace profile slot [ slot # | all ]

show ip trace profile slot [ slot # | all ] vpn vpn_id
```

LSP Trace Utility

The LSP trace utility is used to trace the LSP call signalling sequence for multipoint-to-point Label Switched Paths (MPT LSPs), point-to-point Label Switched Paths (point-to-point LSPs), and multicast Label Switched Paths (LSPs).

You can use trace filters to select desired trace information dynamically from massive trace statements. The LSP trace filtering utility enables you to use trace filters to select only the LSP trace information that you need. This utility is command-driven, and is available only at the switch console.

When you use LSP trace filtering, you specify *cookie* values for the destination address, LSP ID, and exclusions. After setting up the filters, you can start the trace.

The `ctr` command sets up the filters, and the `tr` command starts the trace. To use these commands, you must enter debug mode on the switch by issuing the `enable debug` command and supplying the correct password.

ctr Command Syntax

The syntax of the `ctr` command is:

```
ctr [appid] [card] [ctype] [dest_node] [mptid][exclusion]
```

The following parameters are used with the `ctr` command:

- *`appid`* is always 16 for LSPs.
- *`card`* specifies the slot number of the card used by the LSP.
- *`ctype`* specifies any combination of numbers from zero to six, where:
  - 0 (zero) specifies that one through five should be turned on.
  - 1 specifies an MPT LSP.
  - 2 specifies a point-to-point LSP.
  - 3 specifies a multicast LSP.
  - 4 specifies an unknown LSP type (the default).
  - 5 specifies all the error traces (the default).
  - 6 specifies that the trace must match the user's requirements.
- *`dest_node`* specifies the IP address of the destination (leaf) node, in hexadecimal format.
- *`mptid`* specifies the ID of the LSP, which you can obtain by issuing the `show mpt all command.`
- *`exclusion`* is a hexadecimal number that specifies any combination of the following values:
  - 0x00000000 (the default) specifies no exclusion (everything included).
  - 0x00000001 specifies that Hello PDUs/messages are excluded.
  - 0x00000002 specifies that LSP enable/disable manager traces are excluded.
  - 0x00000003 specifies that LSP conversion traces are excluded.

tr Command Syntax

The syntax of the `tr` command is:

```
tr [ appid ] [ level ] [ output ] [ card ]
```

The following parameters are used with the `tr` command:

- *`appid`* is always 16 for LSPs.
- *`level`* specifies the trace level from one to eight. The level specifies the amount of information that you want in the trace, from the least information (one) to the most information 8 (eight). Zero (0) disables the trace.
- *`output`* specifies how the trace statements are output as follows:
  - 0 (No output)
  - 1 (Output to console)
  - 2 (Debug)
  - 3 (Output to both console and debug)
- *`card`* specifies the slot number of the card used by the LSP.

Sample LSP Trace Output

Example 1:

```
cheetah## ctr 16 8 1 0x01020304 1 0

cheetah## tr 16 8 1 8
```

In this example, the trace output displays all of the trace statements that apply to the card in slot 8, according to the following criteria:

- The LSP type is MPT LSP (1).
- The IP address of the leaf node is 1.2.3.4 (0x01020304).
- The LSP ID is 1.
- The trace statements are output to console if the filter cannot determine whether the user’s requirements are met.

Example 2:

```
cheetah## ctr 16 8 12 0x0a0b0c0d 1 1

cheetah## tr 16 8 1 8
```

In this example, the trace output displays all of the trace statements that apply to the card in slot 8, according to the following criteria:

- The LSP types are MPT LSP (1) and point-to-point LSP (2).
- The IP address of the leaf node is 10.11.12.13 (0x0a0b0c0d).
- The LSP ID is 1.
- Ignore all of the hello messages or PDUs (1).
- The trace statements are output to console if the filter cannot determine whether the user's requirements are met.

Example 3:

```
cheetah## tr 16 8 126 0x0a0b0c0d 2 3
```

```
cheetah## tr 16 8 1 8
```

In this example, the trace output displays all of the trace statements that apply to the card in slot 8, according to the following criteria:

- The MPT type is MPT LSP (1) and point-to-point LSP (2).
- The IP address of the leaf node is 10.11.12.13.
- The MPT ID is 2.
- The trace statements associated with the hello PDU/messages and MPT LSP manager are ignored.
- The trace statements are shown if they match the user's requirements (this is specified by the 6 at the end of 126).

Label Switched Paths

This section describes methods of understanding label switched paths (LSPs). These methods include understanding LSP diagnostic messages and examining the data path of an LSP.

Terms relating to LSPs have changed.Table B-1shows the new LSP terminology.

Table B-1. Terminology Changes

Old Term	New Term
Reverse Multipoint-to-Point Tunneling (RMPT)	Multipoint-to-Point Label Switched Path (MPT LSP)
Multipoint-to-Point (MPT) Point-to-Point Connection	Point-to-Point Label Switched Path (LSP)
Forward Multipoint-to-Point Tunneling (FMPT)	Multicast Label Switched Path (LSP)

▶ The LSP commands used in the Command Line Interface (CLI) use the `MPT` keyword. For example, `show mpt all`.

Overview

IP Navigator uses LSPs as a means of forwarding IP traffic over switched paths (no intermediate IP lookups) through the Lucent network. LSPs are a unique feature provided by Lucent.

Depending on the type of LSP, an LSP can allow:

- Multiple nodes to share the same circuit for transmission to a single destination. This type of LSP is called a *multipoint-to-point LSP (MPT LSP)*.
- A pair of nodes to share a point-to-point connection. This connection overrides a single root-to-leaf connection that would otherwise be part of a unicast LSP. This type of LSP is called a *point-to-point LSP*.
- A single node to use one circuit for transmission to multiple destinations. This type of LSP is called a *multicast LSP*. It is used to transmit IP multicast traffic.

Both the MPT LSP and multicast LSP networks are *rooted* at the switch. A root is a standard circuit endpoint that is created at initialization time on every CP card in the B-STDX 8000/9000 and every SP card in the CBX 500. Traffic flow occurs from the leaves to the root on MPT LSPs, and from the root to the leaves on multicast LSPs. On point-to-point MPT connections, traffic is bi-directional, since each node configured on a point-to-point LSP connection acts as both a leaf and a root.

► See "What Are Label Switched Paths?" on page 12-2 for more information about LSPs.

LSP Call Signalling

LSPs use a process called *call signalling* to set up virtual circuits and to determine the label switched path. Virtual circuits are used to forward IP traffic over label switched paths. A label switched path is a path with a circuit on it that allows data to be switched instead of forwarded hop-by-hop.

The following list describes how the switch processes call signalling for a cell LSP:

1. As each cell IOM initializes on a CBX 500 root switch, the cell IOM sets up connections to each forwarding engine on each frame IOM in the switch. These connections allow LSP data that is received on the cell trunk to be forwarded to the appropriate forwarding engine.
2. When OSPF determines that a node is reachable, it informs the LSP component running on the SP. The LSP component then requests OSPF for a route to the leaf node.
3. The SP sends a Call Setup to the cell IOM. The LSP component on the cell IOM allocates a VPI and informs the hardware that the VPI is an LSP VPI.
4. A Call PDU is then sent to the next switch as follows:
   a. At a transit node, the LSP component on the cell IOM receives the Call Setup and forwards the Call Setup to the next cell IOM.
   b. At a leaf node, the LSP component on the cell IOM receives the Call Setup and forwards the Call Setup to the SP.
5. The SP receives the Call Setup and sends a confirmation back to the LSP root indicating which FE with IP interfaces it has on its node.
6. At the root node, the LSP component on the SP receives the Confirm message and allocates a RID for each FE with IP interfaces at the leaf.
7. The LSP component on the SP informs each FE at the root about the RIDs that it has allocated and sends a Call Setup to each FE at the leaf.
8. At the leaf node, the LSP component on the cell IOM receives the call setup and forwards each Call Setup to the appropriate FEs.

LSP Call Forwarding

LSPs use a process known as *call forwarding* to forward IP data over label switched paths.

The CBX 500 call forwarding functions for a cell circuit as follows:

1. At the leaf node, the FE determines that data can be forwarded on an LSP. The FE segments the IP packet into ATM cells and inserts the MPT VCI into the cell header. The ATM cell is then sent across the switching fabric to the cell IOM.
2. At the transit node, the LSP cells are switched on virtual paths until they reach the LSP root.
3. At the root node, if the LSP cell is received on a cell IOM, the IOM determines the cell is an LSP based on its VPI. The IOM examines the first five bits of the VCI to determine which FE will receive the ATM cell. The VPI bits are used to separate streams of multiple LSPs.
4. At the root node, if the LSP cell is received on the frame IOM, the cells are reassembled based on the RID (found in the remaining 11 bits of the VCI). Once the packet is reassembled, another IP lookup is performed and the data is forwarded out of the physical interface.

MPT LSP Information and Restrictions

This section describes the MPT LSP network configuration requirements and known restrictions.

MPT LSP Configuration Prerequisites

MPT LSPs must be enabled at the switch and at least one IP interface must exist for a switch to act as an MPT LSP root or a leaf.

An IP interface is not required if the switch is a boundary node. For example, a boundary node can be one of the following:

- OSPF area border router (ABR)
- Transit 9000 switch within a cell network
- Frame/cell domain border router
- Optimum trunk border router

VNN OSPF Domains

MPT LSPs cannot traverse multiple VNN OSPF domains (areas). However, point-to-point LSPs and multicast LSPs can traverse multiple VNN OSPF domains.

Within a single VNN OSPF domain, MPT LSPs are switched. However, as soon as traffic reaches an area boundary (that is, an area border router), a routing lookup must be made and the MPT LSP traffic must be switched to a new LSP.

Switch Domains

MPT LSPs and point-to-point LSPs are only established between switches in the same domain. Multicast LSPs can be established between switches in different domains.

There are two types of switch domains. A *cell* domain is a path that traverses direct ATM trunks and ATM OPTimum trunks. A *frame* domain is a path that traverse direct frame trunks and frame OPTimum trunks.

A switch that only belongs to one domain cannot add a switch from a different domain to its LSP. To traverse different domains, a boundary switch that belongs to both domains must act as an intermediary. Each endpoint switch connects to the boundary switch via a separate LSP, and the boundary switch performs an IP lookup when routing traffic between the endpoints.

OPTimum Trunks

The following criteria applies to OPTimum trunks:

- OPTtimum Trunk VPI attributes must not overlap. For instance, the VPC VPI Start and Stop range and the Transit LSP and point-to-point LSP VPI Start and Stop values must not overlap with the OPTimum Trunk VPI, the LSP VPI value, or the OPTimum trunk value of a different lport on the same physical port.
- MPT LSPs must use even VPI values between 2-30 and point-to-point LSPs must use odd VPI values between 1-2047.
- The OPTimum Trunk range cannot overlap with the values that are configured for Virtual UNIs on the same physical port or the OPTimum trunk value of a different lport on the same physical port.
- The OPTimum Trunk VPI values do not have to match when data is forwarded from the root to the leaf and the trunk traverses an intermediate ATM network.

  One exception to this rule is when the leaf is a CBX 500 and the root is a B-STDX 8000/9000.

- Switches that act as OPTimum cell trunk endpoints cannot also be point-to-point LSP endpoints.

  – A switch configured as an OPTimum cell trunk endpoint can act as a transit switch for point-to-point LSPs.

  – If you configure a switch as both an OPTimum cell trunk endpoint and as a point-to-point LSP endpoint, the point-to-point LSP will fail to establish.

- On the B-STDX 8000/9000 and CBX 500, the logical port OPTimum trunk VPI value that is configured in the NMS, is used for VCC connections that are needed for IP network management and VC manager control.

Parallel Paths

An MPT LSP or point-to-point LSP will not establish when the path between two nodes contain an IP interface with a lower administrative cost than the configured trunks. If the IP interface is not running OSPF, the trunk will be used because OSPF is preferred. For example, if a parallel IP link and trunk exists between two nodes, the LSP will not establish because the IP link offers the shortest path.

Traffic Descriptors

The MPT LSP Committed Information Rate (CIR) specifies the rate in Kbps at which MPT LSPs transfer data, averaged over a minimum increment of time. In addition, this value reserves bandwidth for MPT LSPs, which the switch originates.

The CIR value is based on the available line rate. The minimum CIR is 100. CIR is used as the minimum cell rate (MCR) of the available bit rate (ABR) LSP.

Disabling MPT LSP

Disabling MPT LSP on a switch prevents the node from acting as a root or a leaf to an MPT LSP.

To enable or disable MPT LSPs on a switch:

1. Select the switch on the network map.
2. From the Administer menu, choose Lucent IP Parameters ⇒ Set IP Parameters. The Set IP Parameters dialog box appears.
3. In the Set IP Parameters dialog box, set the MPT LSP option to *disable* to disable MPT LSPs or *enable* to enable MPT LSPs.

Point-to-Point LSP Information and Restrictions

The following criteria applies to point-to-point LSPs:

- A point-to-point LSP can traverse different OSPF areas and VPN areas.
- A point-to-point LSP does not support CBR QoS Class.
- The GX 550 cannot be a transit node for a point-to-point LSP in a frame domain.
- A point-to-point LSP in a cell domain will not establish if one of the transit nodes is a B-STDX 8000/9000 switch or if a cell optimum trunk is directly connected to a border node.

▶ Point-to-point LSPs are operational when MPT LSPs are disabled on a switch.

VNN OSPF Domains

Point-to-point LSPs can cross VNN OSPF domains.

Switch Domains

Point-to-point LSPs can not cross cell/frame domains.

OPTimum Trunks

The OPTimum trunk restrictions that apply to MPT LSPs also apply to point-to-point LSPs. See “OPTimum Trunks” on page B-22 for a description of these restrictions.

Multicast LSP Information and Restrictions

Disabling multicast LSP on a switch will cause multicast data to be forwarded hop-by-hop if the data is going over a frame trunk:

- If data is going over a cell trunk, only the VPN multicast packets are sent hop-by hop.
- All other multicast data is discarded.

VNN OSPF Domains

Multicast LSPs can cross OSPF VNN domains.

Switch Domains

Multicast LSPs can cross cell/frame domains.

LSP Terms

Many of the following terms are used in the command line display output and can be used in a troubleshooting discussion.

Table B-2. LSP Terms/Acronyms

Term/Acronym	Description
Cell Domain	A network configuration of pure direct ATM trunks and ATM Optimum trunks. MPT LSPs and point-to-point LSPs do not cross cell/frame domain boundaries. Multicast LSPs do cross cell/frame domains boundaries.
Child Virtual Circuit	On the CBX 500 leaf node, the child virtual circuit (vc) receives data on the ingress card.
FE	A forwarding engine (FE) on an Ethernet or Frame IOM2. FEs reassemble cells and perform IP lookups. There are two FEs in each of the CBX 500 Ethernet and Frame cards, one FE on the CBX 500 Channelized DS3 card, and one FE on the B-STDX 8000/9000 Ethernet card.
FE ID	The forwarding engine identifier (FE ID) is a five bit field that identifies the forwarding engine at the root. Four bits of the FE ID indicate which slot the frame card is in and the other bit indicates which FE is on the frame card (zero or one).
Frame Domain	A network configuration of direct frame trunks and frame OPTImum trunks. MPT LSPs and point-to-point LSPs do not cross cell/frame domain boundaries. Multicast LSPs do cross cell/frame domains boundaries.
Leaf	An LSP leaf is the source of data that is forwarded over the LSP. A leaf can either be on a B-STDX 8000/9000 CP or a forwarding engine (FE) on a Frame IOM on the CBX 500.
LSP VCI	The LSP VCI is made up of the 11 bit RID and the 5 bit FE ID. This information provides the source identifier and destination ID.
mptID	The LSP virtual circuit identifier is the value shown in the VC column of the `show mpt vcarray` command. Similarly referred to as fmID.
mptTport	A mptTport maintains state for downstream (nodes further away from the root) connections from the point of view of the parent.
Parent Virtual Circuit	On the CBX 500 leaf node, the parent virtual circuit (vc) forwards data on the egress card.
Parent Function	The parent function on a CBX 500 aggregates cell streams into a common VP and delivers them upstream to the root.

Table B-2. LSP Terms/Acronyms (Continued)

Term/Acronym	Description
RID	The reassembly identifier (RID) is used when forwarding data to the source (root). The RID allows the root to know which source FE/node a cell came from when cells are interleaved on the same default connection between the frame card and the cell card.
Root	An LSP root is the destination switch of data that is forwarded over the LSP. A root is anchored on the SP or CP and is created at switch initialization. The root tracks the state of other nodes or FEs on the network that belong to the LSPs.
RVc	The RVc found in the `show mpt vcarray` command output is the remote VC in the LSP chain.
RVC vctype	On the B-STDX, a RVC virtual circuit (vc) is used to send data to a CBX 500 or receive data from either a CBX 500 or a B-STDX 8000/9000.
tleaf	An LSP tleaf represents the source endpoint node and forwarding engine (FE) state of the leaves.
vcarray	A vcarray is the set of virtual circuit entries for the specified node or card.
VCC	The B-STDX 8000/9000 switch uses a virtual channel connection (VCC) for each LSP crossing an optimum trunk.
vcentry	A vcentry is a data structure used to store information about a virtual circuit.
vcID	The vcID is the internal local VC identifier of the LSP circuit used to forward the IP data.
vctype	The vctype is a parent, child, FE, RVC, root or a leaf. A FE type is a forwarding engine vcentry type on a leaf or root 500.
VPC	The CBX 500 switch uses a virtual path connection (VPC) for each LSP.
VPI	The CBX 500 switch uses a virtual path identifier (VPI) for each LSP crossing an optimum trunk.

Examining a LSP Data Path

You can examine the data path of a LSP to diagnose IP forwarding problems.

For a MPT LSP and a point-to-point LSP the data path travels from the leaf switch to the root switch. For a multicast LSP, the data path travels from the root switch to the leaf switch. A leaf switch is also referred to as a leaf node and a root switch is referred to as a root node.

The user should be familiar with the network topology in order to identify switches and interfaces in the following sections.

Examine the Data Path of an MPT LSP

This section describes how to examine the data path of a MPT LSP from a CBX 500 ingress node, through a CBX 500 transit node, to a CBX 500 root node. The switches are connected by cell OPTimum trunks. LSP tracing begins at the ingress switch that is closest to the data source. The data source is the workstation or router that transmits IP data to an IP destination located on the other side of the network cloud. The ingress and transit switches are leaf switches (nodes) and the egress switch is a root switch (node). IP data enters the network at the ingress switch. Each switch along the data path is called a transit switch. IP data exits the network at the egress switch.

Figure B-1 shows the MPT LSP data path network for this section.

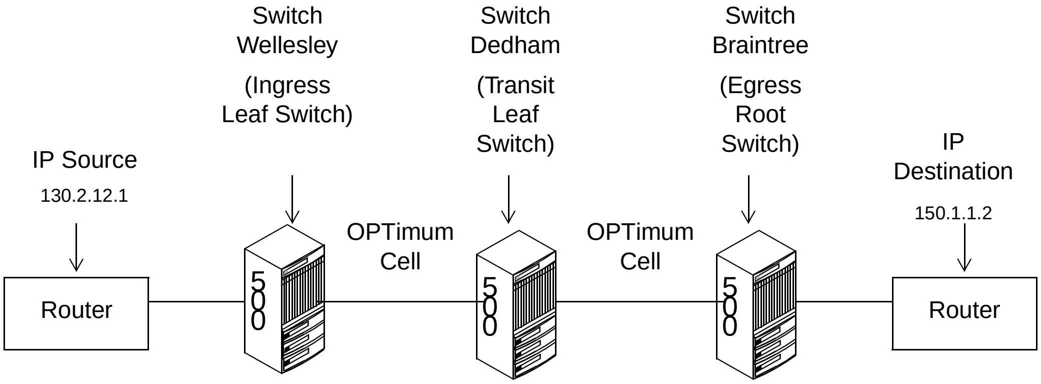

LSP data path is from the leaf node to the root node.

Figure B-1. MPT LSP Data Path Network

This section describes the steps to trace a MPT LSP:

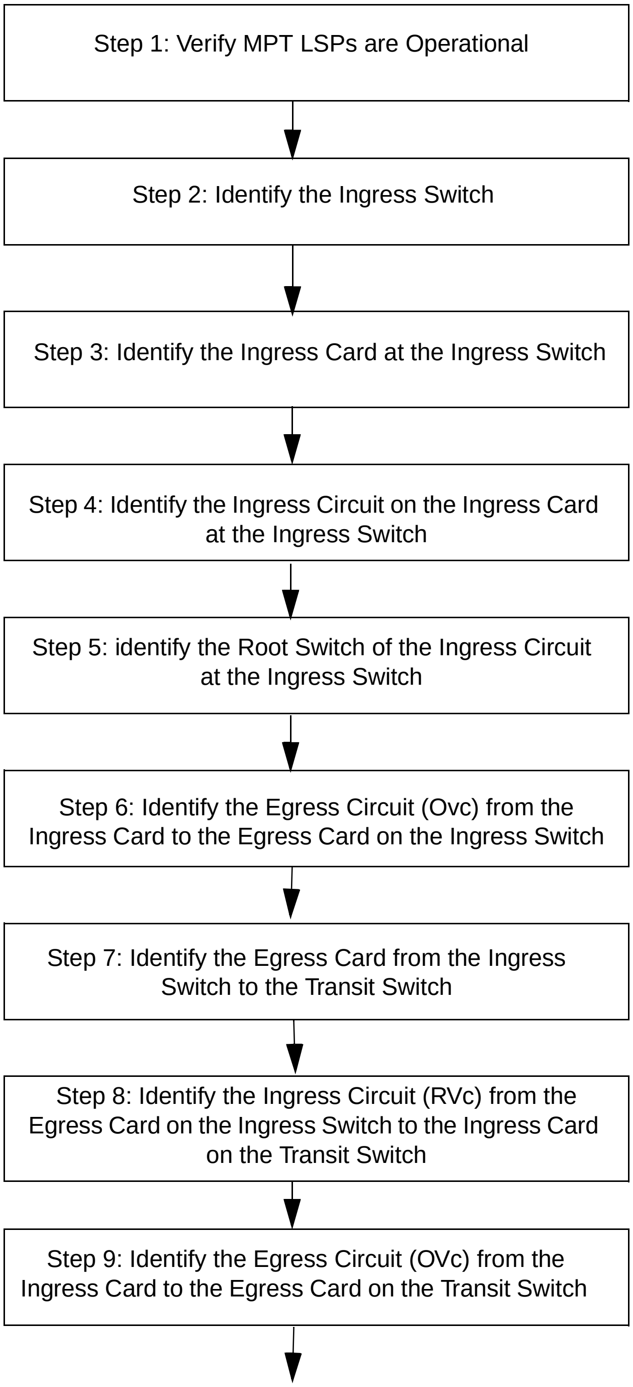
Step 1: Verify MPT LSPs are Operational
Step 2: Identify the Ingress Switch
Step 3: Identify the Ingress Card at the Ingress Switch
Step 4: Identify the Ingress Circuit on the Ingress Card at the Ingress Switch
Step 5: identify the Root Switch of the Ingress Circuit at the Ingress Switch
Step 6: Identify the Egress Circuit (Ovc) from the Ingress Card to the Egress Card on the Ingress Switch
Step 7: Identify the Egress Card from the Ingress Switch to the Transit Switch
Step 8: Identify the Ingress Circuit (RVc) from the Egress Card on the Ingress Switch to the Ingress Card on the Transit Switch
Step 9: Identify the Egress Circuit (OVc) from the Ingress Card to the Egress Card on the Transit Switch

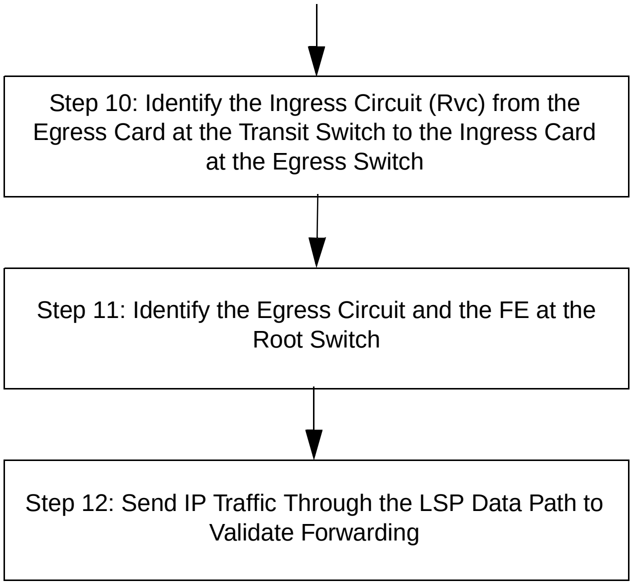
Step 10: Identify the Ingress Circuit (Rvc) from the Egress Card at the Transit Switch to the Ingress Card at the Egress Switch
Step 11: Identify the Egress Circuit and the FE at the Root Switch
Step 12: Send IP Traffic Through the LSP Data Path to Validate Forwarding

Step 1: Verify MPT LSPs are Operational

Use the `show mpt all` command to verify that the MPT LSP state is operational and that each LSP is active between all switches in the data path from the IP source to the IP destination.

In order for the switch to process MPT LSPs, the MPT LSPs value for the switch must be enabled. MPT LSPs are operational when MTP LSPs are enabled and the system completes initializing after the CP/SP or system has been reset. If MPT LSPs are suspended, then the switch has either recently been initialized (it takes a few seconds for MPT LSPs to become operational) or the MPT LSP state has been disabled at the switch.

If MPT LSP connections are inactive, there may be a signalling problem. Note the status of the RecvState and SendState in the `show mpt all` command output. See "LSP Call Signalling" on page B-20 for an explanation of call signalling. See "LSP Connection Failure Reasons" on page B-39 for a list of failure reasons.

```
wellesley> show mpt all

MPT Identifier: 1  [OPERATIONAL]    Type: UNICAST (Multipoint-Point Tunnel)    OSPF Area: 0.0.0.1
Node            IOP/FE RID  Flags  RecvState SendState SendMpt LastFail(Node/Port/Reason)

150.202.83.13   CP/SP  -NA- 0x121006 ACTIVE    ACTIVE    1    None/0/NONE
150.202.83.13   4 /0   8    0x121006 ACTIVE    ACTIVE    1    None/0/NONE
150.202.83.13   7 /0   9    0x121006 ACTIVE    ACTIVE    1    None/0/NONE
150.202.83.13   7 /1   11   0x121006 ACTIVE    ACTIVE    1    None/0/NONE
150.202.83.10   CP/SP  -NA- 0x121006 ACTIVE    ACTIVE    1    None/0/NONE
150.202.83.10   4 /0   3    0x121006 ACTIVE    ACTIVE    1    None/0/NONE
150.202.83.10   5 /0   4    0x121006 ACTIVE    ACTIVE    1    None/0/NONE
150.202.83.10   7 /0   5    0x121006 ACTIVE    ACTIVE    1    None/0/NONE
1150.202.83.10   5 /1   7    0x121006 ACTIVE    ACTIVE    1    None/0/NONE
150.202.83.10   7 /1   8    0x121006 ACTIVE    ACTIVE    1    None/0/NONE
1150.202.83.20   CP/SP  2    0x120806 ACTIVE    ACTIVE    20   None/0/NONE
```

Step 2: Identify the Ingress Switch

Use the `traceroute` command at the IP source to identify the Lucent switch that is closest to the IP source and functions as the ingress switch to the network. This switch functions as an ingress node into the network cloud. In Figure B-1 on page B-27, this is Switch Wellesley.

Step 3: Identify the Ingress Card at the Ingress Switch

At switch Wellesley, use the `show ip interfaces` command to identify the card that connects the switch to the IP source. The Pport column shows the ingress card and port information.

This example shows card 6 has a direct connection to the IP source.

```
wellesley> sho ip interfaces

IpAddr              Lport Pport    Card       MTU   ARP   IARP  OPER  ADMIN

130.2.12.12/24      60     6.1    TFETHER-4   1500  ENA   DIS   UP    ENA


wellesley> sho arp

IpAddr        LinkType     HwAddr           State          EntryType

130.2.12.1    ethernet     0050d197b039     Complete       Dynamic
```

Step 4: Identify the Ingress Circuit on the Ingress Card at the Ingress Switch

At switch Wellesley, use the `show ip route` *`destination_ip slot_# fe`* command to perform a route lookup of the destination IP address and to determine the ingress circuit on the ingress card:

- Enter a complete host IP address for *`destination_ip`*. In this example, you would enter 189.75.50.1 which is the IP address of the destination host.
- Enter the slot number for *`card`* that represents the ingress card identified in "Step 3: Identify the Ingress Card at the Ingress Switch".
- Enter the forwarding engine if the card is a CBX 500 Ethernet or Frame DS3 card.

In this example, the ingress circuit is 159:

```
wellesley> show ip route 150.1.1.2 6 1

Forwarding Engine : 1

                 Network:         150.1.1.0
                 Mask:            255.255.255.0
                 Flags:
                 Hop-by-hop:      [*, 45]
                 Switched:        [0x0/159, TTL: 1]
```

> If the ingress or egress card is a cell trunk, you must perform the route lookup using another card and lport that supports IP Routing. See "Logical Ports That Support IP Routing on the B-STDX 8000/9000" on page B-66 and "Logical Ports Supporting IP Routing on the CBX 500" on page B-67.
>
> Use the `show system` command to locate one of these types of cards. To specify a forwarding engine, you enter the FE identifier after the slot number.

The route lookup display output shows important hop-by-hop and switched information. If a switched circuit is not shown in the table, then the IP data is forwarded hop-by-hop. In the hop-by-hop row, the '*' indicates there is not an ARP entry for the route (use the `show arp` command to view the arp table) and the next hop route is probably a cell trunk. The '45' indicates the egress lport identifier. In the switched row, the '159' value indicates the LSP virtual circuit. The TTL indicates the total number of hops the LSP takes to reach the destination node. The LSP circuit is always 1 hop, from ingress to egress. The transit switches are not counted as hops.

Step 5: Identify the Root Switch of the Ingress Circuit

At switch Wellesley, use the `show mpt rootnodes` command to identify the root switch (node) and root egress card of the ingress circuit identified in the previous step. See "show mpt rootnodes" on page B-55 for a complete description of the command. Note that the command syntax varies for the CBX 500 and the B-STDX.

The command requires the following information:

- Specify the card identified in "Step 3: Identify the Ingress Card at the Ingress Switch".
- Specify forwarding engine one (1) or two (2) Enter the forwarding engine if the card is a CBX 500 Ethernet or Frame DS3 card.
- Specify the LSP virtual circuit identified in "Step 4: Identify the Ingress Circuit on the Ingress Card at the Ingress Switch".

In a complex network, the IP data path can traverse many MPT LSPs because of border router conditions. If the root switch terminates at a border router then the user needs to perform another IP lookup to determine the next egress circuit. Border router conditions are described in "VNN OSPF Domains" on page B-22.

In this example, the root node is 150.202.83.13 (switch Braintree). The egress card to the IP destination on the root node is card 7. The feNum value of 2 represents forwarding engine 2. A value of 1 will represent forwarding engine 1.

```
Wellesley> show mpt rootnodes 6 1 159

MPT ckt/node Info: Card 6 FE 1

VcId    RootNode        Slot    feNum  mptID   vpn
159     150.202.83.13    7        2      1      0
```

Step 6: Identify the Egress Circuit (OVc) from the Ingress Card to the Egress Card on the Ingress Switch

At switch Wellesley, use the `show mpt rsvcarray` *`ingress_card`* command to determine the egress circuit from the ingress card to the egress card. Specify the card that was identified in "Step 3: Identify the Ingress Card at the Ingress Switch".

In the vcarray display output, find the circuit (identified in "Step 4: Identify the Ingress Circuit on the Ingress Card at the Ingress Switch") in the VC column. On the same line, find the egress circuit in the OVc column.

A VCI value is shown if the data is forwarded over a cell OPTimum trunk. The value is used to calculate the card and forwarding engine on the root switch that the data is going to. This information is also available in the `show mpt rootnodes` command used in step 5.

- Convert the decimal VCI value to binary.
- The FEID is contained in the upper five bits of the sixteen bit VCI.
- Of these five bits, the upper most bit represents the FE and the lower four bits represents the slot.
- If the upper most bit (bit 16 of VCI) is zero, the data is destined to FE 1, otherwise it is destined to FE 2.
- Adding a one to the value of the lower four bits of the FEID (bits 15-12), identifies the slot number that the data is destined to.

In the following example, the VCI is 45061.

10110 is the upper five binary bit value of the decimal value 45061.

The upper most bit '1' represents FE 2, the lower four bits plus a value of one represents card 7.

In the following example, the OVc (egress circuit) is 908. The VCType is FE. The egress card that has this egress virtual circuit is slot 10, port 2. The OPrt (45) is the logical port identifier of the egress card. .

```
wellesley> show mpt rsvcarray 6

VC   VPN  Type  VCType State    DNde OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FEID OutFrames InFrames
159  0    RMPT  FE     ACTIVE   530d 908  1    0    45061  10  2    3954  45   5    N/A-0     1
```

Step 7: Identify the Egress Card from the Ingress Switch to the Transit Switch

There are two methods to identify the egress card. You can simply use the IOP value shown in the previous command or you can use the next two commands to identify the egress card. The advantage of using the following command is that these commands identify the next switch and ingress card in the signalling path to the root switch.

The `show mpt spath` command identifies all of the switches in the signalling path from ingress switch, Wellesley, to the root switch, Braintree. Enter the *`node_IP_address`* of the switch identified in "Step 5: Identify the Root Switch of the Ingress Circuit".

The egress card is derived from the logical port identifier of the first node/lport pair shown in the `show mpt spath` display output. Use the `show lport attributes` command to identify the logical port. The display output of this command will show the remote node and remote interface.

In this example, the first node is 150.202.83.12 (switch Dedham), and the lport identifier is for slot 10, port 2. Since the root node has been identified as switch Braintree, Dedham is a transit node in the signalling path to the root node.

```
wellesley> show mpt spath 150.202.83.13

MPT in VPN 0:
  List of node/interface pairs in path to node 150.202.83.13:
  150.202.83.12/45,   150.202.83.10/156
  Path characteristics:
  PURE_CELL,   FOR_IP,   FE_ARCH,   JADEM1_PATH,   OSPF_PATH_REG


wellesley> show lport att 45

Slot:           10
Port:            2
Interface:       45
Data Rate:       40640000

Trunk Status:    Full             Trunk Overhead:  5%
Remote Node:          150.202.83.10  Remote Interface:        114

       Trunk              Out BW        Out BW
       Oversub.:          Avail.        Alloc.
Qos1     100%               91057          4792
Qos2     100%               91057             0
Qos3     100%               91057             0
Qos4     100%               91057             0
Administrative Status:          Up   Operational Status:           Up
```

Step 8: Identify the Ingress Circuit (RVc) from the Egress Card on the Ingress Switch to the Ingress Card on the Transit Switch

At switch Wellesley, use the `show mpt vcarray` *`egress_card`* command to identify the ingress circuit from the egress card on the ingress node to the ingress card at the switch Dedham. Enter a value for *`egress_card`* that matches the value identified in either Step 6 or Step 7.

In this example, the RVc is 98. The VCType is PRNT.

```
wellesley> show mpt vcarray 10

VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FEID Node            OutFrames InFrames
Discard
908  0    RMPT  PRNT   ACTIVE   98   0    1    18   0      0   2    45    0    0    150.202.83.13      4 0
```

Step 9: Identify the Egress Circuit (OVc) from the Ingress Card to the Egress Card on the Transit Switch

At transit switch Dedham (150.202.83.10), identified in Step 7 as the Remote Node, use the `show lport attributes` *`interface`* command to identify the ingress card to the transit switch. The interface value is the Remote Interface value identified in step 7.

Use the `show mpt vcarray` *`ingress_card`* command with the remote interface value to identify the egress circuit from the ingress card to the egress card on the switch.

Alternatively at switch Dedham, use the `show vnn trunks` command to determine the ingress card.

In this example, the `show vnn` and `show lport` commands are used to determine the ingress card which is slot 8, port 2, lport identifier 114. The `show mpt vcarray` command is used to identify the egress circuit is 141. The OPrt is 156. The `show lport command` is used to identify OPrt 156 as the egress card which is card 8, port 1. Note that the CBX 500 VCType is Child.

```
dedham> show vnn trunks

switch/lport  switch/lport  fbw3/0      rbw3/0      delay    cost            area            comments
83.10/114     83.12/45      37703       37703       2        50              0.0.0.1


Dedham> sho lport att 114

Slot:           8
Port:              2
Interface:       114
Data Rate:       40640000

Trunk Status:    Full      Trunk Overhead:           5%
Remote Node:           150.202.83.12 Remote Interface:          45

       Trunk              Out BW        Out BW
       Oversub.:          Avail.        Alloc.
Qos1     100%               90727          4792
Qos2     100%               90727             0
Qos3     100%               90727             0
Qos4     100%               90727           330
Administrative Status:          Up   Operational Status:           Up


dedham> show mpt vcarray 8

VC    VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FEID Node           OutFrames     InFrames
Discard
98    0    RMPT  CHILD  ACTIVE   908  141  1    18   0       8   2    114   156  0    150.202.83.13   N/A-0         N/A-0


dedham> show lport att 156


Slot:        8
Port:           1
Interface:       156
Data Rate:       40640000

Trunk Status:    Full             Trunk Overhead:  5%
Remote Node:           150.202.83.13  Remote Interface:          68

       Trunk              Out BW        Out BW
       Oversub.:          Avail.        Alloc.
Qos1     100%               90947          4792
Qos2     100%               90947             0
Qos3     100%               90947             0
Qos4     100%               90947           110
Administrative Status:          Up   Operational Status:           Up
```

Step 10: Identify the Ingress Circuit (RVc) from the Egress Card at the Transit Switch to the Ingress Card at the Egress Switch

At switch Dedham, use the `show mpt vcarray` *`egress_card`* command to identify the circuit value from the egress card at the transit switch to the ingress card at the egress switch Braintree. This command also uses card 8, as in step 9, since card 8 is the ingress and egress card to switch Dedham.

In this example, the RVc is 80. The VCType is PRNT. LPort 156, is egress card 8, port 1. Oprt 68 is ingress card 5, port 1, at the egress switch.

```
dedham> show mpt vcarray 8

VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FEID Node            OutFrames
InFrames   Discard
141 0    RMPT  PRNT   ACTIVE  80   0    1    18   0      0 2      156   68   0    150.202.83.13   N/A-0
N/A-0
```

Step 11: Identify the Egress Circuit and the FE at the Root Switch

At switch Braintree, use the `show lport attributes` *`interface`* command with the value from the OPrt column from step 10 to determine the ingress card number.

Use the `show mpt vcarray` *`ingress_card`* to identify the egress circuit from the ingress card to the SP.

In the following example for card 5, the egress circuit is 117. A value of one (1) in the IOP column indicates the circuit has been passed to the primary SP in slot 1. A value of 2 will indicate the primary SP is in slot 2. The OPrt value of 4093 is a ‘dummy’ circuit for the SP. The ingress card VCType is Child..

▶ At the root node, when an LSP cell has been received on a cell IOM, in this case card 5, the IOM determines the cell is an MPT LSP based on its VPI. The IOM then examines the first five bits of the VCI to determine which card and FE the data is going to. This value matches the VPI value identified in “Step 6: Identify the Egress Circuit (OVc) from the Ingress Card to the Egress Card on the Ingress Switch”.

Use the `show mpt vcarray` command on the SP to identify the egress circuit from the ingress card that terminates at the SP.

In the example for the SP, the VC value matches the OVc value of ingress card 5. The Lport value is also 4093. The VCType is Root

```
braintree> show mpt vcarray 5

VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI IOP PPrt LPort OPrt  FEID Node           OutFrames     InFrames
Discard
80   0    RMPT  CHILD ACTIVE  141  117  1   18   0   1   1    68   4093 0    150.202.83.13   N/A-0         N/A-0

braintree> show mpt vcarray

VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI IOP PPrt LPort OPrt FEID Node            OutFrames     InFrames
Discard
117  0    RMPT  ROOT   ACTIVE    0    0    1     18   0    0   1    4093  0    0    150.202.83.13      0            0
0
```

Step 12: Send IP Traffic Through the LSP Data Path to Validate Forwarding

Use the `ping` command to verify connectivity from the IP source to the IP destination and then use the `show mpt vcarray` command and `show ip forwarding statistics` command to examine counters and statistics.

For each identified card and virtual circuit along the data path, use the `show mpt vcarray` command to check the InFrames and OutFrames counters of each virtual circuit. ‘n/a’ in these fields indicate that statistics are not available for the selected card.

Also, check the IP forwarding statistics on each identified card in the data path to determine if packets are being forwarded by MPT, hop-by-hop, or being dropped due to an error condition. The IP forwarding statistics counters are further described in “show ip forward statistics” on page B-66.

Use the `show ip interfaces` command to identify the slot and port number on the egress card to the IP destination. Use the `show arp` command to verify the ARP resolution is complete.

```
braintree## show ip interfaces

IpAddr              Lport Pport    Card       MTU   ARP   IARP  OPER   ADMIN

150.1.1.1/24        87    7.4      TFETHER-4  1500  ENA   DIS   UP     ENA

braintree## sho arp

IpAddr      LinkType    HwAddr          State         EntryType

150.1.1.2   ethernet    006097594bb3    Complete      Dynamic
```

LSP Connection Failure Reasons

LSP connection failure reasons are shared by MPT LSPs and point-to-point LSPs. Use the `show mpt all` command to view the MPT failure reasons. The Node/Port/Failure shows switch information that is used to isolate the point of failure.

LSP errors can occur because the LSP network configuration guidelines are not followed or because of system memory allocation issues.

Table B-3. LSP Connection Failure Reasons

Field	MPT LSP	Point -to- point LSP	Description
Table Legend: • = Supported			
AbrNoAltOption		●	The defined path and alternate path failed at the area border router.
ADD_RVC	●	●	VNN could not add a reassembly virtual circuit (RVC). This is a resource error.
ALLOC_CD	●	●	VNN could not allocate a VCC. This is a resource error.
ALLOC_VP	●	●	VNN could not allocate a VPC. This is a resource error.
ATMFR	●	●	LSP is cell-only and the port is frame. This is a cell/frame boundary condition.
CNF_TIMEOUT	●	●	During the LSP calling signalling, a confirmation timer is set to expire after a period of time after a CALL PDU is sent from the root to the leaf. If a confirmation response is not received by the root from a leaf after this period of time, the timer will expire to end the call sequencing. This may occur during a LSP call failure when no return response is received including a LSP Reject PDU.
DEAD	●	●	LSP hello packets are not being received.
DISABLE		●	The point-to-point LSP is disabled.
FRATM	●	●	LSP is frame-only and the port is cell. This is a cell/frame boundary condition.

Table B-3. LSP Connection Failure Reasons (Continued)

Field	MPT LSP	Point -to- point LSP	Description
GROOM	●	●	A better path exists in the network and MPT signalling is trying to create the circuit over the new path.
ILGL_PORT	●	●	The port cannot support the LSP signal.
IMPURE	●	●	A shorter path exists but it is a cell/frame (mixed) path.
INACTIVE	●	●	An inactive trunk was found.
INV_TRUNK	●	●	An invalid trunk was found in the path.
IPROUTED	●	●	OSPF found a better IP-routed path.
MAX_PORT	●	●	The port number exceeds the maximum.
MGMT_ONLY	●	●	There can be only a management trunk in the path.
MULTI9000CELL	●	●	9000 cannot be a cell-only transit node.
NONE	●	●	The MPT is functional.
NO_PORT	●	●	The port does not exist.
NOBANDWIDTH	●	●	There is a bandwidth allocation failure.
NODETYPE	●	●	Node type identification in progress. This is a normal condition.
NOPATH	●	●	The route lookup failed because a route does not exist.
NORESOURCE	●	●	There is a memory allocation resource failure.
NORIDS	●	●	There are no more RIDs available at this node.

Table B-3. LSP Connection Failure Reasons (Continued)

Field	MPT LSP	Point-to-point LSP	Description
OPTIPTROOT		●	The optimum trunk border node cannot be a root for a point-to-point MPT.
OPTIRMT	●		The optimum trunk MPT VPI is not available. Check configuration to ensure a VPI is properly configured.
OPTITRNSPTVPI		●	The optimum trunk transit node point-to-point MPT VPI is not available. Check configuration to ensure a VPI is properly configured.
OPTITRNSRMTVPI	●		The optimum trunk transit node MPT VPI is not available. Check configuration to ensure a VPI is properly configured.
PATH_CLR	●	●	There is an OSPF Path Clear condition.
PATH_REG	●	●	There has been a failure to register a path with OSPF.
PTPTDISABLE		●	Point-to-point LSPs are disabled at the switch.
RMGR_SUSPEND	●		LSPs are disabled at the switch.
RT_LOOKUP	●	●	There is a route lookup failure.
RVC_DEAD	●	●	The Reassembly Virtual Circuit (RVC) is not receiving hello messages.
TD_CHANGE	●	●	The tleaf traffic descriptor changed.
TPORT_CALLING	●	●	VNN detected an MPT port calling error.
TPORT_DEAD	●	●	There is a dead MPT port.
TRKDOWN	●	●	A trunk is down.
UNKNOWN	●	●	The error condition that the switch reported does not match any of the error conditions in this table.

Table B-3. LSP Connection Failure Reasons (Continued)

Field	MPT LSP	Point-to-point LSP	Description
VC_CALLING	●	●	VNN detected a circuit endpoint calling error.
VCEXHAUST	●	●	There are no more VC entries left to be allocated. VC entries are used during call signalling for each VC along the path. There is a fixed amount of resources. Once these resources are allocated, MPTs are unable to setup.
VPICHANGE	●	●	Optimum trunk has a configured conflicting VPI.

LSP Path Failure Reasons

An LSP path is a circuit path consisting of a sequence of node identifiers and outbound interface indexes at switches along the established circuit.

The path failure reasons are shown in the display output of the `show mpt spath` or `show mpt dpath` command. The spath shows signalling path characteristics and the dpath shows data path characteristics for a specified node IP address.

You can also view Point-to-point LSP path parameters from the network map by selecting the appropriate switch object. From the Monitor menu, select Lucent IP Objects ⇒ Show Point-to-Point LSP. The Show All Point-to-point LSP Connections dialog box Fields" appears.

Table B-4. LSP Path Failure Reasons

Field	Description
CONFIRMTIMEOUT	During the LSP calling signalling, a confirmation timer is set to expire after a period of time after a CALL PDU is sent from the root to the leaf. If a confirmation response is not received by the root from a leaf after this period of time, the timer will expire to end the call sequencing. This may occur during a LSP call failure when no return response, including a MPT Reject PDU, is received.
DEAD	LSP HELLO packets are no longer being received.
GROOMING	A better OSPF path exists in the network and signalling is trying to use the new path.
IMPUREPATH	A shorter path of mixed cells and frames exists.
NONE	No problem exists.
PATHCLEAR	OSPF is notifying LSP that the path is no longer preferred.
PATHREGISTER	The path is not registered with OSPF.
ROUTELOOKUP	The route lookup failed.
RVCDIED	The RVC is not receiving LSP HELLO messages.
TPCALLING	A mptTport is calling. This is expected behavior.
TPDEAD	The connection is dead.
TRUNKDOWN	A trunk in the LSP path is down.
VCCALLING	A VC_ENTRY is calling. This is expected behavior.

LSP Flag Characteristics

LSP Flags show characteristics about a LSP leaf (MPT LSP, point-to-point LSP and multicast LSP). A leaf represents the data source. Use the `show mpt all` command to view the LSP flag.

Flags contain a bitmap representation. To interpret a bitmap, convert the hex value to binary.

Example: In flag 0x121006, the 6 is 0110 in binary. The second bit is 010 which represents 2, pure cell. The third bit is 100 which represents 4, MPT LSP leaf.

Table B-5. LSP Flags

Hex	Field	MPT LSP	Point-to-point LSP	Description
Table Legend: • = Supported				
0x00000001	PURE FRAME	●	●	The leaf is a pure frame circuit.
0x00000002	PURE CELL	●	●	The leaf is a pure cell circuit.
0x00000004	FOR_IP	●	●	The leaf is a MPT LSP leaf.
0x00000008	IPROUTED	●	●	A hop-by-hop path exists (not a LSP). This error will be displayed for an inactive state.
0x00000010	NON_FRAME_NODE	●	●	The LSP leaf is on a CBX 500 or GX 550 switch.
0x00000020	NON_VP_CELL	●	●	B-STDX 8000/9000 is a transit node in a pure cell path. This is an invalid configuration.
0x00000040	NEG_PATH		●	The path has insufficient bandwidth.
0x00000080	PT_PT_CONN		●	The leaf is a point-to-point MPT leaf.
0x00000200	DEFINED_PATH		●	A point-to-point defined path is configured.

Table B-5. LSP Flags

Hex	Field	MPT LSP	Point -to- point LSP	Description
0x00000400	ALT_PATH_OPT		●	A point-to-point alternate path is configured.
0x00000800	NODE_ARCH	●	●	The LSP leaf is on a B-STD 8000/9000 switch.
0x00001000	FE_ARCH	●	●	The LSP leaf is either on a CBX 500 or a GX 550.
0x00002000	USE_DEFPATH		●	The point-to-point defined path is being used.
0x00004000	IS_GARNET	●	●	The LSP leaf is on a GX 550 switch.
0x00008000	VC_EXHAUSTED		●	All virtual circuits are exhausted due to a resource error condition.
0x00010000	ENABLE_DEFPATH		●	The point-to-point defined path is enabled.
0x00020000	JADEM1_PATH	●	●	The minimum switch software version along the LSP path is Jade M1.
0x00040000	ABR_TRY_ALT		●	The defined path and alternate path are enabled. The define path fails and the alternate path is selected.
0x00080000	OSPF_PATH_REF	●	●	Register path when call is requested. This is expected call setup behavior.
0x00100000	OSPF_PATH_REG	●	●	Register path when call is completed. This is expected call setup behavior.
0x00200000	FWD			The leaf is a multicast LSP leaf.

LSP Console Commands

LSP command line interface commands are used for monitoring and debugging all types of LSPs. LSPs are debugged from a signalling and data standpoint. You may need to take a combination of the two debugging techniques to fully diagnose an LSP problem.

> The `show ip route` and `show ip forwarding` commands are useful for debugging LSPs. These commands are described on page B-64.

show mpt all

This command displays the MPT LSP and point-to-point LSP leafs established from a root switch. The LSPs are grouped by Type and OSPF Area.

MPT LSPs are operational when MPT LSPs are enabled and are suspended when LSPs are disabled. Point-to-point LSPs remain operational when MPT LSPs are disabled. The MPT LSP state is managed through Naviscore.

Syntax:

```
show mpt all
```

Example:

:

```
wellesley> show mpt all

MPT Identifier: 1  [OPERATIONAL]   Type: UNICAST (Multipoint-Point Tunnel)   OSPF Area: 0.0.0.0
Node            IOP/FE RID  Flags  RecvState SendState SendMpt LastFail(Node/Port/Reason)

150.202.83.8    CP/SP  -NA- 0x120805 ACTIVE    ACTIVE    11    None/0/NONE
150.202.83.13   CP/SP  -NA- 0x121006 ACTIVE    ACTIVE    8     None/0/NONE
150.202.83.13   4 /1   3    0x121006 ACTIVE    ACTIVE    8     None/0/NONE
150.202.83.13   7 /1   8    0x121006 ACTIVE    ACTIVE    8     None/0/NONE
150.202.83.13   9 /1   4    0x121006 ACTIVE    ACTIVE    8     None/0/NONE
150.202.83.13   10/1   5    0x121006 ACTIVE    ACTIVE    8     None/0/NONE
150.202.83.13   7 /2   6    0x121006 ACTIVE    ACTIVE    8     None/0/NONE
150.202.83.13   9 /2   7    0x121006 ACTIVE    ACTIVE    8     None/0/NONE
150.202.83.13   10/2   2    0x121006 ACTIVE    ACTIVE    8     None/0/NON
```

```
Teluride> show mpt all

MPT Identifier: 1  [SUSPENDED]    Type: UNICAST (Multipoint-Point Tunnel)
Node            IOP/FE Flags  RecvState SendState SendMpt LastFail(Node/Port/Reason)
```

The following table describes the fields in the `show mpt all` command.

Table B-6. show mpt all Command Fields

Field	Description
Node	IP address of the switch (node) that has a leaf attached to this root.
IOP/FE	The leaf can either be a CP on a B-STDX 8000/9000 or a FE on a CBX 500. The FE is displayed as slot/FE.
RID	A unique identifier assigned to each new leaf added to a CBX 500 MPT LSP or point-to-point LSP root.
Flags	A flag describes the type of LSP. Flags are described in "LSP Flags" on page B-44.
RecvState	This state determines whether the LSP root is ready to accept data over an LSP. States are `active`, `inactive`, `calling`, `retry`, `wcinact`, or `wcdel`.
SendState	This state determines whether the LSP leaf is ready to send data over an LSP. States are `active`, `inactive`, `calling`, `retry`, `wcinact`, or `wcdel`.
SendMpt	The LSP identifier that data will be sent over.
LastFail (Node/Port/Reason)	Failure reasons are described in "LSP Path Failure Reasons" on page B-43.

show mpt ckt

This command displays general LSP circuit information for the specified card and circuit.

Syntax:

```
show mpt ckt card-number vc-ID
```

Example:

```
wellesley> show mpt ckt 15 345

 MPT Circuit Info: Card: 15 VC: 345
Type   VPI   VCI   SVPI  SVCI   Out_VC  Out_IOP Out_Port
PARENT 642   0     642   0      0       0       1
Rvc Info:
VcId    Node            IOP/FE  Vpi   Vci   SVpi  SVci
Tport Info:
OutPort  IOP/FE  OutVc   OutIop
130      9  /2   381     9
21       9  /1   380     9
```

The following table describes the fields in the `show mpt ckt` command.

Table B-7. show mpt ckt Command Fields

Field	Description
Type	Child or Parent
VPI	Virtual Path Identifier
VCI	Virtual Circuit Identifier
SVPI	Spare Ingress VPI
SVCI	Spare Ingress VCI
Out_VC	Egress Virtual Circuit
Out_IOP	Egress Line Card
Out_Port	Egress port
Receiver Information	
VcID	Virtual Circuit Identifier
Node	Node Identifier
IOP/FE	Line Card/Forwarding Engine
VPI	Virtual Path Identifier
VCI	Virtual Channel Identifier
SVPI	Spare Virtual Path Identifier
SVCI	Spare Virtual Channel Identifier
Tport Information	
OutPort	Egress Port
IOP/FE	Line Card/Forwarding Engine
OutVC	Egress Virtual Circuit
OutIOP	Egress Line Card

show mpt dpath

This command displays the data path (IP data flows from leaf to root) for the selected node IP address and is executed from a leaf node.

Each node and logical lport in the data path to the specified node is shown. The node IP address is a value determined from the `show ip route` or `show mpt rootnodes` display output. The mpt_id is determined from either the MPT identifier value in the `show mpt all` table or the MPT column in the `show mpt vcarray` table. The vpn_id (VPN Identifier) value is determined from the `show mpt vcarray` table.

Syntax:

```
show mpt dpath [node-IP-address | node-IP-address mpt_id]
```

Example:

```
wellesley## show mpt dpath 150.202.83.26

MPT in VPN 0:
  List of node/interface pairs in path to node
150.202.83.26:
  150.202.83.26/27,  150.202.83.24/19, 150.202.83.6/34
```

show mpt error

This command displays error counters for MPT LSPs, point-to-point LSPs, and multicast LSPs.

The Global MPT counters represent an aggregate count for all LSPs. All other counters represent individual error counters for each LSP type. RMPT is a MPT LSP, PTPTMPT is a point-to-point LSP, and FMPT is a multicast LSP.

The default is the active CP/SP unless a specific card value is selected.

`Syntax:`

```
show mpt error {card-number}
```

`Example:`

```
Global MPT Exception Info: Card: 1

 errorChecksum 0
 errorConvert 0
 lastMidInvalid 0
 errorMidInvalid 0
 msgUnknownIn 0
 errorDVCPDUHLOIn 0
 errorDVCmsgHLOIn 0
 errorDVCPDUHAKIn 0
 errorDVCmsgHAKIn 0
 errorDVCPDUCLRIn 0
 errorDVCmsgCLRIn 0
 errorDVCPDUNAKIn 0
 errorDVCmsgNAKIn 0
 errorDVCPDUDISIn 0
 errorDVCmsgDISIn 0
 errorDVCmsgCLRIOP 0
 errorDVCPDUREJIn 0
 errorDVCmsgREJIn 0
 errorDVCPDURVCHLOIn 0
 errorDVCmsgRVCHLOIn 0
 errorDVCPDUTLHLOIn 0
 errorDVCmsgTLHLOIn 0

More...
```

Example: (continued

```
 RMPT Exception Info: Card: 1

 errorOutLink 0
 errorDVC 0
 errorDVCChain 0
 errorOptVP 0
 errorOptVPInUse 0
 errorOptVPNoChild 0
 errorOptVPNoItrChild 0
 errorOptVPNoParent 0
 errorFWDVCDestAlloc 0
 errorFWDVCChildAlloc 0
 errorFWDVCAllocParent 0
 errorSEGVCAlloc 0
 errorVCAlloc 0
 errorMPTVCAlloc 0
 errorDestPort 0
 errorTPABitMapProxy 0
 errorDead 0
 errorVcra 0
 errorParentVpiZero 0
 errorRvcVpiZero 0
 errorRvcVpiZero1 0
 errorRvcVpiZero2 0
 errorOSPFPathLookupFailed 0
 errorVctFreeATMAlloc 0
 errorVctFreeATMNotAlloc 5
 errorVctFreeATMNotFree 0
 errorVctFreeRateQ 0


 errorIPMGRTearDown 0
 errorIPMGROnChanged 0

PTPTMPT Exception Info: Card: 1

 errorOutLink 0
 errorDVC 0
 errorDVCChain 0
 errorOptVP 0
 errorOptVPInUse 0
 errorOptVPNoChild 0
 errorOptVPNoItrChild 0
 errorOptVPNoChild 0
 errorOptVPNoItrChild 0
 errorOptVPNoParent 0
 errorFWDVCDestAlloc 0
 errorFWDVCChildAlloc 0
 errorFWDVCAllocParent 0
 errorSEGVCAlloc 0
 errorVCAlloc 0
 errorMPTVCAlloc 0
 errorDestPort 0
 errorTPABitMapProxy 0
 errorDead 0
 errorVcra 0
 errorParentVpiZero 0
 errorRvcVpiZero 0
 errorRvcVpiZero1 0
 errorRvcVpiZero2 0
 errorOSPFPathLookupFailed 0
 errorVctFreeATMAlloc 0
 errorVctFreeATMNotAlloc 0
 errorVctFreeATMNotFree 0
 errorVctFreeRateQ 0

 More...
```

Example: (continued)

```
FMPT Exception Info: Card: 1

 errorOutLink 0
 errorDVC 0
 errorDVCChain 0
 errorOptVP 0
 errorOptVPInUse 0
 errorOptVPNoChild 0
 errorOptVPNoItrChild 0
 errorOptVPNoParent 0
 errorFWDVCDestAlloc 0
 errorFWDVCChildAlloc 0
 errorFWDVCAllocParent 0
 errorSEGVCAlloc 0
 errorVCAlloc 0
 errorMPTVCAlloc 0
 errorDestPort 0
 errorTPABitMapProxy 0
 errorDead 1
 errorVcra 0
 errorParentVpiZero 0
 errorRvcVpiZero 0
 errorRvcVpiZero1 0
 errorRvcVpiZero2 0
 errorOSPFPathLookupFailed 0
 errorVctFreeATMAlloc 0
 errorVctFreeATMNotAlloc 16
 errorVctFreeATMNotFree 0
 errorVctFreeRateQ 0


 errorAddMcastVc 0
 errorFMPTFwdObjOut 0
 errorFMPTFwdObjIn 0
 errorFWDLnkDummyNull 0
 errorFWDIOM 0
 errorfmptRxConMak 0
 errorfmptleafproxy 0
 errorConnId 0
 errorFMPTDis 0
 errorFmptNotInitialized 0
 errorFmptSendCid 0
 errorFmptSendBD 0
```

The following table describes the fields in the `show mpt error` command.

Table B-8. show mpt error Command Fields

Field	Description
AddMcastVc	add_vc_mid fails because the list is greater than 8.
Checksum	Checksum error exists.
ConnId	A connection identifier is not available.
Convert	PDU conversion error exists.
Dead	The dead timer has fired and killed a VC entry.
DestPort	A destination port cannot be found.

Table B-8. show mpt error Command Fields (Continued)

Field	Description
DVC	Called VC is not available.
DVCChain	Egress VC does not match.
FMPTDis	A multicast LSP request is disabled.
FMPTFwdObjIn	SVC forward object is missing.
FMPTFwdObjOut	RVC forward object is missing.
fmptleafproxy	There is an unknown port at the multicast LSP leaf.
fmptRxConMak	A destination port cannot be found.
FmptSendBD	No buffer multicast LSP root.
FmptSendCid	A CID error has occurred sending IP multicast into multicast LSP root.
FWDIOM	A IOM2 does not exist for destination FE .
FWDLnkDummyNull	The dummy link for multicast LSP FE is null at destination slot.
FWDVCAllocParent	A forward parent VC does not exist.
FWDVCChildAlloc	A forward child VC does not exist at the child leaf.
FWDVCDestAlloc	A forward VC does not exist at the destination.
IPMGROnChanged	The status has changed during the grooming process.
IPMGRTearDown	The IP Manager has torn down the LSP.
lastMidInvalid	MID is invalid.
MidInvalid	MID is not in vcra vcarray.
MPTVCAlloc	A VC does not exist at mpt_vct_alloc.
NoMPTLeaf	A leaf is not found for a LSP HELLO acknowledgement.
NoMPTRoot	A root is not found for a LSP HELLO acknowledgement.

Table B-8. show mpt error Command Fields (Continued)

Field	Description
NotInitialized	Event is ignored because LSP is not yet initialized.
OptVP	VPI not found on optimum trunk.
OptVPInUse	VPI is in use.
OptVPNoChild	VPI configuration problem exists.
OptVPNoItrChild	VPI configuration problem exists.
OptVPNoParent	VPI configuration problem exists.
OSPFathLookupFailed	OSPF path lookup failed on LSP leaf.
OutLink	No link is available.
ParentVpiZero	Detected a VPI of 0 on the parent.
SEGVCAlloc	A segmentation VC does not exist at the VC.
TPABitMapProxy	The tport array cannot find proxy for bitmap.
Unknown Message	An unknown message has been received by the switch.
VCAlloc	Unable to allocate a VC.
Vcra	Not in vcarray.
VctFreeATMAlloc	Attempt to free VC entry but an ATM connection ID is still allocated.
VctFreeATMNotAlloc	ATM connection ID is not allocated when the vc_entry freeing bit is set.
VctFreeATMNotFree	ATM connection ID cannot be freed when the vc_entry bit is set.
VctFreeRateQ	VC is still in rate Q on vct_free.
VpiZero	Deteced a VPI of 0 on RVC.

show mpt rootnodes

This command displays the root node of a circuit. The command is executed from a leaf node. In the display output, `VcId` shows the virtual circuit of the leaf and `slot` and `feNum` shows where the root is located.

On the CBX 500 only, enter a value for card and forwarding engine (`feid`). The card must be an Ethernet or a Frame DS3 type. The `feid` is one (1) for forwarding engine 1 and two (2) for forwarding engine 2.

The vcid is optional. When a virtual circuit (vc) is specified, only information about this particular vc is displayed; otherwise, information about all vcs is displayed.

For the B-STDX, use the following syntax:

```
show mpt rootnodes {vc-id}
```

For the CBX, use the following syntax:

```
show mpt rootnodes card-number fe-id {vc-id}
```

The following CBX 500 example shows the rootnode of vcid 210 is 150.202.83.4.

```
Byfield83_3> show mpt rootnodes 14 1

MPT ckt/node Info: Card 14 FE 1

VcId   RootNode        Slot      FE    mpt    vpn
210    150.202.83.4     3        1      1      0
211    150.202.83.4     4        1      1      0
212    150.202.83.4     12       1      1      0

Byfield83_3> show mpt rootnodes 14 1 210

MPT ckt/node Info: Card 14 FE 1

VcId   RootNode        Slot      FE    mpt    vpn
210    150.202.83.4     3        1      1      0
```

The following B-STDX example shows the rootnode of vcid 48 is 150.202.83.12.

```
andover> show mpt rootnodes

MPT ckt/node Info:

VcId   RootNode        Slot    feNum  mptID   vpn
43     150.202.83.9    CP/SP    N/A     1      0
798    150.202.83.11   CP/SP    N/A     9      167
48     150.202.83.12    6        0      1      0
48     150.202.83.12    7        0      1      0

Andover> show mpt rootnodes 48

MPT ckt/node Info:

VcId   RootNode        Slot    feNum  mptID   vpn
48     150.202.83.12    6        0      1      0
```

show mpt signal

This command displays call signalling information for all types of LSPs. Call signalling is an action which occurs when LSP processing is being established. This action is referred to as call setup. This is a useful command for debugging LSP signalling problems.

The default is the active CP/SP unless a specific card value is selected.

Syntax:

```
show mpt signal {card-number}
```

Example:

```
wellesley> show mpt signal 15

 Global MPT Signalling Info: Card: 15, FEBitMap 0x1e001e0, IPFEBitMap 0x1e001e0

                  IN MESSAGE   OUT MESSAGE       IN LNK       OUT LNK
   GR-TL HLO          2762        17109         12469         8837


 RMPT Signalling Info: Card: 15


            IN MESSAGE   OUT MESSAGE      IN LNK       OUT LNK
  Call               491          150          153          491
  Cnfrm              149          215          215          149
  Hello            11199        27460         3115        11199
  Hello Ack        27453        11199        11199         3112
  Clear              293           92           96          293
  NAK                  0          102          102           10
  Disrupt              0           13           13            9
  Reject               0           91           91            3
  RVC Hello            0            0            0            0
  RVC HAK              0            0            0            0
  TL Hello          7199         4801         4801          713
  IPTlHello          812            0            0            0
  ClrIOP               0            0            0            0

  Rmgr                 0            0


 PtPt-MPT Signalling Info: Card: 15


            IN MESSAGE   OUT MESSAGE      IN LNK       OUT LNK
  Call               463          341          341          463
  Cnfrm              131          166          166          131
  Hello            10738        85139        10606        10734
  Hello Ack        85112        10733        10733        10606
  Clear              243           13           56          243
  NAK                  0          104          104           75
  Disrupt              0           14           11            3
  Reject               0           97           97            0
  RVC Hello            0            0            0            0
  RVC HAK              0            0            0            0
  TL Hello         21817         2250         2250         2098
  IPTlHello         2381            0            0            0
  ClrIOP               0            0            0            0

More ...
```

Example: (continued)

```
FMPT Signalling  Info: Card: 15


            IN MESSAGE   OUT MESSAGE      IN LNK       OUT LNK
  Call               173          298          298          173
  Cnfrm               88           88           88           88
  Hello            46134        76009        43572        46134
  Hello Ack        76003        46134        46134        43568
  Clear               80           24           34           80
  NAK                  0           31           31            7
  Disrupt              0           25           23           27
  Reject               0           35           35            0
  RVC Hello            0            0            0            0
  RVC HAK              0            0            0            0
  TL Hello         17961        10580        10580        12654
  IPTlHello         4905            0            0            0
  ClrIOP               0            0            0            0
```

The following table describes the fields in the `show mpt signal` command:

Table B-9. show mpt signal Command Fields

Field	Description
Clear	Number of clear messages and PDUs sent and received on the card.
ClrIOP	Number of cleared IOP messages on the switch.
Disrupt	Number of disrupt messages and PDUs sent and received on the card.
FEBitMap	A bitmap of active FEs at the switch.
Fwd Call	For multicast LSPs, the number of outgoing call messages, incoming messages, and incoming and outgoing PDUs send and received on the link.
Fwd Cnfrm	For multicast LSPs, the number of outgoing confirm messages, incoming messages, and PDUs sent and received on the link.
GR-TL HLO	Number of group tleaf HELLO messages and PDUs sent and received on the card.
Hello	Number of hello messages and PDUs sent and received on the card.
Hello Ack	Number of hello acknowledgement messages and PDUs sent and received on the card.
IPFEBitMap	A current bitmap on the switch of FEs that have IP interfaces.
IP-TL HLO	Number of IP tleaf HELLO messages and PDUs sent and received on the card.
NAK	Number of NAK messages and PDU’s sent and received on the card.

Table B-9. show mpt signal Command Fields (Continued)

Field	Description
PtP Call	For point-to-point LSPS, the number of outgoing call messages, incoming messages, and incoming and outgoing PDUs sent and received on the link.
PtP Cnfrm	For point-to-point LSPs, the number of outgoing confirm messages, incoming messages, and PDUs sent and received on the link.
Reject	Number of reject messages and PDUs sent and received on the card.
RMGR	Number of RMGR messages sent and received on the switch.
RMT Call	For MPT LSPs, the number of outgoing call messages, incoming messages, and incoming and outgoing PDUs sent and received on the link.
RMT Cnfrm	For MPT LSPs, the number of outgoing confirm messages, incoming messages, and PDUs sent and received on the link.
RVC HAK	Number of RVC HAK message and PDUs sent and received on the card.
RVC Hello	Number of RVC hello messages and PDUs sent and received on the card.
TL-Hello	Number of tleaf HELLO messages and PDUs sent and received on the card.

show mpt spath

This command displays MPT LSP and point-to-point LSP signalling path characteristics for the specified node address. Each node and logical port in the signalling path to the specified node is displayed. The signalling path is from the *root* to the *leaf* and the command is executed at the root switch. The node IP address is the internal switch IP address of the leaf.

You can also use the `show mpt dpath` command that shows the data path from the *leaf* to the *root*. See “show mpt dpath” on page B-49 for a description of the command.

This is a useful command for debugging LSP signalling problems by identifying the nodes in the signalling path.

Path characteristics are described in “LSP Flags” on page B-44.

Syntax:

```
show mpt spath node-ip-address
```

Example:

```
Wellesley> show mpt spath 150.202.83.13

MPT in VPN 0:
  List of node/interface pairs in path to node 150.202.83.13:
  150.202.83.12/71,   150.202.83.28/13,   150.202.83.30/14
  Path characteristics:
  PURE_CELL,   FOR_IP,   FE_ARCH,   JADEM1_PATH,   OSPF_PATH_REG

MPT Pt-Pt in VPN 0:
  List of node/interface pairs in path to node 150.202.83.13:
  150.202.83.12/71,   150.202.83.28/13,   150.202.83.30/14
  Path characteristics:
  PURE_CELL,   PT_PT_CONN,   FE_ARCH,   JADEM1_PATH,   OSPF_PATH_REG

MPT Pt-Pt in VPN 167:
  Path characteristics:
  PT_PT_CONN,   FE_ARCH
```

show mpt vcarray

This command displays a list of virtual circuits for all types of LSPs. A vcarray is a set of virtual circuit entries. The *summary* keyword, used on a frame card, restricts forwarding engine RVCs from being displayed. Otherwise, all RVC vctypes will be displayed. On a CBX 500, only the first FE send circuit is displayed in the vcarray display output.

The default is the active CP/SP unless a specific card value is selected.

Syntax:

show mpt vcarray [*card-number* **|** *node-IP-address* **|** *card-number node-IP-address* **|** summary *card-number*]

Examples of the show mpt vcarray command:

```
wellesley> show mpt vcarray

VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI   IOP PPrt LPort OPrt FE OutFrames InFrames Discards
113  167  PMPT  ROOT   ACTIVE   0    0    3    0    0     0   1    4093  0    0  0         0          0
121  0    PMPT  ROOT   ACTIVE   0    0    2    0    0     0   1    4093  0    0  0         0          0
122  0    RMPT  ROOT   ACTIVE   0    0    1    0    0     0   1    4093  0    0  0         0          0
127  0    PMPT  ROOT   ACTIVE   0    0    4    0    0     0   1    4093  0    0  0         0          0
130  0    RMPT  LEAF   ACTIVE   0    150  1    0    0     7   1    4093  62   0  0         0          0
131  0    RMPT  ROOT   ACTIVE   0    0    12   0    0     0   1    4093  0    0  0         0          0
473  133  FMPT  LEAF   ACTIVE   0    1574 32770 0    0     9   2    4093  65   0  0         0          0
474  0    RMPT  LEAF   ACTIVE   0    1577 1    0    0     9   2    4093  65   0  0         0          0


wellesley> show mpt vcarray 9

VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FE OutFrames InFrames Discard
134  169  FMPT  ROOT   ACTIVE   0    139  32769 0   0      0   8    4042  0    0  N/A-0     N/A-0
1574 133  FMPT  PRNT   ACTIVE   169  0    32770 0   0      0   2    65    0    0  N/A-0     N/A-0
1577 0    RMPT  PRNT   ACTIVE   173  0    1     0   0      0   2    65    0    0  31        N/A-0
1601 0    RMPT  CHILD  ACTIVE   199  122  1     0   0      1   2    65    4093 0  N/A-0     N/A-0


Asbury83_4> show mpt vcarray 150.202.83.2

VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FE OutFrames InFrames Discard
116  0    RMPT  LEAF   ACTIVE   0    78   1    0    0      10  1    4093  30   0  0         0        0
118  0    FMPT  LEAF   ACTIVE   0    481  2    0    0      10  1    4093  30   0  0         0        0


Rowley83_1> show mpt vcarray 11 150.202.83.2
VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FE OutFrames InFrames Discard
41   0    FMPT  PRNT   ACTIVE   35   35   2    0    0      0   1    20    0    0  N/A-0         N/A-
43   0    RMPT  PRNT   ACTIVE   37   0    1    0    0      0   1    20    0    0      0         N/A-


Millwood> show mpt vcarray 5
VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FE     OutFrames InFrames   Discard
40   0    FMPT  PRNT   ACTIVE   85   5    2    0    43     0   1    6     0    1      0            2467
41   0    FMPT  CHILD  ACTIVE   480  6    3    0    42     1   1    6     4093 1  N/A-0         N/A-0
42   0    FMPT  PRNT   ACTIVE   89   7    3    0    45     0   1    6     0    1      0            2530
43   0    RMPT  CHILD  ACTIVE   481  4    1    0    0      1   1    6     4093 1  N/A-0         N/A-0
44   0    RMPT  RVC    ACTIVE   481  4    1    0    48     1   1    6     4093 1      0            0
45   0    RMPT  RVC    ACTIVE   481  4    1    0    49     1   1    6     4093 1      0            0
46   0    RMPT  RVC    ACTIVE   481  4    1    0    50     1   1    6     4093 2      0            0


Millwood> show mpt vcarray summary 5
VC   VPN  Type  VCType State    RVc  OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FE     OutFrames InFrames Discard
40   0    FMPT  PRNT   ACTIVE   85   5    2    0    43     0   1    6     0    1      0            2495
41   0    FMPT  CHILD  ACTIVE   480  6    3    0    42     1   1    6     4093 1  N/A-0         N/A-0
42   0    FMPT  PRNT   ACTIVE   89   7    3    0    45     0   1    6     0    1      0            2558
43   0    RMPT  CHILD  ACTIVE   481  4    1    0    0      1   1    6     4093 1  N/A-0         N/A-0
```

The following table describes the fields in the `show mpt vcarray` command:

Table B-10. show mpt vcarray Command Fields

Field	Description
Discard	Total number of discarded frames
FE	Forwarding engine identifier
Inframes	Total number of incoming frames
IOP	Card that contains a vcentry
Lport	Logical port identifier of the card that contains the vcentry
Mpt	Value maps to the MPT Identifier for the endpoint
OPrt	Lport of the child virtual circuit from the leaf point of view or lport of a parent virtual circuit from a root point of view
OutFrames	Total number of outgoing frames
OVc	Outgoing VC
PPrt	Physical port identifier of the card (IOP)
RVc	Remote VC
State	Active, Inactive
Type	MPT LSP, Point-to-point LSP, Multicast LSP
VCI	Virtual Circuit Identifier
VCI	Virtual Channel Identifier
VCType	Child, Parent, Root, Leaf, FE, RVC
VPI	Virtual Path Identifier
VPN	Virtual Private Network

show mpt svcarray and show mpt rsvcarray

The `svcarray` command displays a list of virtual circuits used to *send* IP data. The `rsvcarray` command displays a list of virtual circuits used to *send* and *receive* IP data. A vcarray is a set of virtual circuit entries.

Use this command on cards with forwarding engines. These types of cards use circuits for sending and receiving IP data. See Table B-15 on page B-68 for a list of cards with forwarding engines.

This is a useful command for debugging LSP signalling problems by identifying the circuits used to send and receive IP data.

Use the `show mpt svcarray` and `show mpt rsvcarray` command to debug LSP signalling problems.

Syntax:

`show mpt svcarray [`*`card-number`*` | `*`card-number node-ip-address`*`]`

`show mpt rsvcarray [`*`card-number`*` | `*`card-number node-ip-address`*`]`

To determine the card value, use the show mpt vcarray command and then select a card number from the IOP column.

Example of the `show mpt rsvcarray` command:

```
dedham> show mpt rsvcarray 4
VC   VPN  Type  VCType State    DNde OVc  Mpt  VPI  VCI    IOP PPrt LPort OPrt FE     OutFrames      InFrames
139  0    RMPT  FE     ACTIVE   530c 79   10   0    10242  11  2    3970  161  1  N/A-0              0
140  0    RMPT  FE     ACTIVE   530c 79   10   0    12290  11  2    3970  161  1  N/A-0              0
155  0    RMPT  FE     ACTIVE   530c 79   10   0    14338  11  2    3970  161  1  N/A-0              0
156  0    RMPT  FE     ACTIVE   530c 79   10   0    16386  11  2    3970  161  1  N/A-0              0
```

The following table describes the fields in the `show mpt vcarray svcarray` and `show mpt rsvcarray` command.

Table B-11. show mpt svcarray and show mpt rsvcarray Command Fields

Field	Description
DNde	Remote destination node
FE	FE identifier
InFrames	Statistical count for incoming frames.
IOP	Egress IOP card

Table B-11. show mpt svcarray and show mpt rsvcarray Command Fields (Continued)

Field	Description
LPort	Logical port
Mpt	MPT identifier
Oprt	Outgoing port
OutFrames	Statistical count for outgoing frames.
OVc	Outgoing virtual circuit
PPrt	Physical port
State	State of the LSP. For example, Invalid, inactive, retry, calling,
Type	Parent, Child, FE
VC	Virtual Circuit
VCI	Virtual Circuit Identifier
VCType	Parent, Child, FE
VPI	Virtual Path Identifier
VPN	Virtual Private Network

IP Forwarding

Lucent IP switching uses a combination of IP lookups and virtual circuit/paths to improve network performance.

Routing Tables

The routing protocols running on the CP/SP are responsible for building a routing table from information learned through their respective routing updates. To maximize IP forwarding throughput, the routing tables are distributed to the IOPs so that routing table lookups are a local operation.

The following routing tables are maintained on the CP/SP:

Unicast – The unicast routing table contains a list of unicast destination addresses. These addresses include the IP address that identifies the next hop to reach a network, the state and cost of the IP route, the logical port being reported on, the age (in seconds) of the route advertisement (RIP only), and a label switch path circuit indicator.

Multicast – The multicast routing table contains a list of multicast source IP addresses. These addresses include multicast destination group IP addresses, upstream neighbor interface, multicast LSP identification number, and number of outgoing interfaces.

How an IP Packet is Forwarded

IP packets are checked at the ingress and egress ports of each switch in the Lucent network. The following describes what happens when a Lucent switch receives an IP packet.

1. Check if received packet matches a configured Packet Filter. If packet matches filter, accept or discard packet according to filter.
2. Datagrams that specify IP Header Options are automatically forwarded to the CP, which implements a full IP protocol stack capable of handling the options.
3. Check if received packet matches a configured Next Hop Resolution Protocol (NHRP) profile. If packet matches profile, forward according to profile.
4. Check if received packet matches a forwarding policy. If packet matches policy, forward according to policy.
5. Perform an IP lookup on the packet’s destination address. If the forwarding entry shows that the destination is a local switch IP address (SNMP, ping, etc.) forward to the CP/SP for processing.
6. If the forwarding entry contains an LSP circuit, forward it over the LSP to the appropriate egress switch in the cloud.

7. If the forwarding entry does not contain a circuit entry, forward it using the next-hop entry.

Route Protocol Priority

Routing protocols are assigned a route priority number, which is used by the CP/SP/NP to choose the best available route. Use the `show ip route` command to display the routing table.

- The higher numbered route has the higher priority. For example, a VNNE1 route, which has a route priority of 3 has priority over a IBGP route, which has a route priority of 2.
- If two different routes with the same protocol priority exist to the same destination IP address, the protocol with the lower cost is chosen. For example, cost is determined by comparing the administrative cost of an OSPF interface to the hop count (number of hops) of a RIP interface.

The following table describes the route protocol priority.

Table B-12. Route Protocol Priority

Route Priority	Route Protocol
0	Indirect
0	None
2	EBGP
2	IBGP
2	OSPFE2
2	RIP
2	VNNE2
3	OSPFE1
3	Static
3	VNNE1
10	OSPFIA
10	VNNIA
11	OSPF
11	VNN
12	Direct

IP Forwarding Console Commands

show ip forward statistics

This command displays Layer 2 (datalink) and Layer 3 (IP) counters and statistics for IOP modules that support IP routing.

IP statistics counters increment when a specific condition is met. For example:

- The RcvIP counter increments when the switch receives an IP packet.
- The FwdCP counter increments when an IP packet is forwarded to the CP/SP/NP.
- The ttl_exceed counter increments when the TTL of an IP packet has been exceeded.

If an expected counter does not increment, troubleshooting begins by determining if links are up, if ARP entries exist, if correct routing entries exist and are being forwarded based on configured route maps, and where the IP packets are being dropped.

Syntax:

```
show ip forward statistics [card card-number fe
forwarding-engine {1|2}]
```

A forwarding engine value is required for any card that a forwarding engine. Table B-15 lists cards with forwarding engines.

Table B-13 lists the logical ports and card types that support IP routing on the B-STDX 8000/9000. Table B-14 lists the logical ports and examples of card types that support IP routing on the CBX 500.

Table B-13. Logical Ports That Support IP Routing on the B-STDX 8000/9000

Logical Port	Card Types	Encapsulation	Address Resolution
FR UNI-DCE FR UNI-DTE FR NNI	Frame cards[a]	RFC1490	InARP (RFC1293) ARP (RFC1490)
PPP	Frame cards[a]	PPP	N/A
ATM UNI DTE ATM UNI DCE	Frame cards[a]	RFC 1483	InATMARP
ATM UNI DTE ATM UNI DCE	ATM cards[b]	RFC 1483	InATMARP

Table B-13. Logical Ports That Support IP Routing on the B-STDX 8000/9000 (Continued)

Logical Port	Card Types	Encapsulation	Address Resolution
Ethernet	2-port Ethernet card	IEEE SNAP Ethernet II	ARP
IP VPN Cloud	N/A	N/A	ARP

[a] Frame Card Examples = UIO, 4-T1, 4-E1, DSX-10, HSSI, Ch DS3, Ch DS 3/1/0, 12-E1

[b] ATM Card Examples = ATM CS DS3, ATM CS E3, ATM OC3

Table B-14. Logical Ports Supporting IP Routing on the CBX 500

Logical Port	Card Types	Encapsulation	Address Resolution
FR UNI-DCE FR UNI-DTE FR NNI	6-Port DS3 FR/IP card 4-Port Channelized DS3/1 FR/IP card	RFC 1490	InARP ARP
PPP	6-Port DS3 FR/IP card 4-Port Channelized DS3/1 FR/IP card	PPP	N/A
ATM UNI DTE ATM UNI DCE	ATM cards with an IP Server PVC connection	RFC 1483	InATMARP
Ethernet	4-port Ethernet card	IEEE SNAP Ethernet II	ARP
IP VPN Cloud	N/A	N/A	ARP

Table B-15. Card Types with Forwarding Engine

Switch	Card Type	Number of Forwarding Engines
B-STDX 8000/9000	2-port Ethernet 10/100 Base-T	1
CBX 500	4-port Ethernet 10/100 Base-T	2
CBX 500	6-port DS3 Frame/IP	2
CBX 500	4-port Channelized DS3	1

This example shows IP forwarding statistics for a card with a forwarding engine.

```
Byfield83_3> show ip forward statistics card 13 fe 1

IP forwarding stats for Frame Engine: 1

FE_1:Inbound Layer-2 Stats/Errors:
  if_id       =         0, trk_pid_val=         0, rcv_arp    =        55
  bad_arp     =         0, bad_snap   =         0, bad_llc    =         0
  bad_eth_v2  =         0, no_iplport =         0, bad_trk_pid=         0
  bad_eth_type=         0, bad_mcast  =         0, discards   =         0
  no_olnk     =         0, Fwd SFPK     =        55

FE_1:Inbound Layer-3 Errors:
  ip_vers     =         0, ip_hdr_short=         0, ip_trunc_pkt   =         0
  ip_chksum   =         0, ip_filter   =         0, rte_reject     =         0
  lnk_bcst    =         0, ip_bc       =         0, ip_lbc         =         0
  lock_discard=         0, ip_lnktype  =         0, inact_fwd_flags=         0
  no_fbufs    =         0, bad_prtcn_id=         0, invalid_mpt    =         0
  no_lnk      =         0, no_iplport  =         0, no_lnkproxy    =         0
  icmp_thtl   =         0, iparp_thtl  =         0, ip_lookup      =         0
  rte_indirect=         0, ttl_xceed   =         0, discard        =         0
  mc_tree_pend=         0, mc_wrong_if =         0, no_mc_src      =         0
  mc_ttl_xceed=         0, postcp_thtl =         0, unicast_in_tunnel=         0

FE_1:Inbound Layer-3 Informational Counters:
  Rcv IP    =       553, Fwd CP =       553, Fwd Opts=         0
  Fwd HBH   =         0, Fwd MPT=         0, Fwd Frag  =         0
  Rcv Tunnel=         0, Fwd Mcast =         0, Fwd Tun=         0
  Fwd FMPT  =         0, Rx Fmptitr=         0, Tx Fmptitr=         0

FE_1:Inbound IP Server statistics:
  IP Svr Tx =         0

FE_1:Outbound Layer-2 Stats/Errors:
  if_id       =         0, trk_pid_val=         0, rcv_arp    =         0
  bad_arp     =         0, bad_snap   =         0, bad_llc    =         0
  bad_eth_v2  =         0, no_iplport =         0, bad_trk_pid=         0
  bad_eth_type=         0, bad_mcast  =         0, discards   =         0
  no_olnk     =         0, Fwd SFPK   =         0

FE_1:Outbound Layer-3 Errors:
  ip_vers     =         0, ip_hdr_short=         0, ip_trunc_pkt   =         0
  ip_chksum   =         0, ip_filter   =         0, rte_reject     =         0
  lnk_bcst    =         0, ip_bc       =         0, ip_lbc         =         0
  lock_discard=         0, ip_lnktype  =         0, inact_fwd_flags=         0
  no_fbufs    =         0, bad_prtcn_id=         0, invalid_mpt    =         0
  no_lnk      =         0, no_iplport  =         0, no_lnkproxy    =         0
  icmp_thtl   =       347, iparp_thtl  =         0, ip_lookup      =         0
  rte_indirect=         0, ttl_xceed   =         0, discard        =         0
  mc_tree_pend=       746, mc_wrong_if =       960, no_mc_src      =         0
  mc_ttl_xceed=         0, postcp_thtl =         0, unicast_in_tunnel=         0

FE_1:Outbound Layer-3 IP Server Counters:
  Rcv IP  =   1142947, Fwd CP =       803, Fwd Opts=         0
  Fwd HBH =         0, Fwd MPT=         0, Fwd Frag  =         0
  Rcv Tunnel=         0, Fwd Mcast =   2274890, Fwd Tun=         0
  Fwd FMPT  =         0

FE_1:Outbound Layer-3 Informational Counters:
  SW Ctl  =       283, SW hbh =       344, SW ckt  =         0
  SW ipsvr=         0, SW Mcast=         0, Tx HBH =       344
  Tx MC ifcs =         0, Tx MC Nbrs =         0, Tx MC Trks =         0
  SW FMPTitr=   1142947, Tx FMPTitr=         0
```

This example shows IP forwarding statistics for a card without a forwarding engine. In this example, the card is a HSSI.

```
Rowley83_1> show ip forward statistics card 11

Ingress Counters:
  FwdIcmp     =       0, FwdIcmpThtl =       0
  postMcast   =       0, postOK      =       0, postThrottld =       0, postTotal      =       0
  fwd_cp      =     215, fwd_hbh     =   92194, fwd_mpt      =       0, fwd_itr        =       0
  options     =       0, split       =       0, dct          =       0

Egress Counters:
  ntu_ctl     =    1026, ntu_ckt     =   62055, ntu_itr      =       0, ntu_hbh        =     118
  mpt_all     =       0, mpt_gr      =       0, mpt_am       =       0, mpt_rdr        =       0
  ntu_hbh_ctl =     118, tx_hbh      =     130
```

The following table describe the fields shown in the `show ip forward statistics` command.

Table B-16. show ip forward Command Fields

Fields	Description
INGRESS COUNTERS	
FwdIcmp	ICMP post attempts
FwdIcmpThtl	Number of throttled ICMP packets sent to the CP/SP.
postMcast	Multicast post attempts
postOK	Posted to background
postThrottld	Discarded due to foreground/background throttle
postTotal	Total post attempts (OKs and failures)
fwd_cp	Datagrams forwarded to CP
fwd_hbh	Packet is forwarded hop-by-hop
fwd_mpt	Packet is forwarded on a LSP
fwd_itr	Packets received on an intermediate LSP node
options	IP datagrams specifying header options
split	IP datagrams with headers split across BTUs
dct	IP data is sent hop-by-hop to Direct Cell Trunk
EGRESS COUNTERS	
ntu_ctl	Datagram is passed to background (control)
ntu_ckt	Packet is forwarded on circuit
ntu_itr	Packet is forwarded on intermediate circuit
ntu_hbh	Packet is forwarded hop-by-hop
mpt_all	Forward all packets on an LSP.
mpt_gr	Forward all green frames
mpt_am	Forward all amber frames

Table B-16. show ip forward Command Fields (Continued)

Fields	Description
mpt_rdr	Forward all red frames
ntu_hbh_ctl	Forward control packets hop-by-hop
tx_ hbh	Packets are transmitted hop-by-hop
FE [1|2] INBOUND LAYER 2 STATISTICS & ERRORS	
if_id	Lport interface number
trk_pid val	IP packet is discarded because the PID in the trunk header is not set for LSP.
rcv_arp	Total number of ARP packets received
bad_arp	Total number of bad ARP packets received
bad_snap	Total number of packets with bad SNAP header
bad_llc	Total number of packets with bad LLC header
bad_eth_v2	Total number of packets with ethertype Version 2
no_iplport	No IP lport associated with IP interface.
bad_trk_pid	Invalid protocol in the trunk header.
bad_eth_type	Ethernet type is neither ARP or IP.
bad_mcast	Packet is neither an ethernet broadcast or a multicast packet.
discards	Discards occur for other reasons than those shown above.
no_olnk	No outgoing link.
Fwd SFPK	Total number of packets forwarded to SPFK
FE [1|2] INBOUND LAYER 3 ERRORS	
ip_vers	Packet is discarded because IP header contains incorrect IP version (not V4).
ip_hdr_short	Packet is discarded because IP header is too short (less than the expected 20 bytes) .

Table B-16. show ip forward Command Fields (Continued)

Fields	Description
ip_trunk_pkt	Packet is discarded because IP header is truncated.
ip_chksum	Packet is discarded because IP Header checksum is bad.
ip_filter	Packet is discarded because a filter has been configured to selectively block IP packets.
rte_reject	Packet is discarded because a filter has been configured to selectively block routes.
lnk_bcst	A link broadcast packet has been received. This packet is discarded because a link broadcast packet is not forwarded beyond the local segment.
ip_bc	A broadcast packet has been received. This packet is discarded because broadcast packets are not forwarded.
ip_lbc	Subnet broadcast not forwarded.
lock_discard	This counter normally increments when a process lock occurs. Counter is specific to IOM modules.
ip_lnktype	IP data has attempted to be forwarded on a bad or unsupported link type
inact_fwd_flags	This counter increments when a forwarding lookup error occurs. For example, LMI may be down at one end of the link, the link state may be changing, or there may be a configuration error at the IP interface.
no_fbufs	This is a memory resource error. This error may occur during unusually intensive IP fragmentation reassembly processing or during extensive forwarding of IP multicast traffic.
bad_prtcn id	This error occurs when the protcon doesn't know where to send the packet.
invalid_mpt	Invalid LSP errors
no_lnk	No link associated with ifNum
no_iplport	No IP Lport associated with ifNum

Table B-16. show ip forward Command Fields (Continued)

Fields	Description
no_lnkproxy	No link proxy
icmp_thtl	ICMP throttle discard
iparp_thtl	IPARP throttle discard
ip_lookup	No IP route entry exists
rts_indirect	Indirect route lookup error
ttl_xceed	IP ttl exceeded
discard	Layer 3 discards
mc_tree_pend	Cache entry in pending state
mc_wrong_if	Multicast packet arrived on wrong interface
no_mc_src	No multicast source entry in routing table
postcp_thtl	Post throttle disards
unicast_in_tunnel	Unicast packet in tunnel
mc_ttl_xceed	Multicast TTL exceeded
FE [1|2] INBOUND LAYER 3 INFORMATION COUNTERS	
Rcv IP	IP packet is received. This counter is an aggregate of all IP packets received on a frame engine. Note that IOM2 ports share forwarding engines.
Fwd CP	IP packet is forwarded to CP/SP. For example, this counter will increment when the switch receives a PING request which contains a destination IP address local to the switch.
Fwd Opts	This counter is increments when a packet containing a header with IP Options is received.
Fwd HBH	IP packet is forwarded hop-by-hop if a LSP is not available. This counter increments when a packet arrives on a FE and is forwarded to another FE or a frame card.
Fwd MPT	IP packet is forwarded on a LSP. If LSP is unavailable, packet will go hop-by-hop.

Table B-16. show ip forward Command Fields (Continued)

Fields	Description
Fwd Frag	Fragmented IP packets are forwarded. Fragmentation occurs when the source MTU is larger than the egress MTU. It is advisable to avoid fragmentation since the fragmentation/reassembly process is CP/SP intensive.
Rcv Tunnel	IP packet is received on tunnel.
Fwd Mcast	IP multicast packet is received.
Fwd Tun	IP packet is forwarded on tunnel.
Fwd FMPT	IP packet is forwarded on multicast LSP.
Rx Fmptitr	IP packet is received on multicast LSP.
Tx Fmptitr	IP packet is transmitted on multicast LSP.
FE [1|2] INBOUND IP SERVER STATISTICS	
IP Svr Tx	IP server local transmission
FE [1|2] OUTBOUND LAYER 2 STATISTICS & ERRORS	
if_id	Lport interface number
trk_pid val	IP packet is discarded because the PID in the trunk header is not set for LSP.
rcv_arp	Total number of ARP packets received
bad_arp	Total number of bad ARP packets received
bad_snap	Total number of packets with bad SNAP header
bad_llc	Total number of packets with bad LLC header
bad_eth_v2	Total number of packets with ethertype Version 2
no_iplport	No IP lport associated with IP interface.
bad_trk_pid	Invalid protocol in the trunk header.
bad_eth_type	Ethernet type is neither ARP or IP.
bad_mcast	Packet is neither an ethernet broadcast or a multicast packet.

Table B-16. show ip forward Command Fields (Continued)

Fields	Description
discards	Discards occur for other reasons than those shown above.
Fwd SFPK	Total number of packets forwarded to SPFK
FE [1|2] OUTBOUND LAYER 3 ERRORS	
ip_vers	Discard packets with incorrect IP version (not V4)
ip_hdr_short	IP header is too short (less than the expected 20 bytes)
ip_trunk_pkt	IP header is truncated
ip_chksum	IP Header checksum is bad
ip_filter	Blocked filter packets
rte_reject	IP route enter REJECT flag is set
lnk_bcst	Link broadcast packet
ip_bc	Broadcasts not forwarded (BC flag is set)
ip_lbc	Subnet broadcasts not forwarded
lock_discard	Route table is locked
ip_lnktype	Bad or unsupported link type
inact_fwd_flags	IP packet is discarded because circuit flow is turned off.
no_fbufs	There are no FMBUFs available
bad_prtcn id	This error occurs when the protcon doesn't know where to send the packet.
invalid_mpt	Invalid MPT errors
no_lnk	No link associated with ifNum
no_iplport	No IP Lport associated with ifNum
no_lnkproxy	No link proxy
icmp_thtl	ICMP throttle discard
iparp_thtl	IPARP throttle discard
ip_lookup	No IP route entry exists

Table B-16. show ip forward Command Fields (Continued)

Fields	Description
rts_indirect	Indirect route lookup error
ttl_xceed	IP ttl exceeded
discard	Layer 3 discards
mc_tree_pend	Cache entry in pending state
mc_wrong_if	Multicast packet arrived on wrong interface
no_mc_src	No multicast source entry in routing table
postcp_thtl	Post throttle disards
unicast_in_tunnel	Unicast packet in tunnel
mc_ttl_xceed	Multicast TTL exceeded
FE [1|2] OUTBOUND LAYER 3 IP SERVER COUNTERS	
Rcv IP	IP packet is received. This counter is an aggregate of all IP packets received on a frame engine. Note that IOM2 ports share forwarding engines.
Fwd CP	IP packet is forwarded to CP/SP. For example, this counter will increment when the switch receives a PING request which contains a destination IP address local to the switch.
Fwd Opts	This counter is increments when a packet containing a header with IP Options is received.
Fwd HBH	IP packet is forwarded hop-by-hop if a LSP is not available. This counter increments when a packet arrives on a FE and is forwarded to another FE or a frame card.
Fwd MPT	IP packet is forwarded on a LSP. If LSP is unavailable, packet will go hop-by-hop.
Fwd Frag	Fragmented IP packets are forwarded. Fragmentation occurs when the source MTU is larger than the egress MTU. It is advisable to avoid fragmentation since the fragmentation/reassembly process is CP/SP intensive.

Table B-16. show ip forward Command Fields (Continued)

Fields	Description
Rcv Tunnel	IP packet is received on tunnel.
Fwd Mcast	IP multicast packet is received.
Fwd Tun	IP packet is forwarded on tunnel.
Fwd FMPT	IP packet is forwarded on multicast LSP.
FE [1|2] OUTBOUND LAYER 3 INFORMATIONAL COUNTERS	
SW Ctl	Control packets are transmitted hop-by-hop.
SW hbh	Packets are transmitted hop-by-hop.
SW ckt	Packets are transmitted by a LSP.
SW ipsvr	Packets are transmitted by a IP Server.
SW Mcast	Multicast packets are transmitted.
Tx HBH	Packets transmitted hop-by-hop.
Tx MC ifcs	Multicast packets are transmitted to PPP and ethernet interfaces.
Tx MC Nbrx	Multicast packets are transmitted to ATM and Frame neighbors.
Tx MC Trks	Multicast packets are transmitted to trunks.
SW FMPTitr	Packets received by multicast LSP.
Tx FMPTitr	Packets transmitted by multicast LSP.

TCP/IP Statistics

Lucent switches use the Transmission Control Protocol (TCP) for various switch applications. These include TELNET, Network Time Protocol (NTP), Accounting Server, and Bulk Statistics Collector.

Use the `show tcp` command to display the TCP/IP statistics on a Lucent switch.

Syntax:

```
show tcp
```

Example:

```
Byfield83_3> show tcp

TCP Statistics:
 Connections:
              0 Attempted                     1 Accepted
              1 Established                   0 Dropped
              0 Closed                        0 Embryonic drop
             65 Rtt timed                    65 Rtt updated
              7 Delayed acks                  0 Timeout drop
              0 Retransmit timeout            0 Persist timeout
              0 Keepalive timeout             0 Keepalive probe sent
              0 Keepalivedrops

 Sent:
             74 Total pkts                   66 Total data pkts
           5309 Total bytes                   8 Ack only packets
              0 Window probes                 0 URG only packet
              0 Update only pkt               0 Control (SYN/FIN/RST) pkt
              0 Data packets                  0 Data bytes
                 retransmitted                   retransmitted

 Received:
            124 Total packets                64 In sequence
            102 Total bytes                   0 Bad cksum
              0 Bad offset                    0 Short packet
              0 Bytes after window            0 Pkts with data after window
              0 Pkts after close              0 Window probes
              0 Duplicate acks                0 Acks for unsent data
             65 Ack packets                5310 Bytes acked
              0 Duplicate packets             0 Duplicate bytes
              0 Out of order pkts             0 Out of order bytes
              0 Window update                 0 Pkts with duplicate data
              0 Duplicate bytes in partially duplicate packets
```

The following table describes the fields for the `show tcp` command.

Table B-17. show tcp Command Fields

Field	Description
TCP CONNECTIONS	
Attempted	Total number of initiated connections.
Accepted	Total number of SYNs (synchronize sequence numbers used to establish connection) received in LISTEN state.
Established	Total number of connections established actively or passively at switch.
Dropped	Total number of dropped connections (after SYN received).
Closed	Total number of connections closed (includes drops).
Embryonic drop	Total number of embryonic connections dropped before a synchronize sequence numbers (SYN) is received.
Rtt Timed	Total number of segments for which TCP tried to measure Rtt (round trip time).
Rtt Updated	Total number of times Rtt estimators updated.
Delayed acks	Total number of delayed ACKS (acknowledgement number) sent.
Timeout drop	Total number of dropped connections in retransmission timeout.
Retransmit timeout	Total number of retransmit timeouts.
Persist timeout	Total number of times persist timer expires.
Keepalive timeout	Total number of times keepalive timer or connection-establishment timer expire.
Keepalive probe sent	Total number of keepalive probes sent.
Keepalive drops	Total number of connections dropped in keepalive (established or awaiting SYN).
SENT	
Total pkts	Total number of packets sent in sequence.

Table B-17. show tcp Command Fields (Continued)

Field	Description
Total data ptks	Total number of packets sent.
Total bytes	Total number of bytes sent in sequence.
Ack only packets	Total number of ACK (acknowledgement number) packets sent.
Window probes	Total number of window probe packets sent.
URG only packet	Total number of packets sent with URG-only (data length = 0).
Update only pkt	Total number of window update-only packets sent (data length = 0).
Control (SYN/FIN/RST) pkt	Total number of (SYN/FIN/RST) packets sent (data length = 0).
Data packets retransmitted	Total number of data packets retransmitted.
Data bytes retransmitted	Total number of data bytes retransmitted.
RECEIVED	
Total packets	Total number of packets received.
In sequence	Total number of packets received in sequence.
Total bytes	Total number of received bytes.
Bad cksum	Total number of packets received with checksum errors.
Bad offset	Total number of packets received with invalid header length.
Short packet	Total number of packets received that are too short.
Bytes after window	Total number of bytes received beyond advertised window.
Pkts with data after window	Total number of packets received with some data beyond advertised window.
Pkts after close	Total number of packets received after connection is closed.
Window probes	Total number of window probe packets received.

Table B-17. show tcp Command Fields (Continued)

Field	Description
Duplicate acks	Total number of duplicate ACKs received.
Acks for unsent data	Total number of ACKs for unsent data.
Ack packets	Total number of received ACK packets.
Bytes acked	Total number of bytes ACKed by received ACKs.
Duplicate packets	Total number of duplicate ACKs received.
Duplicate bytes	Total number of bytes received in completely duplicate packets.
Out of order pkts	Total number of out-of-order bytes received.
Out of order bytes	Total number of out-of-order bytes received.
Window update	Total number of received window update packets.
Pkts with duplicate data	Total number of packets received with completely duplicate bytes.
Duplicate bytes in partially duplicate packets	Total number of duplicate bytes in part-duplicate packets.

UDP Statistics

Lucent switches use the User Datagram Protocol (UDP) protocol for various switch applications. These include software download, PRAM Sync, Switch Console Dump Function, Radius authentication, Trap/Fault Manager and the SNMP agent.

The UDP statistics increment for datagrams received and forwarded at the switch.

Syntax:

```
show udp
```

Example:

```
Byfield83_3> show udp

UDP Statistics:
           40579 total input packets,      40035 total output packets
 Errors:       0 pkt shorter than header,      0 bad checksum
               0 data len larger than pkt,     0 no socket on port
               0 no socket for broadcast,      0 dropped socket full
               0 other
```

The following table described the fields shown in the `show udp` command.

Table B-18. show udp Command Fields

Field	Description
Total input packets	Total number of received packets.
Total output packets	Total number of transmitted packets.
ERRORS	
Pkt shorter than header	Received datagrams with packet shorter than header.
Bad checksum	Received datagrams with checksum error.
Data len larger than pkt	Received datagram with data length larger than packet.
No socket on port	Received datagram with no process on destination port.
No socket for broadcast	Received broadcast/multicast datagrams with no process on destination port.
Dropped socket full	Received datagrams not delivered because input socket is full.

Index

NavisCore IP Navigator Configuration Guide
Customer Comments

Please take time to fill out this questionnaire so that we can do our best to meet your documentation needs. Then fax or e-mail your comments to our Technical Publications Dept. Your opinions are of great value to us!

FAX: (978) 692-1510 (Attn: Tech Pubs)
E-MAIL: cspubs@lucent.com

✍ What tasks did you perform using this guide? ______________________________

✍ Did you install the hardware/software? ______________________________

✍ If you were having trouble performing a task, were you able to find the information you needed? Was the index useful? ______________________________

✍ Were the examples and illustrations helpful for performing tasks? If not, how can they be improved? ______________________________

✍ Was there any information you needed that was not in the manual? If so, how can we best deliver that information to you? ______________________________

✍ What did you like/not like about the manual? ______________________________

✍ Do you have any other comments about the manual? ______________________________

Page **Description of Error**

________ ______________________________
________ ______________________________
________ ______________________________

Name ____________________ Company ____________________
Mailing Address ______________________________
Phone ____________________ E-mail address ____________________
Fax No. ____________________

✂ Cut Here