Configuring IP Services for B-STDX User’s Guide

Ascend Communications, Inc.

Product Code: 80056
Revision C
December 24, 1997

ASCEND COMMUNICATIONS, INC. END-USER LICENSE AGREEMENT

ASCEND COMMUNICATIONS, INC. 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 ASCEND SWITCH(ES) CONTAINING THE SOFTWARE, AND BEFORE USING THE ACCOMPANYING USER DOCUMENTATION. OPENING THE PACKAGE(S) OR USING THE ASCEND 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, ASCEND 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 ASCEND, AND IT SUPERSEDES ANY PRIOR PROPOSAL, REPRESENTATION OR UNDERSTANDING BETWEEN THE PARTIES.

1. License Grant. Ascend 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 Ascend’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 Ascend’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. Ascend’s Rights. You agree that the Software and the User Documentation are proprietary, confidential products of Ascend or Ascend'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 Ascend or Ascend'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 Ascend. Ascend may terminate this License Agreement upon the breach by you of any term hereof. Upon such termination by Ascend, you agree to return to Ascend the Program and all copies or portions thereof. Termination of this License Agreement shall not prejudice Ascend's rights to damages or any other available remedy.

5. Limited Warranty. Ascend warrants, for your benefit alone, for a period of 90 days from the date of shipment of the Program by Ascend (the “Warranty Period”) that the program diskettes in which the Software is contained are free from defects in material and workmanship. Ascend 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 Ascend of any warranties made under this Agreement. EXCEPT FOR THE WARRANTIES SET FORTH ABOVE, THE PROGRAM IS LICENSED “AS IS”, AND ASCEND 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. Ascend’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 Ascend for the use of the Program. In no event shall Ascend be liable for any indirect, incidental, consequential, special, punitive or exemplary damages or lost profits, even if Ascend has been advised of the possibility of such damages.

7. Proprietary Rights Indemnification. Ascend 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 Ascend 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 Ascend to direct the defense and settlement of the claim, and (c) provide Ascend with the authority, information, and assistance that Ascend 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 Ascend’s written consent. In any action based on such a claim, Ascend 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 Ascend, 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, Ascend 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 Ascend 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 Ascend 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 Ascend 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

12 Monitoring IP Objects

A CLI Commands

List of Figures

List of Tables

About This Guide

Configuring IP Services for B-STDX is a task-oriented guide that describes, step-by-step, the process for configuring and monitoring an IP interface. This guide is intended for users who will be accessing the CascadeView/UX NMS to configure and monitor IP interfaces in a network.

What You Need To Know

As a reader of this guide, you should be familiar with basic UNIX operating system commands and know how to use a mouse. You should possess a working knowledge of Internet routing architectures and IP routing procedures. This guide assumes that you have installed the Cascade switch hardware as well as the CascadeView/UX NMS. In addition, you should be familiar with the configuration process for a Frame Relay or ATM logical port within a B-STDX network.

Refer to the following guides for more information:

B-STDX 8000/9000 Hardware Installation Guide

Network Management Station Installation Guide

Network Configuration Guide for B-STDX/STDX

Customer Comments

Customer comments are welcome! Please fill out the Customer Comment Form located at the back of this guide and return it to us.

How To Use This Guide

This guide is intended for those users who are using Release 3.0 of the CascadeView/UX NMS to configure and monitor an IP interface on a Cascade 8000/9000 switch.

Before you read this guide, read the Software Release Notice (SRN) that accompanies the software.

Read	To Learn About
Chapter 1	An overview of the product.
Chapter 2	How to create an IP interface, establish a BGP session, and set up a route map. This chapter provides a variety of different configuration examples in tutorial format.
Chapter 3	Configuring an IP interface on a logical port.
Chapter 4	Configuring packet filters.
Chapter 5	Configuring an IP QoS PVC, defining an IP QoS flow profile, assigning an IP QoS PVC flow profile to a logical port.
Chapter 6	Configuring static ARP entries.
Chapter 7	Configuring the Routing Information Protocol (RIP).
Chapter 8	Configuring the Border Gateway Protocol (BGP).
Chapter 9	Configuring OSPF at the logical port and switch.
Chapter 10	Configuring static routes.
Chapter 11	Specifying the parameters required in order to establish filters, access lists, route maps, and routing policies.
Chapter 12	Retrieving status information about Cascade IP objects.
Appendix A	Using the CLI commands.
Appendix B	PRAM upload supported objects.

Related Documents

This section lists the related Cascade and third-party documentation that may be useful to reference.

Cascade

- *CascadeView/UX Network Management Station Installation Guide* (product code 80014)
- *CascadeView/UX Network Configuration Guide* (product code 80017)
- *Diagnostic and Troubleshooting Guide for B-STDX* (product code 80018)
- *B-STDX 8000/9000 Hardware Installation Guide* (product code 80011)

Third Party

- *RFC 1700 ASSIGNED NUMBERS. J. Reynolds, J. Postel. October 1994*

Conventions

This guide uses the following conventions to emphasize certain information, such as user input, screen options and output, and menu selections. For example:

Convention	Indicates	Example
`Courier Bold`	User input on a separate line.	**`eject cdrom`**
[*bold italics*]	Variable parameters to enter.	**[*your IP address*]**
`Courier Regular`	Output from a program.	`Please wait...`
Boldface	User input in text.	Type **cd install** and ...
Menu => Option	Select an option from the menu.	CascadeView => Logon
Blue border surrounding text	Notes and warnings.	See examples below.
Italics	Book titles, new terms, and emphasized text.	*CascadeView/UX Network Management Station Installation Guide*

Provides helpful suggestions or reference to materials not contained in this manual.

Warns the reader to proceed carefully in order to avoid equipment damage or personal harm.

Common Acronyms

This document references the acronyms described in the following list

Table 1. List of Common Acronyms

Acronym	Description
ARP	Address Resolution Protocol
AS	Autonomous System
ASE	Autonomous System External
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
CIDR	Classless Inter-Domain Routing
DVMRP	Distance Vector Multicast Routing Protocol
FTP	File Transfer Protocol
IARP	Inverse Address Resolution Protocol
ICMP	Internet Control Message Protocol
IFMP	Ipsilon Flow Management Protocol (RFC 1953)
IGMP	Internet Group Multicast Protocol
IP	Internet Protocol
LSA	Link State Advertisement
MOSPF	Multicast Open Shortest Path First
MPOA	Multi Protocol Over ATM
MPT	Multipoint to Point Tunneling
NBMA	Non-Broadcast Multi-Access
NHRP	Next Hop Routing Protocol

Table 1. List of Common Acronyms (Continued)

Acronym	Description
OSPF	Open Shortest Path First
PIM	Protocol Independant Multicast
PNNI	Private Network to Network Interface
PPP	Point-to-Point
PVC	Permanent Virtual Circuit
RARP	Reverse Address Resolution Protocol
RIP	Routing Internet Protocol
SVC	Switched Virtual Circuit
TCP	Transmission Control Protocol
TDP	Tag Distribution Protocol
UDP	User Datagram Protocol
VNN	Virtual Network Navigator

1

Overview

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

About IP Switching

IP Switching technology allows Ascend’s multiservice WAN switching platforms to assume the characteristics and role of an IP router. The main difference between Ascend’s IP switching and traditional IP routing is that in the core of the Ascend network, IP packets are switched instead of routed. In other words, instead of examining IP headers at each hop, Ascend switches examine the IP header only at the ingress and egress routers to the Ascend 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, adding IP switching allows service providers to optimize data traffic flow by eliminating 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.

Core routers are still required for connection LAN-based servers.

Ascend's Implementation of IP Switching

Ascend 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 CascadeView/UX for Ascend'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 switch on the edge of a WAN to run standard IP routing protocols with B-STDX switches on the edge of the Ascend cloud forwarding packets based on the IP address of the frames. Both frame relay and ATM interfaces are supported. Inside the cloud, CBX 500's 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 obtained via standard routing protocols running on the switch. Inside the cloud, OSPF is used as the Interior Gateway Protocol (IGP). The Border Gateway Protocol (BGP) is used as the Exterior Gateway Protocol (EGP), learning routes to networks in other autonomous systems (AS). You can use RIP, OSPF, or static routing on the links to the CPE Routers to learn what networks are reachable through them.

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:

- TCP/IP
- OSPF
- RIP-2
- BGP-4

Interior Gateway Protocols (IGP)

RIP, RIP-2 and OSPF are internal 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 Ascend switches.

Exterior Gateway Protocols (EGP)

BGP is a protocol that exchanges routing information between autonomous systems (ASs). An AS 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 Protocol (IGP), which advertises routes within the same autonomous system, such as over OSPF and RIP. 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. Interconnection between ASs typically do not use interior gateway protocols, instead they use protocols that are classified as exterior gateway protocols (EGP), such as BGP.

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 (IARP)
- 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. In a future release, IP Navigator will support Internet Group Multicast Protocol (IGMP), Multicast OSPF (MOSPF), and Distance Vector Multicast Routing Protocol (DVMRP) for multicasting.

Exchanging Route Table Information

In any routed network, routers learn about the view 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. IP Navigator uses Ascend's Virtual Network Navigator (VNN) for this purpose. Essentially, VNN is OSPF with Ascend-proprietary extensions.

Each router-enabled switch in the Ascend network communicates with every other router-enabled switch. Route tables are maintained in each switch and a master table is maintained in the switch's control processor (CP). 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.

Multipoint-to-Point Tunneling

IP Navigator uses Multipoint-to-Point Tunnels (MPT) to interconnect the router and switches. MPT allows multiple circuits to share the same circuit to transmit to a single destination. MPTs are a unique feature provided by Ascend. An MPT can be thought of as the inverse of the point-to-multipoint virtual circuit used to allow the sending of packets from one source to multiple destinations. Connections of this type are commonly used in multicast applications such as video distribution. The MPT reverses the direction of the information flow.

On adding a new switch to an IP Navigator network, the switch establishes MPTs to all other switches running IP Navigator in the network. This provides every switch with a path to the new switch.

MPTs are established using the best route available through the network when a circuit is established. The 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 the service provided on the link.

Establishing MPT Circuits

Switches in an IP Navigator network automatically establish MPT circuits upon startup. However, you must configure one MPT parameter. For more information on this, refer to “Configuring IP Parameters” on page 9-23. VNN calculates the best path from each node to the new switch using the extended version of OSPF. These paths are used to form the MPT. Major network changes cause VNN to recalculate the MPTs. If configured to do so, VNN will continually monitor the MPTs and recalculate them based on performance. VNN treats the MPTs the same way it treats any other connection. No special configuration is required for MPT monitoring.

About MPT Paths

The reverse multipoint tunnel (MPT) enables IP Navigator to switch connectionless protocols (e.g. IP) across the Ascend switch network. MPTs provide an efficient, fault tolerant, high performance protocol switching layer that is scalable to 400 switches in a network. MPTs run across CBX 500s and B-STDX 9000s, which run over direct or optimum trunks.

The reason for using MPTs over point-to-point tunnels is because MPTs need less circuits than point-to-point tunnels. MPTs 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.

Figure 1-1 displays an example of an MPT.

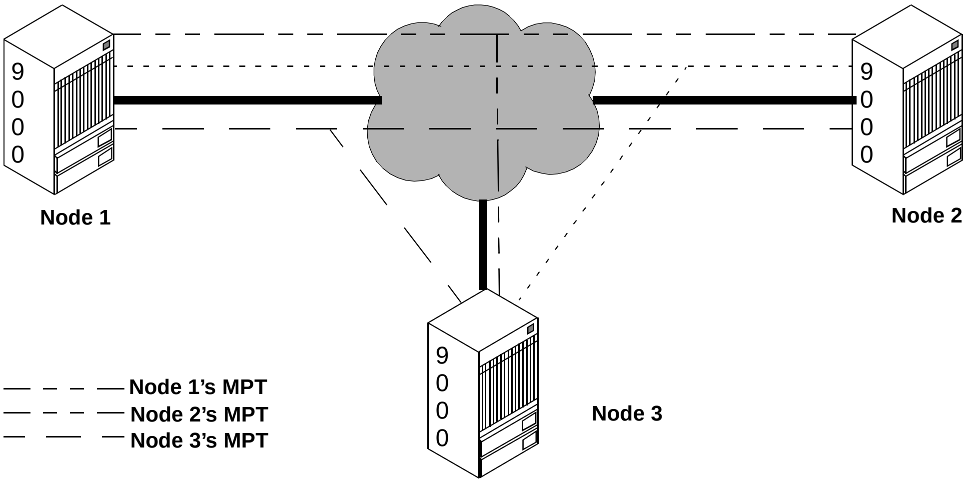

Figure 1-1. MPT Example

In Figure 1-1, the solid lines represent trunks into the Ascend network. Node 1 establishes its MPT first, followed by nodes 2 and 3. When node 1 establishes its MPT, nodes 2 and 3 can forward packets back to node 1, thereby reversing the direction of the MPT. The same idea applies to nodes 2 and 3 when they establish their MPTs. The number of circuits is equal to the number of nodes.

Every switch maintains one MPT circuit network and contains a root, which

- keeps track of MPT nodes
- adds and deletes nodes
- keeps nodes alive

A root is a standard VC_ENTRY that is created at initialization time on every CP card in the 9000 and every SP card in the 500.

Switch Domains

There are two types of switch domains:

Cell Domain — Paths that traverse direct ATM trunks.

Frame Domain — Paths that traverse pure direct frame trunks.

A switch can belong to multiple domains, however, the domains must be 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.

OSPF determines how MPTs connect two switches in different domains. The following factors apply when determining MPTs:

- A switch that only belongs to one domain cannot add a switch from a different domain to its MPT. MPTs are only established between switches in the same domain.
- If the shortest path between two switches in the same domain traverses a different domain, the switches cannot add each other to their MPTs.

MPTs use cell and frame domains to address limitations in switching ATM cells. Reassembling cells on these boundary switches require more resources than currently available on Ascend cell cards. However, a compromise can be reached, whereby large networks can reassemble cells but with additional processing required for data switched across cell/frame cards.

Establishing End-To-End QoS

The basic service offered by IP Navigator is considered *best effort*. IP traffic is transmitted with the lowest level of priority. If this best effort service is considered unacceptable, IP Navigator enables you to provision virtual circuits for a specific route. When you configure a circuit for a specific route, you can assign a desired QoS. As with all services offered on Ascend's switch platforms, the QoS is guaranteed on an end-to-end basis.

B-STDX Switching Functions

In the B-STDX, IP switching functions are divided between the control processor (CP) and the I/O processor (IOP) modules. Functions that collect, maintain, and distribute management and control information reside on the CP. These include routing protocols, ARP, SNMP, Telnet, TFTP, FTP, and configuration. The IOPs only contain enough information to forward all unicast and most multicast datagrams without CP intervention. The CP maintains a global view of the switch, while each IOP contains only enough information to make local forwarding decisions for data packets.

Physical and Logical Ports

You can configure any physical port on a B-STDX switch for IP routing. On these ports, the IP header of received frames is examined. The routing table, built by the native routing protocols, is used to determine the egress switch/port within the B-STDX cloud. IP user traffic is carried over automatically created virtual paths (VPs) connecting all switches in an autonomous system (AS). When the egress switch/port is determined, the IP data is forwarded over the corresponding MPT to that switch.

Packets received on a port are forwarded based on the contents of their IP header. Data encapsulation and address resolution are determined by the type of logical port.

Table 1-1 outlines the B-STDX logical ports that support IP routing.

Table 1-1. Logical Ports Supporting IP Routing

Logical Port	Card Types	Encapsulation	Address Resolution
FR UNI-DCE FR UNI-DTE	Frame cards[a]	RDC1490‘UI	Inverse ARP (RFC 1293 ARP (RFC1490) Static configuration
PPP (new lport type)	Frame cards[a]	None	Static configuration
ATM UNI DTE ATM UNI DCE	Frame cards[a]	RFC 1483	InATMARP ATMARP
ATM UNI DTE ATM UNI DCE	ATM cards[b]	RFC 1483	InATMARP ATMARP

[a] Frame Cards = UIO, 4-T1, 4-E1, DSX-10, HSSI, Ch T3
[b] ATM Cards = ATM UNI Rev C, ATM CS, ATM OC3

Configuration and Management

With the addition of IP Navigator, CascadeView/UX 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 [3]
- BGP-4 MIB [4]
- Routing Table MIB [5]
- RIP v2 MIB [6]

2

Configuration Tutorial

This chapter provides configuration examples for the following tasks:

- Configuring an IP interface
- Configuring a BGP session
- Configuring a BGP neighbor
- Confirming the addition of a BGP neighbor
- Configuring a BGP route map

The examples in this chapter describe the steps required to establish the network shown in Figure 2-1. All procedures assume you are logged on to CascadeView/UX and have already configured a switch for a Frame Relay or ATM logical port connection.

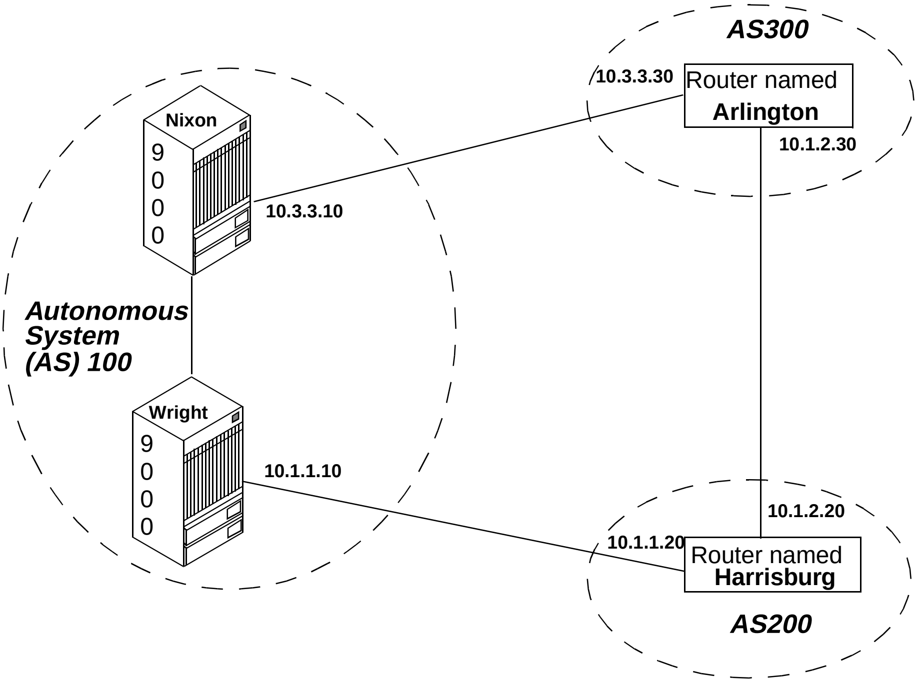

Figure 2-1. Network Used in this Tutorial

Configuring an IP Interface

To establish an IP interface:

1. Select the appropriate switch icon on the network map. For this example, the first connection is from the switch, Nixon, to the router named Arlington.

2. Create a Frame Relay logical port named *fr_to_ARLINGTON* on Nixon and enable it. Refer to the *Network Configuration Guide for B-STDX/STDX* for instructions about creating a Frame Relay logical port.

3. From the Administer menu select Cascade IP Parameters ⇒Set All IP Lports. The Set All IP Lports dialog box appears.

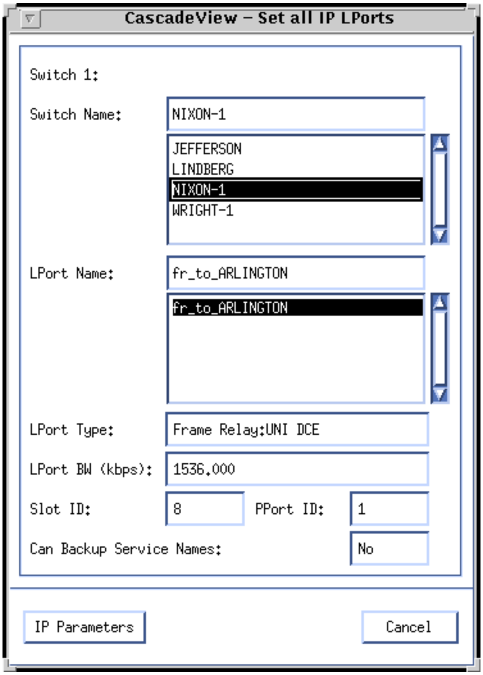

Figure 2-2. Set All IP LPorts

4. Select the Frame Relay logical port named *fr_to_ARLINGTON* and choose IP Parameters. The first of two Set IP Parameters dialog boxes appears.

5. Choose Add IP Lport. The second Set IP Parameters dialog box appears (Figure 2-3).

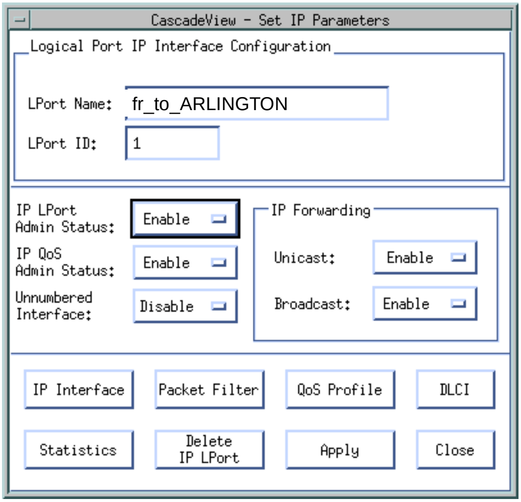

Figure 2-3. Set IP Parameters

6. Choose IP Interface. The Set IP Interface Addresses dialog box appears (Figure 2-4).

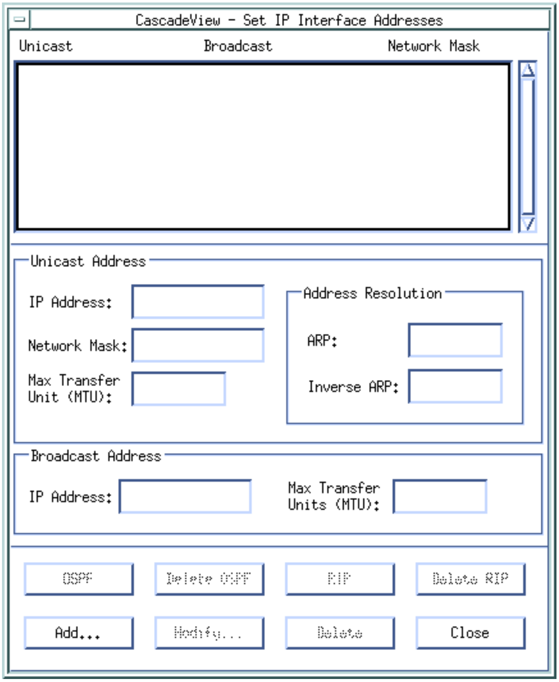

Figure 2-4. Set IP Interface Addresses

7. Choose Add. The Set IP Interface Address (Figure 2-5) dialog box appears.

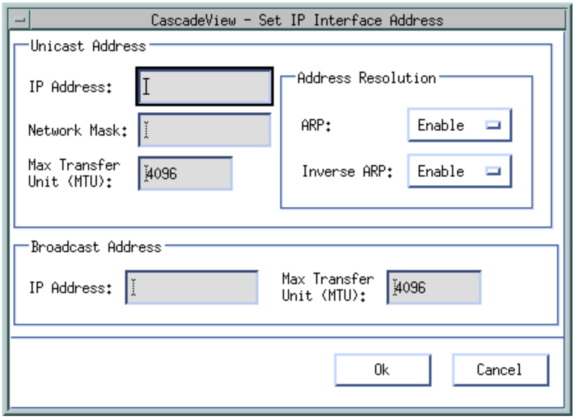

Figure 2-5. Set IP Interface Address

8. Enter the IP address, mask, and broadcast address as follows:
   - IP address: 10.3.3.10
   - Network Mask: 255.255.255.0
   - Broadcast Address: 10.3.3.255
9. Choose OK. The Set IP Interface Addresses dialog box reappears. The system lists the interface that you have just added at the top of the screen.
10. Choose Close to return to the Set IP Parameters dialog box.

Specifying the DLCI Parameter

To specify the DLCI parameter for the IP interface:

1. Choose DLCI from the Set IP Parameters dialog box. The IP Protocol Connection ID dialog box appears (Figure 2-6).

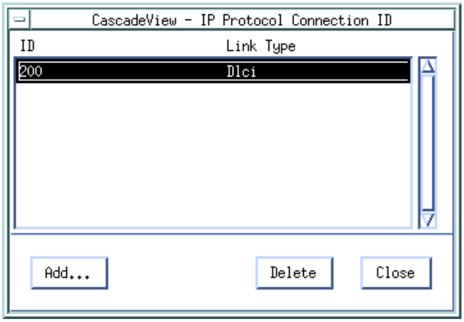

Figure 2-6. IP Protocol Connection ID Dialog Box

2. Choose Add from the IP Protocol Connection ID dialog box. The Set IP Protocol Connection ID dialog box appears.

3. Enter 100 as the DLCI value.

4. Choose OK.

Confirming the Addition of an IP Interface

To confirm the addition of the IP interface:

1. From the console issue the following command:

   show IP interface

   The system displays the following output:

IPAddr	LPort	PPort	Card	MTU	ARP	IARP	OPER	ADMIN	HWADDR
10.3.3.10./24	1	8.1	UIO-8	1500	ENA	ENA	UP	ENA	500

Configuring a BGP Session

To establish a BGP session:

1. Select the Wright switch icon from the network map.
2. Select Cascade IP Parameters ⇒Set All BGP⇒Set All BGP Parameters from the Administer menu. The Set BGP dialog box appears.

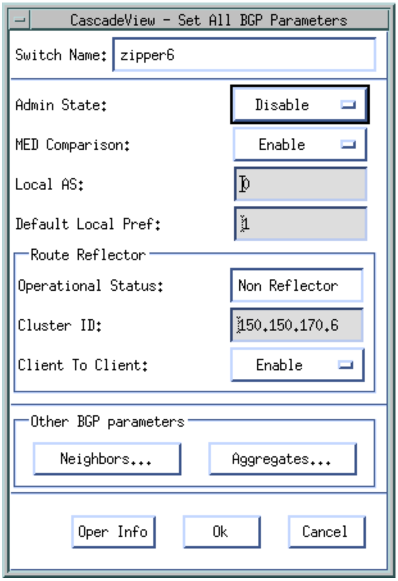

Figure 2-7. Set BGP Dialog Box

3. Set the Admin State value to Enable.
4. Accept the default values for the MED Comparison field.
5. Specify 100 for the Local AS number.
6. Specify the default local preference as 100.
7. Accept the default Route Reflector values.

Configuring a BGP Neighbor

This section describes how to configure the following BGP neighbor connections:

- From the B-STDX Switch (Wright) in AS 100 to a router (Harrisburg) in AS 200
- From the B-STDX Switch (Nixon) to a router (Arlington) in AS 300
- From the B-STDX Switch (Wright) in AS 100 to the second B-STDX Switch (Nixon) in AS 100

See Figure 2-1 for an example of these connections.

To establish a BGP neighbor:

1. Choose Neighbors from the Set All BGP Parameters dialog box. The Set All BGP Neighbors dialog box appears.

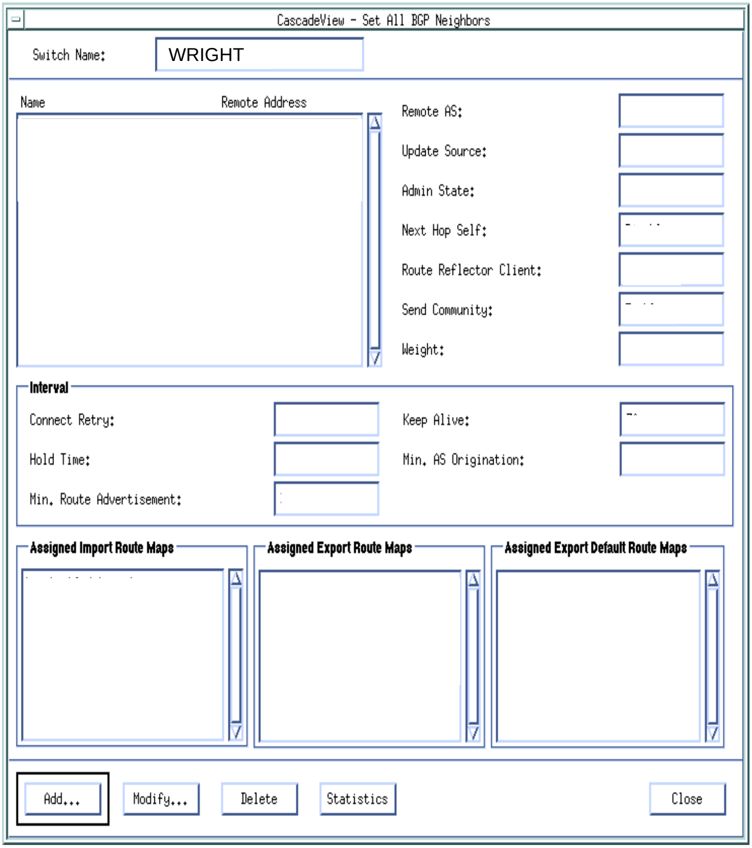

Figure 2-8. Set All BGP Neighbors

2. Choose Add. The Add BGP Neighbor dialog box appears (Figure 2-9).

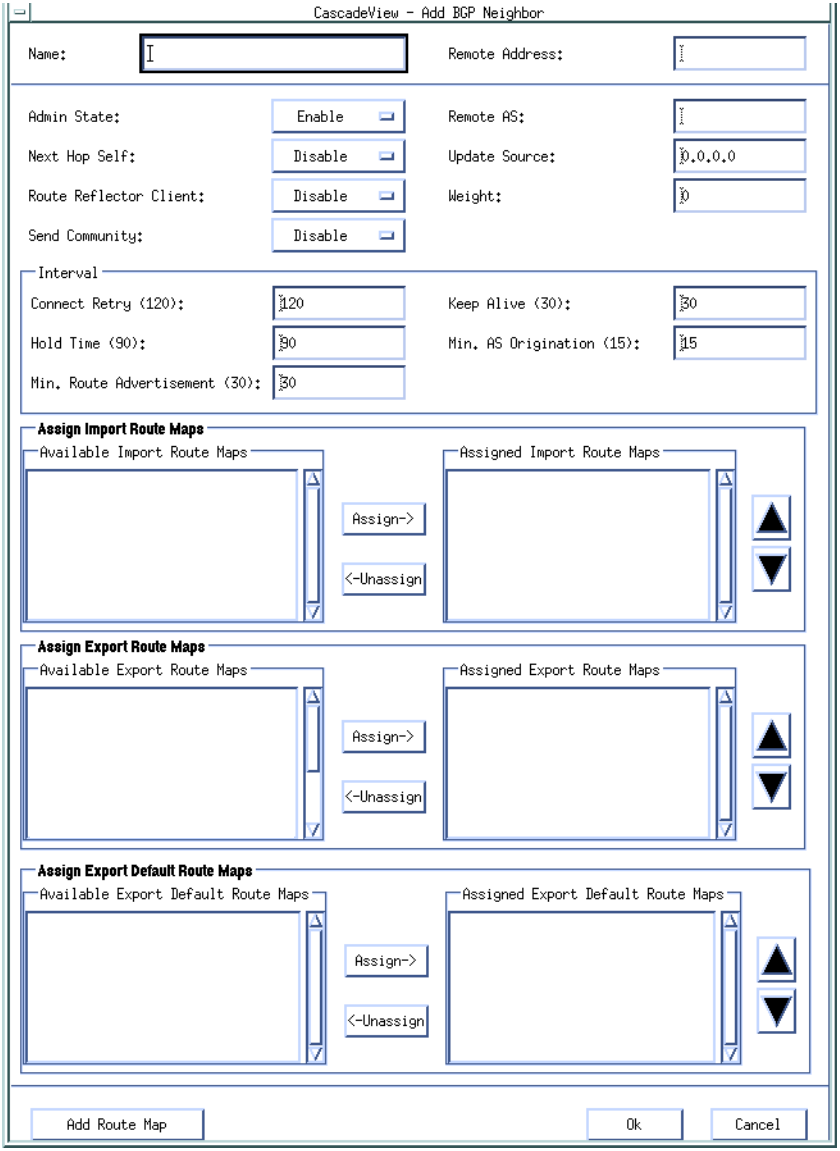

Figure 2-9. Add BGP Neighbor

3. Enter the name of the BGP Neighbor in the Name field.(In the first portion of this example, the name of the BGP Neighbor is *Harrisburg.*)

4. Specify the IP address of the BGP neighbor in the Remote Address field. (In the first portion of this example, the remote address is *10.1.1.20*).

5. Specify the remote AS number in the Remote AS field. (For the first portion of this example, the remote AS number is *200* because Harrisburg is in Autonomous System (AS) 200.

6. Choose OK.

Adding the Connection from Nixon-1to Arlington

1. Configure a BGP session in Nixon-1

2. Configure Arlington as a BGP neighbor. Specify the following values on the Add BGP Neighbor dialog box:

   Name: **Arlington**

   Remote Address: **10.3.3.30**

   Remote AS: **300**

3. Choose OK.

Adding the Connection from Wright to Nixon

1. Configure a BGP session in Nixon

2. Configure Wright as a BGP neighbor. Specify the following values on the Add BGP Neighbor dialog box:

   Name: **Wright**

   Remote Address: **Use the internal address**

   Remote AS: 1**00**

3. Choose OK.

Confirming the Addition of a BGP Neighbor

To confirm the addition of a BGP neighbor:

From the console terminal, issue the following command:

show bgp neighbor

The system displays the output for each BGP neighbor as shown in Figure 2-10.

```
wright## show bgp neighbor

Neighbor: 10.1.1.20

Admin State:             Enabled        Operating State:         Established
Remote AS:               200            Route Reflector Client:  No
Confg Keepalive Time:    30             Calc Keepalive Time:     30
Connect Retry Interval:  120            Configured Hold Time:    90
Established Time:        0:12:52:17     Negotiated Hold Time:    90
Update Elapsed Time:     0:12:52:16     Established Transitions: 1
Messages In:             1541           Updates In:              2
Messages Out:            1546           Updates Out:             0
Send-Community:          Disabled       Weight:                  0
Send-Default-Route:      Disabled
```

Operating state must be established or you do not have a connection with the remote AS

Figure 2-10. Output From the Show BGP Neighbor Command

Configuring a BGP Route Map

Route maps enable you to control and modify routing information and define the parameters that your system uses to redistribute routes between routing protocols and domains. Route maps enable you to alter route information before it is stored in the routing table.

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 can assign the map to an entity as a type of filter. In this example, the following conditions exist:

- The connection from Nixon to Arlington is a DS3 connection
- The connection from Wright to Harrisburg is a DS1 connection

You can control the exit path that is used by altering the local preference attribute. For this example, you will alter the local preference so that all data exits via Nixon.

To do this, you must create a BGP route map that specifies matching criteria for Nixon and then sets the local preference attribute when the match is found. The matching criteria that is used in this example is the last AS value (which for Nixon is currently a value of 300).

Creating the Route Map

A route map acts on various routing protocols. The route map for this example will affect the BGP protocol. In the following example you will create a map that filters routes based on match criteria for a Last AS value of 300. The set criteria that is specified as a result of a match is that the local preference value is set to 500 on all of the filtered routes.

To create the route map:

1. Select Cascade IP Parameters ⇒Set All Route Policies⇒Set All Route Maps from the Administer menu. The Set All Route Maps dialog box appears.
2. Choose Add. The Add Route Map dialog box appears (Figure 2-11).

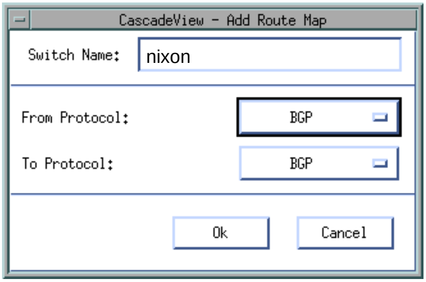

Figure 2-11. Add Route Map

3. Choose OK from the Add Route Map dialog box. The system displays the Add BGP to BGP Route Map dialog box shown in Figure 2-12.

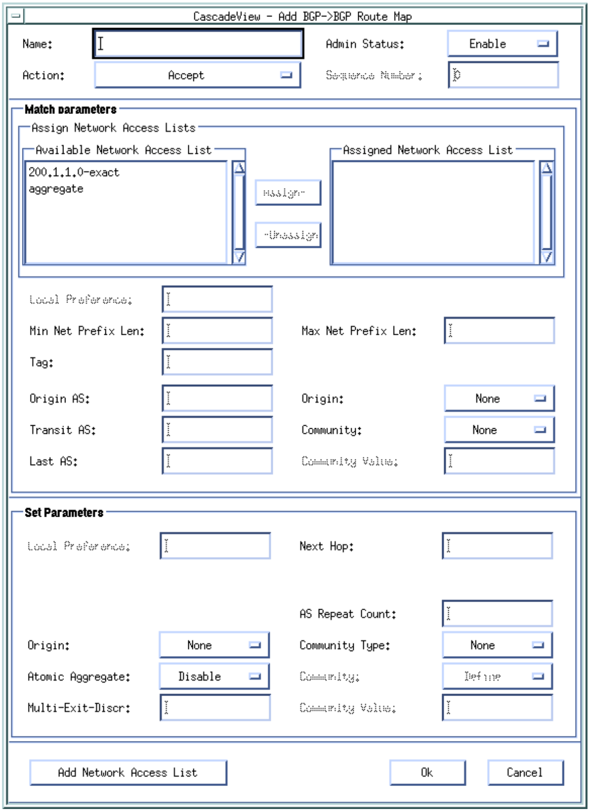

Figure 2-12. Add BGP to BGP Route Map

4. Specify the route map name. For this example use the name *LocalPrefforAS300.*

5. Under Match Parameters, specify a value of 300 for the Last AS field. All other Match Parameter values can be left blank.

6. Under Set Parameters, specify a local preference value of 500.

7. Set the Admin Status field to Enable.

8. Choose OK. The system adds the route map to the list of route maps.

Associating the Route Map with a Neighbor

The next step is to associate the route map with a BGP entity. In this example, you will associate the route map with the BGP neighbor named Arlington.

To associate the route map with a BGP neighbor:

1. Select Cascade IP Parameters ⇒Set All BGP⇒Set All BGP Neighbors from the Administer menu. The Set All BGP Neighbors dialog box appears.
2. Select the neighbor named Arlington.
3. Choose Modify. The Modify BGP Neighbor dialog box appears.

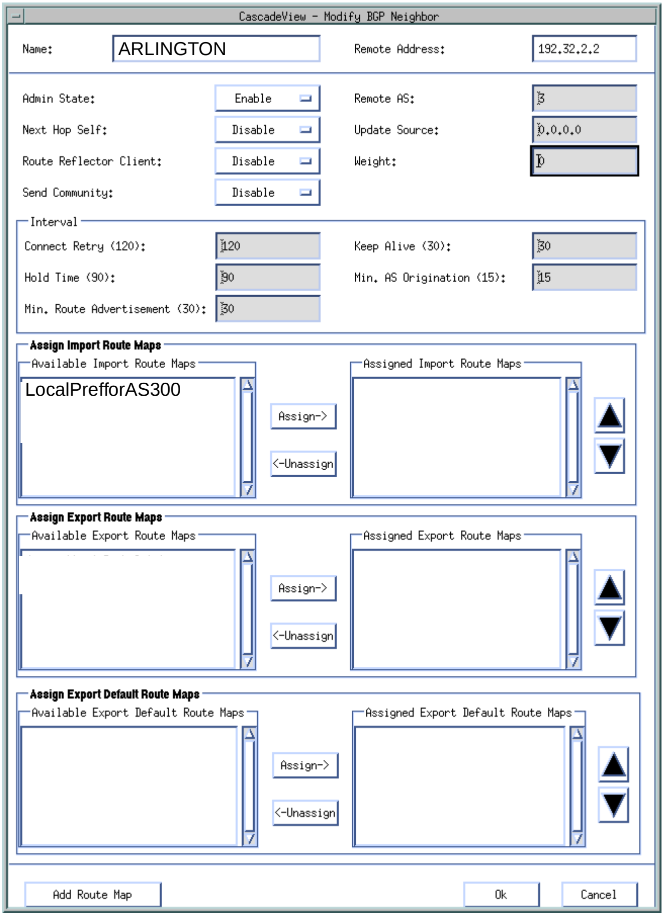

Figure 2-13. Modify BGP Neighbor

4. Select the *LocalPrefforAS300* route map from the Available Route Maps list and choose Assign to move this route map to the Assigned Route Maps list.

5. Choose OK. The system assigns the route map to the IP interface.

Resetting BGP

The BGP protocol only sends route updates when the connectivity state changes.To send a route update you must change the connectivity state to “disable” and then back to “enable” to reset BGP.

You can reset the BGP protocol by using the following console commands from the command line interface:

set bgp state disabled

set bgp state enabled

> *Using the set bgp state enabled/disabled command from the console causes the CP to go out of sync.*

Or you can use the CascadeView/UX NMS as the following section describes.

From the NMS

To reset BGP from the NMS:

1. Select Cascade IP Parameters ⇒Set All BGP_Set All BGP Parameters from the Administer menu. The Set BGP dialog box appears (Figure 2-7).
2. Change the Admin State field to Disable.
3. Choose OK.
4. Redisplay the Set BGP dialog box and change the Admin State field to Enable.
5. Choose OK.

Reviewing the Routing Tables for Modifications

The Nixon and Wright routing table examples shown on this and the following page illustrate the current IP and BGP routing tables for Nixon and Wright. Note that the IP routing table did not change in Nixon. However, the local preference value for the route to network 10.1.2.0.24 through AS 300 has increased from 100 to 500. This indicates that the route map is using the modified local preference value.

Note that Wright's IP routing table has changed so that it now sends traffic destined for network 10.1.2.30/24 through Nixon (peer 150.150.150.1). The BGP routing table now lists the local preference value for route 150.150.150.1 as 500.

New IP and BGP Routing Tables for Nixon

`NIXON## show ip route`

Dest	Next_hop	State	Cost	Lport	Age	Ckt
10.1.1.0/24	*150.150.150.2	BGP	0	2	*	6/36
10.1.2.0/24	10.3.3.2	BGP	1	3	*	
10.3.3.0/24	0.0.0.0	Direct	1	3	*	
10.3.3.1/32	*	Local	1	1cl	*	
150.150.150.1/32	*	Local	0	4072	*	
150.150.150.2/32	150.150.150.2	OSPF	100	2	*	6/36
152.148.230.0/24	0.0.0.0	Direct	1	eth	*	
224.0.0.0/4	*	Local	1	*	*	
224.0.0.0/24	*	Local	1	*	*	
255.255.255.255/32	*	Local	1	*	*	

`NIXON## show bgp route in`

10.1.2.0/24
  nbr id 10.3.3.2 origin 0 as path 300 nexthop 10.3.3.2 med 0 local pref 500
  nbr id 150.150.150.2 origin 0 as path 200 nexthop 10.1.1.2 local pref 100
10.1.1.0/24
  nbr id 150.150.150.2 origin 0 as path nexthop 150.150.150.2 local pref 100

New IP and BGP Routing Tables for Wright-1

show ip route

Dest	Next_hop	State	Cost	Lport	Age	Ckt
10.1.1.0/24	0.0.0.0	Direct	10	1	*	
10.1.1.10/32	*	Local	1	lcl	*	
*150.150.150.2/32	0.0.0.0	Local	1	*	*	
152.148.230.0/24	0.0.0.0	Direct	1	eth	*	
224.0.0.0/4	0.0.0.0	Direct	1	*	*	
224.0.0.0/24	*	Local	1	*	*	
255.255.255.255/32	*	Local	1	*	*	

show bgp route in

10.1.2.0/24

nbr id 150.150.150.1 origin 0 as path 300 nexthop 150.150.150.1 med 0

local pref 500

nbr id 10.1.1.2 origin 0 as path 200 nexthop 10.1.1.2 local pref 100

10.3.3.0/24

nbr id 150.150.150.1 origin 0 as path nexthop 150.150.150.1 local pref 100

3

Configuring IP Logical Ports

This chapter describes how to configure IP services on B-STDX logical ports. For a complete list of the cards and logical port types supported by Ascend, refer to Table 1-1 on page 1-9.

Prior to configuring IP services, verify that the following tasks are complete as described in the *Network Configuration Guide for B-STDX/STDX*.

- ☑ Create a network map
- ☑ Configure the switch parameters
- ☑ Configure the physical port parameters
- ☑ Configure the logical port for Frame Relay or ATM service

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

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)

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

Address Resolution Protocol (ARP) — Is used when an IP address of a given destination is known, but the destination hardware address (DLCI or VPI/VCI) is not.

Inverse Address Resolution Protocol (InARP) — Is used 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 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 IOPs distribution.

Configuring Logical Ports for IP Services

Figure 3-1 illustrates the steps for configuring a logical port for IP.

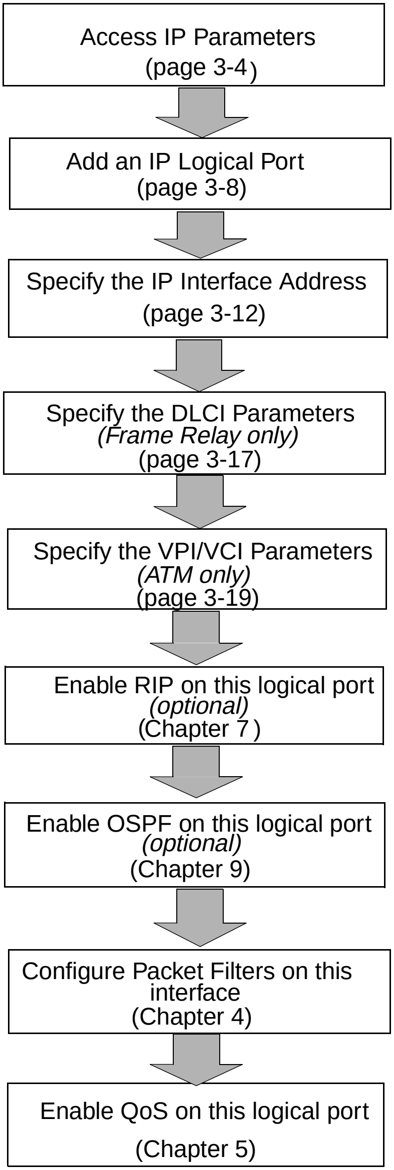

Figure 3-1. IP Logical Port Configuration Process

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 CascadeView/UX Menu. Refer to the following section, "Accessing the Set IP Parameters Dialog Box from the CascadeView/UX Menu" for details about this method of access.
- From the Set All Logical Ports in PPort dialog box. Refer to "Accessing the Set IP Parameters Dialog Box from the Set All Logical Ports Dialog Box" on page 3-6.

Accessing the Set IP Parameters Dialog Box from the CascadeView/UX Menu

To access the Set IP Parameters dialog box from the CascadeView/UX menu:

1. Select the appropriate switch icon from the network map.
2. Select Cascade IP Parameters ⇒Set All IP Lport from the Administer menu. The Set All IP LPorts dialog box appears (Figure 3-2).

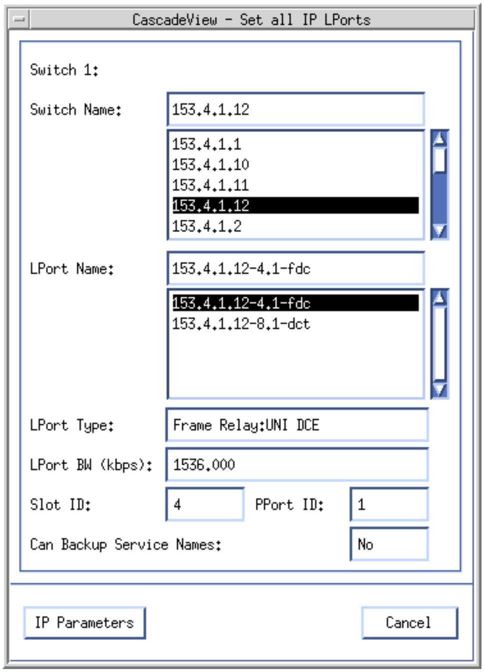

Figure 3-2. Set All IP LPorts Dialog Box

3. Select the LPort name from the list of LPorts.

4. Choose IP Parameters. The Set IP Parameters dialog box appears (Figure 3-3).

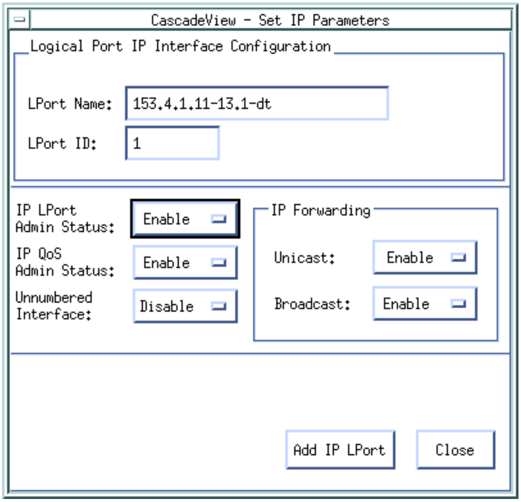

Figure 3-3. Set IP Parameters Dialog Box

5. See “Adding an IP Logical Port” on page 3-8 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 CascadeParameters ⇒Set Parameters. The Switch Back Panel appears.
3. Select the physical port and choose Set Attr... The Set Physical Port Attributes dialog box appears.
4. Choose Logical Port. The The Set All Logical Ports in PPort dialog box appears.

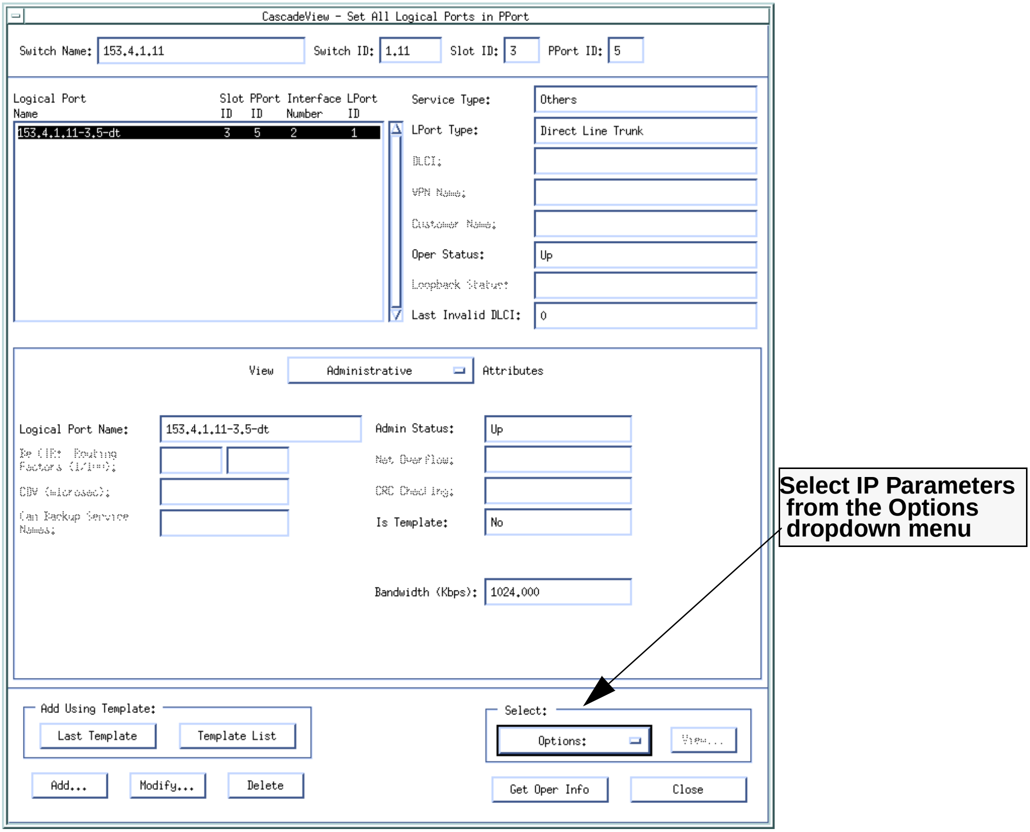

Figure 3-4. Set All Logical Ports in PPort

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

Adding an IP Logical Port

To add an IP logical port:

1. Choose Add IP LPort from the Set IP Parameters dialog box (Figure 3-3 on page 3-6). The second Set IP Parameters dialog box appears (Figure 3-5).

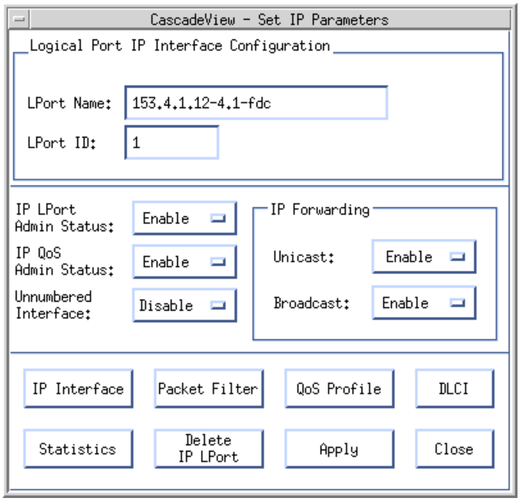

Figure 3-5. Second Set IP Parameters Dialog Box

See Table 3-1 for a description of each of the push buttons on the Set IP Parameters dialog box.

All of the Set IP Parameters command buttons are unavailable if you are configuring an IP logical port that is a trunk.

Table 3-1. Set IP Parameters Buttons

Button	Description
IP Interface	Displays the Set IP Interface Addresses dialog box enabling you to configure the IP interface address. Refer to page 3-12.
Packet Filter	Displays the Assign Logical Port IP Filter dialog box enabling you to specify inbound and outbound packet filters. Refer to Chapter 5, “Configuring Packet Filters” for more details on this function.
QoS Profile	Displays the Associate LPort QoS Profile dialog box enabling you to add and associate Quality of Service profiles. Refer to Chapter 5, “Provisioning IP Quality of Service” for more details on this function.
DLCI	*(For Frame Relay modules only)* Displays the IP Protocol Connection ID dialog box enabling you to specify the Data Link Connection Identifier (DLCI) for the IP logical port. Refer to “Setting the DLCI for Frame Relay Logical Ports” on page 3-17 for more details on this function.
VPI/VCI	*(For ATM modules only)* Displays the IP Protocol Connection ID dialog box enabling you to specify the Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) for the IP logical port. Refer to“Setting the VPI/VCI for ATM Logical Ports” on page 3-19 for more details on this function.
Statistics	Displays the IP Lport Statistics dialog box shown in Figure 12-9 on page 12-13. Refer to “Monitoring IP Logical Port Statistics” on page 12-13 for more information about IP logical port statistics.
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.

Table 3-1. Set IP Parameters Buttons (Continued)

Button	Description
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 3-2.

Table 3-2. 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 a QoS PVC, it is suggested that the Lport Name identify the port as a QoS logical port. When you later use this logical port to associate to the QoS PVC, you will have to select the logical port from a list of Lport Names. Chapter 5, "Provisioning IP Quality of Service" provides details about QoS.
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.
IP QoS Admin Status	Select one of the following options: *Enable* – enables the use of a QoS flow profile for the logical port. *Disable* – disables the use of the QoS flow profile.

Table 3-2. IP Parameter Fields (Continued)

Field	Action/Description
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 address. *Disable* – indicates that this IP logical port is part of a subnet.
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. Refer to “Setting the IP Interface Address” in the following section for details about how to specify a unicast address for an 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. Refer to “Setting the IP Interface Address” in the following section for details about how to specify a unicast address for an IP interface.

3. The next step is to specify the IP interface address for the IP logical port. Refer to 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 (Figure 3-3 on page 3-6) choose IP Interface. The Set Interface Addresses dialog box appears (Figure 3-6).

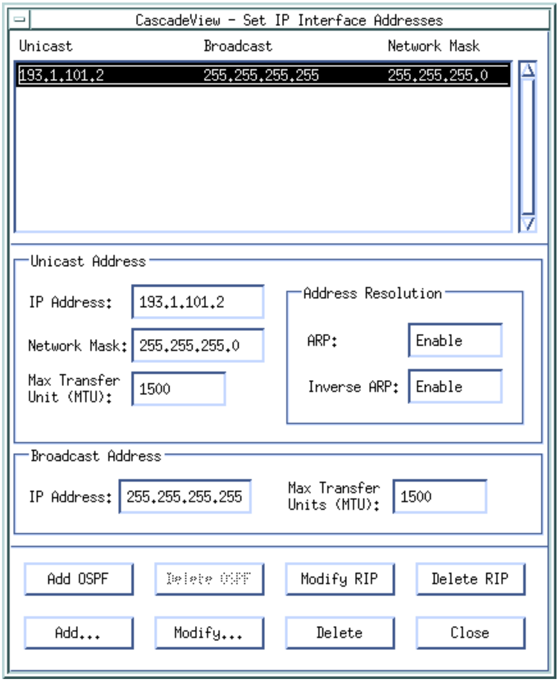

Figure 3-6. Set IP Interface Addresses Dialog Box

> *If you are configuring an IP interface on a trunk logical port, only the Unicast address and Max Transfer Unit (MTU) are configurable. All other fields are displayed as read-only fields and the Broadcast Address information does not display.*
>
> *In addition, all option buttons shown in Table 3-1 are not available when you are configuring a trunk as an IP logical port.*

Table 3-3. Set IP Interface Addresses Buttons

Button	Description
Add OSPF	Displays the Add OSPF Interface dialog box enabling 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 enabling 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 enabling 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 enabling 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 enabling 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 enabling 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 (Figure 3-7).

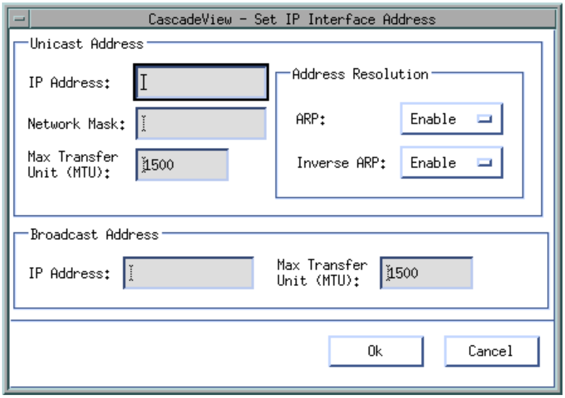

Figure 3-7. Set IP Interface Address Dialog Box

3. Specify the IP Interface Address values described in Table 3-4.

Table 3-4. IP Interface Address Fields

Field	Action/Description
Unicast Address	
IP Address	The IP address for this interface. A maximum of 10 IP addresses can be configured on each IP logical port.
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 is 1500 bytes.
Address Resolution	
ARP	Select one of the following options: *Enable* – Enables the Address Resolution Protocol (ARP). *Disable* – Disables the ARP. Refer to "Address Resolution" on page 3-2 for details.
Inverse ARP	*(Frame Relay Only)* Select one of the following options: *Enable* – enables the Inverse Address Resolution Protocol (InARP). *Disable* – disables the InARP. Refer to "Address Resolution" on page 3-2 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 is 1500 bytes.

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 3-17
- “Setting the VPI/VCI for ATM Logical Ports” on page 3-19

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 0-933.

To specify the DLCI for Frame Relay Logical Ports:

1. From the Set IP Parameters dialog box choose DLCI. The IP Protocol Connection ID dialog box appears (Figure 3-8).

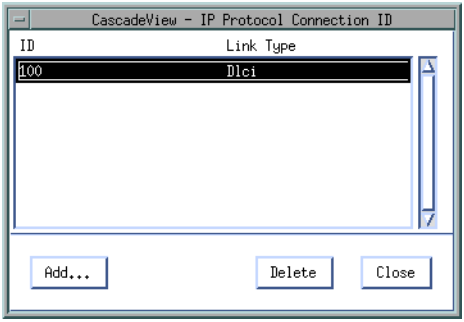

Figure 3-8. IP Protocol Connection ID Dialog Box (For Frame Relay Lports)

2. Choose Add. The Set IP Protocol Connection ID dialog box appears (Figure 3-9).

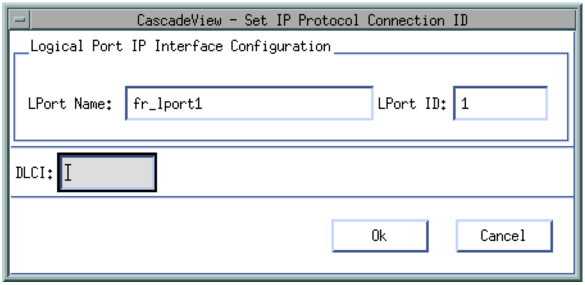

Figure 3-9. Set IP Protocol Connection ID Dialog Box (For Frame Relay LPorts)

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

Table 3-5. IP Protocol Connection ID (For Frame Relay LPorts) Fields

Field	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 0-933: 0-15 – is reserved 16-933 – is available for all link management types

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. The ATM cell header contains both a VCI and a VPI, which gives an ATM cell a unique VCI and associates it with a particular virtual path. Every ATM cell uses these VPI/VCI identifiers.

> *The VPI and VCI are used only for establishing connections between two ATM entities, not the end-to-end connection.*
>
> *On an IP logical port, the number of bits for the VPI value is set to 4. The number of bits for the VCI value is set to 8. This value cannot be changed.*

To specify the VPI and VCI for ATM logical ports:

1. From the Set IP Parameters dialog box choose VPI/VCI. The IP Protocol Connection ID dialog box appears (Figure 3-10).

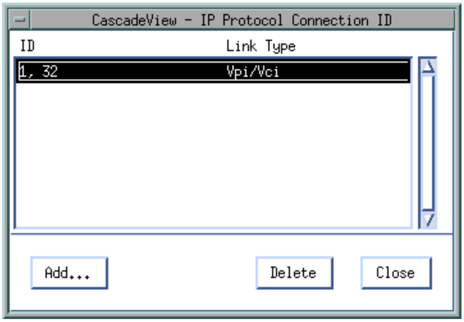

Figure 3-10. IP Protocol Connection ID Dialog Box (ATM LPorts)

2. Choose Add. The Set IP Protocol Connection ID dialog box appears (Figure 3-11).

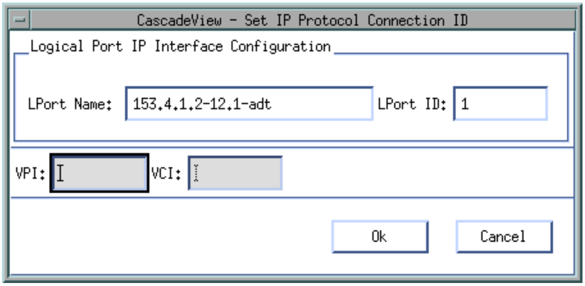

Figure 3-11. Set IP Protocol Connection ID Dialog Box (ATM LPorts)

3. Specify the field values as described in Table 3-6.

Table 3-6. 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 ID number assigned to the selected logical port.
VPI	The virtual path identifier (VPI).
VCI	The virtual channel identifier (VCI).

4

Configuring IP Packet Filters

This chapter describes how to configure and assign IP 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 CascadeView/UX.

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

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 type of service field indicates the packet's priority.

Transport — 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 zero.

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. For outbound filters, the packet goes through a similar process only after the packet is received and routed to an interface.

Packet filtering can trace incoming and outgoing packets when the packets match a filter in the filter list. If a match occurs, the packets are sent to the trace manager. You can enable filtering to count the number of incoming or outgoing packets.

Configuring IP Packet Filters

When you define an IP packet filter, you specify specific 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 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 Cascade IP Parameters ⇒ Set All Packet Filters ⇒ Set All Packet Filters. The Set All Packet Filters dialog box appears (Figure 4-1).

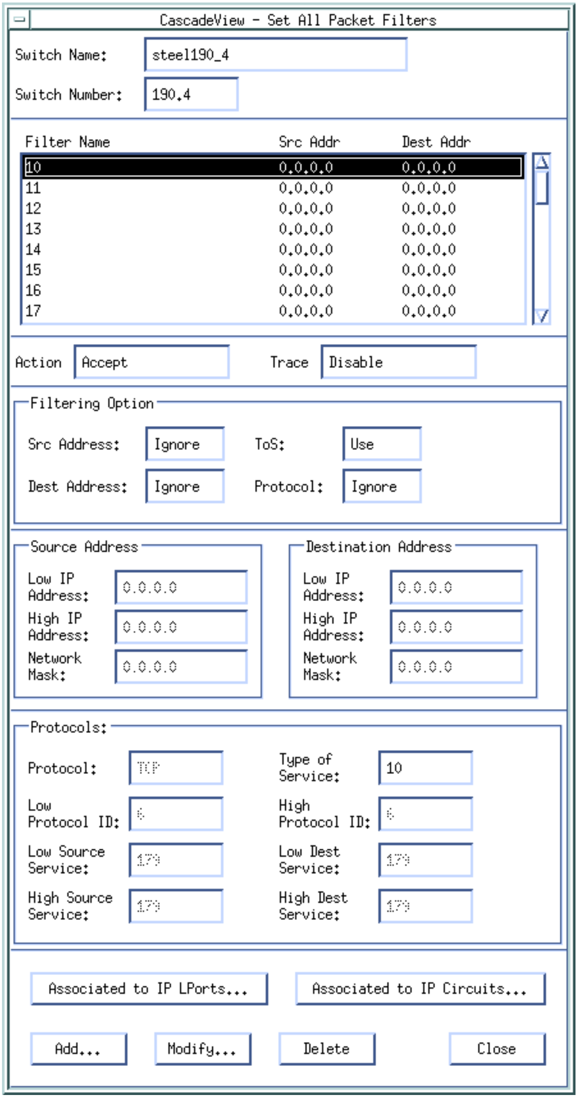

Figure 4-1. Set All Packet Filters Dialog Box

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

Button	Description
Associated to IP LPorts	Displays an IP packet filter's associated logical ports. For more information, see "Viewing an IP Packet Filter's Configuration" on page 4-28.
Associated to IP Circuits	Displays an IP packet filter's associated circuits. For more information, see "Viewing an IP Packet Filter's Configuration" on page 4-28.
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 (Figure 4-2).

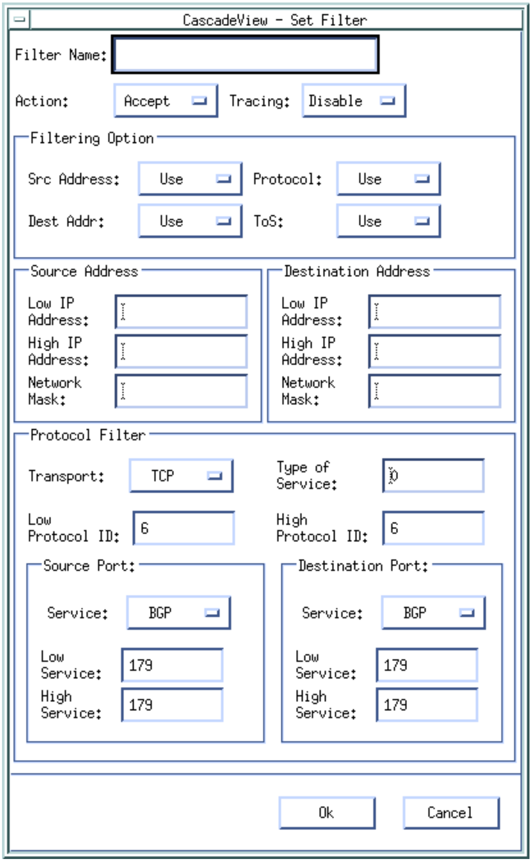

Figure 4-2. Set Filter Dialog Box

4. Complete the fields as described in Table 4-1.

Table 4-1. 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. In addition, the source address fields are grayed out.
Protocol	See "Src Address" field for a description of the Use and Ignore options.
Dest Addr	See "Src Address" field for a description of the Use and Ignore options.
TOS	See "Src Address" field for a description of the Use and Ignore options.

Table 4-1. Set Filter Fields

Field	Action/Description
Source Address	
Low IP address	Enter the low IP address of the node that sends the packet. When you enter this value, you enter a range between this value and the high IP address. If the packet’s source address is between the low and high IP address, 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).* When you enter this value, you enter a range between the low IP address and this value. If the packet’s source address is between the low and high IP address, there is a match.
Network Mask	Enter the Network Mask that applies to the source address.

Table 4-1. Set Filter Fields

Field	Action/Description
Destination Address	
Low IP Address	Enter the low IP address of the node that receives the packet. When you enter this value, you enter a range between this value and the high IP address. If the packet's destination address is between the low and high IP address, there is a match. ***Note:*** *If you want to filter packets going to 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 receives the packet (the default is *high IP address=low IP address).* When you enter this value, you enter a range between the low IP address and this value. If the packet's destination address is between the low and high IP address, there is a match.
Network Mask	Enter the Network Mask that applies to the destination address.

Table 4-1. Set Filter Fields

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.*
Type of Service	Enter a value between 0 and 254. Protocols use the type of service value to specify the packet's priority.

Table 4-1. Set Filter Fields

Field	Action/Description
Low Protocol ID	If you selected Others in the Transport field, enter the low protocol ID. See *RFC 1700* for protocol ID numbers. When you enter this value, you 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. ***Note:*** *To filter packets that have the same protocol ID, specify the protocol ID in the low protocol ID field. You do not have to specify a value in the high protocol ID field.*
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 4-1. Set Filter Fields

Field	Action/Description
Source Port	
Service	If you selected TCP in the transport protocol field, 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* *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.

Table 4-1. Set Filter Fields

Field	Action/Description
Service (Continued)	If you selected UDP in the Transport Protocol field, 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, console, or terminal to indicate that an event occurred. *TFTP* – Trivial File Transfer Protocol *Ignore* – Enables you to filter on all UDP packets. *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. When you enter this value, you enter a range between this value and the high service port number. If the packet's service port number is between the low and high service port numbers, there is a match. ***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 4-1. Set Filter Fields

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 "Service" field description in the source port field section.
Low Service	See "Low Service" field description in the source port field section.
High Service	See "High Service" field description in the source port field section.

5. Choose OK.

6. At the Set IP Filter List dialog box, choose Close.

Packet Filter Configuration Example

The following configuration (filter name: *reject152.148.51.118*) is an example of a filter that restricts packets coming from the source IP address 152.148.51.118.

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Cascade IP Parameters ⇒ Set All Packet Filters ⇒ Set All Packet Filters. The Set IP Filter List dialog box appears (Figure 4-1 on page 4-5).
3. Choose Add. The Set Filter dialog box appears (Figure 4-2 on page 4-7).
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 Source Address field.
   - Select Ignore in the Destination Address, Protocol, and TOS fields.

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

8. In the Low IP Address field for 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.

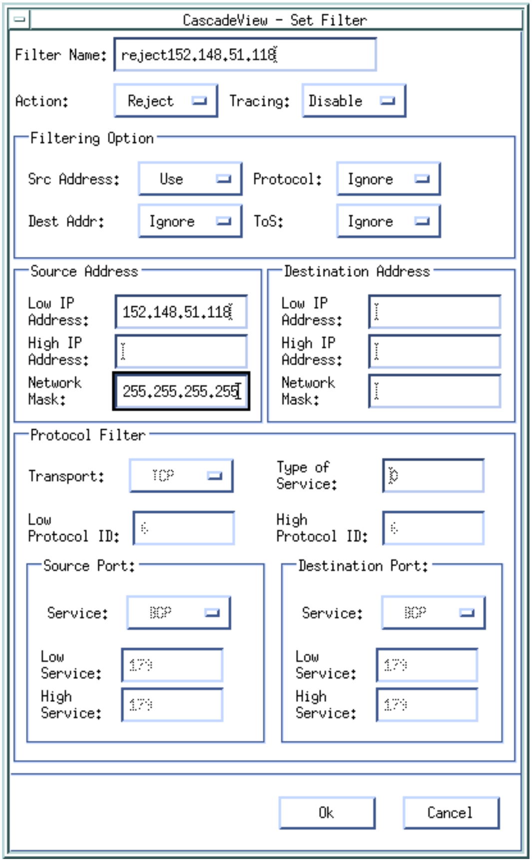

Specify 152.148.51.118 in the Low IP Address field and 255.255.255.255 in the Network Mask field.

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

Figure 4-3. Example 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 Cascade IP Parameters ⇒ Set All Packet Filters ⇒ Set All Logical Port Filters. The Set All Logical Port Filters dialog box appears (Figure 4-4).

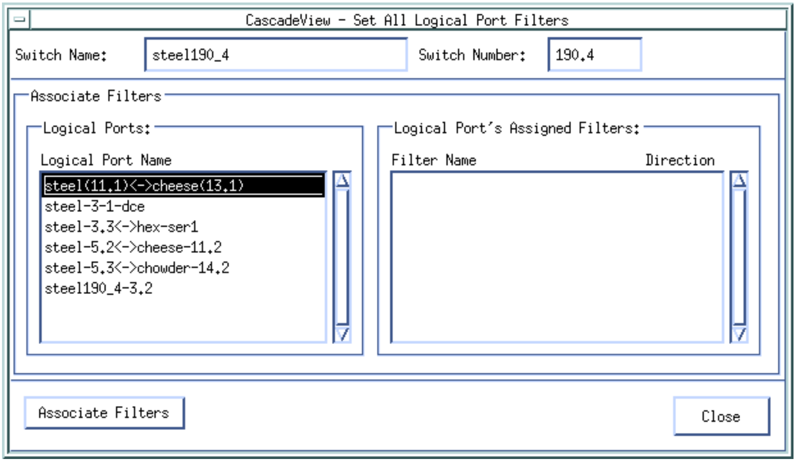

Figure 4-4. Set All Logical Port Filters

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 (Figure 4-5).

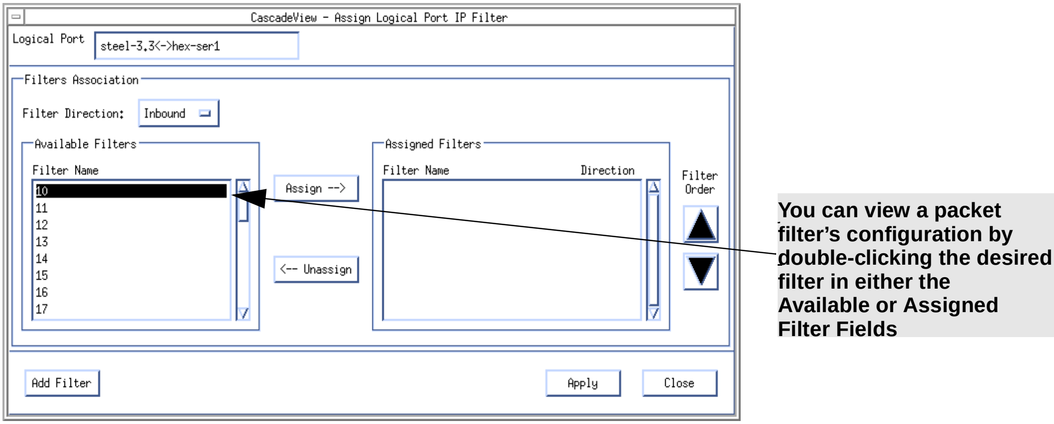

Figure 4-5. Assign Logical Port IP Filter Dialog Box

> *You can view a packet filter's configuration by double-clicking the desired filter in either the Available or Assigned Filters fields. The Set All Packet Filters dialog box (Figure 4-10) appears with the filter's configuration.*

The Assign Logical Port IP Filter dialog box displays the logical port name in the logical port field and provides the following buttons:

Button	Action/Description
Filter Direction	Enables you to indicate the direction (inbound and outbound) in which you want the packets filtered through this logical port.
Available Filters	Lists all available filters.
Assigned Filters	Lists all filters assigned to this logical port.
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.
Filter Direction	Enables you to indicate the direction you want the packets filtered through this logical port.

5. In the Filter Direction field, select either Inbound or Outbound to indicate the direction 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 Cascade IP Parameters ⇒ Set All Packet Filter⇒Set All Host Filters. The Set All Host filters dialog box appears (Figure 4-6).

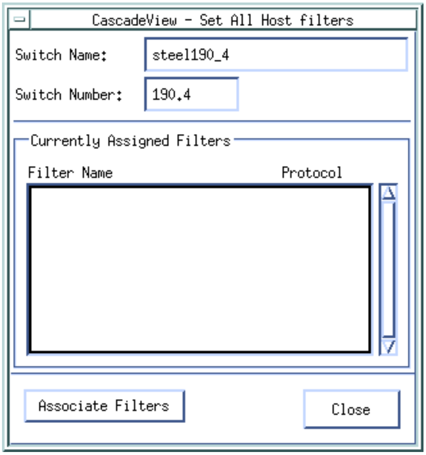

Figure 4-6. Set All Host filters Dialog Box

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

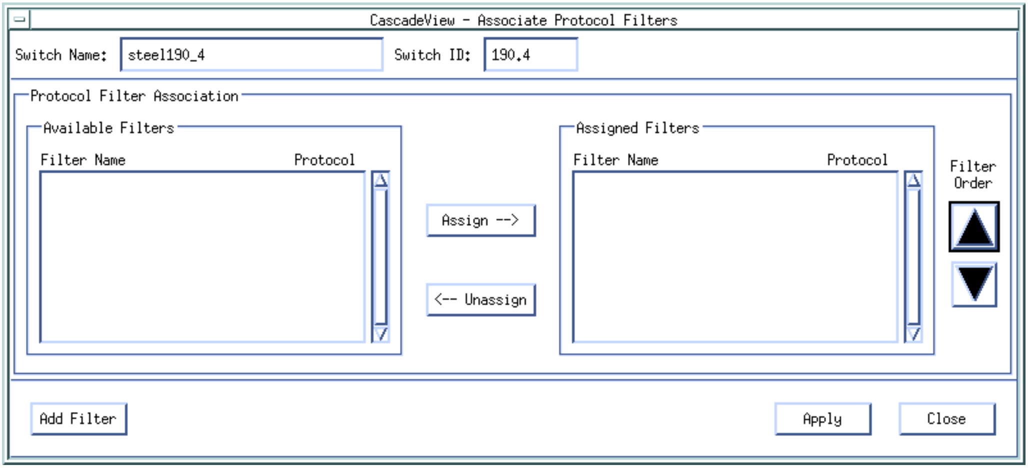

Figure 4-7. Associate Protocol Filters Dialog Box

> *You can view a packet filter's configuration by double-clicking the desired filter in either the Available or Assigned Filters fields. The Set All Packet Filters dialog box (Figure 4-10) appears with the filter's configuration.*

The Associate Protocol Filters dialog box provides the following fields:

Table 4-2. Associate Protocol Filters Fields

Field	Description
Switch Name	Displays the name of the switch.
Switch ID	Displays the switch ID.

The Associate Protocol Filters dialog box provides the following buttons:

Button	Action/Description
Available Filters	Lists all available filters.
Assigned Filters	Lists all filters assigned to this host.
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.

4. From the Available Filters List box, select the filter and choose Assign to assign the IP packet filter to this host.
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”.

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

To assign a filter to the circuit:

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

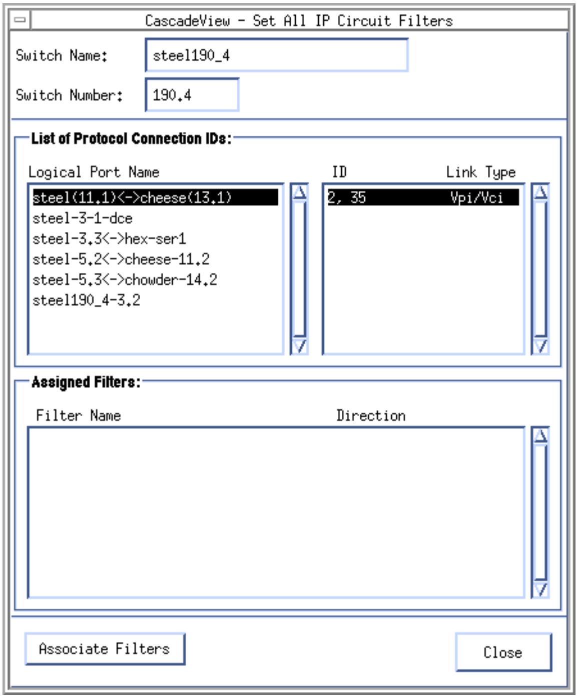

Figure 4-8. Set All IP Circuit Filters Dialog Box

3. Choose Associate Filters. The Associate IP Circuit Filter List dialog box appears.

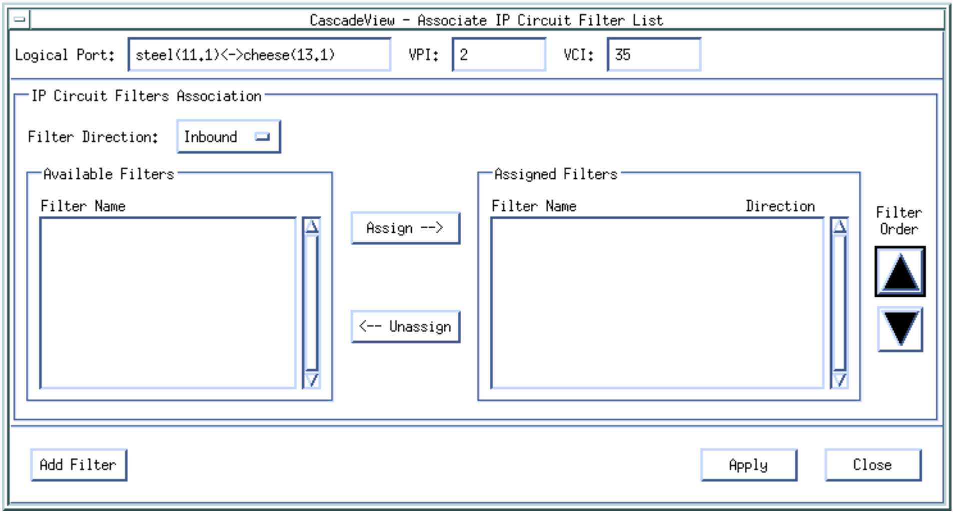

Figure 4-9. Associate IP Circuit Filter List

> *You can view a packet filter's configuration by double-clicking the desired filter in either the Available or Assigned Filters fields. The Set All Packet Filters dialog box (Figure 4-10) appears with the filter's configuration.*

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

Table 4-3. 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.

The Associate Protocol Filter dialog box provides the following buttons:

Button	Action/Description
Filter Direction	Enables you to indicate the direction (inbound and outbound) in which you want the packets filtered through this circuit.
Available Filters	Lists all available filters.
Assigned Filters	Lists all filters assigned to this circuit.
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.

4. From the Available Filters List box, select the filter and choose Assign to assign the IP packet filter to this host.
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. The following is an example of a packet filter assigned to a logical port:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Cascade IP Parameters ⇒ Set All Packet Filters ⇒ Set All Packet Filters. The Set All Packet Filters dialog box appears (Figure 4-10).

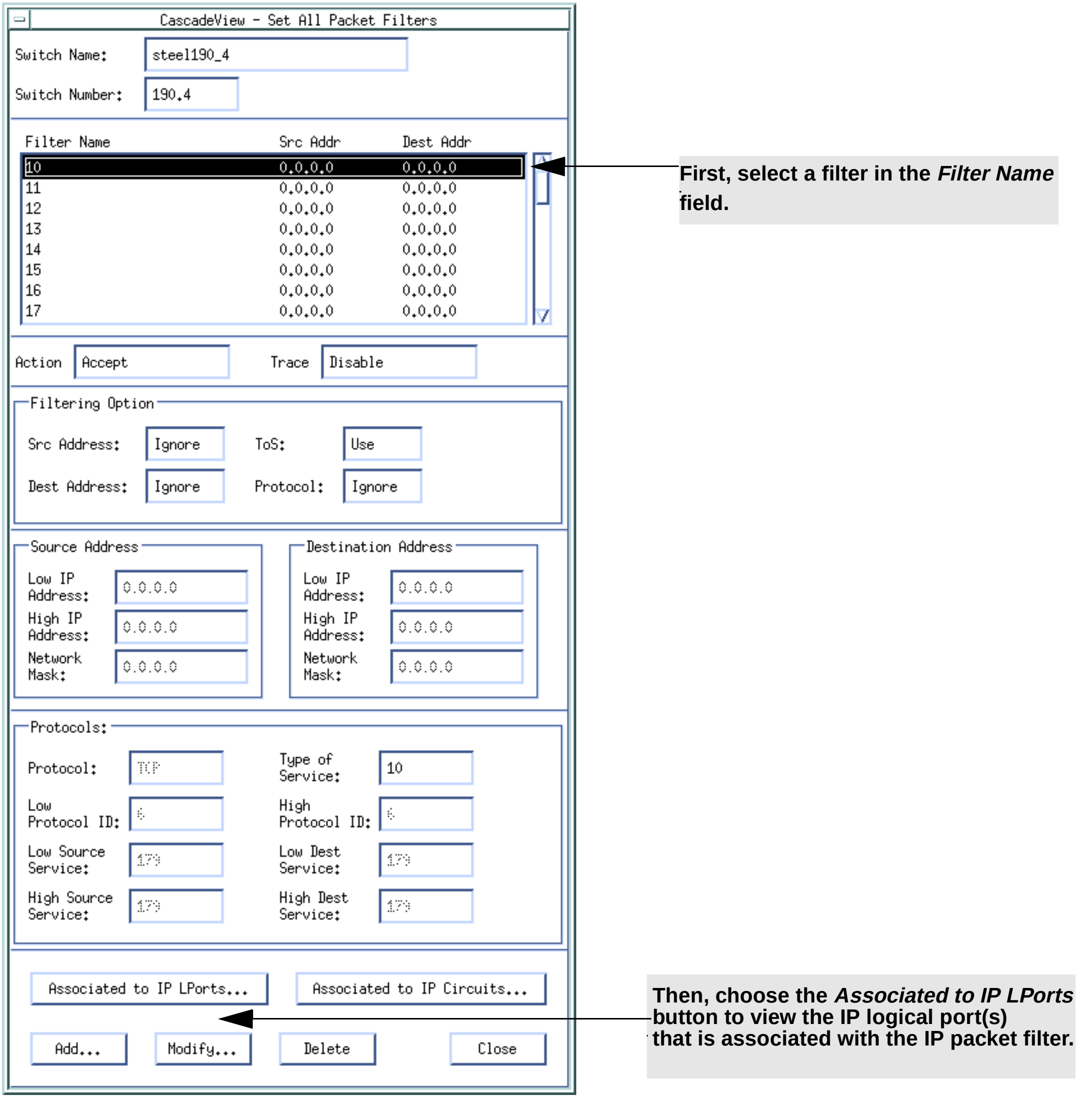

Figure 4-10. Set All Packet Filters Dialog Box

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 (Figure 4-11).

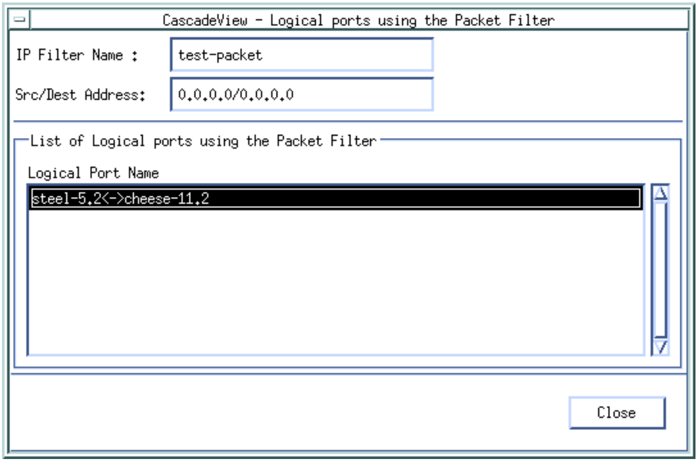

Figure 4-11. Logical ports using the Packet Filter Dialog Box

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

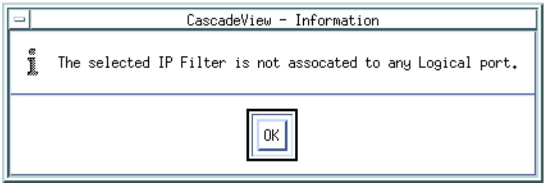

Figure 4-12. Packet Filter Error Message Dialog Box

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 LPort** button.

5

Provisioning IP Quality of Service

About IP Flow Profiles

An IP Flow Profile directs traffic from a source address to a destination address over an IP QoS PVC. When a logical port receives an IP datagram, it compares it with a predefined IP Flow Profile. If there is a match, the logical port forwards the datagram over the specified IP QoS PVC. If the datagram does not match an IP Flow Profile, the logical port forwards it over a MultiPoint-to-Point Tunnel (MPT), which is a best effort PVC. You associate IP Flow Profiles with existing PVCs which you set up with Quality of Service attributes (QoS class, traffic descriptors such as PCR, SCR).

Some important facts about IP QoS and flow profiles include the following:

- One PVC may be shared by multiple flow profiles.
- One flow profile may be shared by multiple logical ports.
- You can assign 100 flow profiles to one switch.
- You can assign 50 flow profiles to one logical port.

- You can modify an IP address in a flow profile with a 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. + anything causes a match.
- You must verify there is enough bandwidth on the PVC for the IP profiles that share the PVC.
- You configure IP filtering before IP QoS.
- You can order the importance of the IP Flow Profiles. The first match wins.
- When you reboot the switch, the PVC takes time to become active.
- The CP maintains the QoS information, but then distributes it to the appropriate IOPs so the CP is not required for proper QoS operation and eliminates a single point of failure.

About IP QoS PVCs

The IP QoS PVC is like an ordinary PVC, except it is a switch-to-switch PVC, not logical port-to-logical port. You can configure an infinite number of PVCs between switches. The IP QoS PVC is defined by its logical port ID and source DLCI. To create the IP QoS PVC, you use a point-to-point connection. The endpoints of the QoS PVC are special logical ports that you configure when you add the first node to the map. This release supports IP QoS Frame Relay circuits only.

Provisioning the QoS Database

This section describes how to:

- Configure an IP QoS PVC
- Define an IP flow profile and attach it to the IP QoS PVC
- Assign an IP flow profile to a logical port

Configuring an IP QoS PVC

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

To configure an IP QoS PVC:

1. From the Administer menu, select Cascade Parameters⇒Set All Circuits⇒Point-to-Point. The Set All PVCs On Map dialog box appears (Figure 5-1).

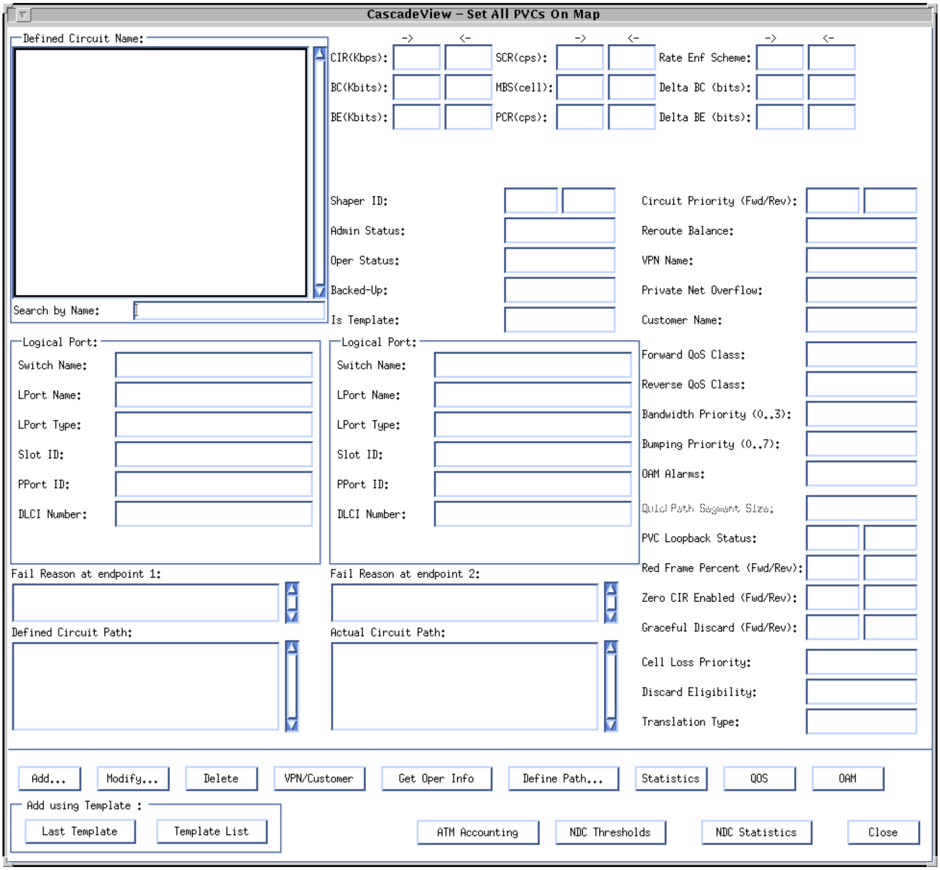

Figure 5-1. Set All PVCs On Map Dialog Box

The Set All PVCs On Map dialog box displays status information for the circuit name you select from the *Defined Circuit Name* list.

You can also use the *Search by Name* field to use wild card characters to search for a specific circuit name. To do this:

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

- To match the ? character, type \?.
- To match the \ character, type \\.

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

Table 5-1. Set All PVCs on Map Buttons

Command Button	Description
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. See the* Diagnostic and Troubleshooting Guide for B-STDX/STDX *for more information.*
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 Diagnostic and Troubleshooting Guide for B-STDX/STDX.
Define Path	Enables you to manually define a circuit path.
Statistics	Displays the summary statistics for the selected circuit. For information, see the *Diagnostic and Troubleshooting Guide for B-STDX/STDX*.
QoS	Displays the Quality of Service 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.

Table 5-1. Set All PVCs on Map Buttons (Continued)

Command Button	Description
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.
ATM Accounting	Accesses the ATM accounting functions for a PVC. For more information, see the *Accounting System Administrator's Guide.*
NDC Thresholds	Displays the configured network data collection (NDC) thresholds for the selected circuit. For more information about these thresholds, see the *Diagnostic and Troubleshooting Guide for CBX 500.*
NDC Statistics	Displays the NDC statistics for the selected circuit. For more information about NDC statistics, see the *Diagnostic and Troubleshooting Guide for CBX 500.*
Close	Exits the dialog box and returns you to the network map.

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

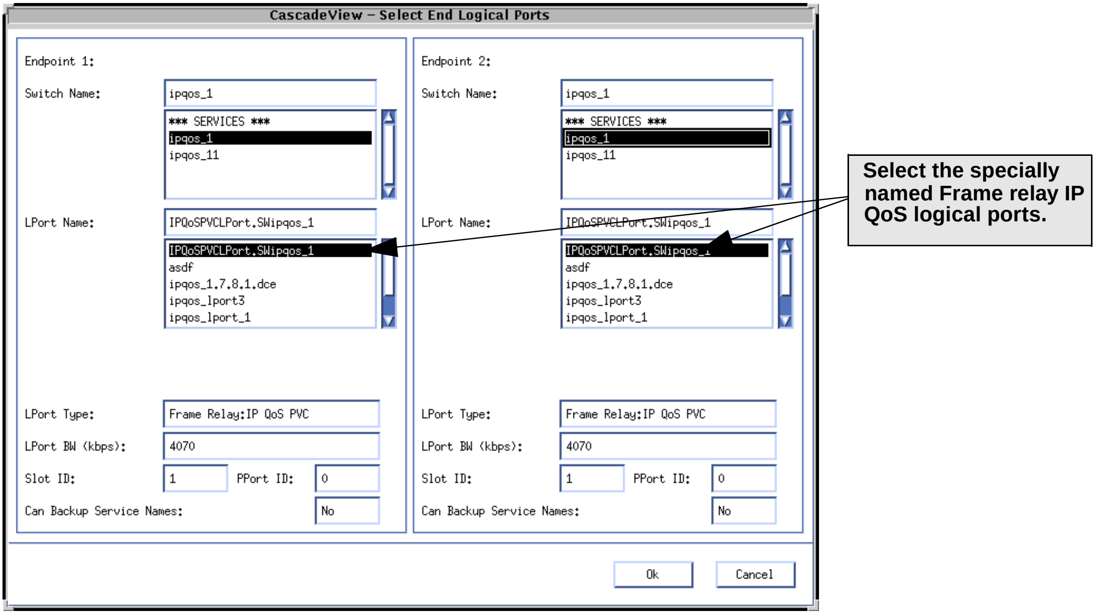

Figure 5-2. Select End Logical Ports Dialog Box

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 Fields

Field	Description
LPort Type	Displays the logical port type for each port in the circuit configuration.
LPort Bandwidth	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.

3. Configure Endpoint 1 and Endpoint 2 as follows:

   a. Select a switch name from the list.

   b. Select the Frame Relay IP QoS logical port (e.g. *IPQoSPVCLPort.SWipqos_1* in Figure 5-2).

4. Choose OK. The Add PVC dialog box appears displaying the current parameters.

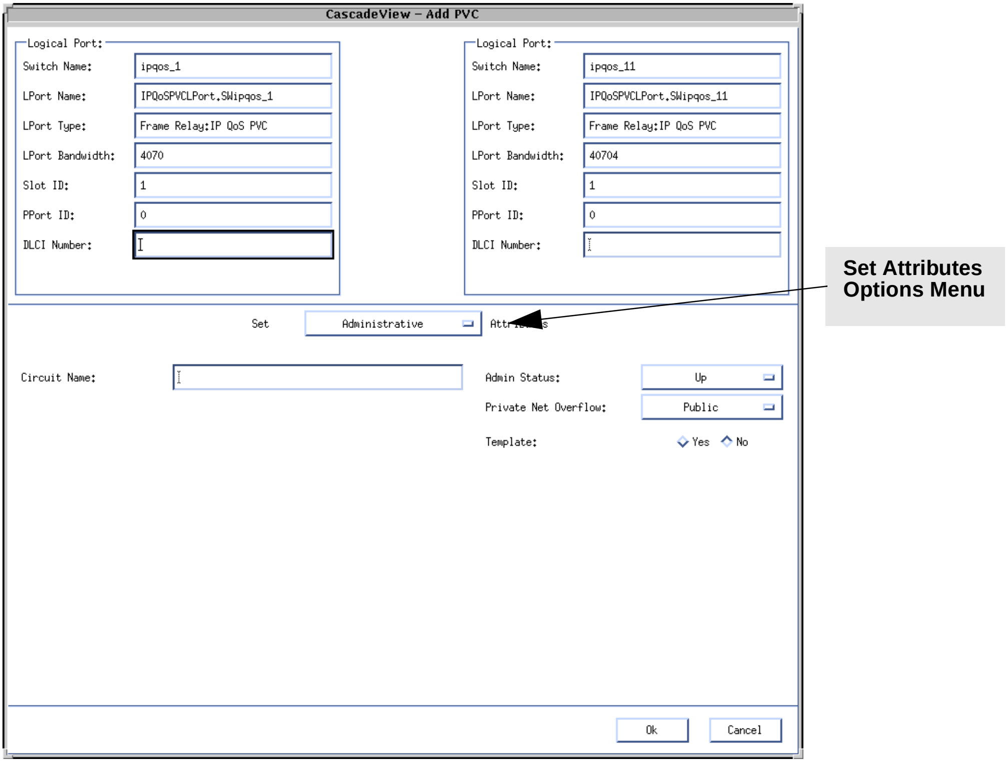

Figure 5-3. Add PVC-Set Administrative Attributes Dialog Box (Frame Relay:IP QoS PVC)

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

Administrative Attributes

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

Table 5-3. Set Administrative Attributes Fields

Field	Action/Description
DLCI Number	Enter a unique DLCI for this logical port. For more information, see Chapter 3, "Setting the DLCI for Frame Relay Logical Ports" on page 3-17.
Circuit Name	Enter any unique, continuous, alphanumeric name for the QoS circuit. Do not use parentheses and asterisks.
Admin Status	Select Up or Down to activate or deactivate. *Up (default)* – Activates the circuit. *Down* – Takes the circuit off-line to run diagnostics such as PVC loopback.
Private Net Overflow	(*For 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 VPN or use public (shared) trunks during overflow conditions. For more information, see the *Network Configuration Guide for B-STDX/STDX*. 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.

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

Field	Action/Description
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.

Traffic Type Attributes

1. Select Traffic Type attributes from the Set Attributes option menu.

   The traffic type fields appear (Figure 5-4).

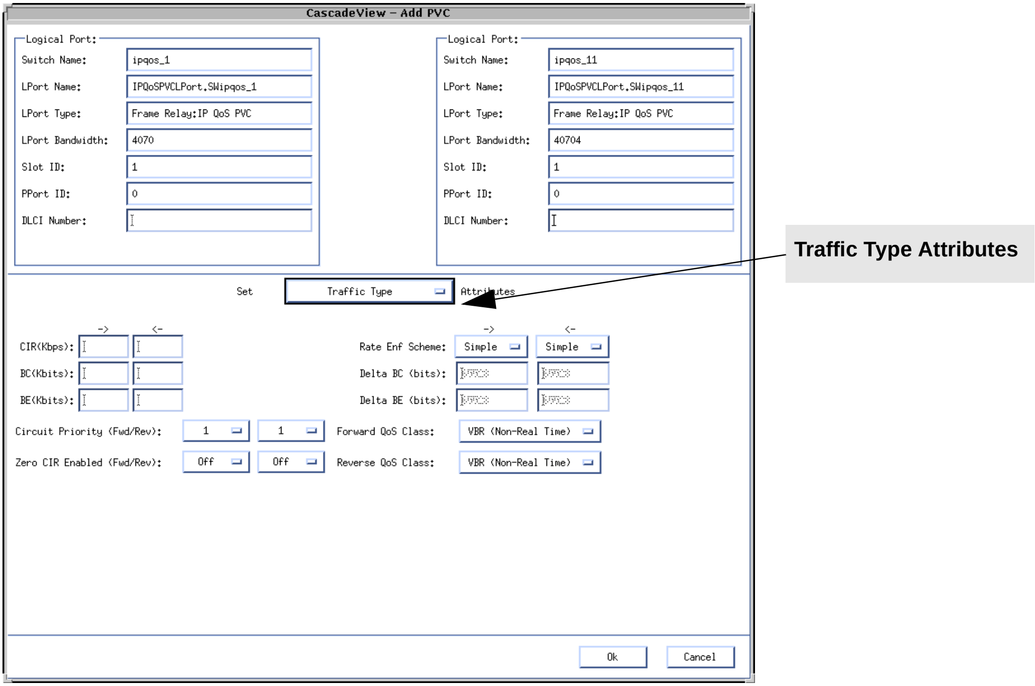

Figure 5-4. Add PVC - Set Traffic Type Attributes Dialog Box (Frame Relay:IP QoS PVC)

2. Complete the required fields described in Table 5-4.

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

Table 5-4. Add PVC - Set Traffic Type Attributes Fields

Field	Action/Description
CIR (Kbps) (Committed Information Rate)	Enter the rate in Kbps at which the network transfers data under normal conditions. Normal conditions see 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 BC/CIR. The network treats this data as “discard eligible” (DE) data.

Table 5-4. Add PVC - Set Traffic Type Attributes Fields (Continued)

Field	Action/Description
Circuit 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). See *Networking Services Technology Overview* for information about congestion control and circuit priority.
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.*
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 (See the* Network Configuration Guide for B-STDX/STDX *for information) is disabled.*

Table 5-4. Add PVC - Set Traffic Type Attributes Fields (Continued)

Field	Action/Description
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.
Forward QoS Class	Ascend currently supports only the *VBR non-real time* class of service value for non-real time applications. This option enables the circuit to transfer large amounts of data over a long period of time using a pre-established ATM connection. Class of service values enable you to prioritize circuit traffic.

Table 5-4. Add PVC - Set Traffic Type Attributes Fields (Continued)

Field	Action/Description
Reverse QoS Class	Ascend currently supports only the *VBR non-real time* class of service value for non-real time applications. This option enables the circuit to transfer large amounts of data over a long period of time using a pre-established ATM connection. Class of service values enable you to prioritize circuit traffic.

User Preference Attributes

1. Select User Preference attributes from the Set Attributes option menu.

   The user preference fields appear (Figure 5-5).

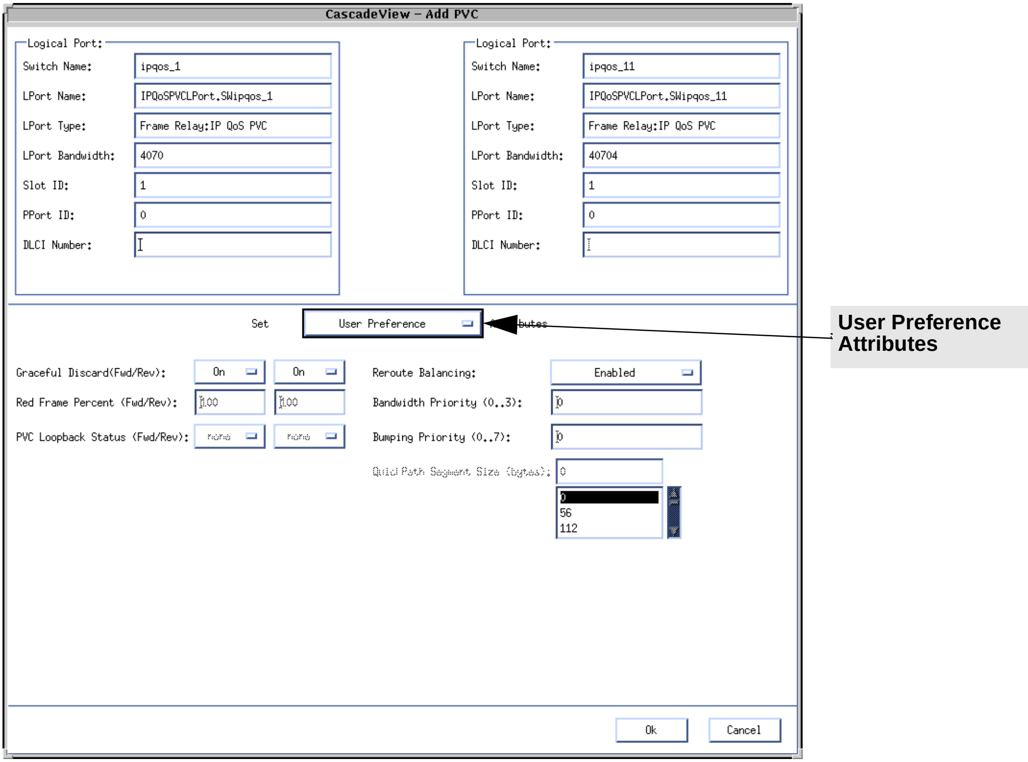

Figure 5-5. Add PVC - Set User Preference Attributes Dialog Box

2. Complete the required fields described in Table 5-5.

Table 5-5. Add PVC - Set User Preference Fields

Field	Action/Description
Graceful Discard (Fwd/Rev)	Select either *On* 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.
Red Frame Percent (Fwd/Rev)	Set this value only if Graceful Discard is set to “On.” For more information on the rate enforcement discard process, see the *Networking Services Technology Overview Guide*. The red frame percent limits the number of red frames the network is responsible to deliver.
PVC Loopback Status (Fwd/Rev)	Displays the current loopback state.

Table 5-5. Add PVC - Set User Preference Fields (Continued)

Field	Action/Description
Reroute Balance	Choose *Enable* to allow this circuit to use reroute tuning. This feature enables the switch to redistribute PVCs across trunks based on OSPF updates and cost metrics. You must first configure the reroute tuning parameters for the selected switch. For more information, see the *Network Configuration Guide for B-STDX/STDX*. If you *Disable* this option, this circuit does not use the reroute tuning parameters.
Bandwidth Priority	Set a value from 0 through 3 where 0 is the default and indicates the highest priority. See the *Network Configuration Guide for B-STDX/STDX* for more information.
Bumping Priority	Set a number from 0 through 7 where 0 is the default and indicates the highest priority. See the *Network Configuration Guide for B-STDX/STDX* for more information.
QuickPath Segment Size (Bytes)	Not supported.

3. 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.

4. (*Optional*) To configure this PVC for a specific VPN and customer, see the *Network Configuration Guide for B-STDX/STDX*.

5. Choose Close to return to the network map.

   Now you can define the IP flow profile for this circuit.

Defining an IP Flow Profile

To define an IP QoS PVC Flow Profile:

1. From the network map select the appropriate switch icon.
2. From the Administer menu, select Cascade IP Parameters ⇒ Set All QoS Profiles ⇒ Set All QoS Profiles. The Set All QoS Profiles dialog box appears (Figure 5-6).

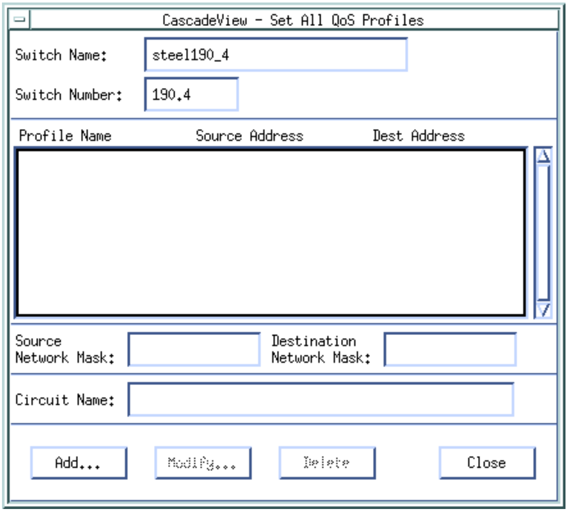

Figure 5-6. Set All QoS Profiles Dialog Box

The Set All QoS Profiles dialog box displays the following buttons.

Button	Description
Add	Enables you to add an IP QoS PVC flow profile
Modify	Enables you to modify an IP QoS PVC flow profile
Delete	Enables you to delete an IP QoS PVC flow profile

3. Choose the Add button. The Add IP QoS PVC Flow Profile dialog box appears (Figure 5-7).

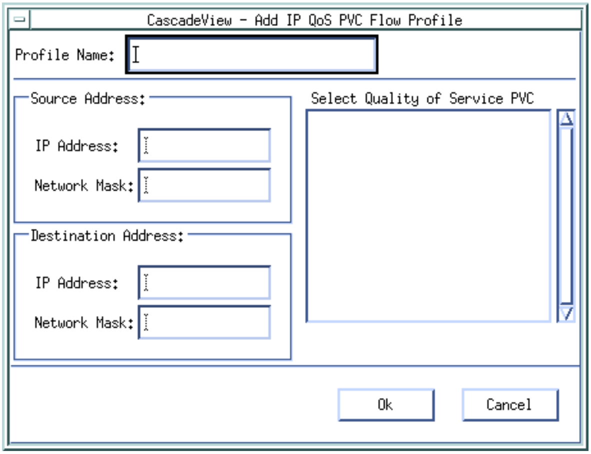

Figure 5-7. Add IP QoS PVC Flow Profile Dialog Box

4. Complete the fields as described in Table 5-6.

Table 5-6. Add IP QoS PVC Flow Profile Fields

Field	Action/Description
Profile Name	Enter a profile name that associates the profile 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.
Select Quality of Service PVC	Select the PVC to which you want to assign the IP QoS PVC flow profile.

You do not have to specify both the destination and source address. However, you must specify at least one.

5. Choose OK.

6. At the Set All QoS Profiles dialog box, choose Close.

Assigning an IP Flow Profile to a Logical Port

> *Before you assign an IP QoS PVC flow profile to the ingress IP logical port, you must enable QoS on the IP logical port. See Chapter 3, "Configuring IP Logical Ports" for information.*

To assign an IP QoS PVC flow profile to a logical port:

1. From the network map select the appropriate switch icon.
2. From the Administer menu, choose Cascade IP Parameters ⇒ Set All QoS Profiles ⇒ Set All Logical Port QoS Profiles. The Set All Logical Port QoS Profiles dialog box appears (Figure 5-8).

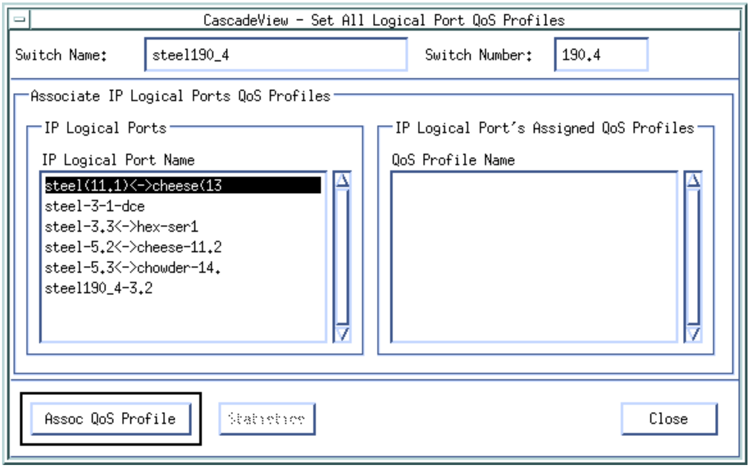

Figure 5-8. Set All Logical Port QoS Profiles Dialog Box

3. Select the IP logical port with which you want to associate an IP QoS PVC flow profile and choose the Assoc QoS Profile button. The Associate LPort QoS Profile dialog box appears (Figure 5-9).

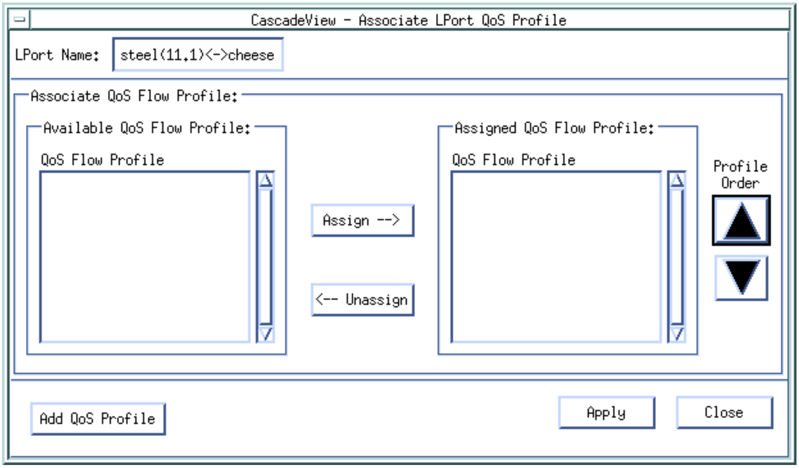

Figure 5-9. Associate Lport QoS Profile Dialog Box

The Associate LPort QoS Profile dialog box displays the logical port name in the LPort Name field and the following buttons:

Button	Action/Description
LPort Name	Displays the name of the logical port.
Available QoS Flow Profile	Lists all current IP QoS PVC flow profiles that are available.
Assigned QoS Flow Profile	Lists all current IP QoS PVC flow profiles assigned to this logical port.
Assign Button	Enables you to assign an IP QoS PVC flow profile to the logical port.
Unassign Button	Enables you to delete an IP QoS PVC flow profile from a logical port.
Profile Order Buttons	Enables you to assign the flow profile’s order of importance. In cases where the IP flow profile matches multiple profiles, the first profile is always used.

Add QoS Profile	Enables you to define additional IP QoS PVC flow profiles. See “Defining an IP Flow Profile” for information.

4. From the Available QoS Flow Profile list box, select the profile to associate with the logical port and choose Assign
5. Choose Apply.

   To create another IP QoS PVC Flow Profile, choose Add QoS Profile and see “Defining an IP Flow Profile” on page 5-18.
6. Choose Close.
7. At the Show IP QoS Filters dialog box, choose Close.

6

Configuring Static ARP/InARP

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 MAC, DLCI, or VPI-VCI addresses. The hardware address you define depends on the link type.

Address Resolution

A node requires the following information in order 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)

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 uses one of the following protocols to resolve the lack of a hardware or IP address.

Address Resolution Protocol (ARP) — When an IP address of a given destination is known but the destination hardware address (DLCI or VPI/VCI) is not.

Inverse Address Resolution Protocol (InARP) — 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 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 IOPs distribution.

Defining a Static ARP Entry

To define a static ARP entry:

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

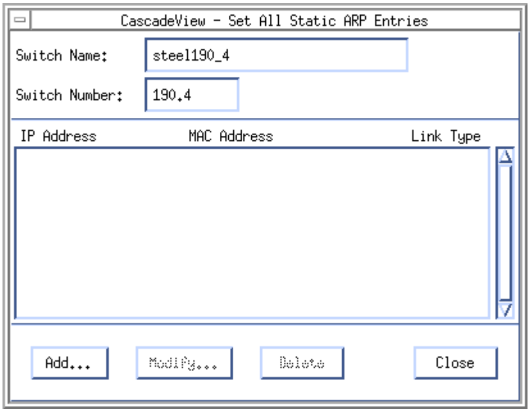

Figure 6-1. Set All Static ARP Entries List Dialog Box

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

Button	Description
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

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

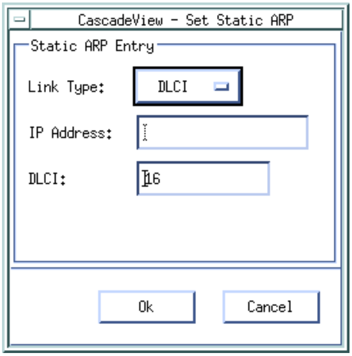

Figure 6-2. Set Static ARP Dialog Box

4. Select the link type and complete the fields described in Table 6-1.

Table 6-1. 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.

Table 6-1. Static ARP Fields (Continued)

Link Type	Field	Action/Description
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.

5. Choose OK. The ARP table entry is created for these addresses.

   You enable ARP/InARP at the Set IP Interface Address dialog box. For more information, see "Setting the IP Interface Address" on page 3-12.

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. RIP is a distance vector protocol, which bases all routing decisions on the path distance.

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, refer to "Adding an IP Logical Port" on page 3-8.
2. Choose Add RIP from the Set IP Interface Addresses dialog box. The Add RIP Interface dialog box appears (Figure 7-1 on page 7-2).

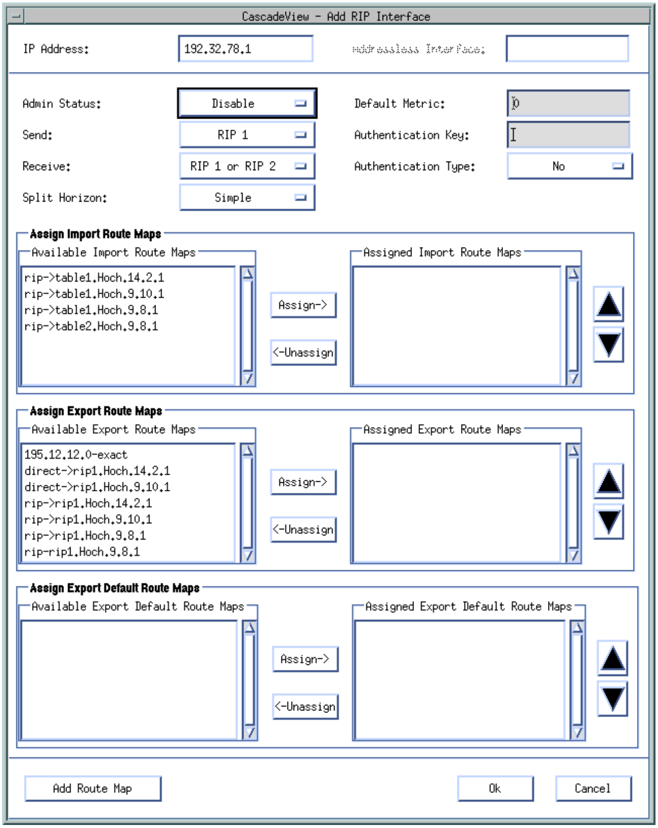

Figure 7-1. Add RIP Interface

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

4. To add a route map for this RIP interface, choose Add Route Map. Refer to Chapter 11, "Configuring Route Maps" for more information on route maps.

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 never 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.*
Receive	Possible values are: *RIP 1*, *RIP2*, *RIP 1 or RIP 2*, or *Disable.*
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. *Poisonous 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.

Table 7-1. RIP Interface Fields

Field	Action/Description
Default Metric	Avariable 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 if you specified a value of *None* 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.
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.
Assign Import Route Maps	

Table 7-1. RIP Interface Fields

Field	Action/Description
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, refer to Chapter 11, "Configuring Route Maps". 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.
Assign Export Route Maps	

Table 7-1. RIP Interface Fields

Field	Action/Description
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, refer to Chapter 11, "Configuring Route Maps". 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 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.
Assign Default Export Route Maps	

Table 7-1. RIP Interface Fields

Field	Action/Description
Available Export Default Route Maps	The export default 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 Maps”. 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 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. For more information about creating route maps, see Chapter 11, “Configuring Route Maps”. To display the parameters for any listed route map, double-click on the map.

8

Configuring BGP Parameters

This chapter describes how to configure Border Gateway Protocol (BGP) parameters on the B-STDX switch, including:

- BGP switch parameters
- BGP neighbors
- BGP aggregates

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 Protocol (IGP), which advertises routes between internal peers within the same autonomous system, such as Open Shortest Path First (OSPF) and Routing Information Protocol (RIP).

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

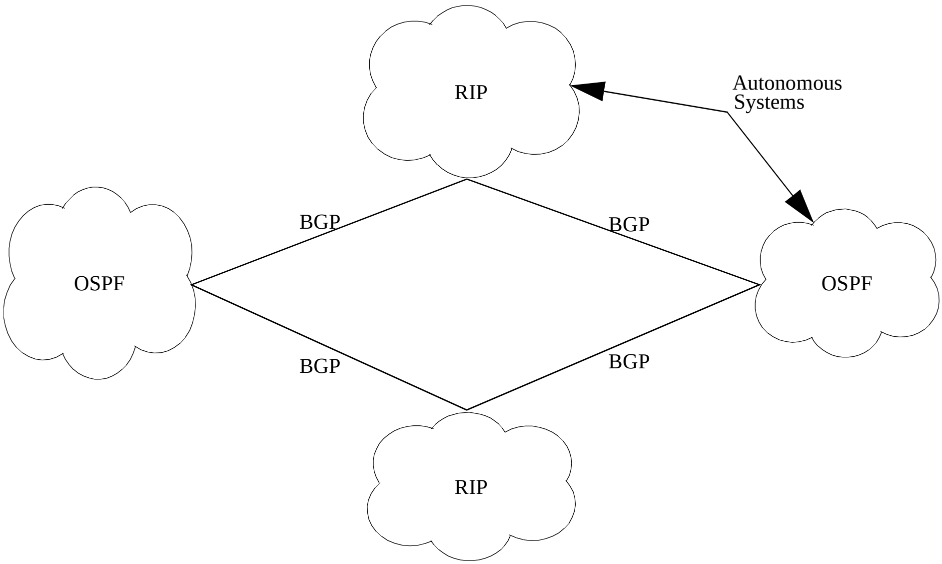

Figure 8-1. Autonomous System Examples

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 updates route information. You use the autonomous system “weight” parameter to determine the path weight and the last route.

BGP peers form a connection between each other, exchanging messages to open and confirm a 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 keeps track of route changes through 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.

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

In a full mesh IBGP, all BGP 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.

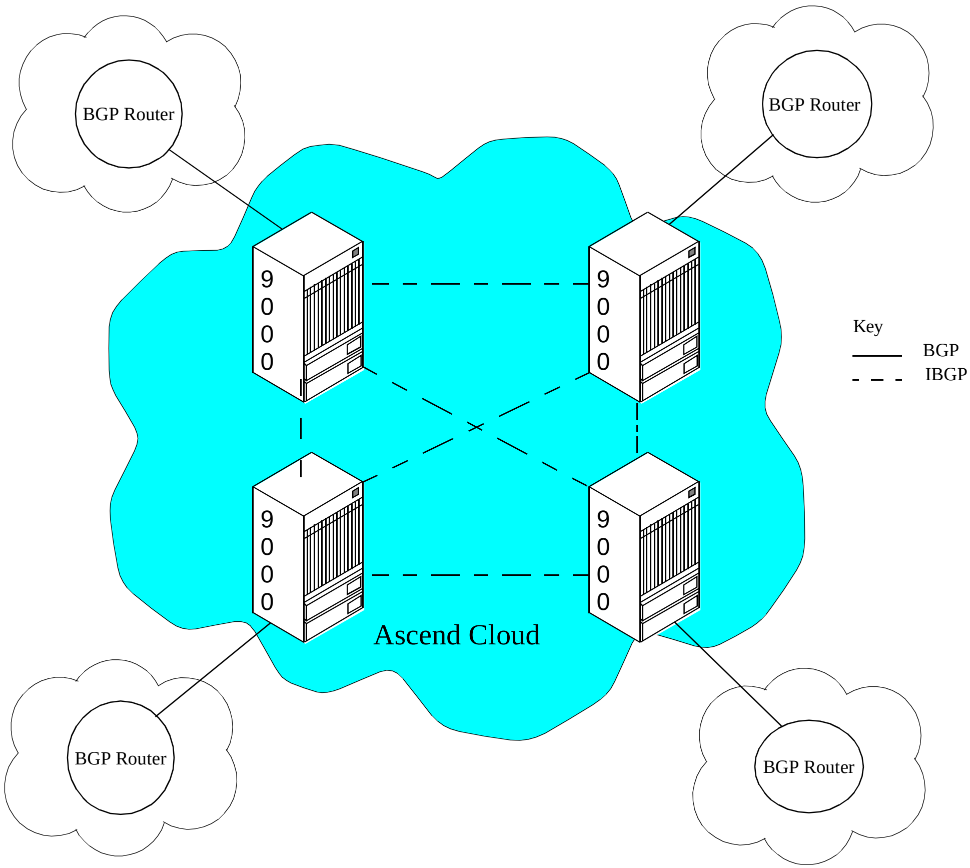

Figure 8-2. Full Mesh Interior Border Gateway Protocol Example

Route reflection 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.

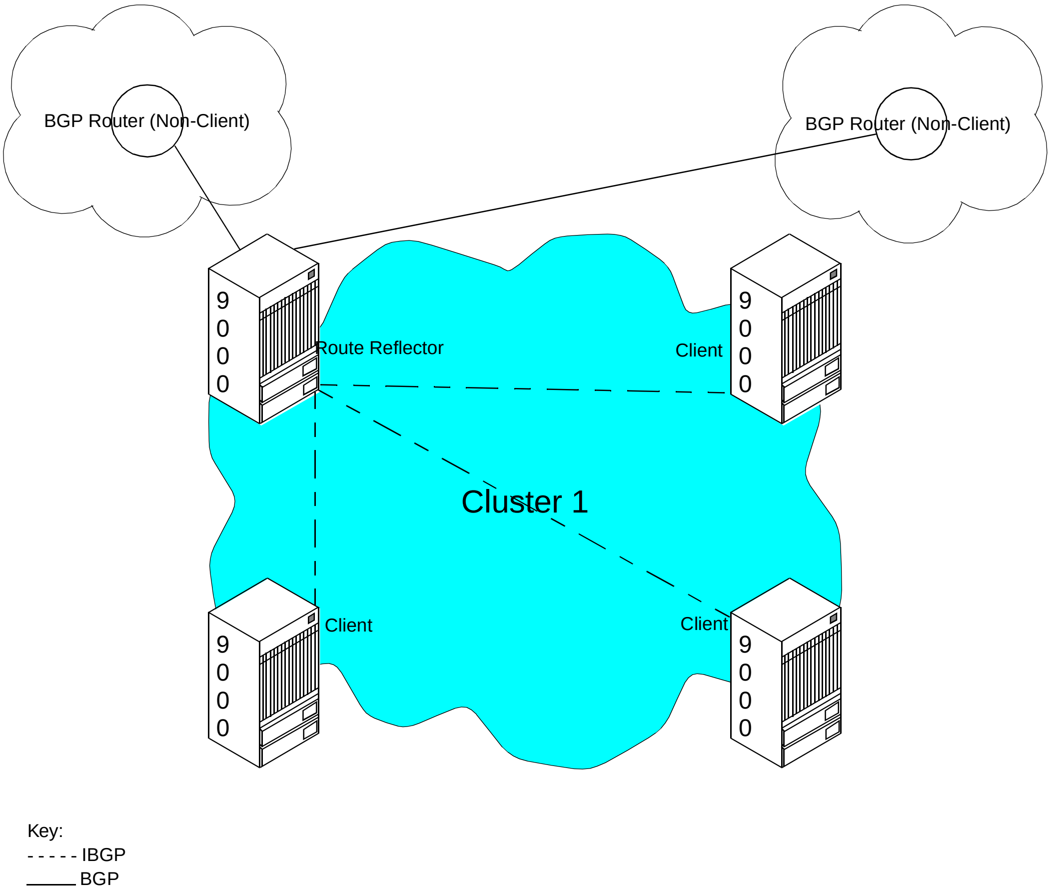

Figure 8-3. Route Reflection Example

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 Aggregates

An aggregate is the combining of specific networks to less specific ones. This reduces 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 are classless (CIDR) and 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 satisfied, BGP forms the aggregate networks and advertises routes to peers.

Configuring BGP

This section describes how to set BGP parameters on the switch level and includes the following tasks:

- Define BGP switch parameters
- Define a BGP neighbor and assign route filters to it
- Define a BGP aggregate

Defining BGP Switch Parameters

To define BGP switch parameters:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Cascade IP Parameters ⇒ Set All BGP ⇒ Set All BGP Parameters. The Set All BGP Parameters dialog box appears (Figure 8-4).

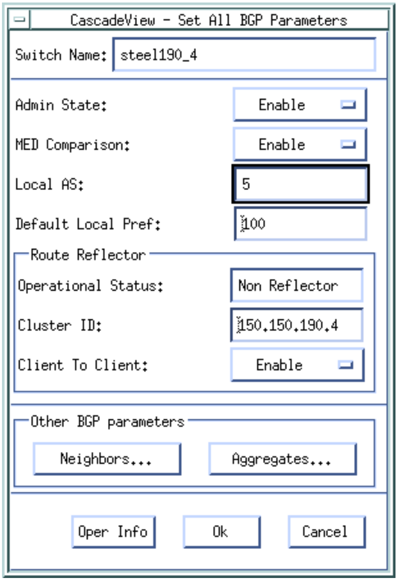

Figure 8-4. Set All BGP Dialog Box

3. Complete the fields described in Table 8-1.

Table 8-1. BGP Switch Parameter Fields

Field	Action/Description
Admin State	Select one of the following options: *Enable* – Allows the selected switch to communicate BGP. *Disable* – Prevents the selected switch to communicate BGP.
Med Comparison	Select one of the following options: *Enable* – Allows you to use 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.
Local AS	Enter a value between 1 and 65535. This parameter is the switch's autonomous system number.
Default Local Pref	Enter a value between 1 and 4294967295 (the default is 100). This value is sent to internal neighbors. 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.

Table 8-1. BGP Switch 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 that contains more than one route reflector. 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 IBGP peers only**)	Select one of the following options: *Enable*– If you enable this parameter, any routes that are received by the selected switch from a client will be sent to all other clients (the default is enable). *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.*
Other BGP parameters	
Neighbors...	Choose this button to define BGP neighbor parameters. For more information, see “Defining a BGP Neighbor and Assigning a Route Filter” on page 8-10.
Aggregates...	Choose this button to define BGP aggregate parameters. For more information, see “Defining a BGP Aggregate” on page 8-20.

4. Choose Oper Info to display a status message about the selected port.

5. Choose OK.

Defining a BGP Neighbor and Assigning a Route Filter

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

To define a BGP neighbor to a switch:

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

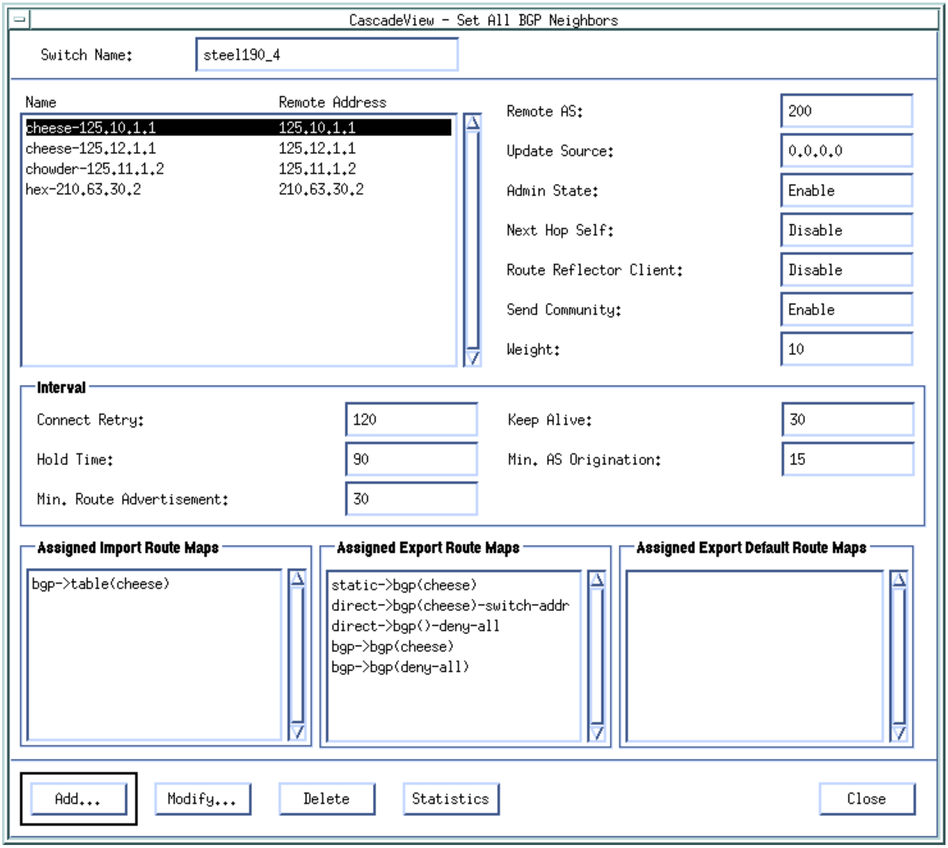

Figure 8-5. Set All BGP Neighbors Dialog Box

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

Button	Description
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 Chapter 12, “Monitoring IP Objects”.

3. Choose Add. The Add BGP Neighbor dialog box appears (Figure 8-6).

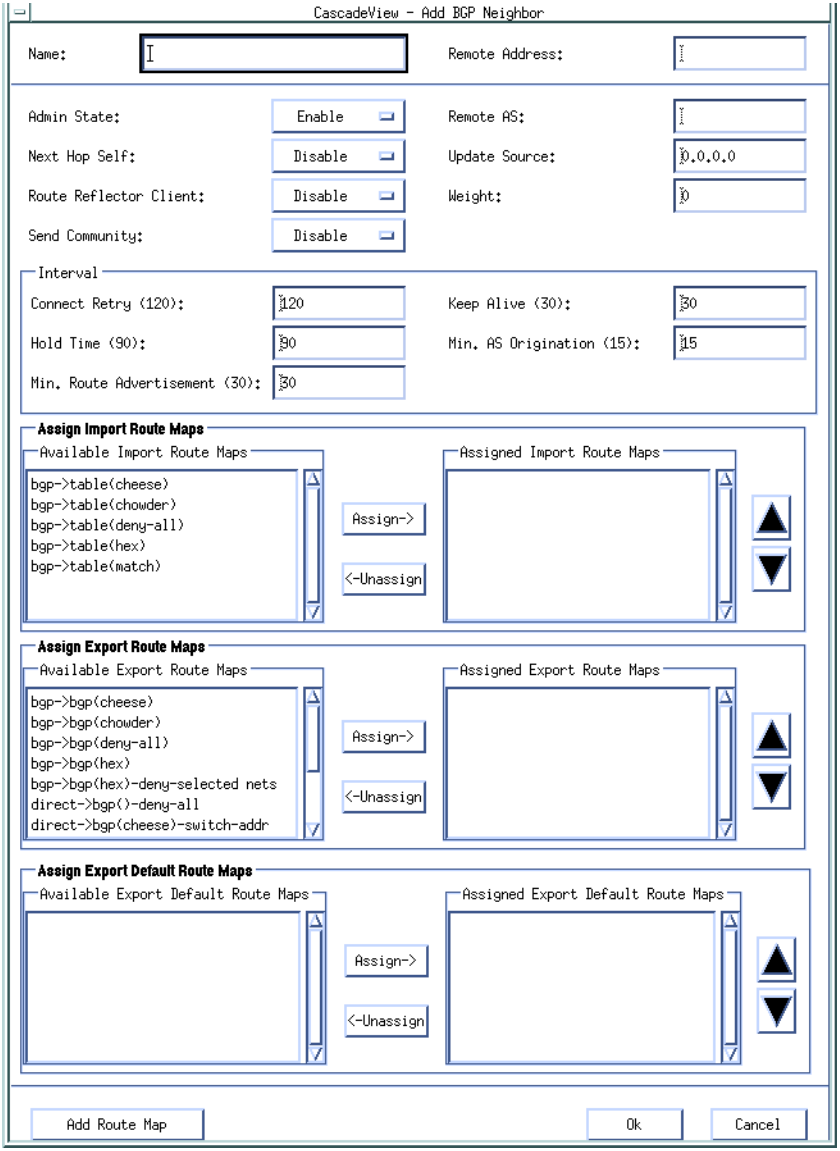

Figure 8-6. Add BGP Neighbor Dialog Box

4. Specify the values as listed in Table 8-2.

Table 8-2. 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* – Activates the connection between the selected switch and this BGP neighbor. *Disable* – Deactivates the connection between the selected switch and this BGP neighbor.
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* – 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.

Table 8-2. BGP Neighbor Fields (Continued)

Field	Action/Description
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* – 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 route.
Send Community	Select one of the following options: *Enable* – Enables you to send community attributes of all updates to this 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* – Disables the sending of community attributes of all updates to this neighbor.

Table 8-2. BGP Neighbor Fields (Continued)

Field	Action/Description
Interval The default value for each interval field in the Add BGP Neighbor dialog box (Figure 8-6) 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.
Min. AS Origination (15)	Enter a value between 1 to 65535 (the default value is 15). The parameter represents the minimum time, in seconds, between route update advertisements belonging to the home AS.

Table 8-2. BGP Neighbor Fields (Continued)

Field	Action/Description
Min. Route Advertisement (30)	Enter a value between 1 and 65535 (the default value is 30). This parameter represents the time, in seconds, that specifies the minimum time between route update advertisements belonging to the external AS.
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 Maps".
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.**

Table 8-2. 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 Maps".
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.**
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 Maps".

Table 8-2. BGP Neighbor Fields (Continued)

Field	Action/Description
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.**
Order Buttons	Enables you to list the route maps in order of importance.

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

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

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

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

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

10. To delete 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.

11. Choose OK.

12. To add an additional network access list, choose Add Route Map. See “Adding Route Maps” on page 11-15 for more information.

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

Defining a BGP Aggregate

To define BGP Aggregates:

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

2. From the Administer menu, choose Cascade IP Parameters ⇒ Set All BGP ⇒ Set All BGP Aggregates. The Set All BGP Aggregates dialog box appears (Figure 8-7).

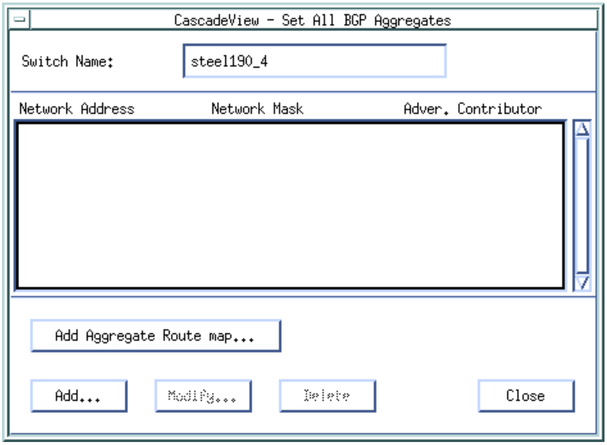

Figure 8-7. Set All BGP Aggregates Dialog Box

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

Button	Description
Add Aggregate Route Map	Enables you to add an Aggregate Route Map. For more information, see Table 11-26 on page 11-53.
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, select Add. The Add BGP Aggregates dialog box appears (Figure 8-8).

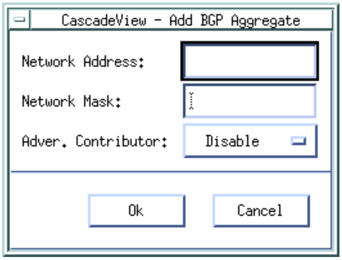

Figure 8-8. Add BGP Aggregate Dialog Box

4. Specify the values as listed in Table 8-3.

Table 8-3. 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. Select OK.
6. At the Set All BGP Aggregates dialog box, choose Close.

9

Configuring OSPF Parameters

This chapter describes how to configure Open Shortest Path First (OSPF) parameters for IP Services.

About OSPF

The Open Shortest Path First protocol is a link-state routing protocol. With link-state routing protocols, routers maintain a link-state database that contains topology current link-state information. This information enables routers to determines the best routes to each destination network in the autonomous system.

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** — Restrict flooding to configured areas, thereby reducing the database size.

The Link-State Database

OSPF “floods” routers with link-state advertisements (LSAs) which contain topology information. Routers store LSAs in the link-state databases. Flooding ensures that all routers have identical databases and the same topology information.

Link-state databases include:

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

Router’s 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.

In addition, designated routers:

- Monitor the health of adjacent routers
- Establish adjacencies

A backup designated router is 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 because many 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.

Despite these benefits, flooding controls reduce the reliability of flooding. Flooding is reliable because many copies of the message travel through the network, ensuring a high probability that one message will arrive safely.

OSPF Areas

As networks grow large, the link-state database grows large as well. This causes problems for the following reasons:

- Increased memory space is consumed
- Generating route tables 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 link-state advertisements (LSAs) to routers that are part of the same area. If an autonomous system 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.

Figure 9-1 describes the concept of areas.

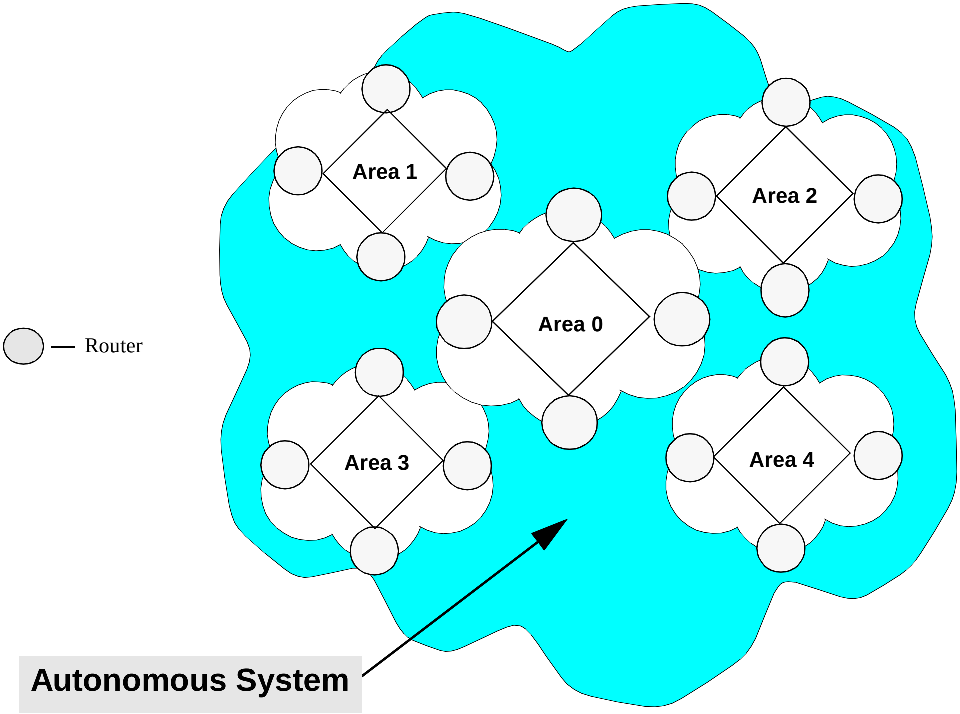

Figure 9-1. OSPF Areas

In Figure 9-1, the autonomous system is divided into 5 areas. Each area represents smaller networks within the autonomous system, and maintains separate link-state databases. Area 0 is the backbone and connects all areas within the autonomous system.

Area Aggregates

Area aggregates consolidate multiple routes (or addresses) within an area (or areas) into one single link-state advertisement. This consolidation enables one advertisement representing a range of 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 1000 interfaces

The Backbone

You do not have to divide the autonomous system into areas so that every router is part of a defined area. Routers that belong to more than one area are referred to as the *backbone*. Backbone routers maintain separate link-state databases of each area to which they belong to.

The backbone is itself an area and is designated as area 0. OSPF requires the backbone to be contiguous to all areas in the autonomous system.

If the backbone breaks, OSPF uses virtual links to patch the backbone. A virtual link is a logical link used by OSPF to connect physically separate portions of the backbone. In addition, virtual links connect areas that are not physically connected to the backbone. Figure 9-2 shows an example of a broken backbone repaired by a virtual link.

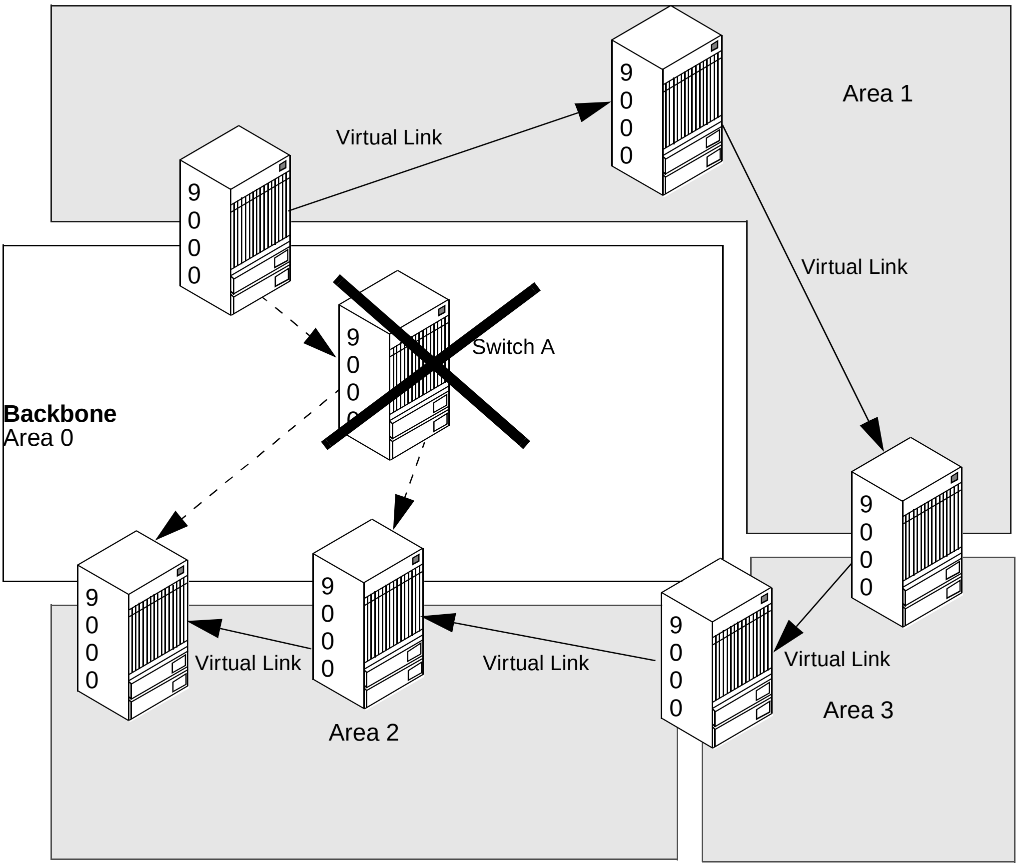

Figure 9-2. OSPF Virtual Linking

In Figure 9-2, if Switch A loses power and goes off-line, you can create a virtual link between area 1 and area 2 to patch the backbone.

About Clustering

Clustering is a way of grouping OSPF areas into subareas. Clustering enables you to use set increments (allows you to use 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 OSPF areas, while performing route aggregation at the *area border switch* or *area border router.* A switch that spans one or more OSPF areas is considered to be an area border switch (ABS). A cluster forms a subset of an OSPF area. A cluster enables additional address aggregation at the ABS and reduces the size of the IP routing table, link-state database, and the number of summary LSAs.

Use clustering only if you plan to do the following:

- Implement OSPF areas using switch software Version 5.0. Note that OSPF areas are supported for IP traffic only.
- Deploy new nodes with the same subnet addresses into multiple OSPF areas (for example, due to a lack of IP addresses)

In CascadeView/UX Version 2.3, you can define an IP address subnet 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 switches are added in that cluster the switch number/host ID in the IP address increments according to the cluster ID. For example, Table 9-1 shows the cluster ID IP-address range using the default IP address, 152.148.50.x.

Table 9-1. Cluster ID and IP Addresses

Cluster ID	IP Address Range
0	152.148.50.1 - 152.148.50.30
1	152.148.50.33 - 152.148.50.62
2	152.148.50.65 - 152.148.50.94
3	152.148.50.97 - 152.148.50.126
4	152.148.50.129 - 152.148.50.158
5	152.148.50.161 - 152.148.50.190

Table 9-1. Cluster ID and IP Addresses (Continued)

Cluster ID	IP Address Range
6	152.148.50.193 - 152.148.50.222
7	152.148.50.225 - 152.148.50.254

Figure 9-3 illustrates the benefits of clustering.

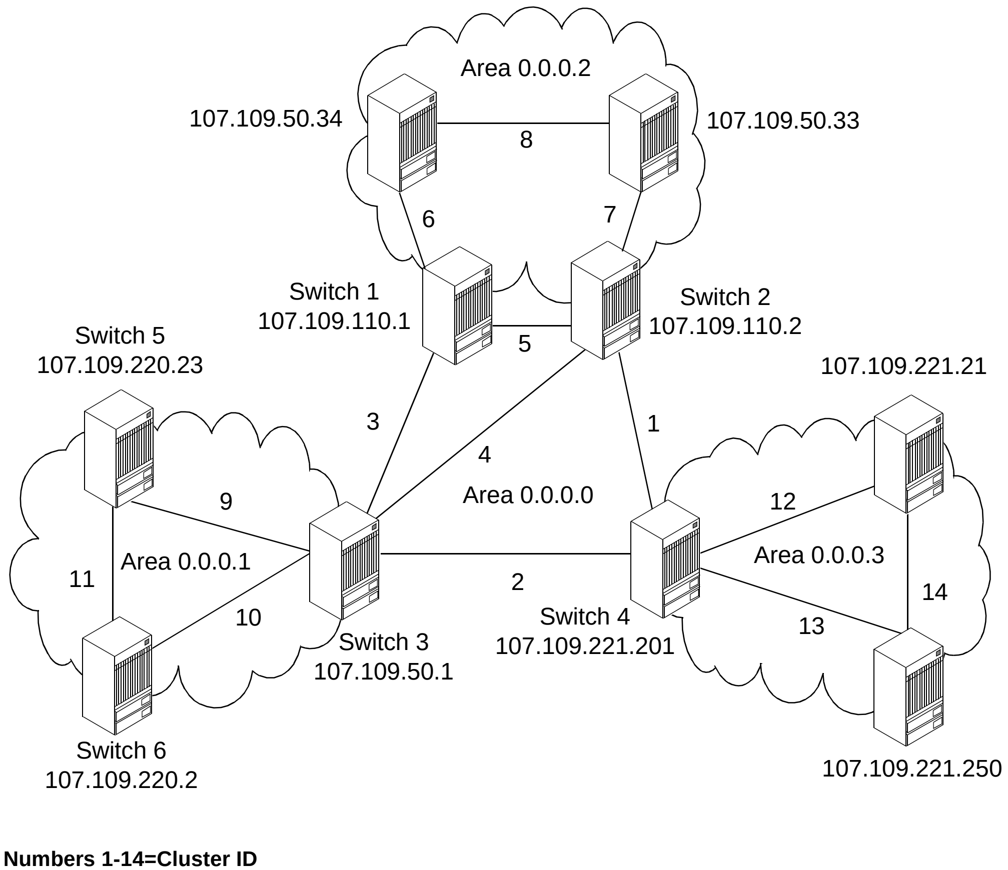

Figure 9-3. OSPF Area Configuration Example

OSPF requires that all OSPF areas be directly attached to the OSPF backbone area. 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 attached to the OSPF 0.0.0.0 backbone area. Packets are forwarded between areas, from the source area, through the backbone, and then into the destination area.

Ascend'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.

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 area border switches. For example, Area 1 is assigned a subnet of 106.105.110.0/*24*. The number 24 specifies a subnet mask of 24, so all IP addresses in the range, 106.105.110.1-254, are sent as a single OSPF summary LSA. In this example, an OSPF summary LSA is sent for each address 152.148.50.33 and 152.148.50.34 without clustering. With clustering, the Ages 106.105.110.1 and 106.105.110.2 are configured for address range 152.148.50.32/27. A single summary LSA is sent for the 152.148.50.32/27 address range.

Address aggregation is not required, but when used, it results in fewer summary-LSAs. Fewer summary LSAs reduce the size of the routing table and OSPF link-state databases. In Figure 9-3, Area 1 routing table would be 106.105.110.0/24, 106.105.50.33, and 106.105.50.34 without OSPF areas. If 152.148.50.0 is designated as cluster 1, the routing tables would have entries for 107.109.11.0/24 and 107-109.50.32/27.

If you anticipate a lack of network IP addresses and the use of a particular subnet address into multiple OSPF areas, you should add new switches to a cluster. Otherwise, you may not need to cluster.

OSPF Routing and Router Classifications

There are two types of routing:

- Intra-area routing
- Inter-area routing

Intra-area routing is routing within an area, and inter-area routing is routing between areas. These types of routing are performed by different classifications of routers, including:

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

Area border routers — Routers with links to more than one area, or between an area and backbone.

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 9-4 shows an example of OSPF routing and router classifications.

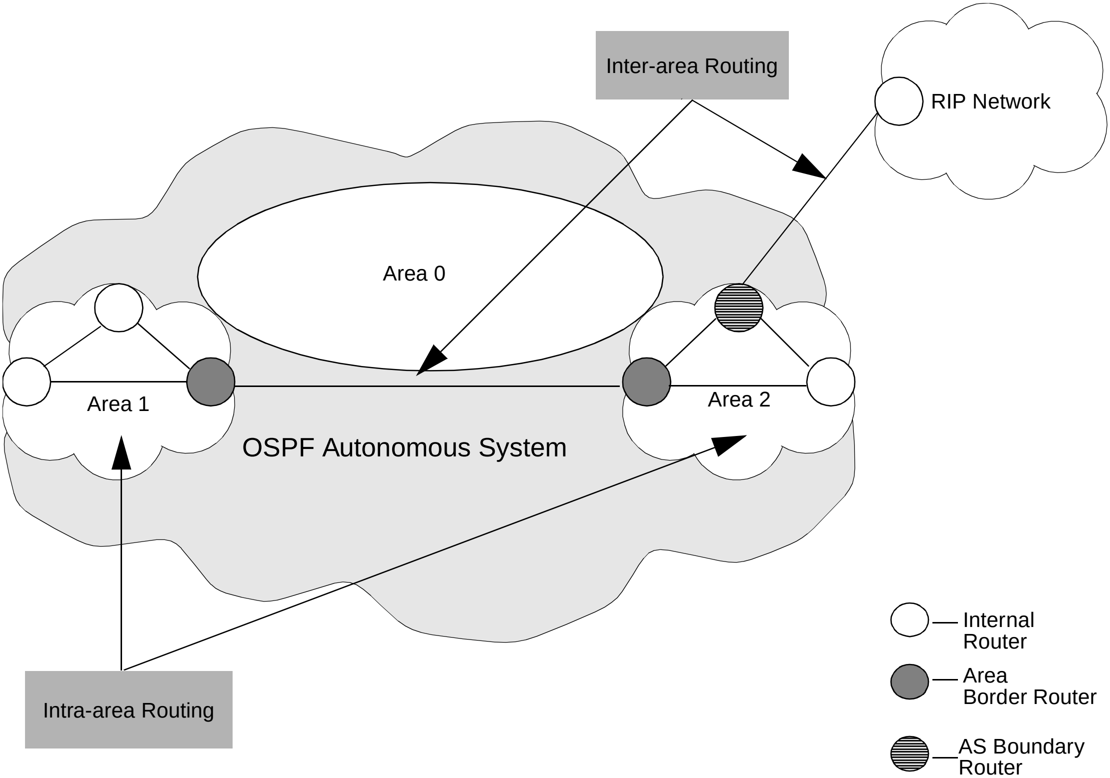

Figure 9-4. Router Classifications

Configuring OSPF

This section describes how to set OSPF parameters and includes the following tasks:

- Configuring OSPF parameters at the logical port
- Configuring OSPF parameters at the switch, including
  - Configuring IP parameters
  - Defining an OSPF neighbor
  - Defining an OSPF area aggregate
  - Defining an OSPF virtual link
  - Configuring OSPF route maps

Configuring OSPF at the Logical Port

To configure OSPF on the logical port:

1. Enable the logical port for IP services as described in "Configuring Logical Ports for IP Services" on page 3-3.
2. Choose Add OSPF from the Set IP Interface Addresses dialog box. The Add OSPF Interface dialog box appears (Figure 9-5).

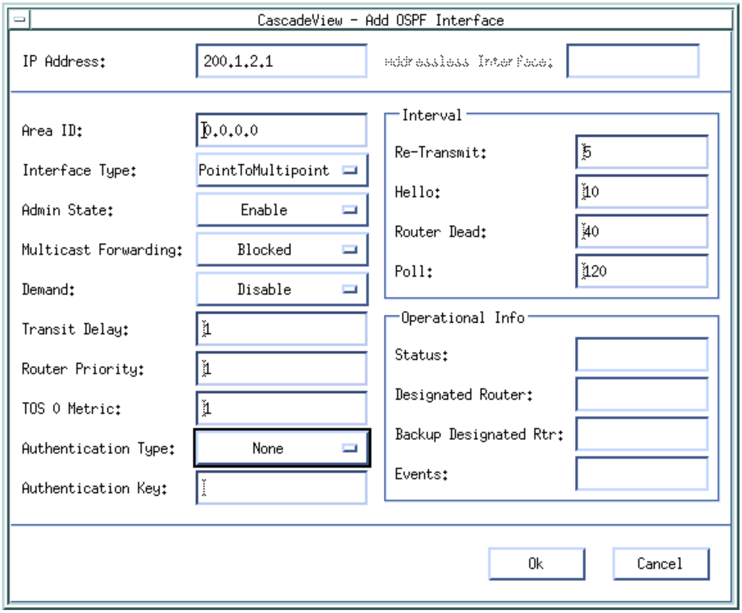

Figure 9-5. Add OSPF Interface

3. Complete the fields described in Table 9-2.

Table 9-2. 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 # or interface #.
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. ***Note***: *Area 1 is reserved for Ascend switches. If you configure the OSPF interface in Area 1, see “Configuring IP Parameters” on page 9-23*.

Table 9-2. Add OSPF Interface Fields (Continued)

Field	Action/Description
Interface Type	Select one of the following options for the logical port interface type: *Broadcast* – 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* – 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-Point* – A point-to-point network joins two routers together. The IP address of the neighboring routers interface is advertised. Hello packets are sent to the neighbor every *Hello Interval* seconds. For more information on Hello, see page 9-19. *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 routers interfaces are advertised. *Virtual Link* – A virtual-link network links areas that are not physically connected to the backbone and patches the backbone if a disconnect occurs in the backbone.
Admin State	Select one of the following options: *Enable* – This parameter allows this interface to communicate OSPF. In addition, this interface can send or receive Hello packets. *Disable* – This parameter prevents this interface from communicating OSPF. In addition, this interface cannot send or receive Hello packets.

Table 9-2. Add OSPF Interface Fields (Continued)

Field	Action/Description
Multicast Forwarding	Not Supported.
Demand	Not Supported.
Transit Delay	Enter a value betweeen 0 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.
Router Priority	Enter a value between 0 and 255. This number identifies the priority of the router associated with this logical port and is used to elect the designated 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 Metric	Enter a value between 1 and 65535. This value specifies the type of service cost. The lowest TOS 0 has the highest priority for routing.

Table 9-2. Add OSPF Interface Fields (Continued)

Field	Action/Description
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* – 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* – Specifies that an encryption method be used, which converts the authentication key to a number. The number is forwarded with the route rather than the actual key.
Authentication Key	Enter an authentication password in this field if you specified either *Simple* or *MD5* as the authentication type. This value is not required if you specified *None* as the authentication type.

Table 9-2. Add OSPF Interface Fields (Continued)

Field	Action/Description
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). The number of seconds between router Hello messages. This is a configurable parameter and controls the frequency of router Hello messages on an interface.
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. 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.
Poll	Enter a value greater than or equal to 0 (the default value for this field is 120). The time, in seconds, between Hello packets sent to an inactive non-broadcast multi-access (NBMA) neighbor.

Table 9-2. Add OSPF Interface Fields (Continued)

Field	Action/Description
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 include: *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 based on the Hello interval every *Hello Interval* seconds. See page 9-19 for details about Hello. *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.

Table 9-2. Add OSPF Interface Fields (Continued)

Field	Action/Description
Status (continued)	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 backup designated router or 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 9-2. 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 OSPF Parameters at the Switch

This section describes how to configure the following OSPF switch parameters:

- IP parameters
- OSPF neighbors
- OSPF area aggregates
- OSPF virtual links

Configuring IP Parameters

> *Configure this parameter only if you configured the switch's OSPF interface in Area 1, which is used only for Ascend switches.*
>
> *You also use Area 1 for routing updates between Ascend switches. This enables switches running switch software prior to 5.0 operate with switches running 5.0. If the interface is connected to a non-Ascend device, you cannot use Area 1.*

To configure IP parameters:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Cascade IP Parameters ⇒ Set IP Parameters. The Set IP Parameters dialog box appears (Figure 9-6).

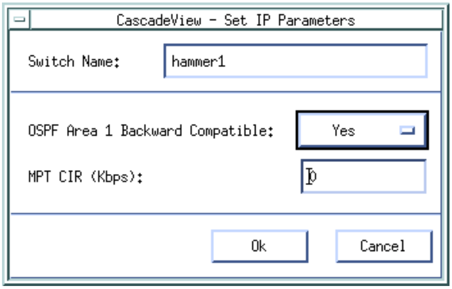

Figure 9-6. Set IP Parameters Dialog Box

3. Complete the fields described in Table 9-3.

Table 9-3. Set IP Parameters Field Descriptions

Field	Description
Switch Name	Displays the name of the switch.
OSPF Area 1 Backward Compatible	Select either Yes or No. If you select Yes, the switch: – can communicate with other Ascend switches running pre-5.0 switch software – can communicate with other Ascend switches running 5.0 switch software, which is set to Yes in this field – **cannot** communicate with other vendor routers If you select No, the switch: – cannot communicate with other Ascend switches running pre-5.0 switch software – can communicate with other Ascend switches running 5.0 switch software, which is set to No in this field – **can** communicate with other vendor routers
MPT CIR (Kbps) (Multi-Point-to-Point Tunneling Committed Information Rate)	Enter the rate in Kbps at which the Multi-Point-to-Point tunnel (MPT) transfers data, averaged over a minimum increment of time. In addition, this value reserves bandwidth for MPTs, which the switch originates. For more information on MPTs, see Chapter 1, "Multipoint-to-Point Tunneling" on page 1-5. ***Note:*** *This value applies to all links in the MPT.*

4. Choose OK.

Defining an OSPF Neighbor

> *You do not have to define OSPF neighbors if you assign OSPF to an interface. OSPF automatically discovers its neighbors through Hello packets. However, if you configure an NBMA network, you must define OSPF neighbors. See the description for NBMA networks on page 9-16.*

To define an OSPF neighbor:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Cascade IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Neighbors. The Set All OSPF Neighbors dialog box appears (Figure 9-7).

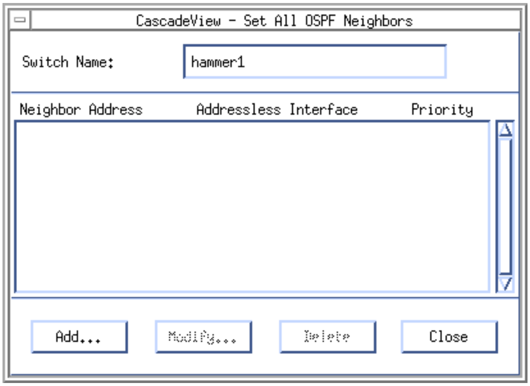

Figure 9-7. Set All OSPF Neighbors Dialog Box

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

Button	Description
Add	Enables you to add an OSPF neighbor.
Modify	Enables you to modify an OSPF neighbor.
Delete	Enables you to delete an OSPF neighbor.

3. At the Set All OSPF Neighbors dialog box, choose the Add button. The Add OSPF Neighbor dialog box appears (Figure 9-8).

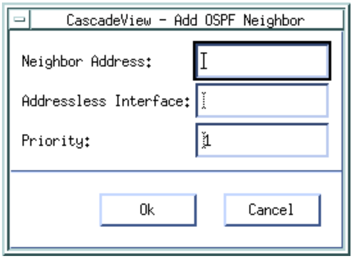

Figure 9-8. Add OSPF Neighbor Dialog Box

4. Specify the values as 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 address links, the address is not 0.0.0.0 but the address of the neighbors interface.
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 # or interface #.
Priority	Enter a value between 0 and 255. 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.

Defining an OSPF Area Aggregate

To define an OSPF area aggregate:

1. From the network map select the appropriate switch icon.
2. From the Administer menu, select Cascade IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Area Aggregates. The Set All OSPF Area Aggregates dialog box appears (Figure 9-9).

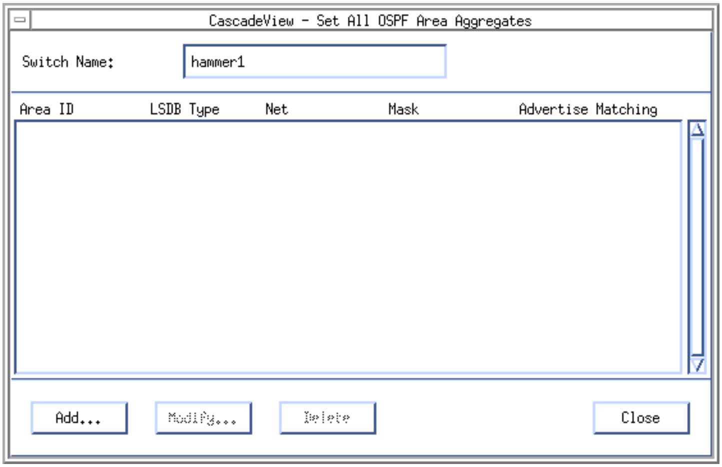

Figure 9-9. Set All OSPF Area Aggregates Dialog Box

The Set All OSPF Area Aggregates dialog box displays the following buttons:

Button	Description
Add	Enables you to add an OSPF area aggregate
Modify	Enables you to modify an OSPF area aggregate
Delete	Enables you to delete an OSPF area aggregate

3. Choose the Add button.

The Add OSPF Aggregates dialog box appears (Figure 9-10).

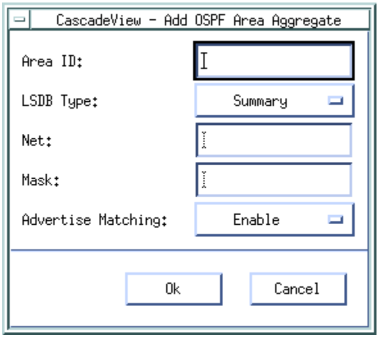

Figure 9-10. Add OSPF Area Aggregates Dialog Box

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

Table 9-5. Add OSPF Area Aggregates Fields

Field	Field/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. ***Note****: Area 1 is reserved for Ascend switches. If you configure the OSPF interface in Area 1, see "Configuring OSPF Parameters" on page 9-13.*

Table 9-5. Add OSPF Area Aggregates Fields (Continued)

Field	Field/Description
LSDB Type	Specify the link state database type to which this address aggregate applies. Options include: *Summary* – 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* – 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.

Defining an OSPF Virtual Link

To define an OSPF virtual link:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Cascade IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Virtual Links. The Set All OSPF Virtual Links dialog box appears (Figure 9-11).

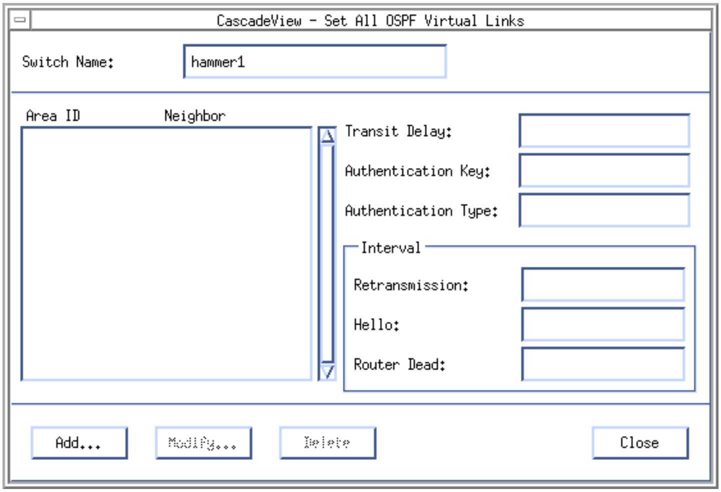

Figure 9-11. Set All OSPF Virtual Links Dialog Box

The Set All OSPF Virtual Links dialog box displays the following buttons:

Button	Description
Add	Enables you to add an OSPF virtual link.
Modify	Enables you to modify an OSPF virtual link.
Delete	Enables you to delete an OSPF virtual link.

3. Choose the Add button. The Add OSPF Link dialog box appears (Figure 9-12).

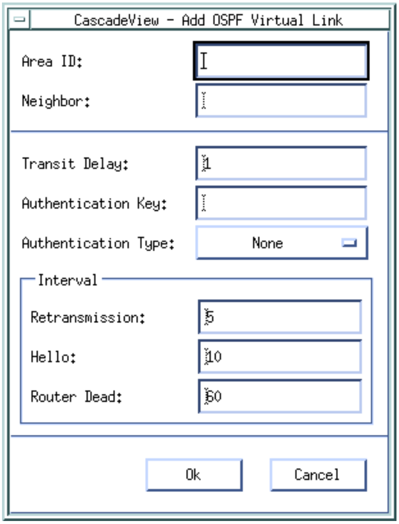

Figure 9-12. Add OSPF Virtual Link Dialog Box

4. Complete the fields described in Table 9-6.

Table 9-6. OSPF Virtual Link Fields

Field	Field/Description
Area ID	Enter the area ID (x.x.x.x) in which you want to locate the neighbor. Area 0.0.0.0 is the network backbone area. Areas are collections of networks, hosts, and routers. The area ID identifies the area. ***Note**: Area 1 is reserved for Ascend switches. If you configure the OSPF interface in Area 1, see "Configuring OSPF Parameters" on page 9-13.*
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 the neighbor's interface.
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 specified either *Simple* or *MD5* as the authentication type. This value is not required if you specified *None* as the authentication type.

Table 9-6. OSPF Virtual Link Fields (Continued)

Field	Field/Description
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* – 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* – Specifies that an encryption method be used, which converts the authentication key to a number. The number is forwarded with the route rather than the actual key.

Table 9-6. OSPF Virtual Link Fields (Continued)

Field	Field/Description
Interval	
Retransmission	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). The number of seconds between router Hello messages. This is a configurable parameter and controls the frequency of router Hello messages on an interface.
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. 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 an OSPF Route Map

Chapter 11, “Configuring Route Maps” provides detailed information about all types of route maps (including OSPF route maps) that you can configure, using IP services. See Chapter 11 before you begin any route map configuration

To configure an OSPF route map from the OSPF parameter menu:

1. From the network map, select the appropriate switch icon.
2. From the Administer menu, select Cascade IP Parameters ⇒ Set All OSPF ⇒ Set All OSPF Route Maps. The Set All OSPF Route Maps dialog box appears (Figure 9-13).

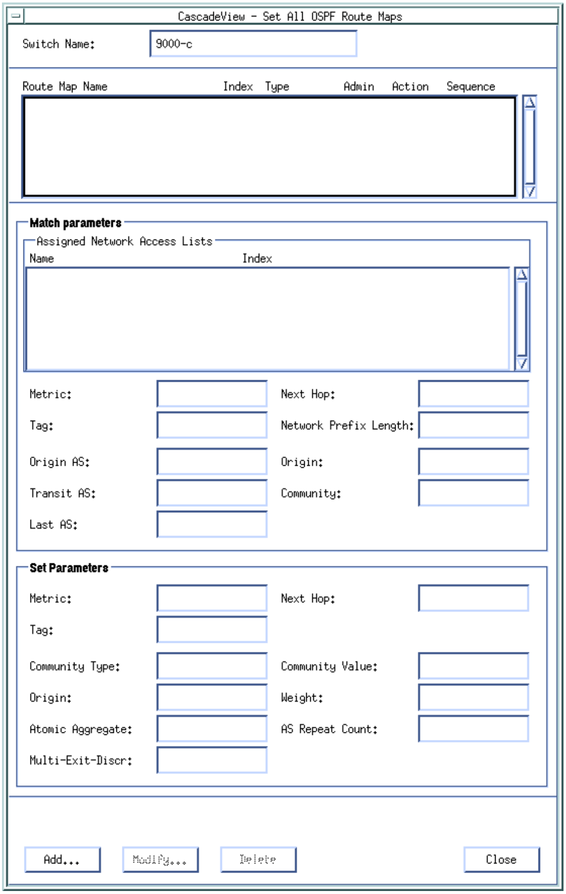

Figure 9-13. Set All OSPF Route Maps

See “About Route Maps” on page 11-3 for detailed information about configuring an OSPF route map and for dialog box field information.

10

Configuring 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 Cascade IP Parameters ⇒ Set All Static Routes. The All Static Route dialog box appears (Figure 10-1).

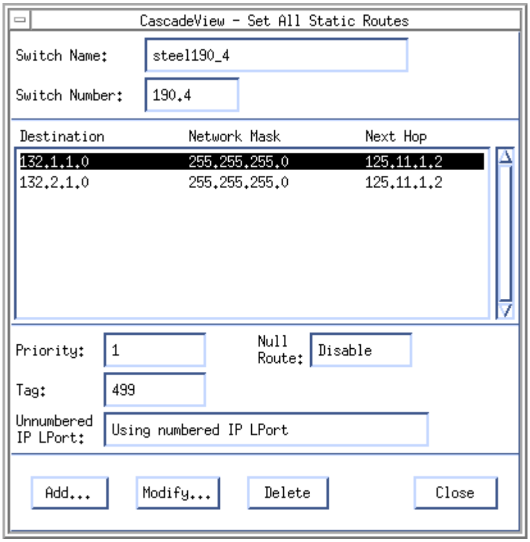

Figure 10-1. Set All Static Route Dialog Box

The Set All Static Route dialog box displays the following buttons.

Button	**Description**
Add	Enables you to add a static route
Modify	Enables you to modify a static route
Delete	Enables you to delete a static route

3. Choose the Add button. The Set Static Route dialog box appears (Figure 10-2).

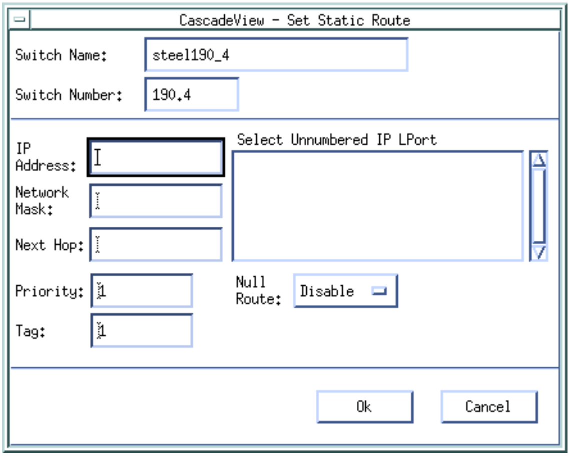

Figure 10-2. Set Static Route Dialog Box.

4. Complete the values as listed in Table 10-1.

Table 10-1. Static Route Fields

Field	Action/Description
Switch Name (read only)	Displays the switch name you want to add a static route to.
Switch Number (read only)	Displays the switch number.
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.
Tag	Enter the tag value, which you use to group multiple static route entries together.
Null Route	Select either *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.

Table 10-1. Static Route Fields

Field	Action/Description
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.

11

Configuring Route Maps

This chapter describes 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. 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 are using multiple route maps for the same neighbor, you can specify the order that IP Navigator uses the specified 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). *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 the traffic or a combination of both the traffic source and destination. A route map c be enabled or disabled administratively.

When you create a route map you specify two routing protocols. These include a Fr Protocol and a To Protocol. The route map specifies how routes are redistributed fr one routing protocol to another. This is done between two different protocols as we as within the same protocol (for example, from BGP to BGP). The route maps are a 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 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 From options include the following options:

- BGP
- OSPF
- RIP
- Static
- Direct
- Aggregate
- Any

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.

The To choices that are available for selection vary depending upon the selected From choice. For example, the routing table option can only be selected if the From choice protocol is BGP or RIP. The possible list of To choices include the following options:

- BGP
- OSPF
- RIP
- Routing Table

What Determines if a Route Map is for Import or Export?

The protocol that you specify 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 with the Routing Table as the To choice are performed before IP Navigator adds the routes to the routing table.

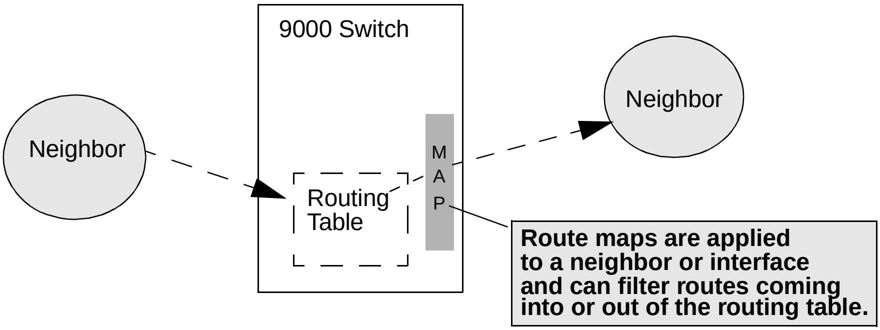

Figure 11-1. Using Route Maps to Filter Routes

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 OSPF routes to the routing table. For this reason you cannot use a route map to specify the acceptance or rejection of specific routes between OSPF and the routing table.

OSPF to OSPF – IP Navigator always advertises OSPF routes to the OSPF routing domain. Link state protocols assume that all routes share the same information. For this reason you cannot use a route map to specify the acceptance or rejection of specific routes being sent to an OSPF neighbor.

NMS Paths to OSPF – IP Navigator always advertises any NMS paths configured as Autonomous System External-Link State Advertisements (ASE-LSAs). (An NMS path is a static route that uses the Network Manageme Station as its destination.) For this reason you cannot use a route map to specify the acceptance or rejection of specific NMS paths to the OSPF protocol.

What Happens if You Don’t Use a Route Map?

If you do not use a route map for route filtering or route redistribution, the followin import and export operations occur by default:

- Routes from all protocols except for EBGP are imported into the routing table b 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 f the IP interface.

When are Route Maps Required?

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 attribute

Protocol Pairs That Require Route Maps

Route maps are also required for each of the following protocol pairs:

Protocol	Description
Static ⇒ OSPF Direct ⇒ OSPF BGP ⇒ OSPF RIP ⇒ OSPF	Route maps are required in order to advertise any Static, Direct, BGP, and RIP routes into the OSPF routing domain. By default IP Navigator does not advertise Static, direct, BGP and RIP routes into the OSPF routing domain.
Static ⇒ RIP Direct ⇒ RIP BGP ⇒ RIP OSPF ⇒ RIP	Route maps are required to advertise any Static, Direct, BGP, and OSPF routes into the RIP routing domain. By default IP Navigator does not advertise Static, Direct, BGP, and OSPF routes into the RIP routing domain.
Static ⇒ BGP Direct ⇒ BGP OSPF ⇒ BGP RIP ⇒ 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, and OSPF routes into the BGP routing domain.
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.

When are Route Maps Not Required?

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

Steps For Configuring a Route Map

Use the following steps to configure a route map:

1. Optionally define the network filters depending on your system’s needs. Refer to “Adding a Network Filter” on page 11-10 for more information.
2. Optionally use the defined network filters to create the network access lists. Refer to “Adding a Network Access List” on page 11-11 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. Refer to “Adding Route Maps” on page 11-15 for more information.
4. Assign the route map to a BGP neighbor or a RIP interface. You assign route maps to BGP interfaces on the Modify BGP Neighbor dialog box. Refer to Chapter 8, “Configuring BGP Parameters” for more information about accessing the BGP functions. You assign route maps to RIP interfaces on the Modify RIP Interface dialog box.Refer to Chapter 7, “Configuring RIP” for more information about accessing the RIP functions.

5. If you are using multiple route maps, specify the order that IP Navigator should use the assigned route maps. You do this by using the arrow buttons on the Modify BGP Neighbor and Modify RIP Interface dialog box. Route maps are used to filter routes on the interface in the order that 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.

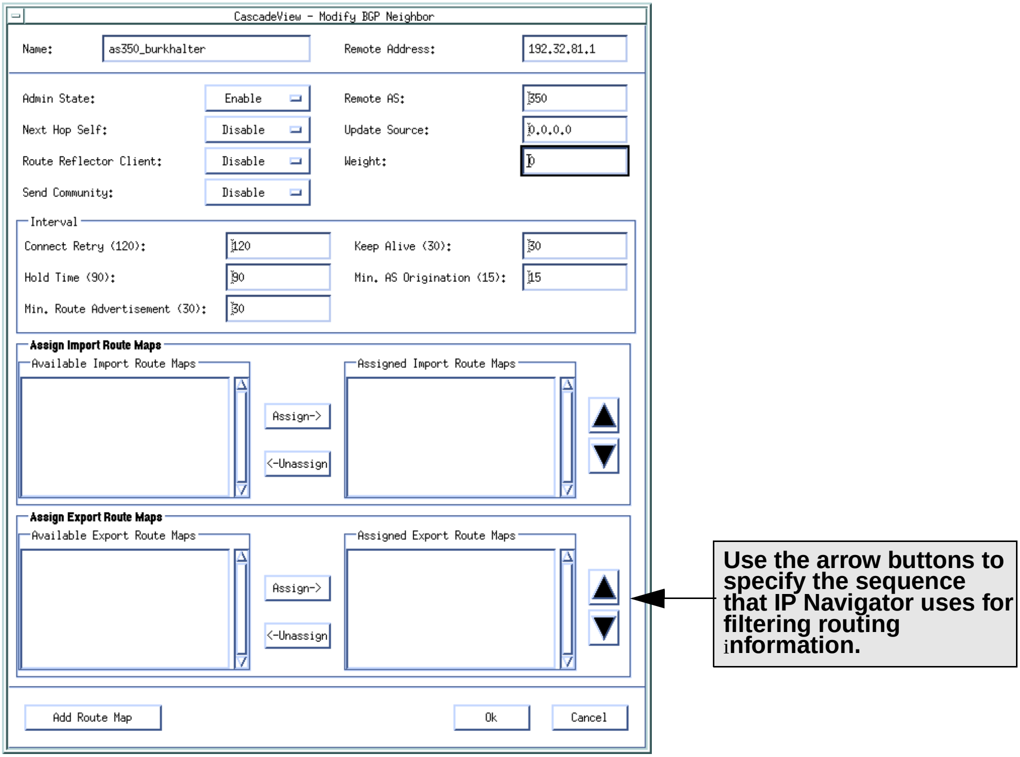

Figure 11-2. Using the Arrow Buttons to Sequence Route Maps

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 Cascade IP Parameters ⇒Set All Route Policies⇒Set All Network Filters. The Set All Network Filters dialog box displays.

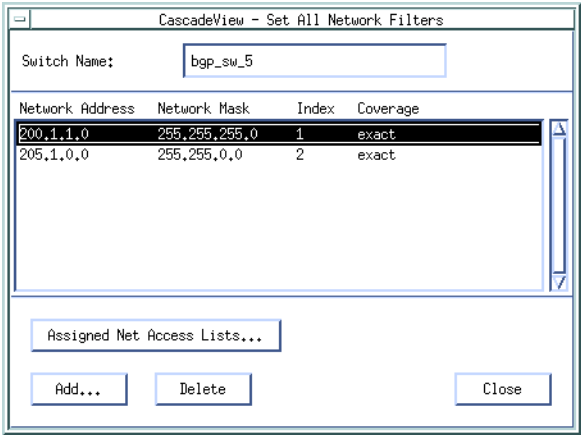

Figure 11-3. Set All Network Filters Dialog Box

Table 11-1 describes each of the Set All Network Filters command buttons.

Table 11-1. Set All Network Filters Command Buttons

Command	Description
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.
Assigned Net Access Lists	Displays a list of any network access lists that use the selected filter.

3. Choose Add. The Add Network Filter dialog box displays.

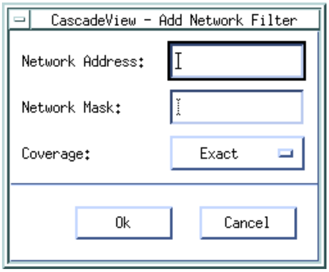

Figure 11-4. Add Network Filter Dialog Box

4. Specify the necessary network filter values listed in Table 11-2.

Table 11-2. Network Filter Fields

Field	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). 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 Cascade IP Parameters ⇒Set All Route Policies⇒Set All Network Access Lists. The Set All Network Access Lists dialog box displays.

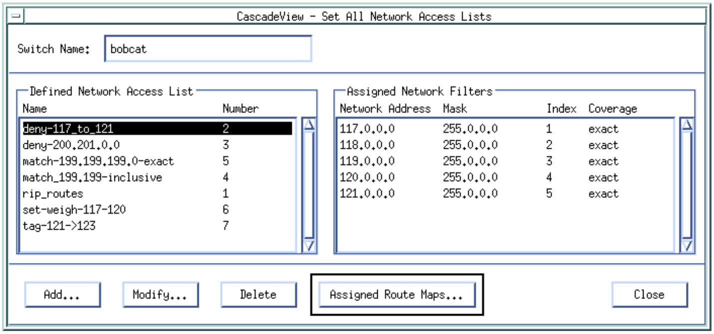

Figure 11-5. Set All Network Access Lists Dialog Box

Table 11-3 describes each of the Set All Network Access Lists command buttons.

Table 11-3. Set All Network Access List Command Buttons

Command	Description
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.
Route Maps	Displays a list of any route maps that use a selected network access list.

3. Choose Add. The Add Network Access List dialog box displays.

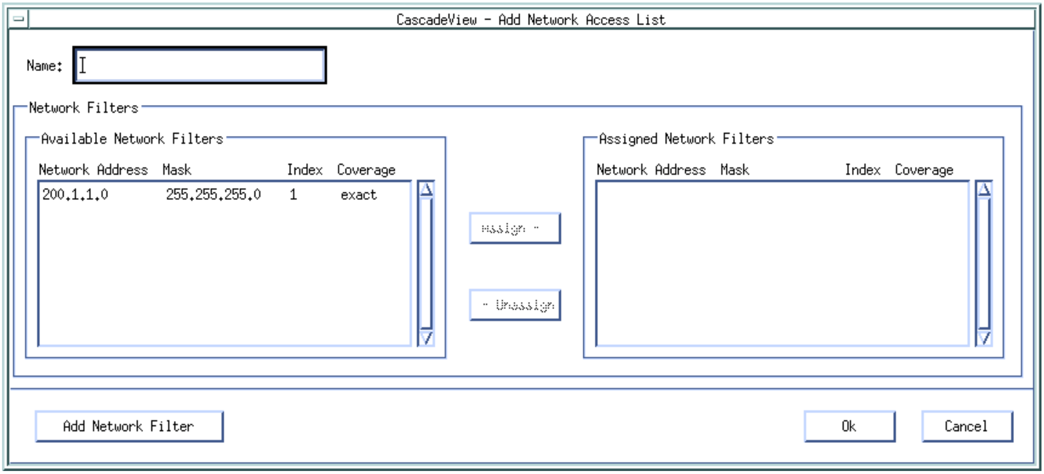

Figure 11-6. Add Network Access List Dialog Box

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 11-4 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 on Figure 11-4 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 include in the network access list. One network access list can include up to 300 network filters.

Table 11-4. Network Access List Fields

Field	Description
Name	Specify a unique network access list name.
Available Network Filters	A list of filters that are available for inclusion in the network access list.
Network Address	The network address for the filter.
Mask	The network mask for the filter.
	The index field is generated by CascadeView 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 CascadeView 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 Cascade IP Parameters ⇒Set All Route Policies⇒Set All Route Maps. The Set All Route Maps dialog box displays.

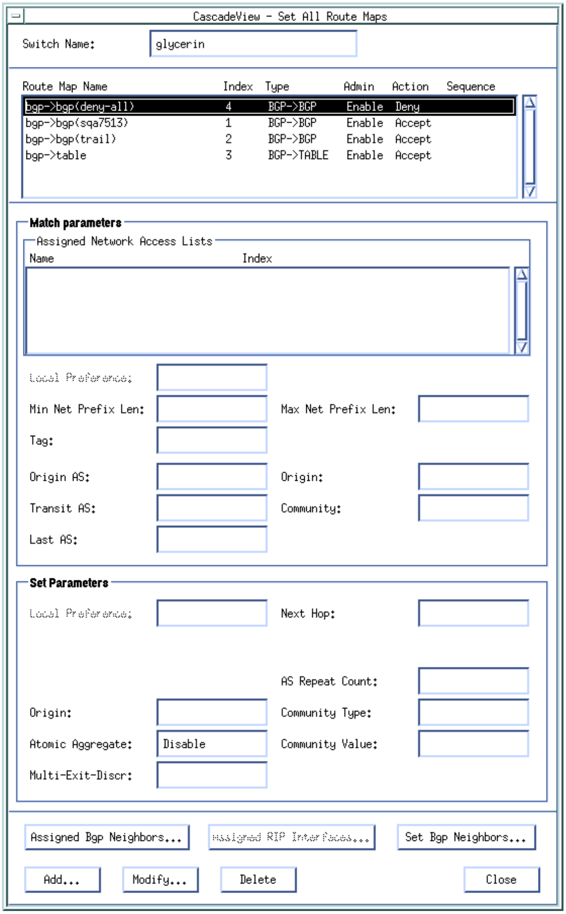

Figure 11-7. Set All Route Maps Dialog Box

Table 11-5 describes each of the Set All Route Maps command buttons. Table 11-6 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 11-9 for a reference to the section of this chapter that describes the Match and Set parameters for each route map type.

Table 11-5. Set All Route Maps Command Buttons

Command	Description
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.
Assigned by BGP Neighbors	List all BGP neighbors that use a selected route map.
Assigned by RIP Interfaces	List all RIP interfaces that use a selected route map.
Set BGP Neighbors	Displays the Set All BGP Neighbors dialog box to enable you to assign a route map to a BGP neighbor.

Table 11-6. Set All Route Maps Common Values

Field	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 CascadeView and is unique within the switch. This field is for internal system use only and cannot be modified.

Table 11-6. Set All Route Maps Common Values (Continued)

Field	Description
Type	Displays the From protocol and To protocol that identify the route distribution type. See Table 11-9 on page 11-21 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 displays.

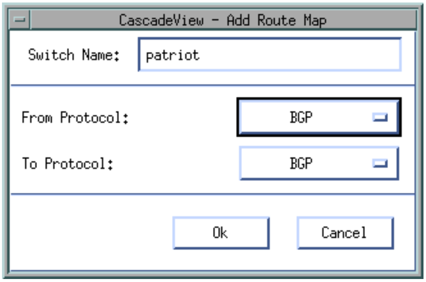

Figure 11-8. Add Route Map Dialog Box

4. Specify the values listed in Table 11-7.

Table 11-7. Route Map Descriptions

Field	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, or ANY.
To Protocol	Specify one of the following values: BGP, OSPF, RIP, or Routing Table. The routing table option can only be selected if the From protocol is either BGP or RIP.

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 11-9.

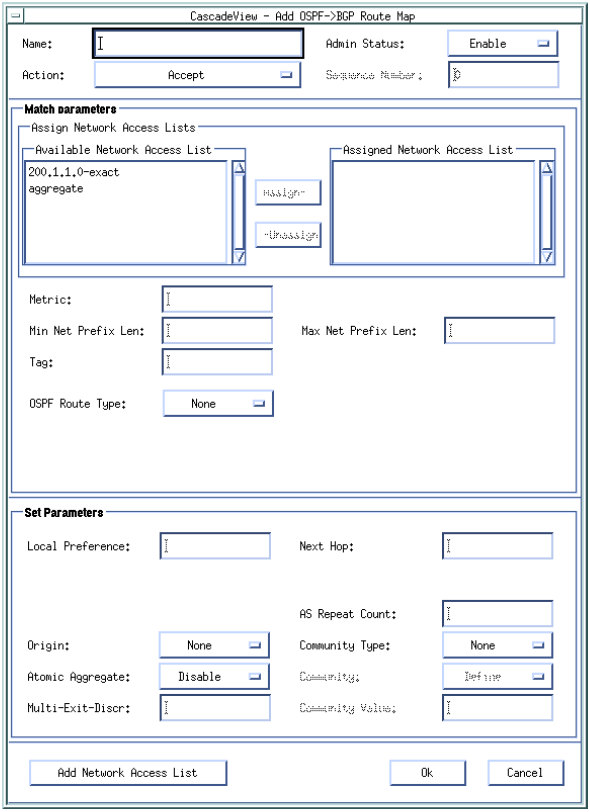

Figure 11-9. Second Add Route Map Dialog Box

6. Specify the Route Map Name, Admin Status, and Action values as described in Table 11-8.

7. Specify the necessary match and set parameters for this route map. If you need to add an access list for inclusion in the route map, choose Add Access Lists. Instructions for adding access lists start on page 11-11.

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 11-9 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 11-10 through Table 11-25 are optional. It is possible to specify a routing policy that uses no match and no set values.

Table 11-8. Add Route Map Fields

Field	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 11-9. Match and Set Parameter Field Descriptions

Route Map Type	See...
BGP to BGP	Table 11-10 on page 11-22
BGP to OSPF	Table 11-11 on page 11-26
BGP to RIP	Table 11-12 on page 11-28
BGP to Routing Table	Table 11-13 on page 11-31
OSPF to BGP	Table 11-14 on page 11-34
OSPF to RIP	Table 11-15 on page 11-37
RIP to RIP	Table 11-16 on page 11-38
RIP to BGP	Table 11-17 on page 11-39
RIP to OSPF	Table 11-18 on page 11-41
RIP to Routing Table	Table 11-19 on page 11-42
Static to OSPF	Table 11-20 on page 11-44
Static to BGP	Table 11-21 on page 11-45
Static to RIP	Table 11-22 on page 11-47
Any or Direct to BGP	Table 11-23 on page 11-48
Any or Direct to OSPF	Table 11-24 on page 11-51
Any or Direct to RIP	Table 11-25 on page 11-52
Aggregate to BGP	Table 11-26 on page 11-53

*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 11-10. BGP to BGP Match and Set Parameter Fields

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 that is 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 that is 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.
Origin AS	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. 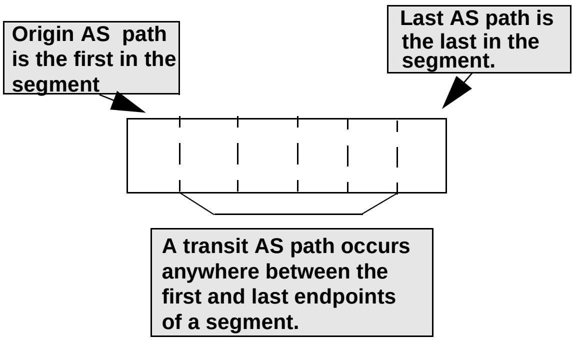

Table 11-10. BGP to BGP Match and Set Parameter Fields (Continued)

Field	Description
Transit AS	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 the figure in the Origin AS description above for further details.
Last AS	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 the figure in the Origin AS description above for further details.
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.

Table 11-10. BGP to BGP Match and Set Parameter Fields (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.**
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** – a new community number is assigned to replace the old value. **Additive** – a community is added to an existing community **None** – no community modification will occur.

Table 11-10. BGP to BGP Match and Set Parameter Fields (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 11-11.BGP to OSPF Match and Set Parameter Fields

Field	Description
Match Parameters	BGP routes can be distributed to OSPF based on matches to the following paramers. 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 that is 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 that is 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.
Origin AS	Specify a match parameter for the Autonomous System (AS) where the route originated.
Transit AS	Specify a match parameter for transit Autonomous System (AS) that is recorded in the route.
Last AS	Specify a match parameter for the last Autonomous System (AS) in the route.
Origin	Specify one of the following values to indicate the BGP origin code:IGP, EGP, Incomplete, or None.

Table 11-11.BGP to OSPF Match and Set Parameter Fields (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 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 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 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 specifies the next hop to reach a network. If you leave this field blank, a default of 0 is used.

Table 11-12.BGP to RIP Match and Set Parameter Fields

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 that is 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 that is 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.
Origin AS	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. 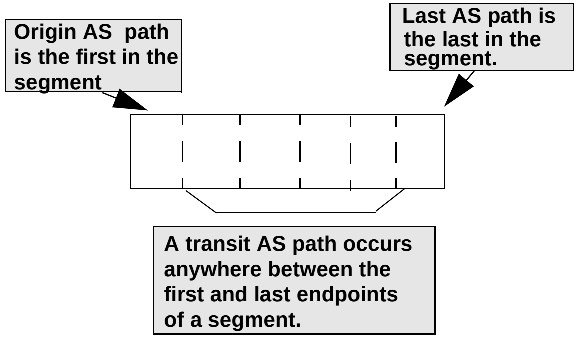

Table 11-12.BGP to RIP Match and Set Parameter Fields (Continued)

Field	Description
Transit AS	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 the figure in the Origin AS description above for further details.
Last AS	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 the figure in the Origin AS description above for further details.
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 Parameter	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.**

Table 11-12.BGP to RIP Match and Set Parameter Fields (Continued)

Field	Description
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.
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 11-13. BGP to Routing Table Match and Set Parameter Fields

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 that is 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 that is greater than this value are not selected.
Origin AS	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. 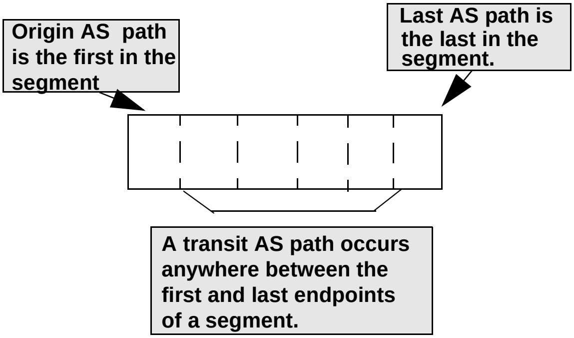
Transit AS	Specify a match parameter for the transit Autonomous System (AS) that is recorded in the route. An AS path value uses the originating, transit, or last AS path to further identify a route. See the figure in the Origin AS description above for further details.

Table 11-13. BGP to Routing Table Match and Set Parameter Fields

Field	Description
Last AS	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 the figure in the Origin AS description above for further details.
Next Hop	Specify the IP address that specifies the next hop to reach a network. Only routes that match this next hop value are selected.
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 Parameter	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.

Table 11-13. BGP to Routing Table Match and Set Parameter Fields

Field	Description
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 is an indication that a route is more preferred. 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.

Table 11-14. OSPF to BGP Match and Set Parameter Fields

Field	Description
Match Parameters	The redistribution of routes from an 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 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 that is 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 that is 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.
OSPF Route Type	Specify one of the following 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. **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 is an indication that a route is more preferred. This value is local to the AS and is exchanged between IBGP peers only. It is not passed to EBGP peers.

Table 11-14. OSPF to BGP Match and Set Parameter Fields (Continued)

Field	Description
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 specifies 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** – indicates that a new community number is assigned to replace the old value. **Additive** – indicates that a community is added to an existing community **None** – indicates that no community modification will occur.

Table 11-14. OSPF to BGP Match and Set Parameter Fields (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 11-15. OSPF to RIP Match and Set Parameter Fields

Field	Description
Match Parameters	The redistribution of routes from 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 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 that is 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 that is 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.
OSPF Route Type	Specify one of the following OSPF Metric Type values Intra, Internal, External-1, External-2, or None.
Set Parameter	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 specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-16.RIP to RIP Match and Set Parameter Fields

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 that is 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 that is 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 specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-17.RIP to BGP Match and Set Parameter Fields

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 that is 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 that is 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 is an indication that a route is more preferred. 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.

Table 11-17.RIP to BGP Match and Set Parameter Fields (Continued)

Field	Description
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 specifies 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** – a new community number is assigned to replace the old value. **Additive** – a community is added to an existing community. **None** – no community modification will occur.
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 11-17.RIP to BGP Match and Set Parameter Fields (Continued)

Field	Description
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. RIP to OSPF Match and Set Parameter Fields

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 that is 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 that is 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.

Table 11-18. RIP to OSPF Match and Set Parameter Fields

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	The 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 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 11-19. RIP to Routing Table

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 that is less than this value are not selected.

Table 11-19. RIP to Routing Table

Field	Description
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length that is 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 specifies 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 11-20.Static to OSPF Match and Set Parameter Fields

Field	Description
Match Parameters	Static routes can be distributed to 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.
Metric	The 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 that is 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 that is 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 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 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.
Next Hop	The IP address that specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-21.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.
Metric	The 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 that is 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 that is 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 for the same destination. A higher local preference value is an indication that a route is more preferred. 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.

Table 11-21.Static to BGP Match and Set Parameters (Continued)

Field	Description
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 specifies 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** – a new community number is assigned to replace the old value. **Additive** – a community is added to an existing community. **None** – no community modification will occur.
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 11-21.Static to BGP 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 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.Static to RIP Match and Set Parameter Fields

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.
Metric	The 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 that is 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 that is 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.

Table 11-22.Static to RIP Match and Set Parameter Fields (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 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 specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-23.Any or Direct to BGP

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.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length that is less than this value are not selected.

Table 11-23. Any or Direct to BGP (Continued)

Field	Description
Max Net Prefix Len	Specify a value from 0 to 32 to indicate the maximum network prefix length. Any routes with a prefix length that is 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.
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 for the same destination. A higher local preference value is an indication that a route is more preferred. 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 specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-23.Any or Direct to BGP (Continued)

Field	Description
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** – indicates that a new community number is assigned to replace the old value. **Additive** – indicates that a community is added to an existing community **None** – indicates that no community modification will occur. 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 11-24.Any or Direct to OSPF 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.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length that is 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 that is 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 OSPF cost. Ifyou leave this field blank, a default value from the routing table is used.
Tag	The tag to be set in the redistributed routes to 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 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.

Table 11-24. Any or Direct to OSPF 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 value of 0 is used.

Table 11-25. Any or Direct to RIP 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.
Min Net Prefix Len	Specify a value from 0 to 32 to indicate the minimum network prefix length. Any routes with a prefix length that is 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 that is 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.**

Table 11-25. Any or Direct to RIP Parameters

Field	Description
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 specifies the next hop to reach a network. If you leave this field blank, a default value of 0 is used.

Table 11-26. Aggregate to BGP

Field	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 for the same destination. A higher local preference value is an indication that a route is more preferred. 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.

Table 11-26.Aggregate to BGP (Continued)

Field	Description
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 specifies 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** – indicates that a new community number is assigned to replace the old value. **Additive** – indicates that a community is added to an existing community **None** – indicates that no community modification will occur. 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.

Table 11-26.Aggregate to BGP (Continued)

Field	Description
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.

12

Monitoring IP Objects

This chapter describes how to retrieve status information about the various Ascend IP objects such as IP logical ports and associated parameters including OSPF, BGP, and RIP. Table 12-1 lists the sections in this chapter that describe each of the Monitor functions.

Table 12-1. Monitor Functions

To Monitor	Refer to...
IP Logical Ports	"Monitoring Logical Ports" on page 12-3
OSPF	"Monitoring OSPF" on page 12-19
BGP	"Monitoring BGP" on page 12-26
Network Filters	"Monitoring Filters, Access Lists, and Route Maps" on page 12-38
Network Access Lists	"Monitoring Network Access Lists" on page 12-40
Route Maps	"Monitoring Route Maps" on page 12-42
Packet Filters	"Monitoring IP Packet Filters" on page 12-44
QoS Profiles	"Monitoring QoS Profiles" on page 12-49
Protocol Filters	"Monitoring Static Routes" on page 12-53
Static Routes	"Monitoring Static Routes" on page 12-53
ARP	"Monitoring Static ARP Parameters" on page 12-54
MPT Path	"Monitoring MPT Path" on page 12-55
RIP2	"Monitoring RIP" on page 12-60
IP Routing Table	"Monitoring the IP Routing Table" on page 12-63

Monitoring Logical Ports

To view an IP logical port:

1. On the network map, select the appropriate switch object.
2. From the Monitor menu, select Cascade IP Objects ⇒ Show All IP LPorts. The Show All IP LPorts dialog box appears (Figure 12-1).

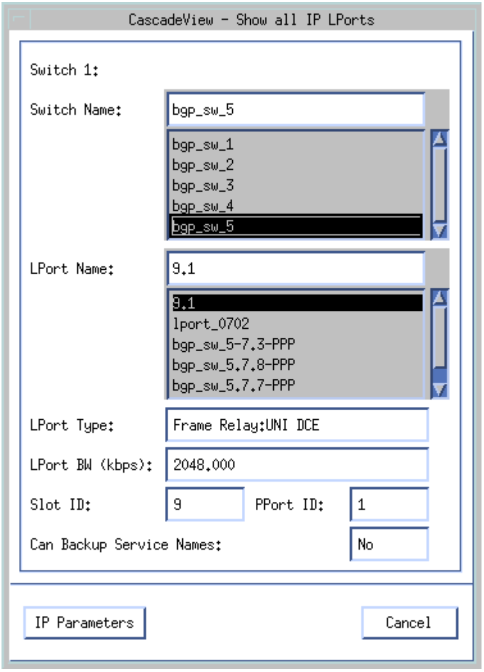

Figure 12-1. Show All IP LPorts

3. Select the IP Lport for which you want to view IP parameters.
4. Choose IP Parameters. The Show IP Parameters dialog box appears (Figure 12-2).

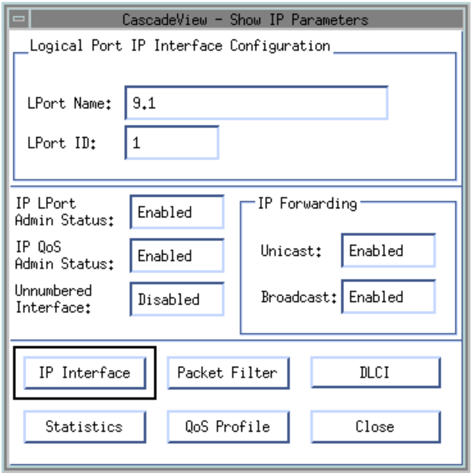

Figure 12-2. Show IP Parameters Dialog Box

Table 12-2 describes the Show IP Parameters buttons. Table 3-2 on page 3-10 describes the IP Parameters fields.

Table 12-2. Show IP Parameters Buttons

Button	Description
IP Interface	Enables you to monitor the IP interface address configuration. Refer to page 3-12 for more details on this function.
Packet Filter	Enables you to display inbound and outbound packet filters. Refer to Chapter 4, "Configuring IP Packet Filters" for more information about packet filters.
DLCI	*(For Frame Relay modules only)*. Enables you to monitor the Data Link Connection Identifier (DLCI) for the logical port. Refer to "Setting the DLCI for Frame Relay Logical Ports" on page 3-17 for more details on this function.
VPI/VCI	*(For ATM modules only)*. Enables you to monitor the virtual path identifier (VPI) and virtual channel identifier (VCI) for the logical port. Refer to"Setting the VPI/VCI for ATM Logical Ports" on page 3-19 for more details on this function.
Statistics	Displays the IP Lport Statistics dialog box shown in Figure 12-9 on page 12-13. Refer to "Monitoring IP Logical Port Statistics" on page 12-13 for more information about IP logical port statistics.
QoS Profile	Enables you to monitor the Quality of Service profiles. Refer to Chapter 5, "Provisioning IP Quality of Service" for more information on QoS profiles.

Monitoring an IP Interface Address

To view an IP interface address:

1. Access the Show IP Parameters dialog box (Figure 12-2) as described in "Monitoring Logical Ports" on page 12-3.

2. Choose IP Interface. The Show IP Interface Addresses dialog box appears (Figure 12-3).

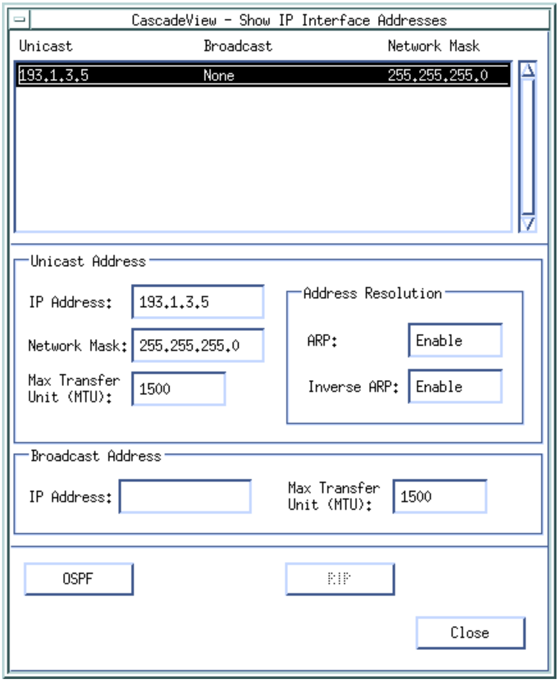

Figure 12-3. Show IP Interface Addresses Dialog Box

Table 12-3 describes the Show IP Interface Addresses buttons. Table 3-4 on page 3-15 describes the IP Interface Address fields.

Table 12-3. Show IP Interface Addresses Buttons

Button	Description
OSPF	Enables you to monitor the OSPF parameters for the logical port.
RIP2	Enables you to monitor the RIP parameters for the logical port.

Monitoring the Logical Port OSPF Interface

To view an OSPF interface for a logical port:

1. Access the Show IP Interface Addresses dialog box (Figure 12-3). If you are not sure how to do this, refer to "Monitoring an IP Interface Address" on page 12-6.
2. Choose OSPF. The Show OSPF Interface dialog box appears (Figure 12-4).

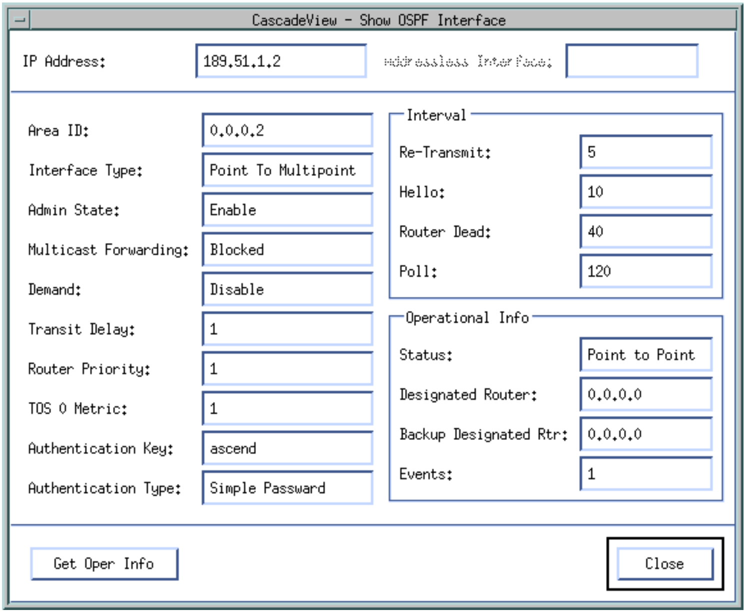

Figure 12-4. Show OSPF Interface Dialog Box

3. Select the Get Oper Info button to display the operational status for the selected Lport in the Status field.

   Table 9-2 on page 9-15 describes the OSPF interface fields.

4. Choose Close.

Monitoring the Logical Port RIP Parameters

To view the RIP parameters for a logical port:

1. Access the Show IP Interface Addresses dialog box (Figure 12-3) as described in "Monitoring an IP Interface Address" on page 12-6.

   Choose RIP2. The Show RIP dialog box appears (Figure 12-5). See Table 7-1 on page 7-3 for a description of each of the RIP fields.

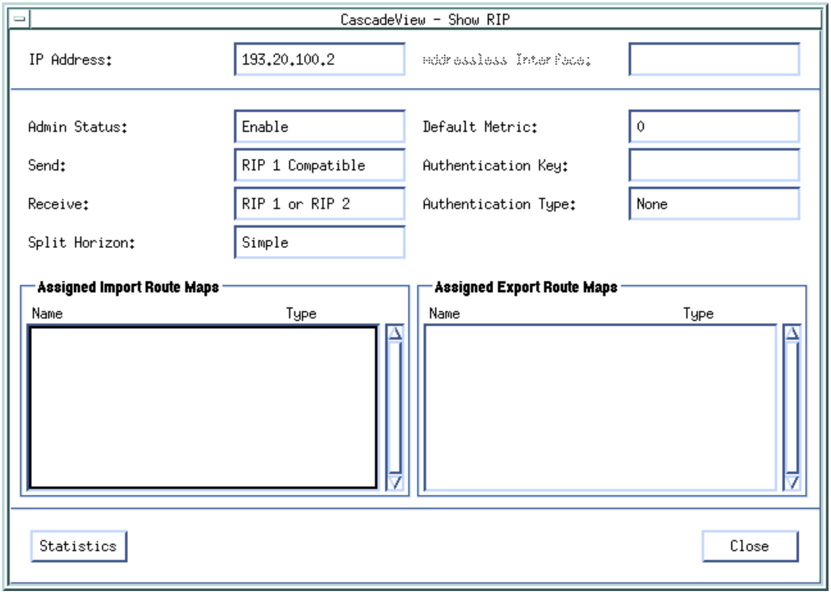

Figure 12-5. Show RIP Dialog Box

2. Use the Statistics option to display statistics for this IP Lport. The RIP2 Statistics dialog box appears (Figure 12-6 on page 12-9).

3. Choose Close.

Monitoring RIP Statistics

To monitor RIP statistics:

1. Access the Show RIP dialog box (Figure 12-4) as described in "Monitoring the Logical Port RIP Parameters" on page 12-8.

2. Choose Statistics. The RIP2 Statistics dialog box appears (Figure 12-6).

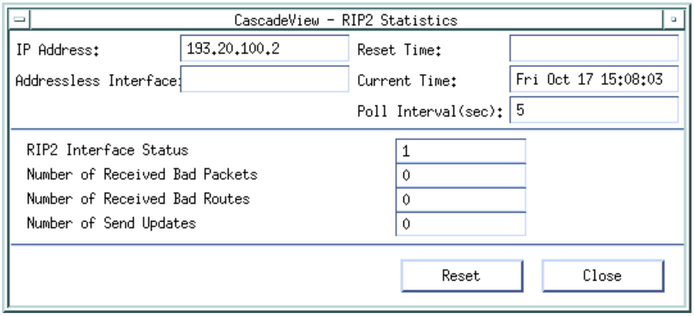

Figure 12-6. RIP2 Statistics Dialog Box

Table 12-4 describes the RIP2 Statistics fields.

3. Choose Reset to clear all statistics counters.

4. Choose Close.

A warm boot clears all of the cumulative statistics and resets the counters.

Table 12-4. RIP2 Statistics

Field	Description
Identifying Fields	
IP Address	Displays the Ethernet IP address of the switch on the indicated subnet. For un-numbered interfaces, the system displays a value of 0.0.0.N where the least significant 24 bits (N) represent the ifIndex for the IP interface.
Addressless Interface	Displays the un-numbered interface address if this IP logical port is not part of a subnet. Because the port does not have a specific address, it uses the router ID address instead.
Reset Time	Displays the time that the Reset command button was last selected to reset counters.
Current Time	Displays the current system time.
Poll Interval (sec)	Displays the time interval for the collection of statistical data. Refer to the *Network Configuration Guide for B-STDX/STDX* for information about setting the poll interval.
Cumulative Statistics	
RIP2 Interface Status	Displays the status of the RIP2 logical port interface.
Number of Received Bad Packets	Displays the total number of RIP response packets received by the RIP process (since the last reset) that were discarded for any reason (for example, a version 0 packet or an unknown command type).
Number of Received Bad Routes	Displays the total number of routes (in valid RIP packets) that were ignored for any reason (for example, an unknown address family or an invalid metric).
Number of Send Updates	Displays the number of triggered RIP updates sent from this interface. This counter does not include full updates that were sent containing new information.

Monitoring Packet Filters

You must be logged on in order to access this function.

To monitor the packet filters that are currently assigned to a logical port:

1. Access the Show IP Parameters dialog box (Figure 12-2). If you are not sure how to do this, refer to "Monitoring Logical Ports" on page 12-3.
2. Choose Packet Filter. The Assign Logical Port IP Filter dialog box appears. The Assigned Filters column lists the packet filters assigned to this logical port. The Available Filters column lists the packet filters that are available for use with this logical port. For more information about packet filters, refer to Chapter 4, "Configuring IP Packet Filters".

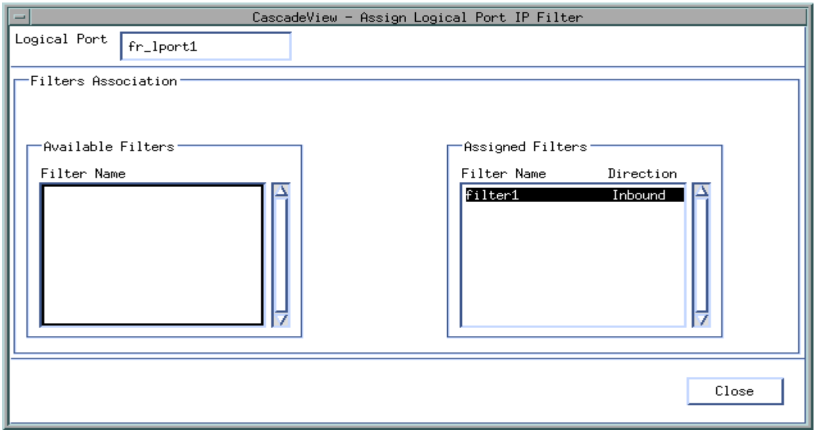

Figure 12-7. Assign Logical Port IP Filter Dialog Box

3. Choose Close to return to the Show IP Parameters dialog box (Figure 12-2).

Monitoring a DLCI Data Link ID

1. Access the Show IP Parameters dialog box (Figure 12-2) as described in "Monitoring Logical Ports" on page 12-3.

2. Choose DLCI. The Show IP Interface Data Link ID dialog box appears (Figure 12-8).

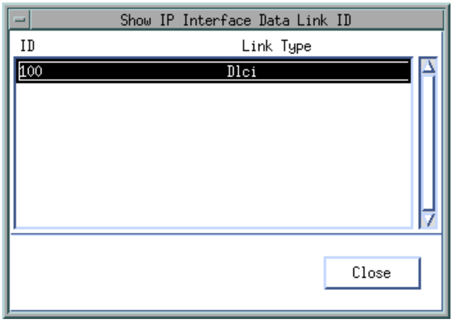

Figure 12-8. Show IP Interface Data Link ID Dialog Box

Table 12-5 describes the Show IP Interface fields.

Table 12-5. Show IP Interface Fields (For Frame Relay Logical Ports)

Field	Description
ID	The DLCI value for this IP interface. Refer to "Setting the DLCI for Frame Relay Logical Ports" on page 3-17 for more information on this field.
Link Type	Displays DLCI to indicate that this logical port uses a DLCI link. For more information about this link type refer to "Setting the DLCI for Frame Relay Logical Ports" on page 3-17.

3. Choose Close to return to the Show IP Parameters dialog box (Figure 12-2).

Monitoring IP Logical Port Statistics

To view IP Logical port statistics:

1. Access the Show IP Parameters dialog box (Figure 12-2). As described in “Monitoring Logical Ports” on page 12-3.
2. Choose Statistics. The IP LPort Statistics dialog box appears (Figure 12-9).

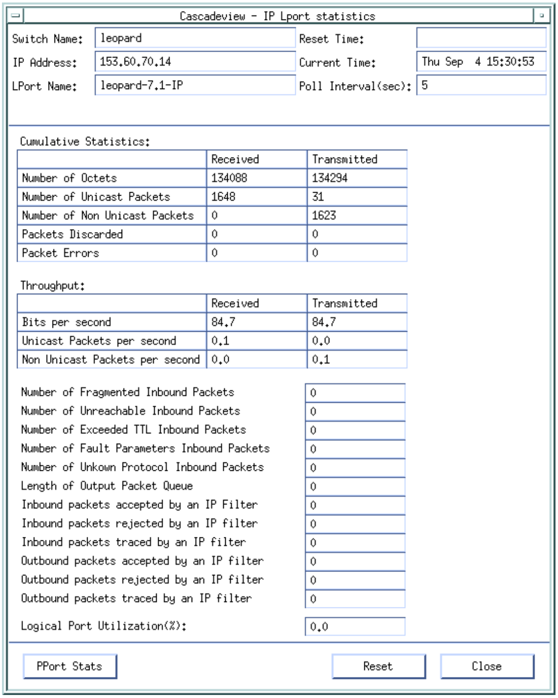

Figure 12-9. IP Lport Statistics Dialog Box

The Logical Port Utilization field measures the traffic queued for transmission on a logical port as a percentage of the logical port speed. It does not measure the amount of logical port bandwidth. Therefore, the Logical Port Utilization field value can exceed 100%.

A warm boot clears all of the cumulative statistics and resets the counters.

The IP LPort Statistics dialog box displays information about the transmission and receipt of data on the logical port. Cumulative Statistics list the number of each type of packet (received, transmitted, or tagged due to an error). Throughput fields display the bits and packets per second for the logical port. Table 12-6 describes the IP logical port statistics fields.

Table 12-6. IP LPort Statistics

Field	Description
IP Address	Displays the Ethernet IP address of the switch.
Addressless Interface	Displays the unnumbered interface address if this IP logical port is not part of a subnet. Because the port does not have a specific address, it uses the router ID as its address.
Reset Time	Displays the time that the Reset command button was last selected to reset counters.
Current Time	Displays the current system time.
Poll Interval (sec)	Displays the time interval for collection of statistics. Refer to the *Network Configuration Guide for B-STDX/STDX* for information about setting the poll interval.
Cumulative Statistics	
Number of Octets	Displays the total number of octets received since the last reset.
Number of Unicast Packets	Displays the total number of unicast packets received since the last reset.
Number of Non-Unicast Packets	Displays the total number of non-unicast packets received since the last reset.
Packets Discarded	Displays the total number of packets (frames) received and transmitted since the last reset.
Packet Errors	Displays the total number of packet errors received since the last reset.

Table 12-6. IP LPort Statistics (Continued)

Field	Description
Throughput Statistics	
Bits per Second	Displays the total number of bits received and transmitted each second.
Unicast Packets per Second	Displays the total number of unicast packets received and transmitted each second.
Non-Unicast Packets per Second	Displays the total number of non-unicast packets received and transmitted each second.
Number of Fragmented Inbound Packets	Displays the number of received IP packets that were fragmented because the packet size was greater than the transmit Maximum Transit Unit (MTU) size.
Number of Unreachable Inbound Packets	Displays the number of inbound packets which had a *Destination Unreachable* condition. The router had no route to the IP destination indicated by the Destination Address (DA) in the IP packet.
Number of Exceeded TTL Inbound Packets	Displays the number of inbound packets which had a *Time-to- Live (TTL) Exceeded* condition. This counter is decremented at each hop. If the counter equals zero, the packet is discarded. The purpose of TTL is to prevent packets from endlessly circulating on the network.
Number of Fault Parameters Inbound Packets	Displays the number of inbound packets which had a *Parameter Errors* condition. A parameter error is any value in the header that either appears incorrectly or cannot be interpreted.
Number of Unknown Protocol Inbound Packets	Displays the number of unknown protocol inbound packets. This counter does not apply to traffic that passes through the router (or Ascend switch). Instead, it applies to traffic destined for any IP entity (such as a router or an end station). This counter is incremented when the IP protocol stack has no packet destination (because there is no IP entity that is interested in this protocol).

Table 12-6. IP LPort Statistics (Continued)

Field	Description
Length of Output Packet Queue	Displays the length of the transmit queue, which indicates the level of congestion on the logical port.
Inbound packets accepted by an IP filter	Displays the number of packets which arrived at the port with inbound IP filters and were accepted because the IP filter matched the packet and specified acceptance of the packet as a result of the match.
Inbound packets rejected by an IP filter	Displays the number of packets which arrived at the port with inbound IP filters and were rejected because the IP filter matched the packet and specified rejection of the packet as a result of the match.
Inbound packets traced by an IP filter	Displays the number of packets which arrived at the port with inbound IP filters and were traced because the IP filter matched the packet and specified packet tracing as a result of the match.
Outbound packets accepted by an IP filter	Displays the number of packets which arrived at the port with outbound IP filters and were accepted because the IP filter matched the packet and specified acceptance of the packet as a result of the match.
Outbound packets rejected by an IP filter	Displays the number of packets which arrived at the port with outbound IP filters and were rejected because the IP filter matched the packet and specified rejection of the packet as a result of the match.
Outbound packets traced by an IP filter	Displays the number of packets which arrived at the port with outbound IP filters and were traced because the IP filter matched the packet and specified packet tracing as a result of the match.

Table 12-6. IP LPort Statistics (Continued)

Field	Description
Utilization Statistic	
Logical Port Utilization	Displays the amount of traffic queued for transmission on a logical port as a percentage of the committed information rate (CIR). This statistic does not measure the amount of bandwidth of the logical port. For this reason, the displayed value can exceed 100%.

3. Choose Reset to clear the statistics counters.
4. Choose PPort Stats to display the physical port statistics for the logical port. For more information about physical port statistics, refer to the *Diagnostic and Troubleshooting Guide for B-STDX/STDX*.
5. Choose Close to return to the Show IP Parameters dialog box.

Monitoring a VPI/VCI Link

To view a VPI/VCI:

1. Access the Show IP Parameters dialog box (Figure 12-2) as described in "Monitoring Logical Ports" on page 12-3.

2. Choose VPI/VCI. The Show IP Interface Data Link ID dialog box appears (Figure 12-8).

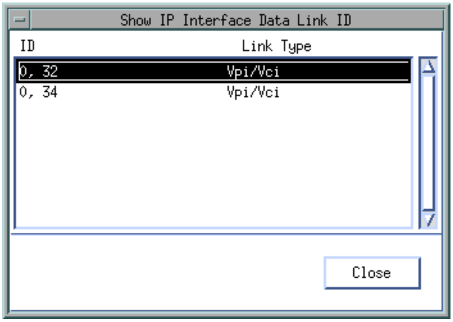

Figure 12-10. Show IP Interface Data Link ID Dialog Box

Table 12-5 describes the Show IP Interface fields.

Table 12-7. Show IP Interface Fields (For ATM Logical Ports)

Field	Description
ID	Displays the VPI/VCI identifier. The virtual path identifier (VPI) is the first number. The virtual channel identifier (VCI) is shown after the comma. Refer to "Setting the VPI/VCI for ATM Logical Ports" on page 3-19 for further information about this field.
Link Type	Displays VPI/VCI to indicate that this logical port uses a VPI/VCI link.

3. Choose Close.

Monitoring OSPF

This section describes how to monitor the following objects:

- IP Parameters
- Neighbors
- Area Aggregates
- Virtual Links

Monitoring IP Parameters

To view IP parameters:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show IP Parameters. The Show IP Parameters dialog box appears (Figure 12-12).

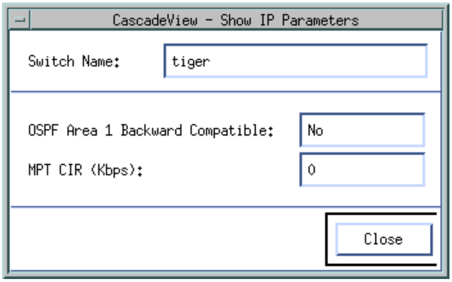

Figure 12-11. Show IP Parameters Dialog Box

Table 9-3 on page 9-24 describes OSPF fields.

3. Choose Close.

Monitoring OSPF Neighbors

To view an OSPF neighbors:

1. From the network map select the appropriate switch object.

2. From the Monitor menu, select Cascade IP objects⇒Show All OSPF⇒ Show All OSPF Neighbors. The Show All OSPF Neighbors dialog box appears (Figure 12-12).

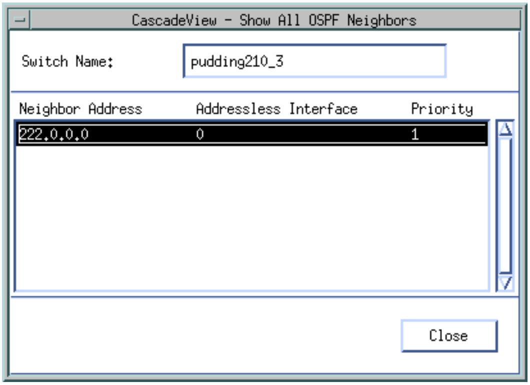

Figure 12-12. Show All OSPF Neighbors Dialog Box

Table 9-4 on page 9-27 describes OSPF Neighbor fields.

3. Choose Close to return to the root map display.

Monitoring OSPF Area Aggregates

To view an OSPF area aggregate:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All OSPF⇒ Show All OSPF Area Aggregates. The Show All OSPF Area Aggregates dialog box appears (Figure 12-13).

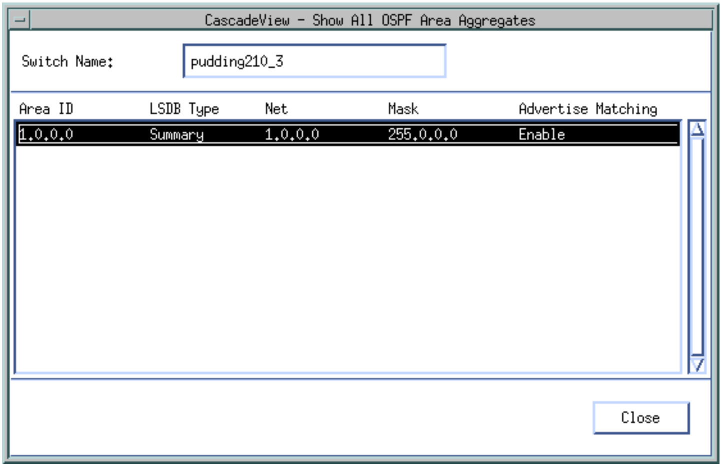

Figure 12-13. Show All OSPF Area Aggregates Dialog Box

Table 9-5 on page 9-29 describes OSPF Area Aggregate fields.

3. Choose Close to return to the root map display.

Monitoring OSPF Virtual Links

To monitor an OSPF switch's virtual links:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All OSPF⇒Show All OSPF Virtual Links. The Show All OSPF Virtual Links dialog box appears (Figure 12-14).

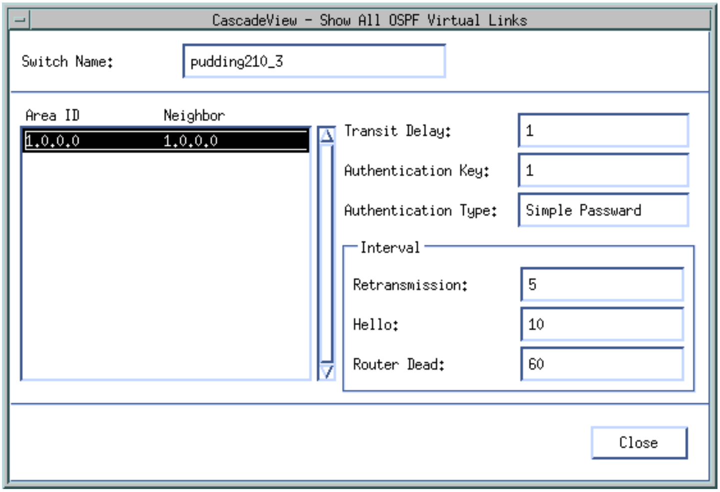

Figure 12-14. Show All OSPF Virtual Links Dialog Box

Table 9-6 on page 9-33 describes OSPF Virtual Link fields.

3. Choose Close to return to the root map display.

Monitoring OSPF Route Maps

To monitor an OSPF route map:

1. From the network map select the appropriate switch object.

2. From the Monitor menu, select Cascade IP objects⇒Show All OSPF⇒Show All OSPF Route Maps. The Show All OSPF Route Maps dialog box appears (Figure 12-15).

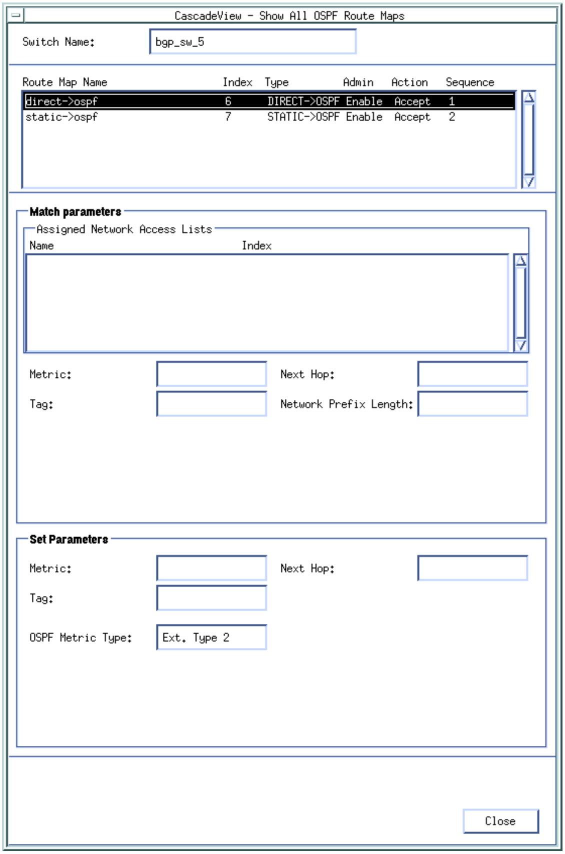

Figure 12-15. Show All OSPF Route Maps

Table 11-6 on page 11-16 describes each of the common fields shown at the top of the Show All Route Maps dialog box.

Note that the Match parameters and Set parameters vary depending on the type of route map that you are defining. Refer to Table 11-9 for information about the Match and Set parameters for each route map type.

3. Choose Close to return to the root map display.

Monitoring BGP

Monitoring BGP Switch Parameters

To view a BGP switch:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All BGP⇒Show All BGP Parameters. The Show All BGP Parameters dialog box appears (Figure 12-16).

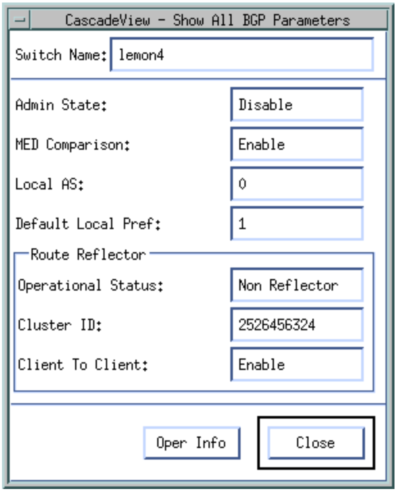

Figure 12-16. Show All BGP Parameters Dialog Box

Table 8-1 on page 8-8 describes the BGP fields.

3. Choose Oper Info to display a status message in the Operational Status field for the selected port.
4. Choose Close to return to the root map display.

Monitoring BGP Neighbors

To view a BGP neighbor:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All BGP⇒Show All BGP Neighbors. The Show All BGP Neighbors dialog box appears (Figure 12-17).

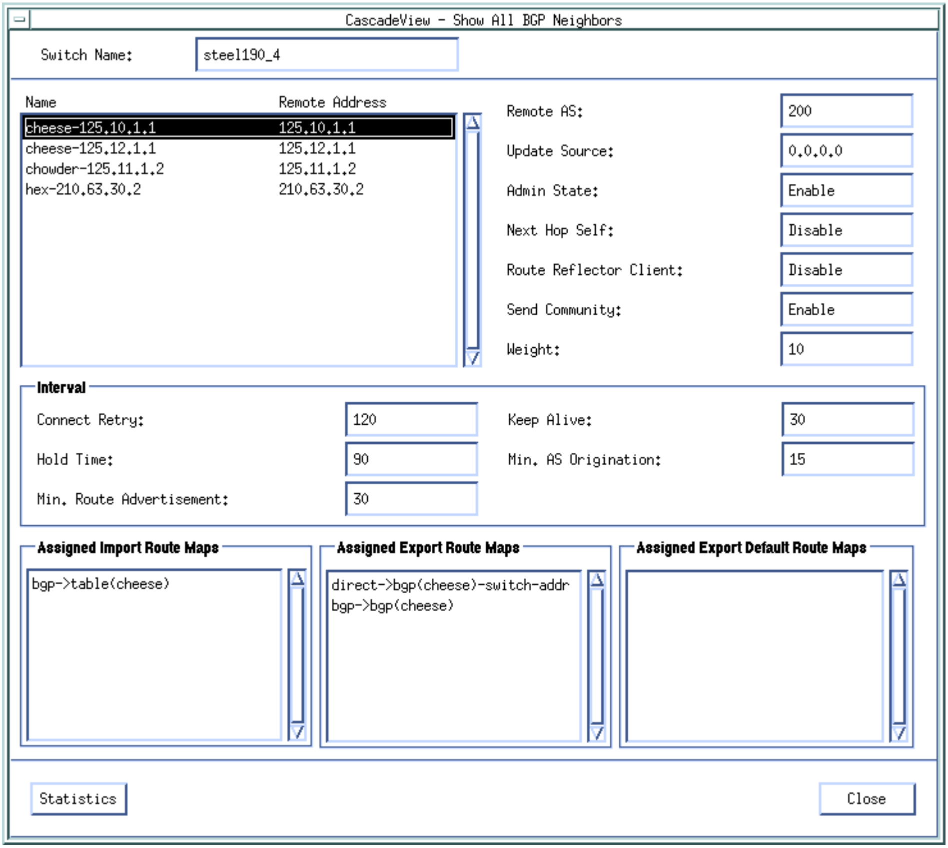

Figure 12-17. Show All BGP Neighbors Dialog Box

Table 8-2 on page 8-14 describes the BGP Neighbor fields.

3. Use the Statistics button to display BGP peer connection statistics. For more information, refer to "Monitoring BGP Connection Statistics" on page 12-28.

4. Choose Close to return to the root map display.

Monitoring BGP Connection Statistics

To view BGP connection statistics:

1. Access the Show All BGP Neighbors dialog box (Figure 12-17) as described in "Monitoring BGP Neighbors" on page 12-27.

2. Choose Statistics. The BGP Peer Connection Statistics dialog box appears (Figure 12-18).

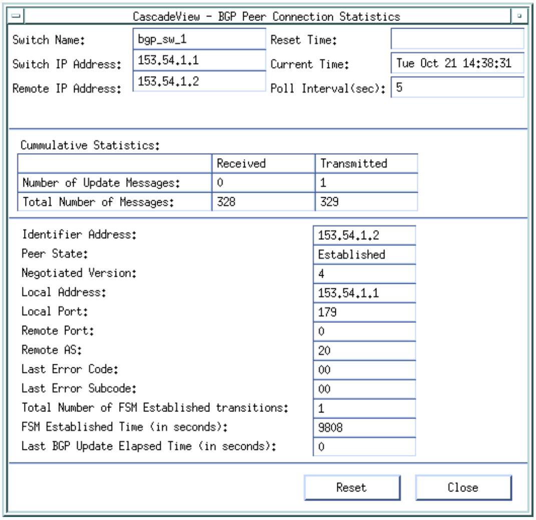

Figure 12-18. BGP Peer Connection Statistics Dialog Box

The BGP Peer Connection Statistics dialog box displays data in separate columns to reflect the transmission and receipt of data on the switch. **Cumulative Statistics** list the number of each type of packet (received, transmitted, or tagged due to an error). Table 12-8 describes each of the BGP Peer Connection Statistics.

Table 12-8. BGP Neighbor Statistics

Field	Description
Switch Name	Displays the name of the switch.
Reset Time	Displays the last time the counters were reset.
Switch IP Address	Displays the IP address of the switch.
Current Time	Displays the current system time.
Remote IP Address	Displays the IP address of the neighbor.
Poll Interval(sec)	Displays the time interval for the collection of statistical data. Refer to the *Network Configuration Guide for B-STDX/STDX* for information about setting the poll interval.
Cumulative Statistics	
Number of Update Messages	Displays the number of BGP update messages this connection transmitted or received.
Total Number of Messages	Displays the total number of messages this connection transmitted to or from a remote peer.
Identifier Address	Displays the 4-byte unsigned integer that indicates the sender's ID.

Table 12-8. BGP Neighbor Statistics (Continued)

Field	Description
Peer State	Displays the BGP peer connection state. Possible values are: *IDLE* — In this state, BGP waits for the operator to initiate a start event. BGP initiates a TCP connection and listens for a connection, which may be initiated by a peer. After BGP is initialized, it advances to the connect state. *CONNECT* — In this state, BGP waits for TCP to complete. If the TCP connection completes, BGP advances to the OPENSENT state. If the TCP connection fails, BGP advances to the ACTIVE state. *ACTIVE* — In this state, BGP tries to establish peering through a TCP connection. If the connection is successful, BGP transmits an OPEN message and advances to the OPENSENT stage. If the TCP connection retry timer expires, BGP restarts the connection timer and retreats to the CONNECT state. *OPENSENT* — In this state, BGP sends an OPEN message to its peer and waits for an OPEN message. Once the message is received, the message is checked for correctness, which includes the router ID, BGP version, AS number, and hold timer. If errors occur, the system sends an error notification message and returns to the IDLE state. If there are no errors, BGP sends keepalive messages. *OPENCONFIRM* — In this state, BGP waits for keepalive messages from its peer. If the message is received, the session advances to the ESTABLISHED state. *ESTABLISHED* — The final state of the BGP connection state.
Negotiated Version	Displays the negotiated version of BGP running between the two peers.
Local Address	Displays the local IP address of this entry’s BGP connection.

Table 12-8. BGP Neighbor Statistics (Continued)

Field	Description
Local Port	Displays the local port that establishes a TCP connection between the BGP peers.
Remote Port	Displays the remote port that establishes a TCP connection between the BGP peers.
Remote AS	Displays the remote AS number.
Last Error Code	Displays the last error code seen by this peer. An error code is part of the notification message which is sent whenever an error occurs. See Table 12-9 for more information on last error codes.
Last Error Subcode	Displays the last error subcode seen by this peer. An error subcode is part of the notification message. The last error subcode provides more specific information about the nature of the error. See Table 12-9 for more information on last error subcodes.
Total Number of FSM Established transitions	Displays the total number of times the BGP Finite State Machine (FSM) transitioned into the established state. The BGP FSM is a process BGP goes through to determine its connection state.
FSM Established Time (in seconds)	Displays the following • How long in seconds this peer has been in the established state. • How long in seconds since this peer was last in the established state. The timer is set to zero when you configure a new peer or you boot the router.
Last BGP Update Elapsed Time (in seconds)	The number of seconds since this peer received the last BGP update message.

Table 12-9. BGP Error Codes

Last Error Code	Last Error Subcode
1 — Message Header Error	1 — Connection Not Synchronized
	2 — Bad Message Length
	3 — Bad Message Type
2 — OPEN Message Error	1 — Unsupported Version Number
	2 — Bad Peer AS
	3 — Bad BGP Identifier
	4 — Unsupported Optional Parameter
	5 — Authentication Failure
	6 — Unacceptable Hold Timer
3 — UPDATE Message Error	1 — Malformed Attribute List
	2 — Unrecognized Well-Known Attribute
	3 — Missing Well-Known Attribute
	4 — Attribute Flags Error
	5 — Attribute Length Error
	6 — Invalid Origin Attribute
	7 — AS Routing Loop
	8 — Invalid NEXT_HOP Attribute
	9 — Optional Attribute Error
	10 — Invalid Network Field
	11 — Malformed AS_path
4 — Hold Timer Expired	NOT applicable

Table 12-9. BGP Error Codes

Last Error Code	Last Error Subcode
5 — Finite State Machine Error (for errors detected by the FSM)	NOT applicable
6 — Cease (for fatal errors besides the ones already listed)	NOT applicable

3. Choose Reset to reset the statistics counters.
4. Choose Close.

Monitoring BGP Aggregates

To view a BGP area aggregate:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All BGP⇒Show All BGP Aggregates. The Show All BGP Aggregates dialog box appears (Figure 12-19).

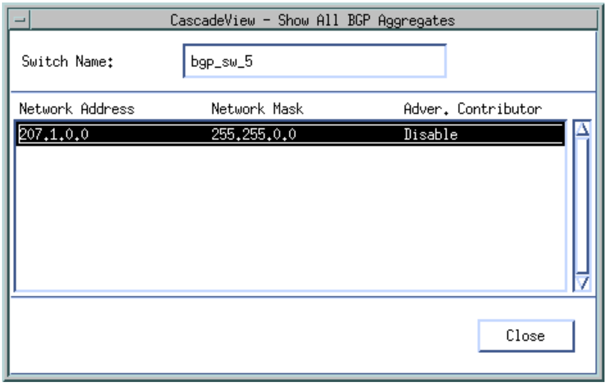

Figure 12-19. Show All BGP Aggregates

Table 8-3 on page 8-22 describes the BGP Aggregate fields.

3. Choose Close to return to the root map display.

Monitoring BGP Path Attributes

BGP attributes are parameters that describe characteristics of a route. The BGP decision process uses these attributes to select the best routes. This section describes BGP switch path attributes.

To monitor a BGP switch's path attributes:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All BGP⇒Show All BGP Path Attributes. The Show All BGP Path Attribute dialog box appears as shown in Figure 12-20.

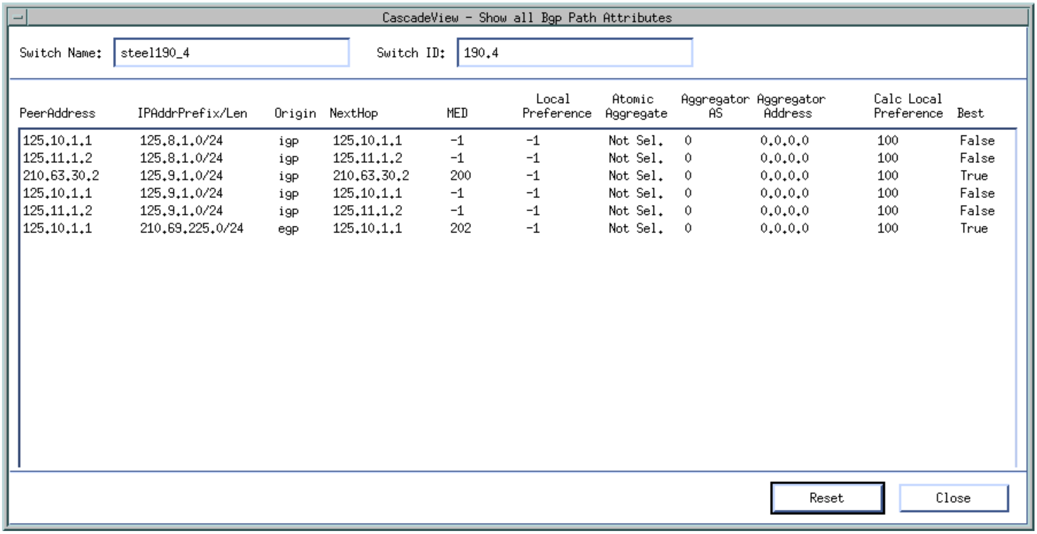

Figure 12-20. Show All BGP Path Attributes Dialog Box

Table 12-10 describes the fields on the Show All BGP Path Attributes dialog box.

Table 12-10. Show All BGP Path Attributes Fields

Field	Description
Peer Address	Displays the IP address of the peer that sent the route information.
IP.Prefix	Displays the IP address prefix in the network layer reachability information field.
IP.Pref Len	Displays the length in bits of the IP address prefix in the network reachability information field. Possible values are between 0 and 32.
Origin	Displays the origin of the path information. Possible values are: • *igp* — Networks learned through interior gateway protocol • *egp* — Networks learned through exterior gateway protocol • *incomplete (3)* — undetermined • *none*
NextHop	Displays the IP address that identifies the next hop to reach a network.
MED	Displays 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 that advertises a MED value of 120 would be preferred over a route with a MED value of 200.
Local Preference	Displays the original BGP speaker’s preference for an advertised route. Possible values are between -1 and 2147483647. A higher local preference indicates the route is preferred. A value of -1 indicates that this attribute does not apply.

Table 12-10. Show All BGP Path Attributes Fields (Continued)

Field	Description
Atomic Aggregate	The BGP speaker uses the atomic aggregate to inform other BGP speakers that the local system selected a less specific route without selecting a more specific route. Possible values are: • Not Selected • Selected
Aggregator AS	Displays the AS number of the last BGP speaker that performed route aggregation. A value of 0 indicates this attribute does not apply.
Aggregator Address	Displays the IP address of the last BGP speaker that performed route aggregation. A value of 0.0.0.0 indicates this attribute does not apply. For more information, see "BGP Aggregates" on page 8-6.
Calc Local Pref.	Displays the calculated local preference for the specified path. This value may be different from the local preference for the specified path. This parameter is configured when you define route maps.
Best	Indicates whether this route was chosen as the best BGP route. Possible values are: *False* — not chosen as best route *True* — chosen as best route

3. Choose Reset to clear each of the statistics counters.
4. Choose Close to return to the root map display.

Monitoring Filters, Access Lists, and Route Maps

This section describes how to monitor the following IP objects:

- Network filters
- Network access lists
- Route maps

Monitoring Network Filters

To view the network filters defined for a node:

1. On the network map, select the appropriate switch object.
2. From the Monitor menu, select Cascade IP Objects ⇒Show Route Policies ⇒Show Network Filters. The Show All Network Filters dialog box appears (Figure 12-21).

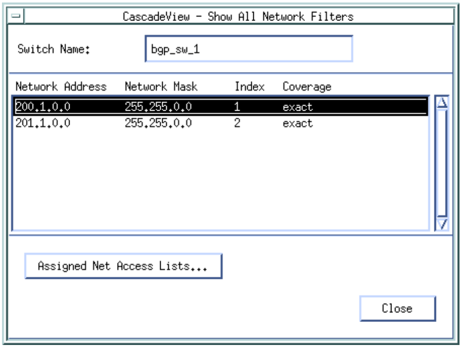

Figure 12-21. Show All Network Filters Dialog Box

Table 11-2 on page 11-11 describes the Network Filter fields.

3. Choose Close to return to the network map display.

Monitoring Assigned Network Access Lists

To view the network access lists assigned to a selected network filter:

1. Choose Assigned Net Access Lists from the Show All Network Filters dialog box shown in Figure 12-21.
2. The Network Access Lists Using a Network Filter dialog box (Figure 12-22) appears.

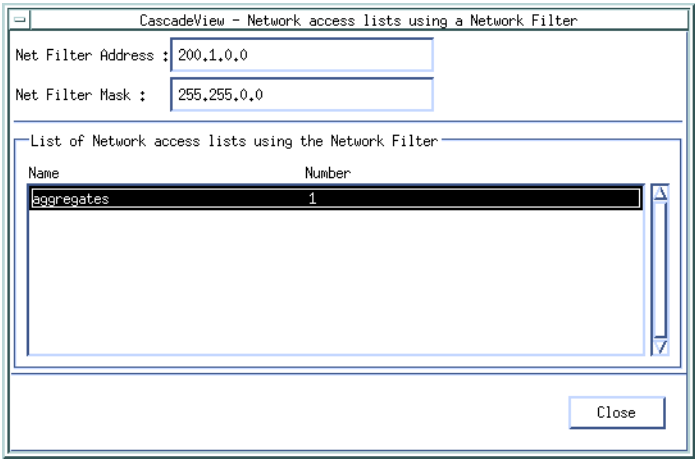

Figure 12-22. Network Access Lists Using a Network Filter Dialog Box

3. Choose Close to return to the network map display.

Monitoring Network Access Lists

To view the network access lists defined for a node:

1. On the network map, select the appropriate switch object.

2. From the Monitor menu, select Cascade IP Objects ⇒Show All Route Policies ⇒Show All Network Access Lists. The Show All Network Access Lists dialog box appears (Figure 12-23).

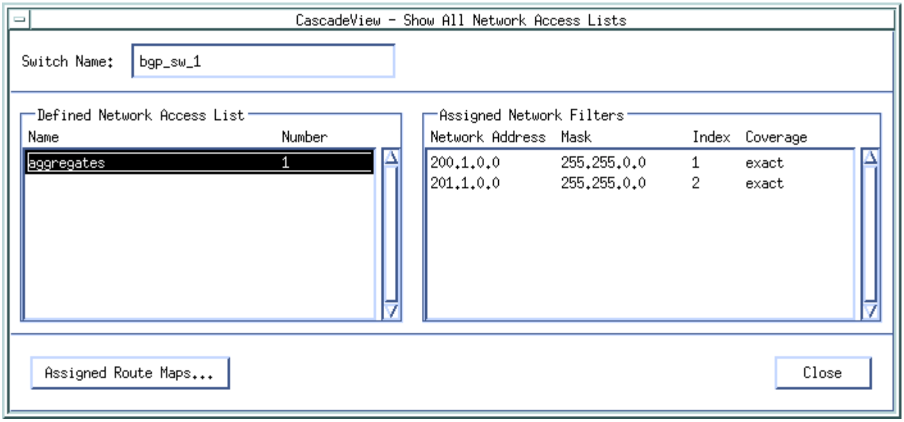

Figure 12-23. Show All Network Access Lists Dialog Box

Table 11-4 on page 11-14 describes of each of the Network Access List fields.

3. Choose Close to return to the network map display.

Monitoring Assigned Route Maps

To view the route maps assigned to a selected network access list:

1. Choose Assigned Route Maps from the Show All Network Access Lists dialog box shown in Figure 12-23.

2. The Route Maps Using Network Access List dialog box appears (Figure 12-24).

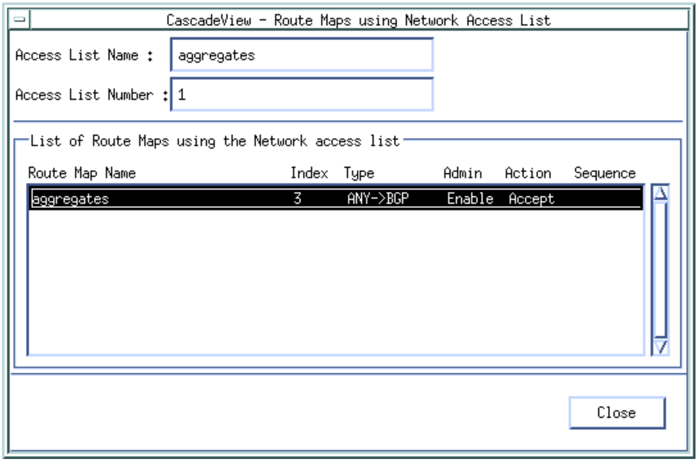

Figure 12-24. Route Maps Using Network Access List Dialog Box

3. Choose Close to return to the network map display.

Monitoring Route Maps

To view the Route Maps defined for a node:

1. On the network map, select the appropriate switch object.
2. From the Monitor menu, select Cascade IP Objects ⇒Show All Route Policies ⇒Show All Route Maps. The Show All Route Maps dialog box appears (Figure 12-25).

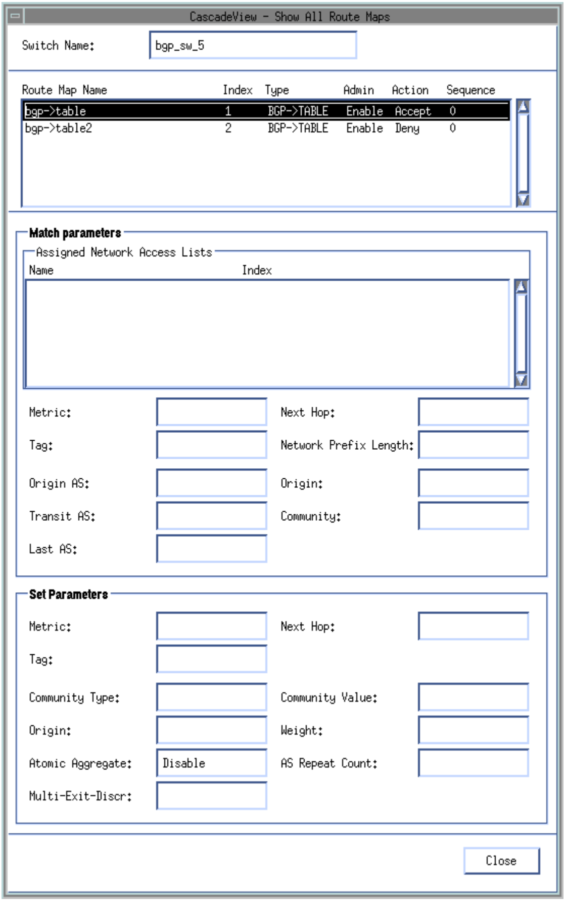

Figure 12-25. Show All Route Maps Dialog Box

Table 11-6 on page 11-16 describes each of the common fields shown at the top of the Show All Route Maps dialog box.

Note that the Match parameters and Set parameters vary depending on the type of route map that you are defining. Refer to Table 11-9 for information about the Match and Set parameters for each route map type.

3. Choose Close to return to the network map display.

Monitoring IP Packet Filters

This section describes how to monitor the following:

- IP packet filters
- IP packet filters assigned to logical ports
- IP packet filters assigned to hosts
- IP packet filters assigned to circuits

Monitoring IP Packet Filters

To view IP Packet filter parameters:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All Packet Filters⇒Show All Packet Filters. The Show All Packet Filters dialog box appears (Figure 12-26).

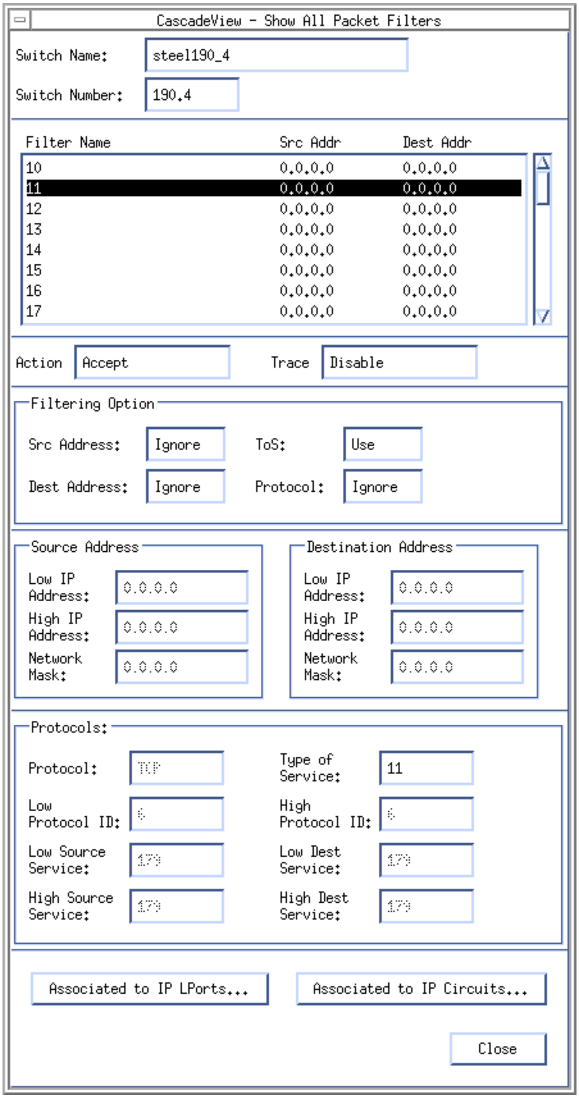

Figure 12-26. Show All Packet Filters Dialog Box

Table 4-1 on page 4-8 describes all packet filter fields. For descriptions of the Associated to IP Lports button and Associated to IP Circuits button, see page 4-6.

3. Choose Close to return to the root map display.

Monitoring IP Packet Filters Assigned to Logical Ports

To view IP packet filters assigned to a logical port:

1. From the network map select the appropriate switch object.

2. From the Monitor menu, select Cascade IP objects⇒Show All Packet Filter⇒Show All Logical Port Filters. The Show All Logical Port Filters dialog box appears (Figure 12-27).

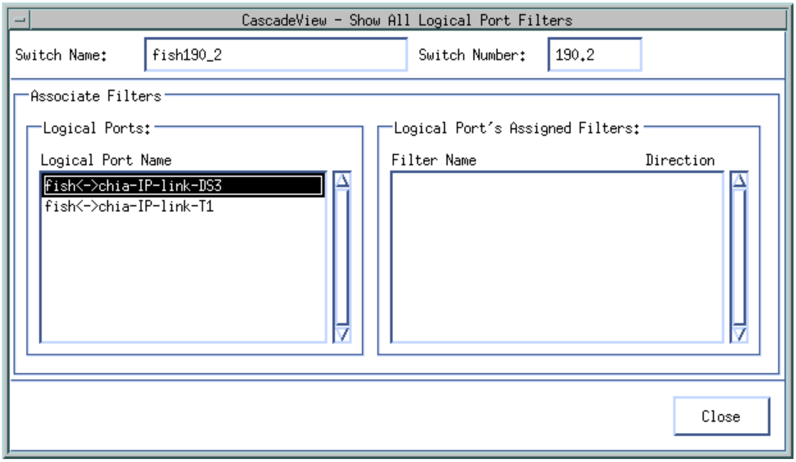

Figure 12-27. Show All Logical Port Filters Dialog Box

Table 4-2 on page 4-22 describes Logical Port IP Filter fields.

3. Choose Close to return to the root map display.

Monitoring IP Packet Filters Assigned to Hosts

To view IP packet filters assigned to hosts:

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

2. From the Monitor menu, select Cascade IP objects⇒Show All Packet Filter⇒Show All Host Filters. The Show All Host Filters dialog box appears (Figure 12-28).

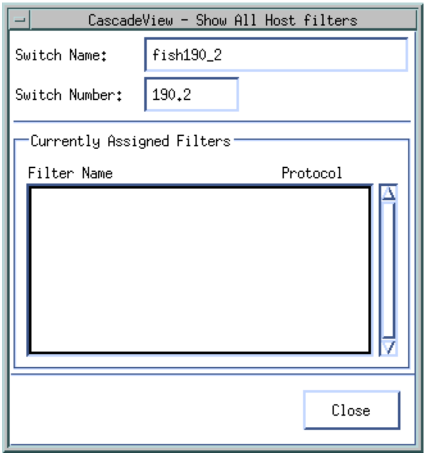

Figure 12-28. Show All Host Filters Dialog Box

Table 4-3 on page 4-26 describes the Show All Host Filters fields.

3. Choose Close to return to the Root Map display.

Monitoring IP Packet Filters Assigned to Circuits

To view IP packet filters assigned to circuits:

1. From the network map, select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All Packet Filter⇒Show All IP Circuit Filters. The Show All IP Circuit Filters dialog box appears (Figure 12-29).

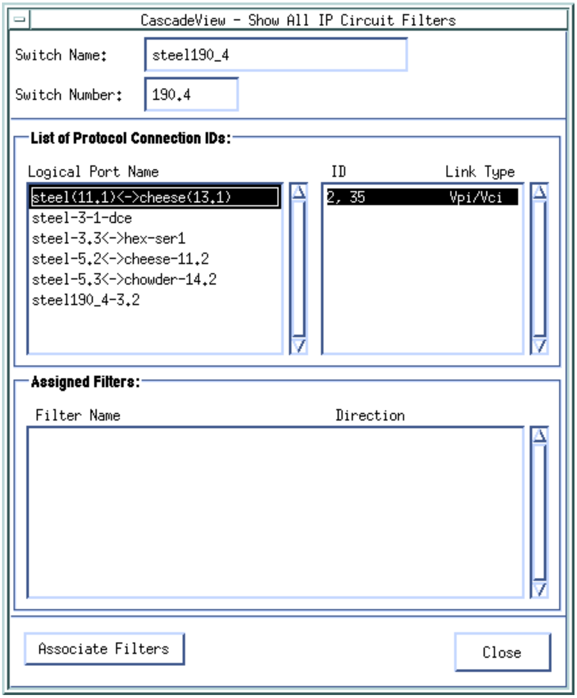

Figure 12-29. Show All IP Circuit Filters Dialog Box

Table 4-3 on page 4-26 describes the Show All IP Circuit Filters dialog box fields.

3. Choose Close.

Monitoring QoS Profiles

This section describes how to monitor:

- QoS flow profiles
- QoS filters assigned to logical ports
- QoS filter statistics

Monitoring QoS Flow Profiles

To monitor IP QoS PVC flow profiles:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All QoS Profiles⇒Show All QoS Profiles. The Show All QoS Profiles dialog box appears (Figure 12-30).

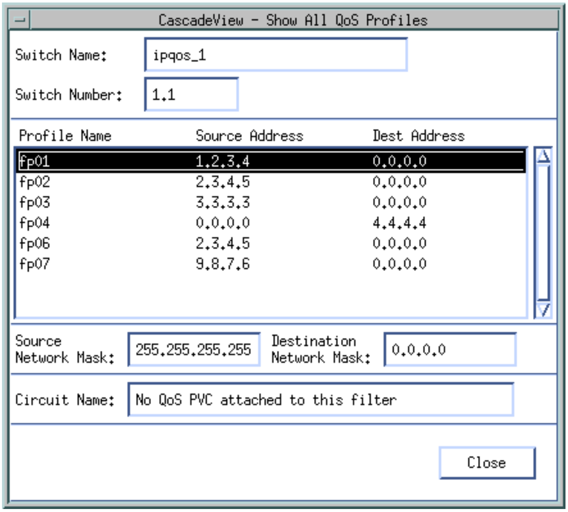

Figure 12-30. Show QoS Flow Profile Dialog Box

Table 5-6 on page 5-20 describes QoS Profile fields.

3. Choose Close to return to the root map display.

Monitoring IP QoS Profiles Assigned to a Logical Port

To view IP QoS profiles assigned to a logical port:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All QoS Profiles⇒Show All Logical Port QoS Profiles. The Show All Logical Port QoS Profiles dialog box appears (Figure 12-31).

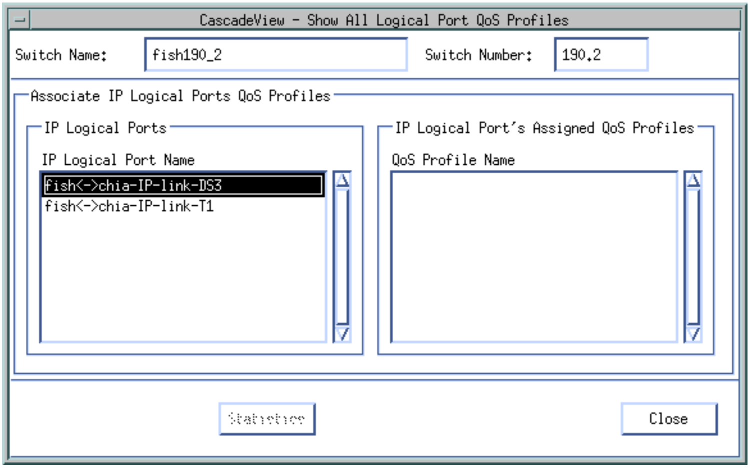

Figure 12-31. Show All Logical Port QoS Profiles Dialog Box

Table 5-6 on page 5-20 describes IP QoS Filter fields.

3. Choose the Statistics button to display LPort QoS filter statistics. For more information, refer to "Monitoring IP QoS Filter Statistics" on page 12-51.
4. Choose Close to return to the root map display.

Monitoring IP QoS Filter Statistics

To view IP QoS filter statistics:

1. Access the Show IP QoS Filter dialog box (Figure 12-32) as described in "Monitoring IP QoS Profiles Assigned to a Logical Port".

2. Choose Statistics. The LPort QoS Filter Statistics dialog box appears (Figure 12-32).

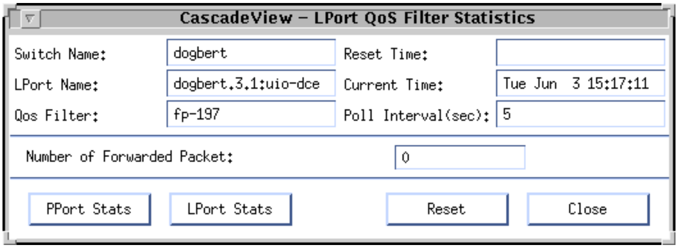

Figure 12-32. LPort QoS Filter Statistics Dialog Box

Table 12-11 describes LPort QoS Filter Statistics fields.

Table 12-11. LPort QoS Filter Statistics

Field	Description
Identifying Fields	
Switch Name	Displays the name of the switch.
Reset Time	Displays the last time the counters were reset.
Lport Name	Displays the name of the logical port.
Current Time	Displays the current system time.
QoS Filter	Displays the name of the QoS profile.
Poll Interval(sec)	The time interval for the collection statistics. Refer to the *Network Configuration Guide for B-STDX/STDX* for information about setting the poll interval.
Number of Forwarded Packet	Displays the number of packets forwarded over the IP QoS PVC that is associated with the IP QoS profile and assigned logical port

3. Choose Reset to clear the statistics counters.
4. Choose PPort Stats to display the physical port statistics and choose LPort Stats to display the logical port summary statistics. For more information, refer to the *Diagnostic and Troubleshooting Guide for B-STDX/STDX*.
5. Choose Close.

Monitoring Static Routes

To view static route parameters:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All Static Routes. The Show All Static Route dialog box appears (Figure 12-33).

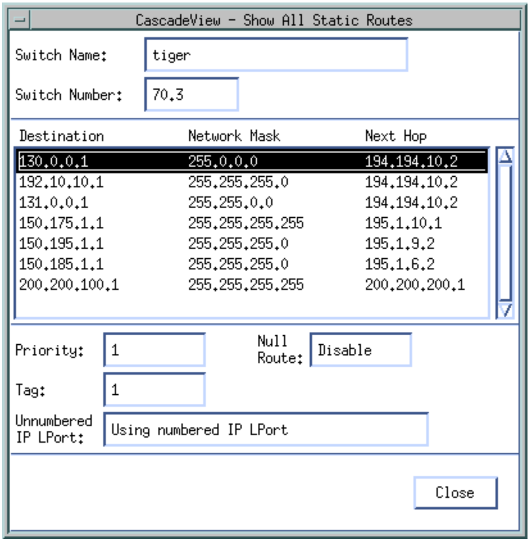

Figure 12-33. Static Route Dialog Box

Table 10-1 on page 10-4 describes Static Route fields.

3. Choose Close to return to the root map display.

Monitoring Static ARP Parameters

To view ARP parameters:

1. From the network map select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show All Static ARP Entries. The Show All Static ARP Entries dialog box appears (Figure 12-34).

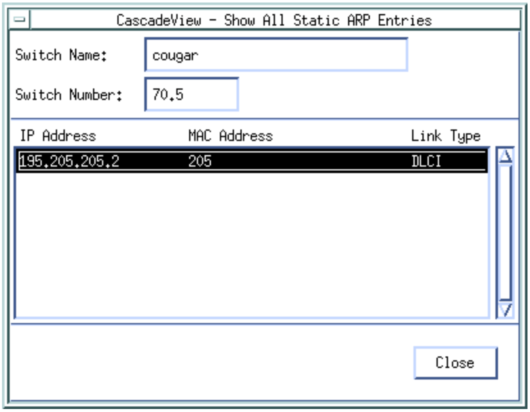

Figure 12-34. Show All Static ARP Entries Dialog Box

Table 6-1 on page 6-4 describes ARP fields.

3. Choose Close to return to the root map display.

Monitoring MPT Path

The reverse multipoint tunnel (MPT) enables IP Navigator to switch connectionless protocols (e.g. IP) across the Ascend switch network. MPTs provide an efficient, fault tolerant, high performance protocol switching layer that is scalable to 400 switches in a network. MPTs run across CBX 500s and B-STDX 9000s, which run over direct or optimum trunks.

The reason for using MPTs over point-to-point tunnels is because MPTs need less circuits than point-to-point tunnels. MPTs 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.

Figure 12-35 displays an example of an MPT.

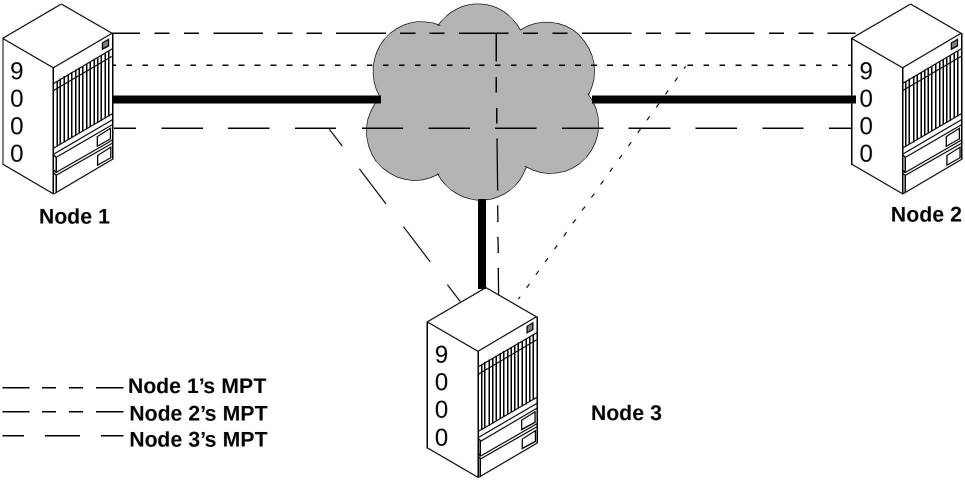

Figure 12-35. MPT Example

In Figure 12-35, the solid lines represent trunks into the Ascend network. Node 1 establishes its MPT first, followed by nodes 2 and 3. When node 1 establishes its MPT, nodes 2 and 3 can forward packets back to node 1, thereby reversing the direction of the MPT. The same idea applies to nodes 2 and 3 when they establish their MPTs. The number of circuits is equal to the number of nodes.

Every switch maintains one MPT circuit network and contains a root, which

- keeps track of MPT nodes
- adds and deletes nodes
- keeps nodes alive

A root is a standard VC_ENTRY that is created at initialization time on every CP card in the 9000 and every SP card in the 500.

Switch Domains

There are two types of switch domains:

Cell Domain — Paths that traverse direct ATM trunks.

Frame Domain — Paths that traverse pure direct frame trunks.

A switch can belong to multiple domains, however, the domains must be 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.

OSPF determines how MPTs connect two switches in different domains. The following factors apply when determining MPTs:

- A switch that only belongs to one domain cannot add a switch from a different domain to its MPT. MPTs are only established between switches in the same domain.
- If the shortest path between two switches in the same domain traverses a different domain, the switches cannot add each other to their MPTs.

MPTs use cell and frame domains to address limitations in switching ATM cells. Reassembling cells on these boundary switches require more resources than currently available on Ascend cell cards. However, a compromise can be reached, whereby large networks can reassemble cells but with additional processing required for data switched across cell/frame cards.

To view MPT path parameters:

1. From the network map, select the appropriate switch object.
2. From the Monitor menu, select Cascade IP objects⇒Show MPT Path. The Show MPT Path dialog box appears (Figure 12-36).

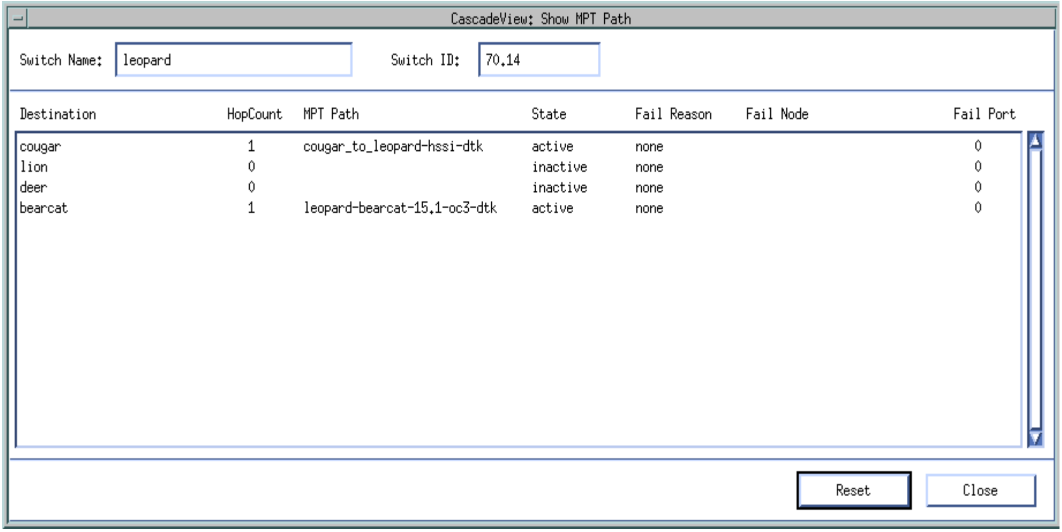

Figure 12-36. Show MPT Path Dialog Box

Table 12-12 describes MPT Path Fields.

Table 12-12. Show MPT Path Fields

Field	Description
Switch Name	Displays the name of the switch.
Switch ID	Displays the switch ID.
Destination	Displays the MPT destination node's switch name.
Hop Count	Displays the MPT destination node path's hop counts.

Table 12-12. Show MPT Path Fields (Continued)

Field	Description
MPT Path	Displays the trunk name that represents the MPT circuit path.
State	Displays the current state of the MPT destination node path. Possible values are: *retry* — The MPT process waits until the local timer expires to call. *calling* — The MPT process waits for the local timer of a message, which rejects or confirms a call. *active* — The MPT process successfully calls and waits for keep-alive events. *inactive* — The MPT process waits for the grooming process (see *grooming* on page 12-59), which activates the MPT leaf node. *wcinact* — The MPT process clears when the call is terminated. *wcdel* — The MPT process calls, then deleted when the calling is over. ***Note:*** *The MPT process maintains the state of each end-to-end path that is sent by OSPF.*

Table 12-12. Show MPT Path Fields (Continued)

Field	Description
Fail Reason	Displays the reason why the MPT path failed. Possible values are *none* — Problems were not encountered. *tpcalling* — An mptTport calls. An mptTport is a data record that describes an MPT ingress port at any given merge point. *vcalling* — A VC_ENTRY calls. *tpdead* —The connection is dead. *routelookup* —The route lookup failed. *confirmtimeout* —The confirm timer expired. *pathclear* — OSPF notifies MPT that the path is no longer preferred. *trunkdown* — The trunk went down. *dead* — Hello packets are no longer received. *grooming* — A better path exists in the network. *pathregister* —The path was not registered with OSPF. *impurepath* — A shorter, mixed cell or frame path exists. *rvcdied* —The RVC died. An RVC is a data structure that reassembles frames into ATM cells.
Fail Node	Displays the last switch associated with the failed MPT path.
Fail Port	Displays the interface # on the failed node that is associated with the failed MPT path.

3. Choose Reset to clear each of the statistics counters.

4. Choose Close to return to the root map display.

Monitoring RIP

Monitoring RIP Peer Information

The B-STDX switch communicates with a RIP peer, such as a remote host or router via the RIP protocol.

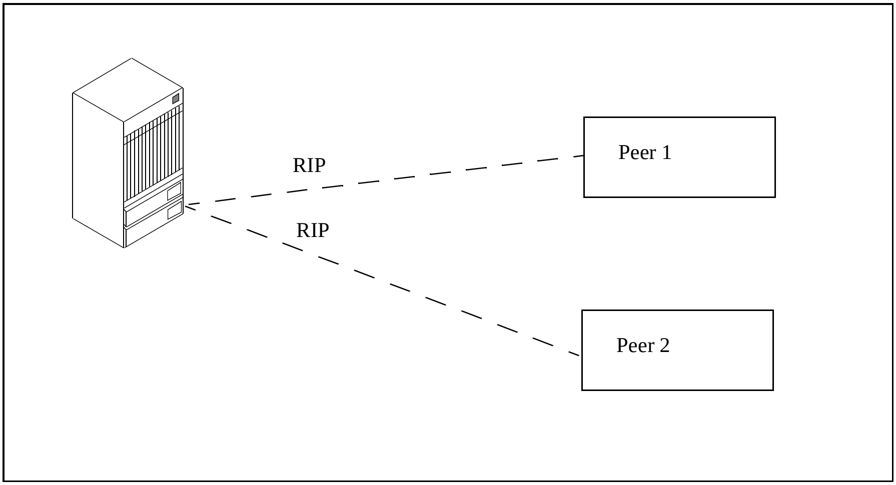

Figure 12-37. RIP Peers

To view the RIP peer information:

1. On the network map, select the appropriate switch object.

2. From the Monitor menu, select Cascade IP Objects ⇒Show Rip2 ⇒Show Rip2 Peers. The Rip Active Peer Table dialog box appears (Figure 12-38).

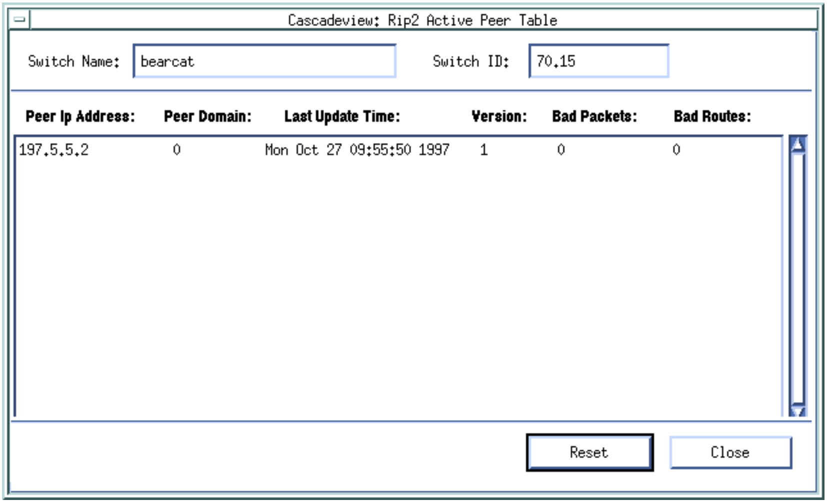

Figure 12-38. RIP Active Peer Table Dialog Box

Table 12-4 describes the RIP Active Peer Table fields.

Table 12-13. RIP Active Peer Table Fields

Field	Description
Peer IP Address	Displays the IP address that the peer uses as its source address. On an unnumbered interface, this address may not be a member of any subnet on the system.
Peer Domain	Displays the value in the Routing Domain field in RIP packets received from the peer. As domain support is lowered, this must be zero.
Last Update Time	Displays the value of *sysUpTime* when the most recent RIP update was last received.
Version	Displays the RIP version number in the header of the last RIP packet received.
Bad Packets	Displays the number of RIP response packets from this peer that were discarded as invalid.

Table 12-13. RIP Active Peer Table Fields (Continued)

Field	Description
Bad Routes	Displays the number of routes from this peer that were ignored because the entry format was invalid.

3. Choose Reset to clear each of the statistics counters.

4. Choose Close to return to the previous screen.

Monitoring RIP Global Counters

To view the RIP peer information:

1. On the network map, select the appropriate switch object.

2. From the Monitor menu, select Cascade IP Objects ⇒Show Rip2 ⇒Show Rip2 Global Counters. The Rip2 Global Counters dialog box appears (Table 12-39).

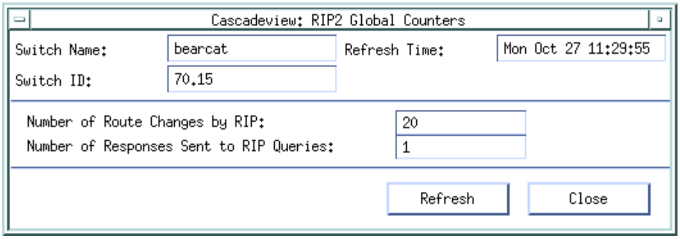

Figure 12-39. RIP2 Global Counters Dialog Box

3. Choose Close to return to the network map.

Monitoring the IP Routing Table

To view the IP routing table:

1. From the network map, select the appropriate switch.
2. From the Monitor menu, select Cascade IP objects⇒Show IP Routing Table. The Show IP Routing Table dialog box appears (Figure 12-40).

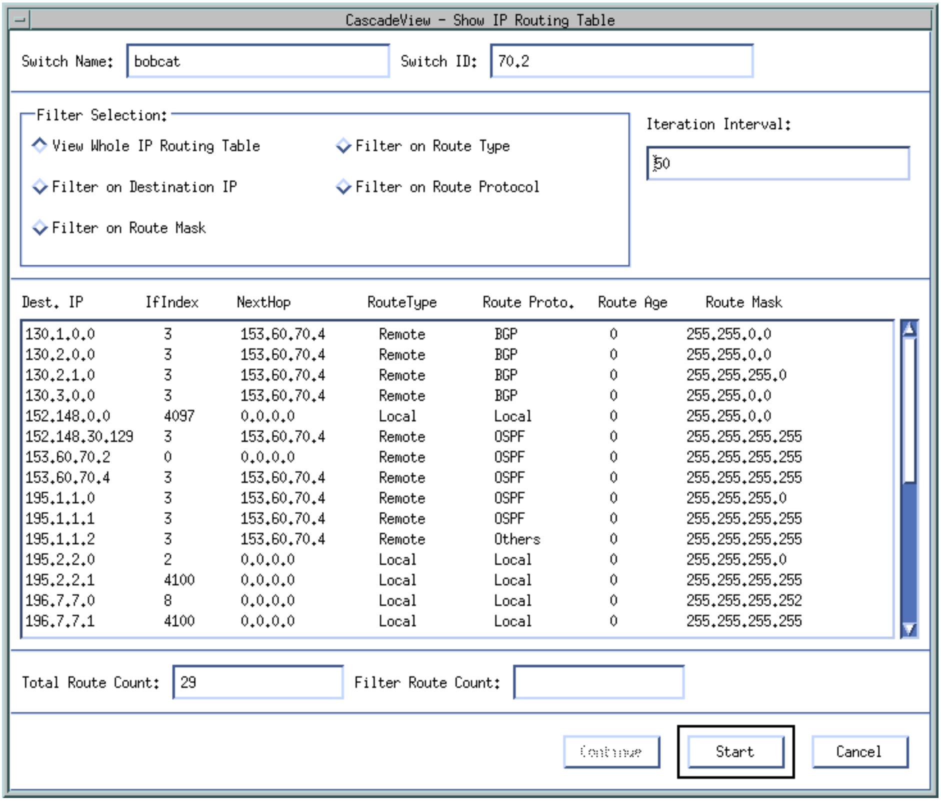

Figure 12-40. Show IP Routing Table Dialog Box

Table 12-14 describes each field in the Filter Selection box. The fields in the filter selection box enable you to view particular routes if the routing table is too large. You can select one of the filter selections to decrease the size of the routing table. This enables you to view routes based on a particular parameter. For example, if you select Filter on Destination IP, you only view the destination IP portion of the routing table. The *View Whole IP Routing Table* filter selection (refer to Table 12-15 on page 12-65) displays all available route table information.

Table 12-14. Filter Selection Fields

Field	Description
View Whole IP Routing Table	Enables you to view the whole routing table.
Filter on Destination IP	Enables you to view destination IP address information in the routing table.
Filter on Route Mask	Enables you to view route mask information in the routing table.
Filter on Route Type	Enables you to view route type information in the routing table.
Filter on Route Protocol	Enables you to view route protocol information in the routing table.

View Whole IP Routing Table

1. Select View Whole IP Routing Table in the Filter Selection box.

   The Show IP Routing Table dialog box displays the fields (Figure 12-40). Table 12-15 describes each of the View Whole IP Routing Table fields.

Table 12-15. View Whole IP Routing Table Fields

Field	Description
Switch Name	Displays the name of the switch.
Switch ID	Displays the ID of the switch.
Iteration Interval	Enter the number of routes you want to view. The default is 50.
Dest. IP	Displays the IP network that is pointed to by this route entry.
IfIndex	Displays the value that identifies the local interface where the next hop is reached.
NextHop	Displays the address of the next system en route on remote routes. Otherwise, the value is 0.0.0.0.
RouteType	Displays the route type. Options include: *Local* — Refers to a route where the next hop is the final destination. *Remote* — Refers to a route where the next hop is not the final destination. *Invalid* — Refers to the entry that is not used to forward IP packets.
Route Proto.	Displays the routing protocol that indicates how the route was learned.
Route Age	Displays the number of seconds since this route was updated or determined to be correct.

Table 12-15. View Whole IP Routing Table Fields (Continued)

Field	Description
Route Mask	Displays the mask that applies to the corresponding entry in the destination IP column.
Total Route Count	Displays the total routes counter.
Filter Route Count	Displays the filter routes counter.
Starting IP Address (with Filter on Destination IP selected)	Enter the starting IP address. When you enter this value, you specify a range between this value and the ending IP address. The range filters out IP addresses that are not in the specified range.
Ending IP Address (with Filter on Destination IP selected)	Enter the ending IP address. When you enter this value, you specify a range between the starting IP address and this value. The range filters out IP addresses that are not in the specified range.
Mask Value (with Filter on Route Mask selected)	Enter a value between 1 and 32. This value filters on mask values between 1 and 32 bits.
Filtering on Routing Type (with Filter on Route Type selected)	Enables you to filter based on the route type. Select one of the following: *Other Route* — not specified by this MIB *Reject Route* — rejected routes *Local Route* — locally configured networks *Remote* — routes learned through routing protocols (e.g. BGP, OSPF, RIP) For more information on filtering on routing type, refer to the *Route Type* description on page 12-65.

Table 12-15. View Whole IP Routing Table Fields (Continued)

Field	Description
Filtering on Routing Protocol (with Filter on Route Protocol selected)	Enables you to filter based on route protocol. Select one of the following protocols: *Local* — not specified *Network Management* — static route *ICMP* — result of ICMP Redirect *EGP* — Exterior Gateway Protocol *GGP* — Gateway-Gateway Protocol *Hello* — FuzzBall HelloSpeak *RIP* — Berkely RIP or RIP-II *IS-IS* — Dual IS-IS *ES-IS* — ISO 9542 *ciscoIgrp* — Cisco IGRP *bbnSpfIgp* — BBN SPF IGP *OSPF*— Open Shortest Path First *BGP* — Bor*der* Gateway Protocol Other

2. Choose the Start button to monitor the routes.
3. Choose the Continue button to continue monitoring the routes.
4. Choose Cancel to return to the root map display.

Filter On Destination IP

1. Select Filter On Destination IP in the Filter Selection box.

   The Show IP Routing Table dialog box displays the fields shown in Figure 12-41. Table 12-15 describes each of the Filter on Destination IP fields.

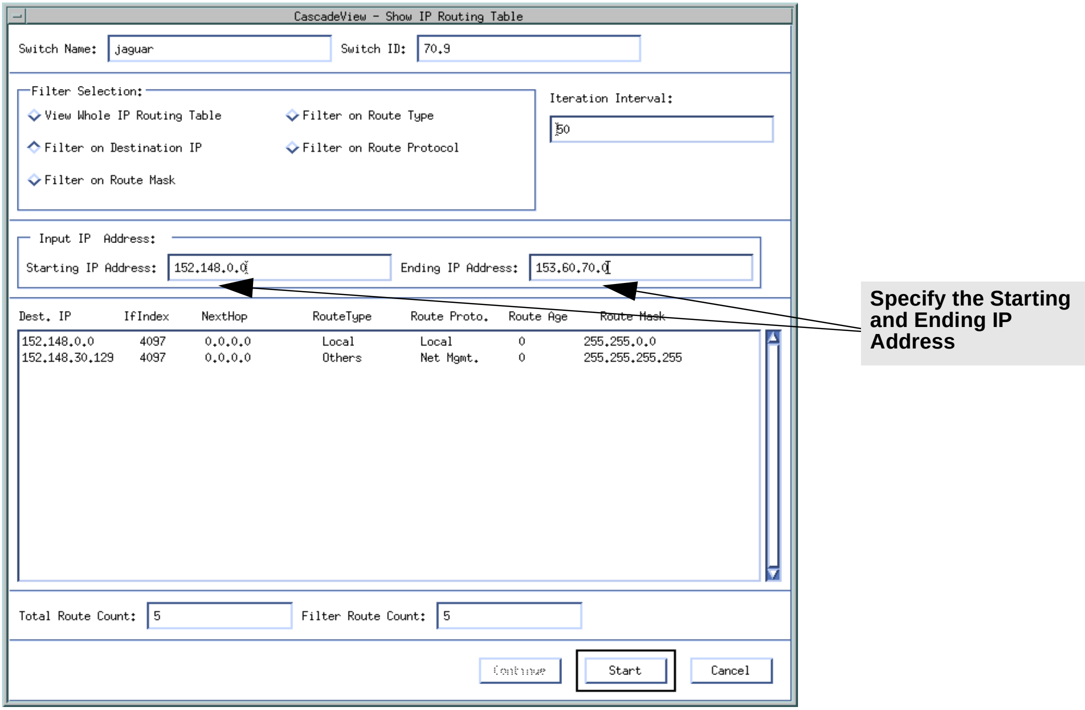

Figure 12-41. Filter On Destination IP Dialog Box

2. In the Iteration Interval field, specify the number of routes you want to view.
3. In the Input IP Address field, enter the starting IP address and the ending IP address.
4. Choose the Start button to monitor the routes.
5. Choose the Continue button to continue monitoring the routes.
6. Choose Cancel to return to the root map display.

Filter On Route Mask

1. Select Filter on Route Mask in the Filter Selection box.

   The Show IP Routing Table dialog box displays the fields shown in Figure 12-42. Table 12-15 describes each of the Filter on Route Mask fields.

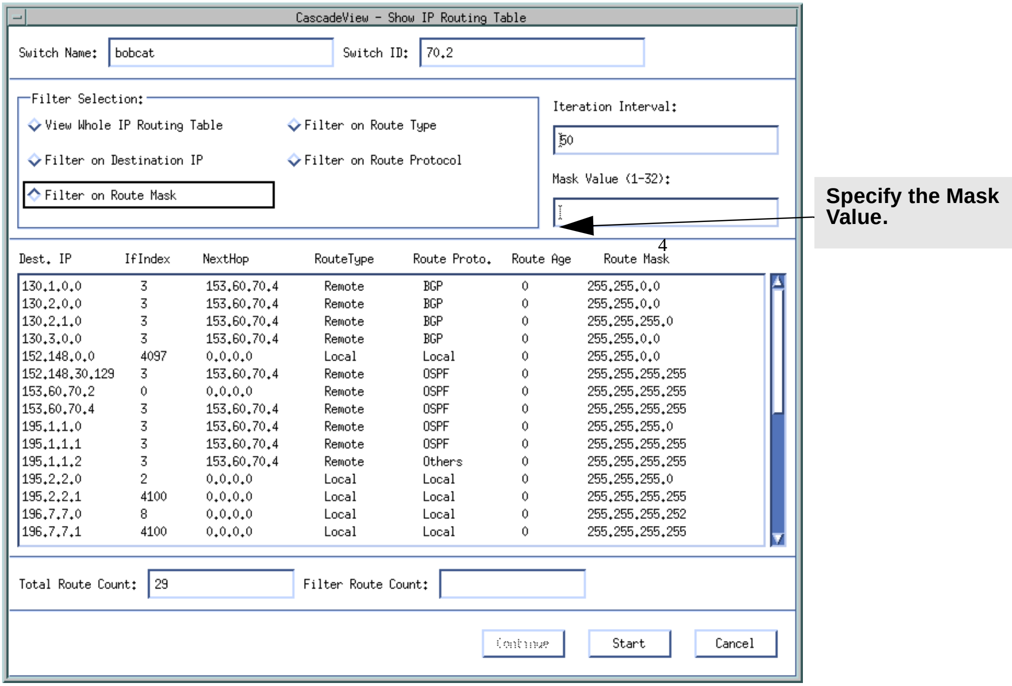

Figure 12-42. Filter On Route Mask Dialog Box

2. In the Iteration Interval field, specify the number of routes you want to view.
3. In the Mask Value field, enter the mask value.
4. Choose the Start button to monitor the routes.
5. Choose the Continue button to continue monitoring the routes.
6. Choose Cancel to return to the root map display.

Filter on Route Type

1. Select Filter on Route Type in the Filter Selection box.

   The Show IP Routing Table dialog box displays the fields shown in Figure 12-43. Table 12-15 describes each of the Filter on Route Type fields.

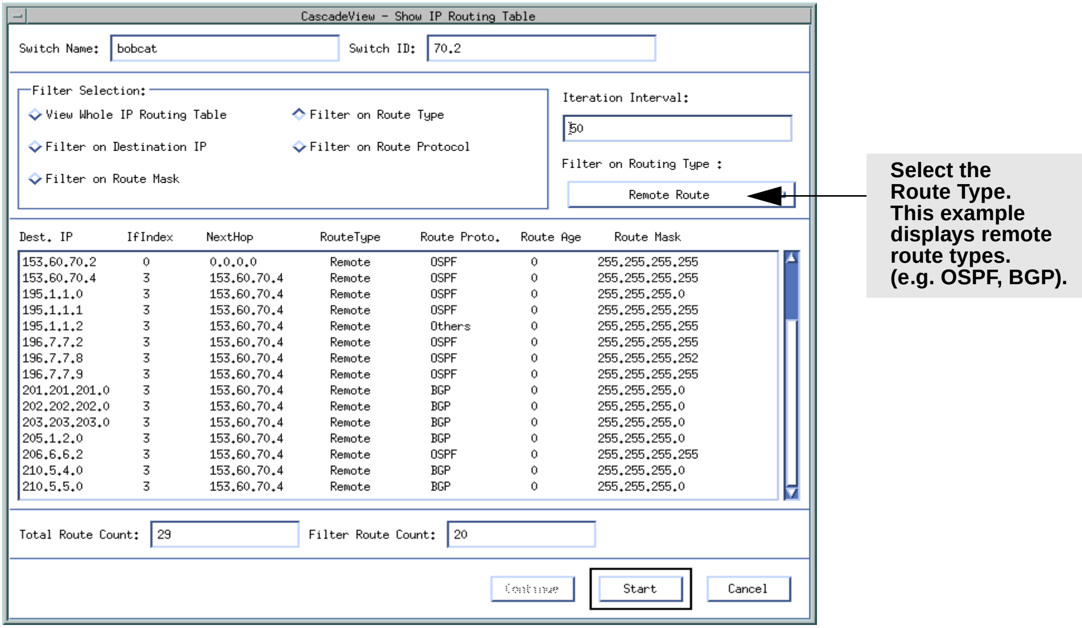

Figure 12-43. Filter On Route Type Dialog Box

2. In the Iteration Interval field, specify the number of routes you want to view.

3. In the Filter on Route Type field, select the route type.

4. Choose the Start button to monitor the routes.

5. Choose the Continue button to continue monitoring the routes.

6. Choose Cancel to return to the root map display.

Filter on Route Protocol

1. Select Filter on Route Protocol in the Filter Selection box.

   The Show IP Routing Table dialog box displays the fields shown in Figure 12-44. Table 12-15 describes each of the Filter on Route Protocol fields.

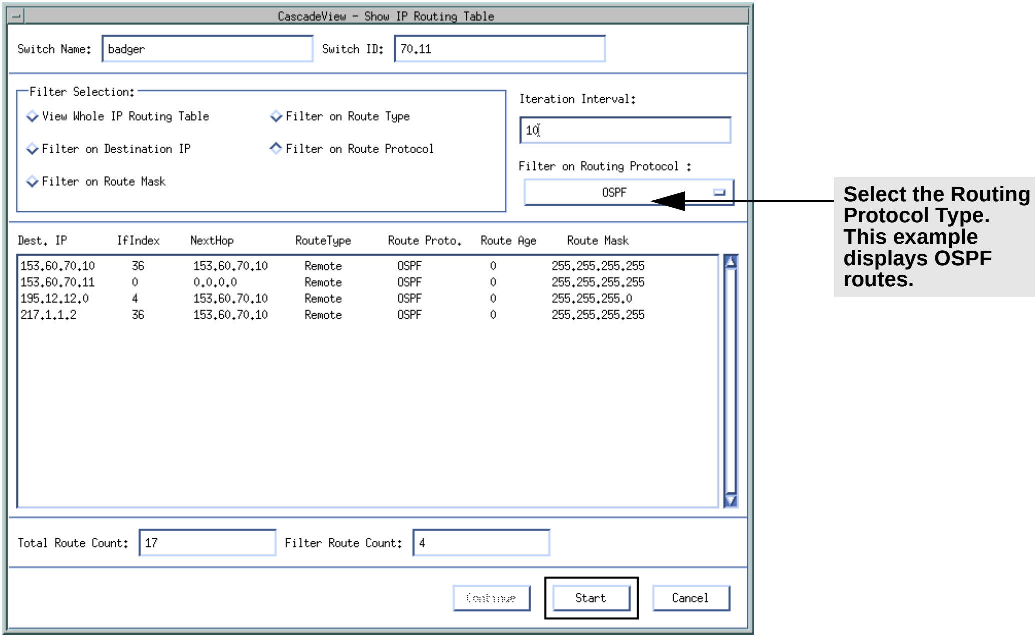

Figure 12-44. Filter On Route Protocol Dialog Box

2. In the Iteration Interval field, specify the number of routes you want to view.

3. In the Filter on Routing Protocol field, select the routing protocol type.

4. Choose the Start button to monitor the routes.

5. Choose the Continue button to continue monitoring the routes.

6. Choose Cancel to return to the root map display.

A

CLI Commands

This appendix provides a listing of the Command Line Interface (CLI) commands that you can enter to perform various tasks on the switch, or to obtain information from the switch. This section describes commands for the B-STDX 8000/9000.

The console remembers the last ten commands that you enter. ^B can be used to recall previous commands as needed.

Conventions

This appendix uses the following conventions to emphasize certain information.

Table A-1. Conventions

Convention	Indicates	Example
Angle brackets <>	user supplied parameters	*<IP address>*
Square brackets []	multiple choice parameters	[enable | disable]

Command Line Interface Help

To determine the correct syntax for a command, type a question mark (?) after the command. For example

```
> show mpt ?

show mpt [statistics|all|path]
```

In addition, CLI provides a brief description of each command.

```
> show bgp ?

  aggregate      Route aggregates

  neighbor       BGP neighbor

  route          Route AS path

  summary        BGP summary
```

CLI Command Summary

This table provides a summary of CLI commands.

Table A-2. Command Line Interface

Commands	Descriptions
get <*oid string*>	SNMP get
next <*oid string*>	SNMP next
set bgp aggregate <*aggregate address*>	Enables you to specify a BGP route summary(page A-17)
set bgp compare_med [enable | disable]	Enables MED comparison (page A-17)
set bgp as <*as number*>	Enables you to specify the local AS (page A-17)
set bgp always_compare_send_bgp_nets [enable | disable]	Enables you to allow BGP networks to be advertised (page A-17)
set bgp client_to_client [enable | disable]	Enables you to set client to client reflection for the peer (page A-18)
set bgp cluster_id <*cluster id*>	Enables you to specify a BGP cluster ID (page A-18)
set bgp def_local_preference <*value*>	Enables you to specify a value for the peer's default local preference (page A-18)
set bgp neighbor <*IP address*> create	Enables you to create a BGP peer instance (page A-18)
set bgp neighbor <*IP address*> delete	Enables you to delete a BGP peer instance (page A-18)
set bgp neighbor <*IP address*> disable	Allows you to disable the bgp neighbor state (page A-19)

Table A-2. Command Line Interface

Commands	Descriptions
set bgp neighbor <*IP address*> enable	Allows you to enable the BGP neighbor's state (page A-19)
set bgp neighbor <*IP address*> hold_interval <*seconds*>	Enables you to specify a BGP neighbor's hold interval time (page A-19)
set bgp neighbor <*IP address*> keepalive_interval <*seconds*>	Enables you to specify a BGP neighbor's keep alive interval time (page A-19)
set bgp neighbor <*IP address*> min_adv_interval <*seconds*>	Enables you to specify a BGP neighbor's minimum advertise interval time (page A-19)
set bgp neighbor <*IP address*> min_orig_interval <*seconds*>	Enables you to specify a BGP neighbor's minimum original interval time (page A-20)
set bgp neighbor <*IP address*> remote_as <*AS number*>	Enables you to specify a BGP neighbor's AS number (page A-20)
set bgp neighbor <*IP address*> retry_interval <*seconds*>	Enables you to specify a BGP neighbor's retry interval time (page A-20)
set bgp neighbor <*IP address*> route_reflector_client [enable | disable]	Enables you to specify a BGP neighbor as a route reflector client (page A-20)
set bgp neighbor <*IP address*> send_community [enable | disable]	Enables you to send community attributes to this BGP neighbor (page A-21)
set bgp neighbor <*IP address*> send_default_route	Enables you to send the BGP default route to the BGP peer (page A-21)

Table A-2. Command Line Interface

Commands	Descriptions
set bgp neighbor <*IP address*> weight <*weight*>	Enables you to specify the weight associated with the BGP neighbor (page A-21)
set bgp network <*network address*>	Enables you to specify the network to be injected into BGP (page A-21)
set bgp state [enable | disable]	Enables you to specify the state of the BGP peer (page A-21)
set iplport <*IP logical port #*> [up | down]	Allows you to enable IP on a logical port (page A-22)
set ospf interface <*IP address*> [up | down]	Enables you to specify an OSPF interface (page A-22)
set ospf interface <*IP address*> area <*IP address*>	Enables you to specify the area the OSPF interface is in (page A-22)
set ospf interface <*IP address*> auth_type [mdfive | simple | none]	Enables you to assign an authentication type (page A-22)
set ospf interface <*IP address*> hello_intvl <*interval*>	Enables you to specify the hello interval (page A-22)
set ospf interface <*IP address*> poll_intvl <*interval*>	Enables you to specify the poll interval (page A-23)
set ospf interface <*IP address*> retransmit_intvl <*interval*>	Enables you to specify the retransmit interval (page A-23)
set ospf interface <*IP address*> router_dead_intvl <*interval*>	Enables you to specify the router dead interval (page A-23)
set ospf interface <*IP address*> router_priority <*priority*>	Enables you to specify the router priority (page A-23)

Table A-2. Command Line Interface

Commands	Descriptions
set ospf interface <*IP address*> tos_metric <*metric*>	Enables you to specify the type of service metric (page A-24)
set ospf interface <*IP address*> transit_delay <*delay*>	Enables you to specify the transit delay (page A-24)
set ospf interface <*IP address*> type [ptp | ptmpt | broadcast | nbmal | virtual link]	Enables you to specify the OSPF logical port interface type (page A-24)
set rip admin <*IP address*> [disable | enable]	Enables you to specify the RIP interface's admin state (page A-24)
set rip authkey <*IP address*> <*auth key*>	Enables you to specify the RIP authentication key (page A-24)
set rip authtype <*IP address*> [disable | simple | mdfive]	Enables you to specify the RIP authentication type (page A-25)
set rip defmetric <*IP address*> <*metric*>	Enables you to specify the RIP default metric (page A-25)
set rip delete <*IP address*>	Enables you to delete RIP from the IP interface (page A-26)
set rip lport <*IP address*> <*IP logical port*>	Allows you to create RIP on the specified IP interface (page A-26)
set rip receive <*IP address*> [disable | one | oneortwo | two]	Enables you to specify RIP receive options (page A-26)
set rip send <*IP address*> [disable | one | onecomp | two]	Enables you to specify RIP send options (page A-26)
set rip splithrz <*IP address*> [disable | simple | poiservs]	Enables you to specify RIP split horizon parameters (page A-26)

Table A-2. Command Line Interface

Commands	Descriptions
set rip version <*IP address*> [one | two]	Enables you to specify the RIP version (page A-27)
show arp	Displays the ARP cache (page A-28)
show bgp aggregate <*network prefix*>	Displays BGP route aggregates (page A-28)
show bgp neighbor	Displays all BGP neighbors (page A-29)
show bgp neighbor <*neighbor address*>	Displays a specific BGP neighbor (page A-30)
show bgp pathdb	Displays BGP path databases (page A-30)
show bgp route in	Displays all BGP routes received by a BGP peer and accepted by its policy (BGP RIB-In) (page A-31)
show bgp route in as <*AS number*>	Displays all BGP routes in the specified AS (page A-32)
show bgp route in community <*community*>	Displays all BGP routes that have this particular community attribute (page A-33)
show bgp route in neighbor <*IP address*>	Displays all BGP routes learned from a specific BGP neighbor (page A-33)
show bgp route in network <*network address*>	Displays all BGP routes in the specified BGP network (page A-33)
show bgp route out	Displays a list of routes sent by a BGP peer using its policies (BGP RIB-Out) (page A-34)

Table A-2. Command Line Interface

Commands	Descriptions
show bgp route summary	Displays the number of BGP routes in RIB-In, RIB-Out, and Rejects. Displays the number of BGP routes sent to a peer, on a peer-to-peer basis (page A-34)
show bgp summary	Displays BGP parameters associated with the configured switch and the switch's BGP neighbor (page A-35)
show card <*slot #*>	Displays the card configuration per slot (page A-36)
show external	Displays System External (ASE) OSPF Autonomous host table (page A-37)
show hardware	The Module Identification Memory (MIM) device allows you to use this command to remotely access your card to determine card type, hardware revision, serial number, manufacturing part number, and product code (page A-38)
show icmp	ICMP statistics (page A-39)
show ip forward statistics card <*card number*>	Displays IP forwarding statistics (page A-40)
show ip forward statistics lport <*logical port #*>	Displays IP forwarding statistics (page A-41)
show IP interface	Displays all IP interfaces (page A-42)

Table A-2. Command Line Interface

Commands	Descriptions
show IP interface slot <*slot #*>	Displays a specific IP interface (page A-42)
show IP interface <*IP lport*>	Displays a specific IP interface (page A-42)
show IP interface <*IP address*>	Displays a specific IP interface (page A-43)
show IP lport	Displays all IP logical ports (page A-44)
show IP lport <*lport #*>	Displays a specific IP logical port (page A-44)
show IP lport slot <*slot #*>	Displays a specific IP logical port (page A-44)
show IP qos pvc	Displays IP QoS pvcs (page A-45)
show IP qos statistics	Displays IP QoS statistics (page A-46)
show ip route	Displays the current IP routing table (page A-47)
show ip route <*ip address*>	Displays an IP route's best match (page A-48)
show ip route <*ip address*> <*net mask*>	Displays an IP route's exact match (page A-48)
show ip route <*ip address*> <*net mask*> all	Displays an IP route's all inclusive matches (page A-49)
show ip route bgp	Displays BGP routes (page A-49)

Table A-2. Command Line Interface

Commands	Descriptions
show ip route direct	Displays direct routes (page A-50)
show ip route ospf	Displays OSPF routes (page A-51)
show ip route rip	Displays RIP routes (page A-52)
show ip route static	Displays static routes (page A-52)
show ip route summary	Displays a summary of the IP routing table (page A-53)
show ip statistics	Displays IP statistics (page A-54)
show fltrbind <*interface #*> <*filter #*>	Displays all filter bindings (page A-54)
show fltrtbl	Displays all filter entries (page A-55)
show lport attributes <*interface #*>	Displays logical port attributes (page A-55)
show lport statistics <*interface #*>	Displays logical port statistics (page A-56)
show mpt all	Displays a list of MPT processes and their current state, the last fail point, and the last failure's result code (page A-56)
show mpt path <*IP address*>	Displays the nodes and outgoing interfaces the MPT path traverses (page A-57)

Table A-2. Command Line Interface

Commands	Descriptions
show mpt statistics	Displays a list of MPTs and the number of nodes attached to them (page A-58)
show ospf adv <*link state type*> <*link state id*> <*adv.router*>	Displays specific link-state advertisements (page A-59)
show ospf adv <*link state type*> <*link state ID*> <*adv.router*> <*area ID*>	Displays specific link-state advertisements (page A-62)
show ospf database	Displays all OSPF link-state databases (page A-63)
show ospf database <*area ID*>	Displays specific OSPF link-state databases (page A-64)
show ospf database absr-summary <*area ID*>	Displays the area border link-state (page A-65)
show ospf database absr-summary <*link state ID*> <*adv.router*>	Displays the area border link-state (page A-65)
show ospf database absr-summary <*link state ID*> <*adv.router*> <*area ID*>	Displays the area border link-state (page A-66)
show ospf database external <*area ID*>	Displays external links (page A-67)
show ospf database external <*link state ID*> <*adv.router*>	Displays external links (page A-68)
show ospf database external <*link state ID*> <*adv.router*><*area ID*>	Displays external links (page A-69)
show ospf database names <*area ID*>	Displays OSPF names (page A-70)
show ospf database names <*link state ID*> <*adv.router*>	Displays OSPF names (page A-71)
show ospf database names <*link state ID*> <*adv.router*><*area ID*>	Displays OSPF names (page A-72)

Table A-2. Command Line Interface

Commands	Descriptions
show ospf database network <*area ID*>	Displays network links (page A-73)
show ospf database network <*link state ID*> <*adv.router*>	Displays network links (page A-74)
show ospf database network <*link state ID*> <*adv.router*> <*area ID*>	Displays network links (page A-75)
show ospf database opaque <*area ID*>	Displays opaque link-state advertisements (page A-76)
show ospf database opaque <*link state ID*> <*adv.router*> <*area ID*>	Displays opaque link-state advertisements (page A-76)
show ospf database opaque <*link state ID*> <*adv.router*>	Displays opaque link-state advertisements (page A-77)
show ospf database router <*area ID*>	Displays router links (page A-78)
show ospf database router <*link state ID*> <*adv.router*>	Displays router links (page A-79)
show ospf database router <*link state ID*> <*adv.router*> <*area ID*>	Displays router links (page A-80)
show ospf database summary <*area ID*>	Displays summary link-state advertisements (page A-81)
show ospf database summary <*link state ID*> <*adv.router*>	Displays summary link-state advertisements (page A-82)
show ospf database summary <*link state ID*> <*adv.router*> <*area ID*>	Displays summary link-state advertisements (page A-83)
show ospf database trunk <*area ID*>	Displays OSPF trunks (page A-84)
show ospf database trunk <*link state ID*> <*adv.router*>	Displays OSPF trunks (page A-85)

Table A-2. Command Line Interface

Commands	Descriptions
show ospf database trunk <*link state ID*> <*adv.router*> <*area ID*>	Displays OSPF trunks (page A-86)
show ospf interface	Displays OSPF interface information (page A-87)
show ospf names	Displays the routing table for resilient names (page A-88)
show ospf namedpath <*type*> <*name*> <*bit length*>	Displays OSPF named path (page A-89)
show ospf namedpath <*type*> <*name*> <*bit length*> <*slot ID*>	Displays OSPF named path (page A-90)
show ospf neighbor	Displays OSPF neighbors (page A-91)
show ospf pathdb <*slot ID*>	Displays the OSPF path database (page A-91)
show ospf pathdb <*switch ID*>	Displays the OSPF path database (page A-92)
show ospf pathdb <*switch ID*> <*interface #*>	Displays the OSPF path database (page A-92)
show ospf pathdb <*switch ID*> <*interface #*> <*slot ID*>	Displays the OSPF path database (page A-92)
show ospf qospath <*IP address*>	Displays OSPF QoS paths (page A-93)
show ospf qospath <*IP address*> <*slot ID*>	Displays OSPF QoS paths (page A-94)
show ospf rtrs	Displays OSPF routers (page A-95)
show ospf statistics <*slot ID*>	Displays OSPF statistics (page A-96)

Table A-2. Command Line Interface

Commands	Descriptions
show ospf statistics <*slot ID*> <*area ID*>	Displays OSPF statistics (page A-97)
show ospf trunk <*switch ID*>	Displays OSPF trunks (page A-98)
show ospf trunk <*interface #*>	Displays OSPF trunks (page A-98)
show ospf trunk <*interface #*> <*slot ID*>	Displays OSPF trunks (page A-99)
show ospf trunk <*qos*>	Displays OSPF trunks (page A-99)
show ospf trunk <*qos*> <*priority*>	Displays OSPF trunks (page A-100)
show ospf vcpath <*IP address*>	Displays OSPF virtual circuit paths (page A-100)
show ospf vcpath <*IP address*> <*slot ID*>	Displays OSPF virtual circuit paths (page A-100)
show ospf vcroute	Displays OSPF virtual circuit routes (page A-101)
show policy interface	Displays the interface association summary (page A-101)
show policy neighbor	Displays the neighbor association summary (page A-102)
show policy netfilter	Displays the network filter summary (page A-102)
show policy netlist	Displays the network filter list summary (page A-103)
show policy ospf	Displays the OSPF route map summary (page A-103)

Table A-2. Command Line Interface

Commands	Descriptions
show policy routemap	Displays the route map summary (page A-103)
show pport attributes <*slot.port*>	Displays physical port attributes(page A-104)
show pport statistics <*slot.port*>	Displays physical port statistics (page A-104)
show pvc attributes <*interface.dlci*>	Displays PVC attributes (page A-105)
show pvc statistics <*interface.dlci*>	Displays PVC statistics (page A-105)
show rip attributes <*IP interface #*>	Displays RIP attributes (page A-108)
show rip attributes all	Displays all RIP attributes (page A-109)
show rip peer <*IP interface #*>	Displays RIP peers (page A-110)
show rip peer all	Displays all RIP peers (page A-110)
show rip statistics <*IP interface #*>	Displays RIP statistics (page A-111)
show rip statistics all	Displays all RIP statistics (page A-112)
show software disk <*slot #*>	Displays disk files for CP ONYX (page A-113)
show software disk all	Displays disk files for CP-ONYX (page A-114)

Table A-2. Command Line Interface

Commands	Descriptions
show software flash <*slot #*>	Displays version information for files on PCMCIA disks (page A-115)
show software flash all	Displays version information for files on PCMCIA disks (page A-116)
show software card <*slot #*>	Displays version information for software running on a card (page A-116)
show software card all	Displays version information for software running on a card (page A-117)
show static	Displays static routes (page A-117)
show system	Displays general system information and status (page A-118)
show tcp	TCP statistics (page A-119)
show udp	UDP statistics (page A-120)
show users	Displays current users logged on the switch through the console or telnet (page A-120)

CLI Command Descriptions

The following pages list each of the command line interface commands in alphabetical order.

set bgp aggregate <*aggregate address*>

Enables you to specify the aggregate network IP address.

```
> set bgp aggregate <aggregate address>
```

set bgp compare_med [enable | disable]

Enables you to allow MED comparison. MED comparison controls the inclusion of multi-exit discriminator in the route selection process. BGP uses MED to communicate preferred path information to external neighbors when the autonomous system has multiple exits to another autonomous system.

```
> set bgp compare_med [enable | disable]
```

set bgp as <*as number*>

Enables you to enter the switch's local autonomous system number. Valid AS numbers are between 1 and 65535.

```
> set bgp as <as number>
```

set bgp always_compare_send_bgp_nets [enable | disable]

Enables you to allow injected BGP networks to be listed in the routing table.

```
> set bgp always_compare_send_bgp_nets [enable | disable]
```

set bgp client_to_client [enable | disable]

Enable this parameter only if the switch is a route reflector. If you enable this parameter, any received routes will be sent to clients.

```
> set bgp client_to_client [enable | disable]
```

set bgp cluster_id <*cluster id*>

Enables you to specify a BGP cluster ID. A cluster is a group of client peers that communicate with one another. The cluster ID specifies the cluster. For the cluster ID, enter the switch's internal IP address.

```
> set bgp cluster_id <cluster id>
```

set bgp def_local_preference <*value*>

Enables you to rank a route according to its importance. This value is sent to internal neighbors. This value is compared to other routes that have the same destination. A higher local preference indicates the route is preferred. Valid local preference values are between 1 and 4294967295.

```
> set bgp def_local_preference <value>
```

set bgp neighbor <*IP address*> create

Enables you to add a BGP neighbor.

```
> set bgp neighbor <IP address> create
```

set bgp neighbor <*IP address*> delete

Enables you to delete a BGP neighbor.

```
> set bgp neighbor <IP address> delete
```

set bgp neighbor <*IP address*> disable

Enables you to disable a BGP neighbor.

```
> set bgp neighbor <IP address> disable
```

set bgp neighbor <*IP address*> enable

Enables you to enable a BGP neighbor.

```
> set bgp neighbor <IP address> enable.
```

set bgp neighbor <*IP address*> hold_interval <*seconds*>

Enables you to specify the time BGP holds before considering this connection to be down if messages are not received from this neighbor. Valid hold interval values are 0, or 3 to 65535.

```
> set bgp neighbor <IP address> hold_interval <seconds>
```

set bgp neighbor <*IP address*> keepalive_interval <*seconds*>

Enables you to specify the time between consecutive keep alive messages sent to this neighbor. Valid keep alive interval values are between 0 and 21845.

```
> set bgp neighbor <IP address> keepalive_interval <seconds>
```

set bgp neighbor <*IP address*> min_adv_interval <*seconds*>

Enables you to specify the minimum time between route update advertisements belonging to the external AS. Valid minimum advertise interval values are between 1 and 65535.

```
> set bgp neighbor <IP address> min_adv_interval <seconds>
```

set bgp neighbor <*IP address*> min_orig_interval <*seconds*>

Enables you to specify the minimum time between route update advertisements belonging to the home AS. Valid minimum original interval values are between 1 and 65535.

```
> set bgp neighbor <IP address> min_orig_interval <seconds>
```

set bgp neighbor <*IP address*> remote_as <*as number*>

Enables you to specify the BGP neighbor's remote AS number. Valid remote AS numbers are between 1 and 65535.

```
> set bgp neighbor <IP address> remote_as <as number>
```

set bgp neighbor <*IP address*> retry_interval <*seconds*>

Enables you to specify the time BGP waits before it tries to connect to this neighbor. Valid retry interval values are between 1 and 65535.

```
> set bgp neighbor <IP address> conn_retry_interval
<seconds>
```

set bgp neighbor <*IP address*> route_reflector_client [enable | disable]

Enables you to specify the neighbor as a BGP route reflector client.

```
> set bgp neighbor <IP address> route_reflector_client
[enable | disable]
```

set bgp neighbor <*IP address*> send_community [enable | disable]

Enables you to allow the sending of community attributes of all updates to this neighbor.

```
> set bgp neighbor <IP address> send community
[enable | disable]
```

set bgp neighbor <*IP address*> send_default_route

Enables you to send the BGP default route to the BGP peer. You configure this parameter if the BGP configured switch does not know a destination, and the BGP peer does.

```
> set bgp neighbor <IP address> send_default_route
```

set bgp neighbor <*IP address*> weight <*weight*>

Enables you to specify the path weight received by the neighbor and applied to every route. Valid path weight values are between 0 and 65535.

```
> set bgp neighbor <IP address> weight <weight>
```

set bgp network <*network address*>

Enables you to specify the network to be injected into BGP.

```
> set bgp network <network address>
```

set bgp state [enable | disable]

Enables you to activate BGP on this switch.

```
> set bgp state [enable | disable]
```

set iplport <*IP address*> [up | down]

Allows you to enable IP on a logical port.

```
> set iplport <ip address> [up | down]
```

set ospf interface <*IP address*> [up | down]

Enables you to specify the OSPF interface status.

```
> set ospf interface <IP address> [up | down]
```

set ospf interface <*IP address*> area <*IP address*>

Enables you to specify the area the OSPF interface is in.

```
> set ospf interface <IP address> area <IP address>
```

set ospf interface <IP address> authtype [mdfive | simple | none]

Enables you to specify the type of authentication that OSPF will use as a security measure to ensure that this logical port and router exchanges information with correct neighbors. See "Authentication Type" on page 9-18 for descriptions of *mdfive*, *simpl*e, and *none*.

```
> set ospf interface <IP address> authtype
[mdfive | simple | none]
```

set ospf interface <*IP address*> hello_intvl <*interval*>

Enables you to specify the number of seconds between router Hello messages. Valid hello interval values are between 1 and 65535.

```
> set ospf interface <IP address> hello_intvl <interval>
```

set ospf interface <*IP address*> poll_intvl <*interval*>

Enables you to specify the time between hello packets sent to an inactive non-broadcast multi-access neighbor. Valid poll interval values are greater than or equal to 0.

```
> set ospf interface <IP address> poll_intvl <interval>
```

set ospf interface <*IP address*> retransmit_intvl <*interval*>

Enables you to specify the time to wait before resending a packet if no acknowledgment is received. Valid retransmit interval values are between 0 and 3600.

```
> set ospf interface <IP address> retransmit_intvl
<interval>
```

set ospf interface <*IP address*> router_dead_intvl <*interval*>

Enables you to specify the number of seconds a router waits to hear a Hello message from a neighbor before the router declares the neighbor down. Valid router dead interval values must be greater than 0.

```
> set ospf interface <IP address> router_dead_intvl
<interval>
```

set ospf interface <*IP address*> router_priority <*priority*>

Enables you to specify the number which identifies the priority of the router associated with this logical port. Valid priority values are between 0 and 255.

```
> set ospf interface <IP address> router_priority <priority>
```

set ospf interface <*IP address*> tos_metric <*metric*>

Enables you to specify the Type of Service 0 Metric. The Type of Service 0 Metric is a cost. The lowest TOS 0 has the highest priority for routing. Valid type of service metric values are between 1 and 65535.

```
> set ospf interface <IP address> tos_metric <metric>
```

set ospf interface <*IP address*> transit_delay <*delay*>

Enables you to specify the estimated number of seconds it takes to transmit a link-state update packet over this interface.

```
> set ospf interface <IP address> transit_delay <delay>
```

set ospf interface <*IP address*> type [ptp | ptmpt | broadcast | nbmal | virtual_link]

Enables you to specify the OSPF logical port interface type. See "Interface Type" on page 9-16 for descriptions of *ptp*, *ptmpt*, *broadcast*, *nbmal*, and *virtual link*.

```
> set ospf interface <IP address>
type [ptp | ptmpt | broadcast | nbmal | virtual_link]
```

set rip admin <*IP address*> [disable | enable]

Enables you to enable or disable RIP on this interface.

```
> set rip admin <IP address> [disable | enable]
```

set rip authkey <*IP address*> <*authentication key*>

Enables you to specify the RIP authentication key.

```
> set rip authkey <IP address> <authentication key>
```

set rip authtype <*IP address*> [disable | simple | mdfive]

Enables you to specify 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. See "Authentication Type" on page 7-4 for descriptions of *disable*, *simple*, and *mdfive*.

```
> set rip authtype <IP address> [disable | simple | mdfive]
```

set rip defmetric <*IP address*> <*metric*>

Enables you to specify the RIP default metric. This value is the metric that is used for the default route entry in RIP updates that originate on this interface.

```
> set rip defmetric <IP address> <metric>
```

set rip delete <*IP address*>

Enables you to delete RIP from the IP interface.

```
> set rip delete <IP address>
```

set rip lport <*IP address*> <*IP logical port #*>

Allows you to create RIP on the specified IP interface.

```
> set rip lport <IP address> <IP logical port>
```

set rip receive <*IP address*> [disable | one | oneortwo | two]

Enables you to specify RIP receive options. See "Receive" on page 7-3 for descriptions of *disable*, *one*, *oneortwo*, and *two*.

```
> set rip receive <IP address>
[disable | one | oneortwo| two]
```

set rip send <*IP address*> [disable | one | onecomp | two]

Enables you to specify RIP send options. See "Send" on page 7-3 for descriptions *disable*, *one*, *onecomp*, and *two*.

```
> set rip send <IP address> [disable | one | onecomp | tw
```

set rip splithrz <*IP address*> [disable | simple | poiservs]

Enables you to specify RIP split horizon parameters. Split horizon is a method for avoiding common situations that require a counting to infinity. See "Split Horizon" page 7-3 for descriptions of *disable*, *simple*, and *poiservs*.

```
> set rip splithrz <IP address> [disable | simple | poiser
```

set rip version <*IP address*> [one | two]

Enables you to specify the RIP version. See "Configuring RIP at the Logical Port" on page 7-1 for descriptions of *one*, and *two*.

```
> set rip version <IP address> [one | two]
```

show arp

Displays Address Resolution Protocol (ARP) entries. ARP entries match IP addresses to specific MAC, DLCI, or VPI-VCI addresses.

```
> show arp

       IpAddr  LinkType  HwAddr        State       EntryType

152.148.30.129  ethernet  08002080629c  Complete     Dynamic

     195.1.1.2  fr-relay           100  Complete    Dynamic

     196.7.7.2  fr-relay           300  Complete    Dynamic

     206.6.6.1       atm      0,    34  Complete    Dynamic

    206.6.6.10       atm      2,    32  Complete    Dynamic
```

show bgp aggregate <*network prefix*>

Displays BGP aggregates. BGP aggregates combine specific BGP networks to less specific BGP networks. This combination reduces the size of the BGP routing table.

> **show bgp aggregate <*network prefix*>**

show bgp neighbor

Displays all BGP neighbors.

```
> show bgp neighbor

Neighbor:  100.100.100.100


Admin State:             Enabled        Operating State:         Idle

Remote AS:               20             Route Reflector Client:  No

Confg Keepalive Time:    30             Calc Keepalive Time:     0

Connect Retry Interval:  120            Configured Hold Time:    90

Established Time:        0:0:0:0        Negotiated Hold Time:    0

Update Elapsed Time:     0:0:0:0        Established Transitions: 0

Messages In:             0              Updates In:              0

Messages Out:            0              Updates Out:             0

Send-Community:          Disabled       Weight:                  0

Send-Default-Route:      Disabled       Update Source:           153.60.70.4


Neighbor:  153.60.70.2


Admin State:             Enabled        Operating State:         Established

Remote AS:               20             Route Reflector Client:  No

Confg Keepalive Time:    30             Calc Keepalive Time:     30

Connect Retry Interval:  120            Configured Hold Time:    90

Established Time:        0:17:46:5      Negotiated Hold Time:    90

Update Elapsed Time:     0:0:0:0        Established Transitions: 1

Messages In:             2136           Updates In:              0
```

```
Messages Out:          2138          Updates Out:            2
Send-Community:        Disabled      Weight:                 0
Send-Default-Route:    Disabled      Update Source:          153.60.70.4
```

show bgp neighbor <*IP address*>

Displays a specific BGP neighbor.

```
> show bgp neighbor 195.1.1.2
Neighbor:  195.1.1.2
Admin State:            Enabled       Operating State:         Established
Remote AS:              10            Route Reflector Client:  No
Confg Keepalive Time:   30            Calc Keepalive Time:     30
Connect Retry Interval: 120           Configured Hold Time:    90
Established Time:       0:22:16:11    Negotiated Hold Time:    90
Update Elapsed Time:    0:22:16:11    Established Transitions: 1
Messages In:            2661          Updates In:              1
Messages Out:           2675          Updates Out:             3
Send-Community:         Disabled      Weight:                  0
Send-Default-Route:     Disabled
```

show bgp pathdb

Displays BGP path databases.

```
> show bgp pathdb
Dir Refcnt Neighbor        Next_hop        Local_pref       MED     AS-path
in  1      153.54.1.5      193.1.2.3       1                0       12 ?
in  1      153.54.1.5      193.1.2.3       1                0       12 i
```

show bgp route in

Displays all BGP routes received by a BGP peer and accepted by its policy. This CLI command is also referred to as Routing Information Base-In (RIB-In).

```
> show bgp route in

Prefix              Neighbor        Next_hop        Local_pref      MED     Weight  TAG
AS-path

11.0.0.0/8         b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 ?

128.1.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

128.2.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

128.3.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

128.4.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

128.5.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

128.6.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

128.7.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

135.1.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

135.2.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

135.3.0.0/16       b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i

152.148.30.129/32  b153.54.1.1      153.54.1.1      100(100)        *       0       *
i

152.148.35.175/32  b153.54.1.1      153.54.1.1      100(100)        *       0       *
i

152.148.35.176/32  b153.54.1.1      153.54.1.1      100(100)        *       0       *
i

153.54.0.0/16      b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 ?

195.218.218.0/24   b153.54.1.5      193.1.2.3       100(100)        0       0       *
12 i
```

show bgp route in as <*AS number*>

Displays all BGP routes in a specified autonomous system.

```
> show bgp route in as 10

130.1.0.0/16 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
130.2.0.0/16 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
130.2.1.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
130.3.0.0/16 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
201.201.201.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
202.202.202.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
203.203.203.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
205.1.2.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
210.5.4.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
210.5.5.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
```

show bgp route in community <*community attribute*>

Displays all BGP routes that have this community attribute in their BGP path attribute.

```
> show bgp route in community <community attribute>
```

show bgp route in neighbor <*IP address*>

Displays all BGP routes that are learned from a specific BGP neighbor.

```
> show bgp route in neighbor 195.1.1.2

130.1.0.0/16 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100

130.2.0.0/16 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100

130.2.1.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100

130.3.0.0/16 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100

201.201.201.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100

202.202.202.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100

203.203.203.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100

205.1.2.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100

210.5.4.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100

210.5.5.0/24 :  nbr id 195.1.1.2  origin 0  as path 10   nexthop 195.1.1.2  med 0   local pref 100
```

show bgp route in network <*network address*>

Displays all BGP routes in the specified BGP network.

```
> show bgp route in network <network address>
```

show bgp route out

Displays a list of routes sent by a BGP peer, using its policies. This CLI command is also referred to as Routing Information Base Out (RIB-Out).

```
> show bgp route out

Prefix              Neighbor        Next_hop        Local_pref      MED     AS-path

150.150.250.3/32    192.32.3.2      192.32.3.1      *               *       350 i

152.148.86.0/24     192.32.3.2      192.32.3.1      *               *       350 i

192.32.1.0/24       192.32.3.2      192.32.3.1      *               *       350 i

192.32.1.1/32       192.32.3.2      192.32.3.1      *               *       350 i

192.32.2.0/24       192.32.3.2      192.32.3.1      *               *       350 i

192.32.2.1/32       192.32.3.2      192.32.3.1      *               *       350 i

192.32.3.0/24       192.32.3.2      192.32.3.1      *               *       350 i

192.32.21.0/24      192.32.3.2      192.32.3.1      *               *       350 i

192.32.21.1/32      192.32.3.2      192.32.3.1      *               *       350 i

192.32.48.0/24      192.32.3.2      192.32.3.1      *               *       350 i

192.32.48.1/32      192.32.3.2      192.32.3.1      *               *       350 i
```

show bgp route summary

Displays the number of BGP routes in *RIB-In*, *RIB-Out*, and *Rejects*. Displays the number of routes sent to a peer, on a peer-to-peer basis.

```
> show bgp route summary

Neighbor         #RIB-in     #RIB-out    #Reject

192.32.2.2       0           0           0

192.32.3.2       17          21          0

192.32.83.2      0           0           0

         Total entries with BGP info: 2502
```

show bgp summary

Displays BGP parameters associated with the configured switch and the switch's BGP neighbor.

```
> show bgp summary

BGP Admin State:   Enabled                 BGP Oper Status:         Up

Router ID:         153.54.1.2

Local AS:          20                      Default Local Preference: 100

Use MED:           Enabled                 Client-to-Client:        Enabled

Route Reflector:   No


Neighbor           V    AS        MsgRcvd      MsgSent      Up/Down        State

153.54.1.1         4    20        3158         3159         1:2:11:50      Established

153.54.1.5         4    20        3164         3162         0:0:49:19      Established
```

show card <*slot #*>

Displays information about the control processor or card. This information includes serial number, hardware, EPROM, software revisions, and status. This command also retrieves statistics for memory and CPU utilization.

```
> show card 3

Serial#:                 0040FB026960    Configured Card Type:        UIO-8
Hardware Revision:           05.00.00    Actual Card Type:            UIO-8
EPROM Revision:           04.02.00.00    Physical Slot:                   3
Software Revision:        05.00.00.00    Logical Slot:                    3

Redund State:                  Active
Card State:                    Active    Memory Available:          2266192
Administrative Status:             Up    CPU Utilization:                2%
Operational Status:                Up
Diagnostics Status:                Up

Packets Received:                6373    Maximum number of PVCs:       2215
Octets Received:               373158    Inactive PVCs:                  19
Packets Sent:                    6478    Free VCs:                     2196
Octets Sent:                   380336
```

show external

Displays OSPF Autonomous System Externals, such as static routes in OSPF.

```
> show external

Destination      Mask             Gateway            Metric  Status  Index  ConnID

152.148.82.11    255.255.255.255  0.0.0.0                 1       1   4097       0

152.148.82.13    255.255.255.255  0.0.0.0                 1       1   4097       0

152.148.82.93    255.255.255.255  0.0.0.0                 1       1   4097       0
```

show hardware

The Module Identification Memory (MIM) device allows you to use this command to remotely access your card to determine card type, hardware revision, serial number, manufacturing part number, and product code. For more information on MIM, contact the Ascend Technical Response Center.

In addition, this command performs the following Cisco command functions: *show iop, show iom, show cp, show sp, and show env status.*

```
> show hardware

Slot        Product Code          MFG. #     H/W Rev.          S/N
1                  11021     810-00152-02           04      22A06238
1  IOA             11023     810-00111-02           02      21A00366
12                 11030     810-00051-02           01      A0000403
12 IOA             11031     810-00139-01           00      21A02518
Card Status
Slot Status   Non-Fatal Error              Time        Fatal Error                 Time
1    ok                                                Internal Error              0:0:0:12
12   marginal PRAM out of Synch            0:0:1:19    Can't write configuration   0:0:4:24
Environment Status
Power Supply 1 - up
Power Supply 2 - down
Fan         Status
1          up
2          up
3          up
4          up
5          up
6          up
```

show icmp

Displays ICMP statistics.

```
> show icmp

ICMP Statistics:

                         Receive        Transmit

Echos:                        38               0

Echo Replies:                  0              38

Dest Unreachables:             0              41

Source Quenches:               0               0

Redirects:                     0               0

Router Advert:                 0               0

Router Solicit:                0               0

Time Exceeded:                 0               0

Parameter Problems:            0               0

Timestamps:                    0               0

Timestamp Replies:             0               0

Info Request:                  0               0

Info Replies:                  0               0

Address Masks:                 3               0

Addr Mask Replies:             0               0

Responses:                38

Errors:                   41    Old IP too short:                0

Old ICMP:                  0    Bad ICMP code:                   0

Pkt too short:             0    Bad checksum:                    0

Bad Length:                0
```

show ip forward statistics card <*slot #*>

Displays IP forwarding statistics.

```
> show ip forward statistics card 14

Layer-2 Errors:
  sap=0, ukn_prot=0, arp_format=0
  arp_lookup=0, arp_invalid=0, discards=0
Layer-3 Errors:
  ip_lport=0[last_id=0]
  ip_vers=0, ip_hdr_short=0, ip_hdr_long=0, ip_chksum=0
  ip_filter=0, ip_lookup=0, ip_reject=0, ip_bc=0
  ip_lbc=0, fmbuf=0, rte_indirect=0, rte_nmgmt=0, ip_ttl=0
  ip_lnktype=0[lastval=0], ip_df=0, discards=0
Layer-3 Counts:
  fg_icmp=0, icmp_throttle=0, fwd_cp=0, fwd_hbh=0
  fwd_mpt=0, options=0, split=0, dct=0
  ntu_ctl=0, ntu_ckt=0, ntu_hbh=0, tx_hbh=0
```

show ip forward statistics lport <*logical port*>

Displays IP forwarding statistics.

```
> show ip forward statistics lport 16

IP-Lport Counts:

  if_inOctets=0, if_inUcastPkts=0, if_inNUcastPkts=0, if_inDiscards=0

  if_inFragmented=0, if_inErrors=0, if_inDestUnreach=0, if_inTtlExceeded=0

  if_inParamErr=0, if_inProtoUnknown=0, if_outOctets=0, if_outUcastPkts=0

  if_outNUcastPkts=0, if_outDiscards=0, if_outErrors=0, if_outQlen=0

  AdminState=1
```

show IP interface

Displays all IP interfaces.

```
> show IP interface

IpAddr            Slot  Card        Lport MTU   MMTU  BMTU  ARP   IARP  OP_STATE  ADMIN_STATE


206.6.6.2/24      14    ATM-IWU     36    1500  N/A   1500  DIS   ENA   UP          ENA
160.160.160.1/16  12    HSSI-2      41    1500  N/A   1500  ENA   ENA   DOWN        ENA
196.7.7.3/24      3     UIO-8       1     1500  N/A   1500  ENA   ENA   DOWN        ENA
195.1.1.1/24      3     UIO-8       2     1500  N/A   1500  ENA   ENA   UP          ENA
```

show IP interface slot <*slot #*>

Displays a specific IP interface.

```
> show IP interface slot 14

IpAddr            Slot  Card        Lport MTU   MMTU  BMTU  ARP   IARP  OP_STATE ADMIN_STATE


206.6.6.2/24      14    ATM-IWU     36    1500  N/A   1500  DIS   ENA   UP          ENA
```

show IP interface <*IP logical port #*>

Displays a specific IP interface.

```
> show IP interface 36

IpAddr              Lport Pport    Card        MTU   ARP   IARP  OPER  ADMIN HWADDR


206.6.6.9/30         36    14.1     ATM-IWU     1500  DIS   DIS   UP    ENA   2,32
206.6.6.2/30         36    14.1     ATM-IWU     1500  DIS   ENA   UP    ENA
```

show IP interface <*IP address*>

Displays a specific IP interface.

```
> show IP interface 206.6.6.9

IpAddr              Lport Pport    Card       MTU   ARP   IARP  OPER  ADMIN HWADDR

1.1.1.1/16           5     8.1       ATMDS3-1   1500  ENA   ENA   UP    ENA

210.11.11.1/24       8     10.3      UIO-8      1500  ENA   ENA   UP    ENA   55

2.2.2.2/8            16    10.6      UIO-8      1500  ENA   ENA   DOWN  ENA
```

show IP lport

Displays all IP logical ports.

```
> show IP lport

    IpLport   Slot   UFWD   BFWD   OP_STATE  ADMIN_STATE

   10          3      ENA    ENA    DOWN       ENA

    5          8      ENA    ENA      UP       ENA

   20         12      ENA    ENA    DOWN       ENA

    8         10      ENA    ENA      UP       ENA

   16         10      ENA    ENA    DOWN       ENA
```

show IP lport <*IP lport #*>

Displays a specific IP logical port.

```
> show IP lport 36

IpLport   Slot   UFWD   BFWD   OP_STATE  ADMIN_STATE

   36         14      ENA    ENA      UP       ENA
```

show IP lport slot <*slot #*>

Displays a specific IP logical port.

```
> show IP lport slot 3

 IpLport   Slot   UFWD   BFWD   OP_STATE  ADMIN_STATE

   10          3      ENA    ENA    DOWN       ENA
```

show IP qos pvc

Displays active IP QoS PVCs. IP QoS PVCs forward datagrams that match IP QoS flow profiles.

```
> show IP qos pvc

Currently active QoS PVCs (by PVC ID):

  207
```

show IP qos statistics <*IP logical port*>

Displays IP QoS PVC statistics.

```
> show IP qos statistics 2

Datagrams forwarded over QoS PVCs for interface # 2:

FlowProfId  SrcIpAddr           DstIpAddr           PvcId          Datagrams

1           1.2.3.4/32          0.0.0.0/0           0                      0

2           2.3.4.0/24          0.0.0.0/0           0                      0

3           3.3.0.0/16          0.0.0.0/0           0                      0
```

show ip route

Displays the IP routing table's current state.

```
> show ip route

 Dest                Next_hop          State    Cost Lport Age Ckt
 130.1.0.0/16         195.1.1.2        EBGP     1    2     *
 130.2.0.0/16         195.1.1.2        EBGP     1    2     *
 130.2.1.0/24         195.1.1.2        EBGP     1    2     *
 130.3.0.0/16         195.1.1.2        EBGP     1    2     *
 152.148.0.0/16       *                Direct   *    eth   *
m152.148.30.129/32    0.0.0.0          Static   1    eth   *
m153.60.70.2/32       153.60.70.2      OSPF     100  4     *     12/38
m153.60.70.4/32       *                Local    *    0     *
 195.1.1.0/24         *                Direct   *    2     *
 195.1.1.1/32         *                Local    *    lcl   *
 196.7.7.0/30         *                Direct   *    1     *
 196.7.7.2/32         *                Local    *    lcl   *
 196.7.7.8/30         *                Direct   *    1     *
 196.7.7.9/32         *                Local    *    lcl   *
 201.201.201.0/24     195.1.1.2        EBGP     1    2     *
 202.202.202.0/24     195.1.1.2        EBGP     1    2     *
 203.203.203.0/24     195.1.1.2        EBGP     1    2     *
 205.1.2.0/24         195.1.1.2        EBGP     1    2     *
 206.6.6.0/30         *                Direct   *    36    *
 206.6.6.2/32         *                Local    *    lcl   *
 210.5.4.0/24         195.1.1.2        EBGP     1    2     *
 210.5.5.0/24         195.1.1.2        EBGP     1    2     *
```

show ip route <*IP address*>

Displays an IP route's best match.

```
> show ip route 195.1.1.0

        Network:        195.1.1.0
        Mask:           255.255.255.0
        Flags:
        Lport:          2
        Next hop:       0.0.0.0
        Tag:            0
```

show ip route <*IP address*> <*net mask*>

Displays an IP route's exact match.

```
> show  ip route 152.148.30.129   255.255.255.255

        Network:        152.148.30.129
        Mask:           255.255.255.255
        Flags:          Management
        Lport:          4097
        Next hop:       0.0.0.0
        Tag:            0x0
```

show ip route <*IP address*> <*net mask*> all

Displays an IP route's all inclusive matches.

```
> show ip route 152.148.30.129   255.255.255.255 all

 Dest                 Next_hop          State    Cost Lport Age Ckt
*152.148.30.129/32    0.0.0.0           Static   1    eth   0
```

show ip route bgp

Displays BGP routes.

```
> show ip route bgp

 Dest                 Next_hop          State    Cost Lport Age Ckt
 152.148.35.175/32     193.1.1.1        EBGP     1    5     *
 152.148.35.176/32     193.1.1.1        EBGP     1    5     *
 201.1.1.0/24          193.1.1.1        EBGP     1    5     *
 201.1.2.0/24          193.1.1.1        EBGP     1    5     *
```

show ip route direct

Displays direct routes.

```
> show ip route direct

Dest                 Next_hop         State    Cost Lport Age Ckt

152.148.35.0/24        *                Direct   *    eth    *

193.1.1.0/24           *                Direct   *    5      *

193.1.1.5/32           *                Local    *    lcl    *

193.1.2.0/24           *                Direct   *    3      *

193.1.2.5/32           *                Local    *    lcl    *

193.1.3.0/24           *                Direct   *    6      *

193.1.3.5/32           *                Local    *    lcl    *

193.1.4.0/24           *                Direct   *    7      *

193.1.4.5/32           *                Local    *    lcl    *
```

show ip route ospf

Displays OSPF routes.

```
> show ip route ospf

 1.0.0.0/8            150.150.170.9    OSPFE2   3     103     *      14/2020
 1.1.0.0/16           150.150.170.1    OSPFE2   1     9       *       8/43
 1.1.1.1/32           150.150.170.1    OSPFE2   1     9       *       8/43
 116.0.0.0/8          150.150.170.3    OSPFE2   2     2       *      10/64
 116.5.5.0/24         150.150.170.1    OSPFE2   1     9       *       8/43
 116.5.5.6/32         150.150.170.1    OSPFE2   1     9       *       8/43
 125.0.0.0/8          150.150.170.3    OSPFE2   3     2       *      10/64
 125.125.12.1/32      150.150.170.1    OSPFE2   1     9       *       8/43
 150.150.0.0/16       150.150.170.9    OSPFE2   3     103     *      14/2020
m150.150.170.1/32     150.150.170.1    OSPF     100   9       *       8/43
m150.150.170.2/32     *                Local    *     0       *
m150.150.170.3/32     150.150.170.3    OSPFIA   100   2       *
m150.150.170.4/32     150.150.170.3    OSPF     200   2       *      10/64
m150.150.170.5/32     150.150.170.3    OSPF     200   2       *
m150.150.170.6/32     150.150.170.1    OSPF     200   9       *       8/43
m150.150.170.7/32     150.150.170.3    OSPF     300   2       *      10/64
m150.150.170.8/32     150.150.170.9    OSPF     200   103     *      14/2020
m150.150.170.9/32     150.150.170.9    OSPF     100   103     *      14/2017
m150.150.201.1/32     150.150.170.3    OSPFIA   301   2       *
m150.150.201.2/32     150.150.170.3    OSPFIA   401   2       *
m150.150.201.3/32     150.150.170.3    OSPFIA   501   2       *
m150.150.201.4/32     150.150.170.3    OSPFIA   351   2       *
```

show ip route rip

Displays RIP routes.

```
> show ip route rip

 Dest                  Next_hop          State     Cost Lport Age Ckt

123.0.0.0/8              189.32.4.203     RIP         2     109    10

189.32.0.0/24            189.32.4.203     RIP         5     109    10

193.1.2.0/24             189.32.4.203     RIP         2     109    10
```

show ip route static

Displays static routes.

```
> show ip route static

 Dest                  Next_hop          State     Cost Lport Age Ckt

m152.148.30.129/32    0.0.0.0             Static    1     eth     *

 205.1.1.0/24          0.0.0.0             Static    1     0       *

 205.1.2.0/24          0.0.0.0             Static    1     0       *
```

show ip route summary

Displays a summary of IP routes.

```
> show ip route summary

Route Source          Route Count

IBGP                  0

EBGP                  10

Direct                12

OSPF                  2

RIP                   0

Static                1



Total Routes          25
```

show ip statistics

Displays IP statistics.

```
> show ip statistics

IP Statistics:

---------

 Datagrams:        13719 Total

                     414 Forwarded                81 Dest unreachable

                       0 Fwd same net          11670 Delivered

                      10 Unknwn Protocol        2906 Locally generated

                       0 Dropped                   0 Reassembled

                       0 Fragmented                0 Output fragmented

                       0 Cannot Fragment           0 Bad Options

                       3 No Route                  0 Bad Version

                       0 Raw IP generated


 Fragments:            0 Received                  0 Dropped

                       0 Timed out


 Errors:               0 Bad checksum              0 Pkt too short

                       0 Not enough data           0 IP header less than data ize

                       0 IP length less than header length
```

show fltrbind <*interface#*> <*filter#*>

Displays all filter bindings.

> **show fltrbind <*interface #*> <*filter#*>**

show fltrtbl

Displays all filter entries

```
> show fltrtbl

FLT A T TOS SOURCE          DESTINATION     PROTO SRC_PORT DST_PORT

--- - - --- --------------- --------------- ----- -------- --------

  1 R D 255       195.1.1.2         0.0.0.0     1      179      179

                  195.1.1.2         0.0.0.0     6      179      179

              255.255.255.0         0.0.0.0
```

show lport attributes <*interface #*>

Displays logical port attributes.

```
> show lport attributes 1

Slot:           3

Port:           7

Interface:      1

Data Rate:      2048000

Maximum Transmission Unit:      8132

Administrative Status:            Up    Operational
Status:         Down
```

show lport statistics <*interface #*>

Displays logical port statistics, including the number of frames transmitted and received. It also provides information about errors, if applicable.

```
> show lport statistics 1



              Receive        Transmit
Frames:       890            825
Octets:       92064          59852
Discards:     0              0
Errors:       0              0
```

show mpt all

Displays the nodes and outgoing interfaces the MPT path traverses.

```
> show mpt all

Node            Flags  RecvState SendState LastFail(Node/Port/Reason)
150.150.170.2   0x6    ACTIVE    ACTIVE    None/0/NONE
150.150.170.3   0x16   INACT     INACT     None/0/NONE
150.150.170.4   0x6    INACT     INACT     150.150.170.3/15/TRUNK_DOWN
150.150.170.5   0x36   INACT     INACT     None/0/NONE
150.150.170.6   0x5    ACTIVE    ACTIVE    None/0/NONE
150.150.170.7   0x5    ACTIVE    ACTIVE    None/0/NONE
150.150.170.8   0x26   INACT     INACT     None/0/NONE
150.150.170.9   0x4    INACT     INACT     None/0/NONE
```

show mpt path <*IP address*>

Displays a list of MPT processes and their current state, the last fail point, and the last failure’s result code.

```
> show mpt path 150.150.170.2

List of node/interface pairs in path to node 150.150.170.2:

150.150.170.1/5

Path characteristics:

PURE_CELL, FOR_IP

MPT sending enabled on slot 8, VC 90 on port 5
```

show mpt statistics

Displays a list of MPTs and the number of nodes attached to them.

```
> show mpt statistics

Source nodes:

Total: 8 Added: 8 Deleted: 0

Destination nodes:

Total: 3 Added: 52 Deleted: 49

Path Leaf Registration:

Added: 7 Deleted: 1 Failed: 0

Path VC Registration:

Added: 7 Deleted: 1 Failed: 45 Cleared: 45
```

show ospf adv <*link state type*> <*link state ID*> <*adv.router*>

Displays link-state advertisements.

For *link state types*, specify one of the numerical values: (1) router link , (2) network link , (3) summary link , (4) ASBR link, and (5) external link. For *link state ID*, specify the IP address of the network or interface you want advertised. For *advertising router* value, specify the IP address of the advertising router.

```
> show ospf adv 1 2.23.160.2 2.23.160.2

            LSA in Area: 0.0.0.0

   LS type:Router-LSA

   ___________________

   LS age:971

   LS options:0x22

   LS ID:2.23.160.2

   Adv rtr:2.23.160.2

   LS Seq #:0x8000052f

   LS Xsum:0x119d (good)

   LS Length:60

   # router type:ASBR

   # interfaces:3

      Link Type:Stub Net

      Link ID:222.222.222.0

      Link Data:255.255.255.0

      TOS 0 Cost:10

      Link Type:Router link
```

show ospf adv <*link state type*> <*link state ID*> <*adv.router*> <*area ID*>

Displays link-state advertisements.

For *link state types*, specify one of the numerical values: (1) router link , (2) network link , (3) summary link , (4) ASBR link, and (5) external link. For *link state ID*, specify the IP address of the network or interface you want advertised. For *advertising router* value, specify the IP address of the advertising router. For *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf adv <link state type> <link state ID>
<adv.router> <area ID>
```

show ospf database

Displays the OSPF link-state database.

```
> show ospf database

   Link-state database for area 0.0.0.2

                ---------------------------- ---------------

Type        ID                Adv-Router        Seq#       Xsum   Age

RTR (10) 1  153.60.70.2       153.60.70.2       0x8000000b 0x6642 148

RTR (10) 1  153.60.70.4       153.60.70.4       0x8000002e 0x8bdc 145

ASE      5  152.148.0.0       153.60.70.2       0x80000007 0x2242 153

ASE      5  152.148.0.0       153.60.70.4       0x80000007 0x322e 999

ASE      5  152.148.30.129    153.60.70.4       0x80000007 0xfbc2 999

ASE      5  185.185.0.0       153.60.70.2       0x80000007 0xb568 153

ASE      5  195.1.1.0         153.60.70.4       0x80000007 0xe0e6 999

ASE      5  195.2.2.0         153.60.70.2       0x80000007 0xb910 153

ASE      5  206.6.6.0         153.60.70.2       0x80000007 0xcde8 153

ASE      5  206.6.6.0         153.60.70.4       0x80000007 0xddd4 999

TRK      14 0.0.0.3           153.60.70.2       0x80000005 0xebf2 151

TRK      14 0.0.0.4           153.60.70.4       0x80000005 0xbfeb 149

TRK      14 0.0.0.9           153.60.70.2       0x80000006 0xaa31 139

TRK      14 0.0.0.35          153.60.70.4       0x80000006 0x767f 137

                # LSAs: 14

                Xsum:   0x909f7
```

show ospf database <*area ID*>

Displays a specific OSPF link-state database.

For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database 0.0.0.2

 Link-state database for area 0.0.0.2

                ---------------------------- ---------------

Type         ID                Adv-Router        Seq#       Xsum   Age

RTR (10) 1   153.60.70.2       153.60.70.2       0x8000000b 0x6642 268

RTR (10) 1   153.60.70.4       153.60.70.4       0x80000030 0x9ba9 65

ASE      5   152.148.0.0       153.60.70.2       0x80000007 0x2242 273

ASE      5   152.148.0.0       153.60.70.4       0x80000007 0x322e 1119

ASE      5   152.148.30.129    153.60.70.4       0x80000007 0xfbc2 1119

ASE      5   185.185.0.0       153.60.70.2       0x80000007 0xb568 273

ASE      5   195.1.1.0         153.60.70.4       0x80000007 0xe0e6 1119

ASE      5   195.2.2.0         153.60.70.2       0x80000007 0xb910 273

ASE      5   206.6.6.0         153.60.70.2       0x80000007 0xcde8 273

ASE      5   206.6.6.0         153.60.70.4       0x80000007 0xddd4 1119

TRK      14 0.0.0.3            153.60.70.2       0x80000005 0xebf2 271

TRK      14 0.0.0.9            153.60.70.2       0x80000006 0xaa31 259

                # LSAs: 12

                Xsum:   0x7e35a
```

show ospf database absr-summary <*area ID*>

Displays the area border link-state.

For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database absr-summary 0.0.0.0

                LSAs in Area: 0.0.0.0

                =====================

Type       ID                 Adv-Router        Seq#        Xsum    Age

ASBR     4  150.150.201.2    150.150.201.1    0x8000000c 0x52ab 411

ASBR     4  150.150.201.2    150.150.201.3    0x80000009 0x4cb2 904

ASBR     4  150.150.201.2    150.150.201.4    0x80000042 0xbf07 342

ASBR     4  150.150.201.3    150.150.201.1    0x80000006 0x405e 538

ASBR     4  150.150.201.3    150.150.201.2    0x80000006 0x4eb3 1008

ASBR     4  150.150.201.3    150.150.201.4    0x80000007 0x1884 375

ASBR     4  150.150.201.4    150.150.201.1    0x80000001 0x5eda 1137

ASBR     4  150.150.201.4    150.150.201.2    0x80000001 0x3acd 1138

ASBR     4  150.150.201.4    150.150.201.3    0x80000001 0x2082 1140
```

show ospf database absr-summary <*link state ID*> <*adv.router*>

Displays the OSPF area border link-state.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router.

```
> show ospf database absr-summary 150.150.201.7 150.150.170.3

                Link-state database for area 0.0.0.0

                ---------------------------- ---------------

Type         ID                Adv-Router       Seq#       Xsum    Age

RTR (2 ) 1   2.23.160.2        2.23.160.2        0x80000689 0xe56d 31

RTR (10) 1   150.150.170.1     150.150.170.1    0x80001a30 0x80dc 753

RTR (10) 1   150.150.170.2     150.150.170.2    0x80001f34 0x7916 638

RTR (10) 1   150.150.170.3     150.150.170.3    0x80000203 0x109d 183

RTR (10) 1   150.150.170.4     150.150.170.4    0x80001161 0x2c73 196

RTR (10) 1   150.150.170.5     150.150.170.5    0x80000ea9 0x7574 177

RTR (10) 1   150.150.170.6     150.150.170.6    0x80001287 0x2715 704

RTR (10) 1   150.150.170.7     150.150.170.7    0x80001295 0xa26e 643
```

show ospf database absr-summary <*link state ID*> <*adv.router*> <*area ID*>

Displays the area border link-state.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router. For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database absr-summary <link state ID>
<adv.router> <area ID>
```

show ospf database external <*area ID*>

Displays external links. External links display destinations beyond the autonomous system.

For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database external 0.0.0.0

                LSAs in Area: 0.0.0.0

                =====================

Type        ID              Adv-Router       Seq#        Xsum    Age

ASE      5   1.0.0.0          150.150.170.2    0x80000002 0xdb0d 121

ASE      5   116.0.0.0        150.150.170.2    0x80000005 0x8af2 1477

ASE      5   116.5.5.0        150.150.170.6    0x80000051 0x3ee8 1531

ASE      5   116.5.5.6        150.150.170.6    0x80000051 0x21f0 1531

ASE      5   123.0.0.0        150.150.170.2    0x80000019 0xe87c 1404

ASE      5   150.150.0.0      150.150.170.2    0x80000008 0xb908 1405

ASE      5   150.150.0.0      150.150.170.8    0x80000004 0x8938 904

ASE      5   152.148.0.0      150.150.170.8    0x8000000a 0x7b40 736

ASE      5   152.148.81.53    150.150.170.4    0x80000340 0x61e0 1600
```

show ospf database external <*link state ID*> <*adv.router*>

Displays OSPF external links. External links display destinations beyond the autonomous system.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router.

```
> show ospf database external 152.148.0.0 153.60.70.4

                              LSA in Area: 0.0.0.2

        LS type:        AS-external-LSA

        ________        ___________

        LS age:         1386

        LS options:     0x0

        LS ID:          152.148.0.0

        Adv rtr:        153.60.70.4

        LS Seq #:       0x80000003

        LS Xsum:        0x2a3a (good)

        LS Length:      36

        Mask:           255.255.0.0

                Cost:           1 (Type 2)

                Forward addr:   153.60.70.4

                Tag:            0
```

show ospf database external <*link state ID*> <*adv.router*> <*area ID*>

Displays external links. External links display destinations beyond the autonomous system.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router. For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database external 200.5.73.0 2.23.160.2 0.0.0.0

                            LSA in Area: 0.0.0.0

LS type:        AS-external-LSA
________        ___________
LS age:         1013
LS options:     0x22
LS ID:          200.5.73.0
Adv rtr:        2.23.160.2
LS Seq #:       0x800000f1
LS Xsum:        0x98d5 (good)
LS Length:      36
Mask:           255.255.255.0
        Cost:           1 (Type 1)
        Forward addr:   0.0.0.0
        Tag:            0
```

show ospf database names <*area ID*>

Displays OSPF names.

For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database names 0.0.0.2

                Link-state database for area 0.0.0.2

                ---------------------------- ---------------

Type        ID                Adv-Router        Seq#       Xsum   Age

RTR (10) 1  153.60.70.2       153.60.70.2       0x8000000a 0x6d43 170

RTR (10) 1  153.60.70.4       153.60.70.4       0x80000009 0x197d 143

ASE      5  152.148.0.0       153.60.70.4       0x80000003 0x3a2a 1519

ASE      5  152.148.30.129    153.60.70.4       0x80000003 0x4be0 1519

ASE      5  152.148.31.0      153.60.70.2       0x80000003 0xd375 1506

ASE      5  185.185.0.0       153.60.70.2       0x80000003 0xbd64 1477

ASE      5  195.1.1.0         153.60.70.4       0x80000003 0xe8e2 1489

ASE      5  195.2.2.0         153.60.70.2       0x80000003 0xc10c 1477

ASE      5  206.6.6.0         153.60.70.2       0x80000003 0xd5e4 1469

ASE      5  206.6.6.0         153.60.70.4       0x80000003 0xe5d0 1482

TRK      14 0.0.0.3           153.60.70.2       0x80000003 0x1bb8 1467

TRK      14 0.0.0.4           153.60.70.4       0x80000003 0xeae6 1465

TRK      14 0.0.0.9           153.60.70.2       0x80000004 0x9052 1453

TRK      14 0.0.0.35          153.60.70.4       0x80000004 0x9b81 1450

                # LSAs: 14

                Xsum:   0x7ee94
```

show ospf database names <*link state ID*> <*adv.router*>

Displays OSPF names.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router.

```
> show ospf database names 152.148.0.0 153.60.70.4

                Link-state database for area 0.0.0.2

                ---------------------------- ---------------

Type        ID                Adv-Router        Seq#       Xsum   Age

RTR (10) 1  153.60.70.2       153.60.70.2       0x8000000a 0x6d43 170

RTR (10) 1  153.60.70.4       153.60.70.4       0x80000009 0x197d 143

ASE      5  152.148.0.0       153.60.70.4       0x80000003 0x3a2a 1519

ASE      5  152.148.30.129    153.60.70.4       0x80000003 0x4be0 1519

ASE      5  152.148.31.0      153.60.70.2       0x80000003 0xd375 1506

ASE      5  185.185.0.0       153.60.70.2       0x80000003 0xbd64 1477

ASE      5  195.1.1.0         153.60.70.4       0x80000003 0xe8e2 1489

ASE      5  195.2.2.0         153.60.70.2       0x80000003 0xc10c 1477

ASE      5  206.6.6.0         153.60.70.2       0x80000003 0xd5e4 1469

ASE      5  206.6.6.0         153.60.70.4       0x80000003 0xe5d0 1482

TRK      14 0.0.0.3           153.60.70.2       0x80000003 0x1bb8 1467

TRK      14 0.0.0.4           153.60.70.4       0x80000003 0xeae6 1465

TRK      14 0.0.0.9           153.60.70.2       0x80000004 0x9052 1453

TRK      14 0.0.0.35          153.60.70.4       0x80000004 0x9b81 1450

                # LSAs: 14

                Xsum:   0x7ee94
```

show ospf database names <*link state ID*> <*adv.router*> <*area ID*>

Displays OSPF names.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router. For the *area ID*, specify the 32-bit number that identifies THE OSPF area.

```
> show ospf database names <link state ID>
<adv.router> <area ID>
```

show ospf database network <*area ID*>

Displays network links. Designated routers generate network links, which describe the set of routers attached to a particular network. Network links are flooded in the area that contains the network.

For the *area ID*, specify the 32-bit number that identifies the OSPF area.

> show ospf database network 0.0.0.2

```
                                 LSAs in Area: 0.0.0.0

                                 =====================

Type        ID                Adv-Router        Seq#       Xsum    Age

NET       2  189.51.101.2      150.150.180.6     0x8000008a 0x556f  852

NET       2  189.51.102.2      150.150.150.52    0x80000003 0xc764  830

NET       2  189.51.104.2      150.150.180.4     0x8000008a 0x10b7  879

NET       2  189.51.108.2      150.150.180.3     0x80000003 0xee5e  889

NET       2  189.51.110.2      150.150.150.52    0x80000090 0x5442  870

NET       2  189.51.112.4      150.150.150.52    0x80000090 0xe860  870

NET       2  189.51.113.2      150.150.180.5     0x80000003 0xcd75  742

NET       2  189.51.114.2      150.150.180.4     0x80000003 0xb094  875

NET       2  189.51.115.1      150.150.150.52    0x80000003 0x5ebf  803

NET       2  189.51.118.2      150.150.180.5     0x80000003 0x96a7  857

NET       2  189.51.128.1      150.150.150.52    0x80000090 0xb3cf  866

NET       2  189.51.130.1      150.150.180.4     0x80000003 0xaf65  867

NET       2  189.51.131.1      150.150.180.4     0x80000003 0xa46f  88

NET       2  189.51.134.1      150.150.180.2     0x80000003 0x7b99  890

NET       2  189.51.134.2      210.38.11.1       0x80000003 0xfbd4  1807

NET       2  189.51.135.1      150.150.180.2     0x80000003 0x70a3  866

NET       2  189.51.135.2      210.38.11.1       0x80000003 0xf0de  1807

NET       2  189.51.137.1      150.150.180.5     0x80000003 0x66a5  472
```

show ospf database network <*link state ID*> <*adv.router*>

Displays network links. Designated routers generate network links, which describe the set of routers attached to a particular network. Network links are flooded in the area that contains the network.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router.

```
> show ospf database network 189.51.137.1 150.150.180.5

                         LSA in Area: 0.0.0.0

        LS type:        Network-LSA
        ________        ___________
        LS age:         559
        LS options:     0x0
        LS ID:          189.51.137.1
        Adv rtr:        150.150.180.5
        LS Seq #:       0x80000003
        LS Xsum:        0xa566 (good)
        LS Length:      32
        Network Mask:   255.255.255.0
                Attached Router:150.150.180.5
                Attached Router:210.38.11.1
```

show ospf database network <*link state ID*> <*adv.router*> <*area ID*>

Displays network links. Designated routers generate network links, which describe the set of routers attached to a particular network. Network links are flooded in the area that contains the network.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router. For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database network  189.51.137.1 150.150.180.5 0.0.0.0

                                   LSA in Area: 0.0.0.0

        LS type:         Network-LSA
        ________         ___________
        LS age:          675
        LS options:      0x0
        LS ID:           189.51.137.1
        Adv rtr:         150.150.180.5
        LS Seq #:        0x80000003
        LS Xsum:         0xa566 (good)
        LS Length:       32
        Network Mask:    255.255.255.0
                Attached Router:150.150.180.5
                Attached Router:210.38.11.1
```

show ospf database opaque <*area ID*>

Displays opaque link-state advertisements.

For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database opaque <area ID>
```

show ospf database opaque <*link state ID*> <*adv.router*>

Displays opaque link-state advertisements.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router.

```
> show ospf database opaque <link state ID> <adv.router>
```

show ospf database opaque <*link state ID*> <*adv.router*> <*area ID*>

Displays opaque link-state advertisements.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router. For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database opaque <link state ID>
<adv.router> <area ID>
```

show ospf database router <*area ID*>

Displays router links. Router links are generated by each router belonging to an area. Router links describe the states of the router's link to the area. They are only flooded within a particular area.

For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database router 0.0.0.0

                LSAs in Area: 0.0.0.0

                =====================

Type       ID               Adv-Router       Seq#       Xsum    Age

RTR (2 ) 1  2.23.160.2        2.23.160.2        0x80000531 0x9913 97

RTR (10) 1  150.150.170.1     150.150.170.1     0x800017ec 0x1c30 441

RTR (10) 1  150.150.170.2     150.150.170.2     0x80001c86 0x9b0b 427

RTR (10) 1  150.150.170.3     150.150.170.3     0x80000cd4 0x5caa 95

RTR (10) 1  150.150.170.4     150.150.170.4     0x80000e86 0x7d63 90

RTR (10) 1  150.150.170.5     150.150.170.5     0x80000c61 0x71ed 95

RTR (10) 1  150.150.170.6     150.150.170.6     0x80001057 0x5d11 1785

RTR (10) 1  150.150.170.7     150.150.170.7     0x80000f8e 0x2c90 96

RTR (10) 1  150.150.170.8     150.150.170.8     0x800012fd 0xba09 96
```

show ospf database router <*link state ID*> <*adv.router*>

Displays router links. Router links are generated by each router belonging to an area. Router links describe the states of the router's link to the area. They are only flooded within a particular area.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router.

```
> show ospf database router  153.60.70.4 153.60.70.4
```

```
                         LSA in Area: 0.0.0.2

        LS type:        Router-LSA
        ________        ___________
        LS age:         325
        LS options:     0x0
        LS ID:          153.60.70.4
        Adv rtr:        153.60.70.4
        LS Seq #:       0x80000009
        LS Xsum:        0x7d19 (good)
        LS Length:      92
        # router type:  ASBR
        # interfaces:   5

                Link Type:      Router link
                Link ID:        153.60.70.2
                Link Data:      206.6.6.2
                TOS 0 Cost:     1
                 Link Type:      Stub Net
                Link ID:        206.6.6.1
                Link Data:      255.255.255.255
                TOS 0 Cost:     1
```

show ospf database router <*link state ID*> <*adv.router*> <*area ID*>

Displays router links. Router links are generated by each router belonging to an area. Router links describe the states of the router's link to the area. They are only flooded within a particular area.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router. For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database router 210.24.2.212 210.24.2.212
0.0.0.0
```

```
                              LSA in Area: 0.0.0.0

        LS type:        Router-LSA

        ________        ___________

        LS age:         29

        LS options:     0x2

        LS ID:          210.24.2.212

        Adv rtr:        210.24.2.212

        LS Seq #:       0x800007be

        LS Xsum:        0x82ed (good)

        LS Length:      60

        # router type:

        # interfaces:   3

                Link Type:      Stub Net

                Link ID:        210.24.2.0

                Link Data:      255.255.255.0

                TOS 0 Cost:     1

                Link Type:      Stub Net

                Link ID:        210.24.1.4

                Link Data:      255.255.255.255

                TOS 0 Cost:     1
```

show ospf database summary <*area ID*>

Displays summary links. Area border routers generate summary links, which describe inter-area routes (routes between areas) to the autonomous system border router

For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database summary 0.0.0.0

                LSAs in Area: 0.0.0.0

                =====================

Type       ID                Adv-Router       Seq#       Xsum    Age

SUMM     3  150.150.170.6    150.150.170.6    0x80000032 0x1260 1764

SUMM     3  150.150.201.1    150.150.201.1    0x80000094 0x6d6e 899

SUMM     3  150.150.201.1    150.150.201.2    0x80000006 0x7094 894

SUMM     3  150.150.201.1    150.150.201.3    0x80000006 0x5649 824

SUMM     3  150.150.201.1    150.150.201.4    0x80000008 0x6ac8 726

SUMM     3  150.150.201.2    150.150.201.1    0x8000000c 0x609e 767

SUMM     3  150.150.201.2    150.150.201.2    0x8000006a 0xb152 193

SUMM     3  150.150.201.2    150.150.201.3    0x80000009 0x5aa5 1260

SUMM     3  150.150.201.2    150.150.201.4    0x80000042 0xcdf9 698
```

show ospf database summary <*link state ID*> <*adv.router*>

Displays summary links. Area border routers generate summary links, which describe inter-area routes (routes between areas) to the autonomous system border router.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router.

```
> show ospf database summary 150.150.201.7    150.150.170.3

                              LSA in Area: 0.0.0.0

          LS type:        Summary-LSA

          ________        ___________

          LS age:         1648

          LS options:     0x0

          LS ID:          150.150.201.7

          Adv rtr:        150.150.170.3

          LS Seq #:       0x80000049

          LS Xsum:        0x8558 (good)

          LS Length:      28

                  Mask:           255.255.255.255

                  TOS 0 Cost:     351
```

show ospf database summary <*link state ID*> <*adv.router*> <*area ID*>

Displays summary links. Area border routers generate summary links, which describe inter-area routes (routes between areas) to the autonomous system border router.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router. For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database summary 150.150.170.3 150.150.170.3 0.0.0.0

                                   LSA in Area: 0.0.0.0

           LS type:        Summary-LSA
           ________        ___________
           LS age:         982
           LS options:     0x0
           LS ID:          150.150.170.3
           Adv rtr:        150.150.170.3
           LS Seq #:       0x80000028
           LS Xsum:        0x2c56 (good)
           LS Length:      28
                   Mask:           255.255.255.255
                   TOS 0 Cost:     0
```

show ospf database trunk <*area ID*>

Displays OSPF trunks.

For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database trunk 0.0.0.0

                LSAs in Area: 0.0.0.0

                =====================

Type       ID              Adv-Router       Seq#       Xsum   Age

TRK      14 0.0.0.1          150.150.170.1    0x80000bb5 0x654d 80

TRK      14 0.0.0.1          150.150.170.6    0x80000031 0x77c4 81

TRK      14 0.0.0.2          150.150.170.1    0x800000be 0xcbe3 37

TRK      14 0.0.0.3          150.150.170.6    0x800000ae 0xaf0a 42

TRK      14 0.0.0.3          150.150.170.8    0x800003aa 0x32d0 225

TRK      14 0.0.0.4          150.150.170.8    0x8000046d 0xf5d4 225

TRK      14 0.0.0.4          150.150.170.9    0x80000489 0xd2ab 226

TRK      14 0.0.0.5          150.150.170.6    0x8000055a 0xdca0 37

TRK      14 0.0.0.5          150.150.170.8    0x80000003 0x1b55 2008

TRK      14 0.0.0.5          150.150.170.9    0x800000a1 0x646d 1744
```

show ospf database trunk <*link state ID*> <*adv.router*>

Displays OSPF trunks.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router.

```
> show ospf database trunk 150.150.150.1 150.150.150.2

LSAs in Area: 0.0.0.1

                                 =====================

Type   ID              Adv-Router       Seq#       Xsum   Age

TRK     0.0.0.2         150.150.150.1    0x80000030 0x915c 646

TRK     0.0.0.14        150.150.150.2    0x8000003b 0x9c45 641
```

show ospf database trunk <*link state ID*> <*adv.router*> <*area ID*>

Displays OSPF trunks.

For the *link state ID*, specify the IP address of the network or interface you want advertised. For the *advertising router* value, specify the IP address of the advertising router. For the *area ID*, specify the 32-bit number that identifies the OSPF area.

```
> show ospf database trunk 0.0.0.103 150.150.170.2 0.0.0.0

                              LSA in Area: 0.0.0.0

        LS type:        Trunk-LSA
        ________        __________
        LS age:         653
        LS options:     0x0
        LS ID:          0.0.0.103
        Adv rtr:        150.150.170.2
        LS Seq #:       0x80000023
        LS Xsum:        0xd3a8 (good)
        LS Length:      48
                Trunk instance: 0x80001f08
                TOS 8 Cost:     1425
                TOS 16 Cost:    4
                TOS 250 Cost:   9
                TOS 253 Cost:   1024
```

show ospf interface

Displays OSPF interface information.

```
> show ospf interface
```

LPort	Interface	Area	Type	State	Cost	NFcnt	DR	BDR
0	0.0.0.0	n/a	Bcast	DR	1	4	n/a	n/a
2	189.51.114.2	0.0.0.0	NBMA	DR	200	1	150.150.180.4	150.150.180.2
3	189.51.104.2	0.0.0.0	NBMA	DR	100	1	150.150.180.4	150.150.180.2
4	189.51.101.1	0.0.0.0	NBMA	DRBack	700	1	150.150.180.6	150.150.180.4
5	189.51.130.1	0.0.0.0	NBMA	DR	200	1	150.150.180.4	210.38.11.1
6	189.51.131.1	0.0.0.0	NBMA	DR	100	1	150.150.180.4	210.38.11.1
7	203.203.203.1	0.0.0.0	p2p	Down	1	0	n/a	n/a
8	202.202.202.2	0.0.0.0	p2p	Down	1	0	n/a	n/a

show ospf names

This command prints a line for every *name* known to the switch. *Names* can be resilient UNI/NNI names, E.164 addresses, and NSAP addresses.

```
> show ospf names
```

```
Type     Cost     Name/Len Primary (Secondaries)

2        0        0x0/0 2/15

3        0        0x0/0 2/15

3        0        0x39018200000000000000000000000000000000/152 1/9

3        0        0x39028200000000000000000000000000000000/152 2/14

3        0        0x392222/24 2/13

3        0        0x39333333333333333333333333333333333333/152 2/15

3        1        0x39989898/32 1/39
```

Attributes

The following list describes each of the attributes for the show ospf names output.

Type — Lists a 1, 2, or 3 to indicate the type of name. Type 1= Resilient UNI/NNI names. Type 2 = E.164 addresses. Type 3 = NSAP addresses (type=3).

Cost — The cost of the path to the switch currently hosting the name (i.e., the name's primary location).

Name/Len Primary — This parameter indicates the name itself along with the primary location for the name. These two parameters are separated by a slash. The name is shown as a hexadecimal string. The length is shown in bits. For example, in the first line of the output shown above, the name is 0x0 and the length is 0 bits.

Secondary — Any secondary locations for the name.

show ospf namedpath <*type*> <*name*> <*bit length*>

Displays the path that a circuit would take when an SVC call is made to the name. Names can be resilient UNI/NNI names, E.164 addresses, and NSAP.

```
> show ospf namedpath 3 0x39998989898 40
```

This example would specify that the system print the path that an SVC which originated on card 8 would take to the NSAP address (type = 3) 0x39998989898/40.

Parameters

The following list describes each of the parameters for the show ospf namedpath console command.

Type — Specify a 1, 2, or 3 to indicate the type of name. Type 1= Resilient UNI/NNI names. Type 2 = E.164 addresses. Type 3 = NSAP addresses (type=3).

Name — This parameter specifies the name. The name is specified as a hexadecimal string. In the example above the name is 0x39998989898. If you do not know the name, refer to "show ospf names" on page A-88 for information on how to obtain the name.

Bit Length — The length of the name is specified in bits. For example, in the sample above, this parameter is 40. If you do not know the length, refer to "show ospf names" on page A-88 for information on how to obtain the length.

show ospf namedpath <*type*> <*name*> <*bit length*> <*slot ID*>

Displays the path that a circuit would take when an SVC call is made to the name. Names can be resilient UNI/NNI names, E.164 addresses, and NSAP.

```
> show ospf namedpath 3 0x39998989898 40 8
```

This example would specify that the system print the path that an SVC which originated on card 8 would take to the NSAP address (type = 3) 0x39998989898/40.

Parameters

The following list describes each of the parameters for the show ospf namedpath console command.

Type — Specify a 1, 2, or 3 to indicate the type of name. Type 1= Resilient UNI/NNI names. Type 2 = E.164 addresses. Type 3 = NSAP addresses (type=3).

Name — This parameter specifies the name. The name is specified as a hexadecimal string. In the example above the name is 0x39998989898. If you do not know the name, refer to "show ospf names" on page A-88 for information on how to obtain the name.

Bit Lenth— The length of the name is specified in bits. For example, in the sample above, this parameter is 40. If you do not know the length, refer to "show ospf names" on page A-88 for information on how to obtain the length.

Slot ID — The card where the SVC originated.

show ospf neighbor

Displays all OSPF neighbors.

```
> show ospf neighbor
```

```
LPort Interface        NbrId            NbrIPAddr        Pri State    Ver  #Rxmt  #LSRq  #DBsum
0     0.0.0.0          0.0.0.3          0.0.0.3          0   Full       10  0      0      0
                       0.0.0.12         0.0.0.12         0   Full       10  0      0      0
                       0.0.0.14         0.0.0.14         0   Full       10  0      0      0
                       0.0.0.16         0.0.0.16         0   Full       10  0      0      0
4     153.60.70.4      153.60.70.2      153.60.70.2      0   Full       10  0      0      0
35    153.60.70.4      153.60.70.2      153.60.70.2      0   Full       10  0      0      0
36    206.6.6.2        153.60.70.2      206.6.6.1        1   Full       2   0      0      0
```

show ospf pathdb <*slot ID*>

Displays OSPF path databases.

```
> show ospf pathdb 1

ID # VCsPath

3.11  170.7/14,170.6/1

3.20  170.1/1

3.30  170.1/1,170.6/9

3.40  170.6/1

3.50  170.4/14,170.7/14,170.6/1

3.60  170.2/9,170.1/1

3.70  170.2/9,170.1/1,170.6/9

3.80  170.9/5,170.2/9,170.1/1

3.90  170.2/9,170.1/2,170.6/9

3.100 170.9/5,170.2/9,170.1/2,170.6/9
```

show ospf pathdb <*switch ID*>

Displays OSPF path databases.

```
> show ospf pathdb <switch ID>
```

show ospf pathdb <*switch ID*> <*interface #*>

Displays OSPF path databases.

```
> show ospf pathdb <switch ID> <interface #>
```

show ospf pathdb <*switch ID*> <*interface #*> <*slot ID*>

Displays OSPF path databases.

```
> show ospf pathdb <switch ID> <interface #> <slot ID>
```

show ospf qospath <*IP address*>

Displays various QoS path statistics.

```
> show ospf qospath 153.60.70.4

   Forward BW (Kbytes):

      Reverse BW (Kbytes):

      Forward QoS (1-4):

      Reverse QoS (1-4):

      Metric (0-3):

      Current Path ID:

      S/W version (2-6):

      Characteristics:

      E-E Delay (milliseconds):

      Private Net ID:

      Void trunk (switch/IFIndex):


      Dest:       0.0.0.0

      Result:     Destination unreachable
```

show ospf qospath <*IP address*> <*slot ID*>

Displays various QoS path statistics.

```
> show ospf qospath 153.60.70.4 8

   Forward BW (Kbytes):

      Reverse BW (Kbytes):

      Forward QoS (1-4):

      Reverse QoS (1-4):

      Metric (0-3):

      Current Path ID:

      S/W version (2-6):

      Characteristics:

      E-E Delay (milliseconds):

      Private Net ID:

      Void trunk (switch/IFIndex):


      Dest:        0.0.0.0

      Result:      Destination unreachable
```

show ospf rtrs

Displays OSPF routers.

```
> show ospf rtrs

Router            State     Cost  TTL

2.23.160.2        Intra     201   2

150.150.170.1     Deleted

150.150.170.2     Intra     100   1

150.150.170.3     Intra     100   255

150.150.170.4     Intra     200   2

150.150.170.5     Intra     200   255

150.150.170.6     Intra     100   1

150.150.170.7     Intra     200   2

150.150.170.8     Intra     300   255

150.150.170.9     Intra     200   1

210.24.2.212      Deleted

210.210.210.44    Intra     101   1
```

show ospf statistics <*slot ID*>

Displays OSPF statistics.

```
> show ospf statistics 7

# switches:           8             # reachable switches: 0
# Dijkstra runs:      43            # Trunks:             0 (0)
Max LSA size:         0             # Stub links:         0 (30)
# LSAs:               0             Database checksum:    0x0
# router-LSAs:        8             # network-LSAs:       18
# AS-external-LSAs:   0             # name-LSAs:          0
# opaque-LSAs:        0
# VC lookups:         0             # VC reroute attempts:0
# successful defaults:0             # specific VC calc.:  0
# QoS failures:       0             # VC unreachables:    0
# VC reroutes:        0
# VC crankbacks       0             # Multipoint cranks:  0
 Max task latency (ms):10            Max lookup time (ms): 0
# OSPF trunk inst ch: 0             # VCMGR trunk inst ch:0
# Bad paths reg:      0             # VCs using old rev:  0
# VCMGR call backs:   0             # VCMGR rpt old inst: 0
# Trunk cost chg neg: 0             # Reg on down trunks: 0
# trunk congestion:   0             # path congestion:    0
# VC congestion:      0
# policy scans:       0             # incremental import: 0
# incremental discrd: 0             # alloc LISTs:        0
# free LISTs:         0             # fwd entries:        570
# alloc Opaques:      0             # returned Opaques:   0
Routing S/W revision: 10            Network S/W revision: 10
```

show ospf statistics <*slot ID*> <*area ID*>

Displays information pertaining to OSPF.

```
> show ospf statistics 7 0.0.0.2

      # switches:           8              # reachable switches: 0
      # Dijkstra runs:      43             # Trunks:             32 (0)
      Max LSA size:         120            # Stub links:         10 (30)
      # LSAs:               104            Database checksum:    0x35dea7
      # router-LSAs:        8              # network-LSAs:       18
      # AS-external-LSAs:   0              # name-LSAs:          0
      # opaque-LSAs:        0
      # VC lookups:         0              # VC reroute attempts:0
      # successful defaults:0              # specific VC calc.:  0
      # QoS failures:       0              # VC unreachables:    0
      # VC reroutes:        0
      # VC crankbacks       0              # Multipoint cranks:  0
      Max task latency (ms):10             Max lookup time (ms): 0
      # OSPF trunk inst ch: 0              # VCMGR trunk inst ch:0
      # Bad paths reg:      0              # VCs using old rev:  0
      # VCMGR call backs:   0              # VCMGR rpt old inst: 0
      # Trunk cost chg neg: 0              # Reg on down trunks: 0
       # trunk congestion:   0              # path congestion:    0
      # VC congestion:      0
      # policy scans:       0              # incremental import: 0
      # incremental discrd: 0              # alloc LISTs:        0
      # free LISTs:         0              # fwd entries:        570
      # alloc Opaques:      0              # returned Opaques:   0
       Routing S/W revision: 10             Network S/W revision: 10
```

show ospf trunk <*switch ID*>

Displays OSPF trunk statistics.

```
> show ospf trunk <switch ID>

sw/prt      sw/prt      fbw2/0      rbw2/0      delay    cost               comments

1.1/7       1.2/7       1900        1900        1        100

1.2/6       1.5/7       1900        1900        52       100
```

show ospf trunk <*interface #*>

Displays OSPF trunk statistics.

```
> show ospf trunk <interface #>

 sw/prt      sw/prt      fbw2/0      rbw2/0      delay    cost
comments

1.1/7       1.2/7       1900        1900        1        100

1.2/6       1.5/7       1900        1900        52       100
```

show ospf trunk <*interface #*> <*slot ID*>

Displays OSPF trunk statistics.

> show ospf trunk <*interface #*> <*slot ID*>

sw/prt	sw/prt	fbw2/0	rbw2/0	delay	cost	comments
1.1/7	1.2/7	1900	1900	1	100	
1.2/6	1.5/7	1900	1900	52	100	

show ospf trunk <*qos*>

Displays OSPF trunk statistics.

> show ospf trunk <*qos*>

sw/prt	sw/prt	fbw2/0	rbw2/0	delay	cost	comments
1.1/7	1.2/7	1900	1900	1	100	
1.2/6	1.5/7	1900	1900	52	100	

show ospf trunk <*qos*> <*priority*>

Displays OSPF trunk statistics.

> **show ospf trunk <*qos*> <*priority*>**

sw/prt	sw/prt	fbw2/0	rbw2/0	delay	cost	comments
1.1/7	1.2/7	1900	1900	1	100	
1.2/6	1.5/7	1900	1900	52	100	

show ospf vcpath <*IP address*>

This command displays various OSPF trunk statistics.

> **show ospf vcpath <*IP address*>**

```
Dest:        0.0.0.0

Result:      Destination unreachable
```

show ospf vcpath <*IP address*> <*slot ID*>

This command displays various OSPF trunk statistics.

> **show ospf vcpath <*IP address*> <*slot ID*>**

```
Dest:        0.0.0.0

Result:      Destination unreachable
```

show ospf vcroute

Displays OSPF virtual circuit routes.

```
> show ospf vcroute

Switch           Cost   fbw            rbw            path

150.150.170.2    100    37762          37762          170.1/5

150.150.170.3    100    138937         138937         170.1/7

150.150.170.4    200    33119          33164          170.1/7,170.3/15

150.150.170.5    200    138937         138937         170.1/7,170.3/21

150.150.170.6    100    20445          20445          170.1/3

150.150.170.7    200    18638/33119    18638/33164    170.1/3,170.6/5

150.150.170.8    300    138937         138937
170.1/7,170.3/21,170.5/14

150.150.170.9    200    20210          20210          170.1/5,170.2/8
```

show policy interface

Displays the interface association summary.

```
> show policy interface

Interface        Dir     MapName                          MapId From    To     Action Nets    Ref Stat Seq


189.32.4.2       export  direct-to-rip                     4     direct rip    accept 0       2   enab 1

                 export  deny-200.0.0.0-ospf-to-rip        8     ospf   rip    deny   3       1   enab 2

                 export  deny-210.0.0.0-ospf-to-rip        3     ospf   rip    deny   1       1   enab 3

                 export  deny-201.0.0.0-ospf-to-rip        9     ospf   rip    deny   2       1   enab 4

                 export  ospf-to-rip                       5     ospf   rip    accept 0       2   enab 5

193.1.3.2        export  direct-to-rip                     4     direct rip    accept 0       2   enab 1

                 export  ospf -> rip                       2     ospf   rip    deny   0       1   disb 2

                 export  ospf-to-rip                       5     ospf   rip    accept 0       2   enab 3
```

show policy neighbor

Displays the neighbor association summary.

```
> show policy neighbor

MapName                          MapId From   To     Action Nets   Ref Stat Seq


Imports for neighbor 193.1.1.1
bgp->table2                      2     bgp    table  deny   1      1   enab 1
bgp->table                       1     bgp    table  accept 0      1   enab 2
Exports for neighbor 193.1.1.1
static->bgp                      3     static bgp    accept 0      1   enab 3
```

show policy netfilter

Displays the network filter summary.

```
> show policy netfilter

Destination        Coverage  Id
200.1.0.0/16       inclusive 1
200.2.0.0/16       inclusive 2
201.1.0.0/16       inclusive 3
201.2.0.0/16       inclusive 4
200.0.0.0/8        inclusive 5
210.0.0.0/8        inclusive 6
```

show policy netlist

Displays the network filter list summary.

```
> show policy netlist

Name                             ListId     Destination        Coverage   Id

inclusive-200.0.0.0              1          200.1.0.0/16       inclusive 1

                                            200.2.0.0/16       inclusive 2

inclusive-201.0.0.0              2          201.1.0.0/16       inclusive 3

                                            201.2.0.0/16       inclusive 4

inclusive-200                    3          200.0.0.0/8        inclusive 5

inclusive-210                    4          210.0.0.0/8        inclusive 6
```

show policy ospf

Displays the OSPF route map summary.

```
> show policy ospf

MapName                          MapId From    To      Action Nets    Ref Stat Seq


rip-to-ospf                      6     rip     ospf    accept 0       1   enab 1

static-to-ospf                   7     static  ospf    accept 0       1   enab 3

direct-to-ospf                   10    direct  ospf    accept 0       1   enab 4
```

show policy routemap

Displays the route map summary.

```
> show policy routemap

bgp->table                       1     bgp     table   accept 0       1   enab 0

bgp->table2                      2     bgp     table   deny   1       1   enab 0

static->bgp                      3     static  bgp     accept 0       1   enab 0
```

show pport attributes <*slot#.port#*>

Displays physical port attributes.

```
> show pport attributes 8.8

             Receive      Transmit
Octets:      302687       992145
Frames:      3680         3689
Discards:    0            0
Errors:      0            0
```

show pport statistics <*slot #.port #*>

Displays physical port statistics. In addition, this command performs the same function as Cisco's *show bus status* command.

```
> show pport statistics 8.1

                         Receive        Transmit
Octets:                   425700          640713
Frames:                     8568            8662
Discards:                      0               0
Errors:                        0               0
Cells:                         0               0
Cell Errors:                   0
Out Discard Cells:                             0
```

show pvc attributes <*interface#.dlci#*>

Displays PVC attributes.

```
   > show pvc attributes 8.8

            Receive     Transmit

Octets:     302687      992145

Frames:     3680        3689

Discards:   0           0

Errors:     0           0
```

show pvc statistics <*interface#.dlci#*>

Displays PVC statistics.

```
> show pvc statistics 8.8

            Receive     Transmit

Octets:     302687      992145

Frames:     3680        3689

Discards:   0           0

Errors:     0           0
```

Receive and Transmit Counters

These statistics are the receive and transmit counters for the following items:

Octets — The number of bytes.

Frames — The number of frames.

DE Frames — The number of Discard Eligible frames.

ODE Frames — The number of Over Discard Eligible frames.

DE Octets — The number of Discard Eligible bytes.

ODE Octets — The number of Over Discard Eligible bytes.

FECN — The number of Forward Explicit Congestion Notification frames.

BECN — The number of Backward Explicit Congestion Notification frames.

Discards — The number of frames discarded (receive only).

Lost Frames — The number of frames lost (transmit only).

DE Frames Lost — The number of lost Discard Eligible frames.

ODE Frames Lost — The number of Over Discard Eligible frames lost.

Lost Octets — The number of lost bytes.

DE Octets Lost — The number of lost Discard Eligible bytes.

ODE Octets Lost — The number of lost Over Discard Eligible bytes.

show rip attributes <*IP address*>

Displays RIP attributes for a specific IP interface.

```
> show rip attributes 152.148.31.121

      RIP Attribute on IP interface [152.148.31.121]

      Admin Status:                  Disable

      Send Host Routes:              Enable

      Split Horizon:                 Simple

      Lport ID:                      4097

      Auth Type:                     None

      RIP Send:                      doNotSend

      RIP Receive:                   rip1OrRip2

      RIP Default Metric:            0

      RIP SRC Address:               152.148.31.121
```

show rip attributes all

Displays RIP attributes for all RIP interfaces.

```
> show rip attributes all

RIP Attribute on IP interface [152.148.86.4]
Admin Status:                 Disable
Send Host Routes:             Disable
Split Horizon:                Simple
Lport ID:                     4097
Auth Type:                    None
RIP Send:                     rip1Compatible
RIP Receive:                  rip1OrRip2
RIP Default Metric:           0
RIP SRC Address:              152.148.86.4
RIP Attribute on IP interface [192.32.7.1]
Admin Status:                 Enable
Send Host Routes:             Enable
Split Horizon:                Simple
Lport ID:                     2
Auth Type:                    None
RIP Send:                     rip1Compatible
RIP Receive:                  rip1OrRip2
RIP Default Metric:           0
RIP SRC Address:              192.32.7.1
show rip peer all
```

show rip peer <*IP address*>

Displays RIP peers for a specific interface.

```
> show rip peer 192.32.10.1

RIP Peer [192.32.10.1]

Peer Lport:                1

Peer Last Update:          699819

Peer Version:              2

# Bad Pkts Received from:  0

# Bad Routes Received from:0
```

show rip peer all

Displays RIP peers for all interfaces.

```
> show rip peer all

      RIP Peer [192.32.10.1]

        Peer Lport:                    1

        Peer Last Update:              699819

        Peer Version:                  2

        # Bad Pkts Received from:      0

        # Bad Routes Received from:    0
```

show rip statistics <*IP address*>

Displays RIP statistics for a specific interface.

```
> show rip statistics 192.32.7.1

          RIP is enabled on IP interface [192.32.7.1]

         # Requests Received:             0

         # Responses Received:            2073

         # Packets Received:              2073

         # Bad Packets Received:          0

         # Bad Routes Received:           0

         # Requests Send:                 0

         # Full Updates Send:             1928

         # Triggered Updates Send:        0

         # Updates Send:                  1928

         # Packets Send:                  1928

RIP is disabled on IP interface [152.148.31.121]
```

show rip statistics all

Displays RIP statistics for all interfaces.

```
> show rip statistics all

 RIP is disabled on IP interface [152.148.86.4]


        RIP is enabled on IP interface [192.32.7.1]


        # Requests Received:           0

        # Responses Received:          2070

        # Packets Received:            2070

        # Bad Packets Received:        0

        # Bad Routes Received:         0

        # Requests Send:               0

        # Full Updates Send:           1926

        # Triggered Updates Send:      0

        # Updates Send:                1926

        # Packets Send:                1926
```

show software disk <*slot #*>

Displays disk files for a specific CP-ONYX card.

```
> show software disk 3

Active CP Software::

Part#         Revision      Size     Description

7000910100    5.00.00.00    970388   CP Application [C3-Q000007FD]
7000900100    4.06.00.00    167849   CP Boot FLASH
                           1138237   Total

7000910200    5.00.00.00    705492   IOPA Application [C3-Q00000391]
7000900200    4.06.00.01    121162   IOPA Boot FLASH
                            826654   Total

7000914600    5.00.00.00    749620   IOPB Application [C3-Q0000034A]
7000904600    4.06.00.02    147556   IOPB Boot FLASH
                            897176   Total

7000914700    5.00.00.00    740197   IOPC Application [C3-Q00000167]
7000904700    5.00.00.00    165490   IOPC Boot FLASH
                            905687   Total
```

show software disk all

Displays disk files for all CP-ONYX cards.

```
> show software disk all

Active CP Software::

Part#         Revision     Size     Description

7000910100    5.00.00.00   898710   CP Application [C3-Q00000743]

7000900100    4.06.00.00   167849   CP Boot FLASH

                          1066559   Total

7000910200    5.00.00.00   705277   IOPA Application [C3-Q0000031F]

7000900200    4.06.00.01   121162   IOPA Boot FLASH

                           826439   Total

7000914600    5.00.00.00   749175   IOPB Application [C3-Q000002D9]

7000904600    4.06.00.02   147556   IOPB Boot FLASH

                           896731   Total

7000914700    5.00.00.00   739810   IOPC Application [C3-Q00000148]

7000904700    5.00.00.00   165490   IOPC Boot FLASH

                           905300   Total
```

show software flash <*slot #*>

Displays software version information for files on PCMCIA disks.

```
> show software flash 3

Active CP Software::


Part#         Revision      Size     Description


7000910100    5.00.00.00    970388   CP Application [C3-Q000007FD]

7000900100    4.06.00.00    167849   CP Boot FLASH

                           1138237   Total


7000910200    5.00.00.00    705492   IOPA Application [C3-Q00000391]

7000900200    4.06.00.01    121162   IOPA Boot FLASH

                            826654   Total


7000914600    5.00.00.00    749620   IOPB Application [C3-Q0000034A]

7000904600    4.06.00.02    147556   IOPB Boot FLASH

                            897176   Total


7000914700    5.00.00.00    740197   IOPC Application [C3-Q00000167]

7000904700    5.00.00.00    165490   IOPC Boot FLASH

                            905687   Total
```

show software flash all

Displays software version information for files on PCMCIA disks. The following example displays flash for all cards.

```
> show software flash all

Active CP Software::

Part#         Revision      Size    Description

 7000910100    5.00.00.00    898710  CP Application [C3-Q00000743]

7000900100    4.06.00.00    167849  CP Boot FLASH

                           1066559  Total

7000910200    5.00.00.00    705277  IOPA Application [C3-Q0000031F]

7000900200    4.06.00.01    121162  IOPA Boot FLASH

                            826439  Total

7000914600    5.00.00.00    749175  IOPB Application [C3-Q000002D9]

7000904600    4.06.00.02    147556  IOPB Boot FLASH

                            896731  Total

7000914700    5.00.00.00    739810  IOPC Application [C3-Q00000148]

7000904700    5.00.00.00    165490  IOPC Boot FLASH

                            905300  Total
```

show software card <*slot #*>

Displays software version information for a specific card.

```
> show software card 3

IOPA Application

Copyright (c) 1991-1997 Ascend Communications

Software Revision: 05.00.00.00

Software ID:       C3-Q00000391

Image generated:   19-Aug-1997.07:49:20
```

show software card all

Displays software version information for all cards.

> show software card all

```
Slot Type         Redund  SW Revision Software ID  Date
1    CP-40        Active  05.00.00.00 C3-Q00000743 12-Aug-1997.08:33:32
3    UIO-8        Active  05.00.00.00 C3-Q0000031F 12-Aug-1997.08:37:30
12   HSSI-2       Active  05.00.00.00 C3-Q000002D9 12-Aug-1997.03:31:56
14   ATM-IWU      Active  05.00.00.00 C3-Q00000148 12-Aug-1997.03:15:49
16   ATMDS3-1     Active  05.00.00.00 C3-Q000002D9 12-Aug-1997.03:31:56
```

show static

Displays static routes.

> show static

```
Dest                   Next_hop             Intf Pri  Tag  Filt Disc Stat
210.11.1.0/24          195.2.1.206          0    1    1    2    2    1
111.1.1.1/32           111.1.1.2            0    1    1    2    2    1
```

show system

This command displays general system information and status.

```
> show system

Switch Name:    schultz

System Desc:    Ascend Communications Corporation B-STDX 9000

Model:          651-09000-00

Location:

Contact:

System State:   Active

Uptime:         47 minutes 30 seconds

Current time:   Mon Aug 25 20:46:07 1997 UTC

Serial Number:  0040FB02F417         Internal IP Addr:   150.150.250.2

Hardware Rev:   07                   Ethernet IP Addr:   152.148.86.4

EPROM Rev:      01.00.00.00          Network Wide Addr:  150.150.0.0

Software Rev:   05.00.00.00          Network Mask:       255.255.0.0

Ethernet Addr:  00:40:FB:02:F4:17

Slot Type            Redund  State    SW Rev      HW Rev      EPROM       Serial#

                      State

1    CP-40           Active  Active   05.00.00.00 07          01.00.00.00 28A29100

4    FT1-4-24        Active  Active   05.00.00.00 06.00.00    04.01.00.00 0040FB028610

5    UT1-4-24        Active  Active   05.00.00.00 03.00.00    04.01.00.00 0040FB02908C

6    FE1-4-30        Active  Active   05.00.00.00 02.00.00    04.03.00.00 0040FB02B421

7    UIO-8           Active  Active   05.00.00.00 05.00.00    04.03.00.00 0040FB02BA23

8    DSX1-10         Active  Active   05.00.00.00 03.00.00    04.05.00.00 0040FB0283F7

10   UE1-4-30        Active  Active   05.00.00.00 03.00.00    04.04.00.00 0040FB02AB59

11   UE1-4-30        Standby Active   05.00.00.00 03.00.00    04.01.00.00 0040FB0299F7

12   UE1-4-30        Active  Active   05.00.00.00 03.00.00    04.01.00.00 0040FB0299E9

13   SFT1-4          Active  Active   05.00.00.00 134.00.00   04.01.00.00 0040FB02

B4BD

16   HSSI-2          Active  Active   05.00.00.00 02.00.00    04.00.00.00 0040FB02054F
```

show tcp

This command displays basic counters on TCP (Transmission Control Protocol) such as the number of TCP connections requested and the number of TCP connections denied.

```
> show tcp

TCP Statistics:
 Connections:            1 Attempted                2 Accepted
                         3 Established              0 Dropped
                         2 Closed                   1 Embryonic drop
                      2075 Rtt timed             2075 Rtt updated
                      1930 Delayed acks             0 Timeout drop
                         0 Retransmit timeout              0 Persist timeout
                         0 Keepalive timeout               0 Keepalive probe sent
                         0 Keepalivedrops

 Sent:             4008 Total pkts              2073 Total data pkts
                  39687 Total bytes             1933 Ack only packets
                      0 Window probes             0 URG only packet
                      0 Update only pkt           2 Control (SYN/FIN/RST) pkt
                      0 Data packets retransmitted            0 Data bytes retransmitted

Received:          4145 Total pkts                   2070 In sequence
                  36915 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
                      2 Duplicate acks                  0 Acks for unsent data
                   2075 Ack packets                 39690 Bytes acked
                      0 Duplicate packets               0 Duplicate bytes
                      1 Out of order packets            0 Out of order bytes
                      0 Window update                 117 Pkts with duplicate data
                      2 Duplicate bytes in partially duplicate packets
```

show udp

This command displays counters on UDP (User Datagram Protocol), such as the number of UDP packets transmitted and received and the number of errors that occurred.

```
> show udp

UDP Statistics:

                30271 total input packets,              18690 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
```

show users

This command displays users currently logged on to the switch through console or telnet.

```
> show users


user          location     time

your name     console      14 minutes 21 seconds
```

B

PRAM Upload Supported Objects

This appendix describes the Upload PRAM feature and supported IP objects.

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 Network Configuration Guide for B-STDX/STDX 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. You can access the following features from the Card PRAM Upload and NMS Synchronization dialog box:

View—Enables you to view the NMS and switch database differences, before performing an upload. Refer to page 12-20 in the *Network Configuration Guide for B-STDX/STDX* for details.

Update NMS Database — Enables you to update the NMS database and decrements the number of records that have been updated. Refer to Refer to page 12-20 in the *Network Configuration Guide for B-STDX/STDX* for details.

The Upload PRAM function requires the following software:

- CascadeView/UX 1.5.x or above.
- Switch Release 4.0.3.xx or 4.1.x or above, where xx is greater than 50 for B-STDX 8000/9000 models.
- Switch Release 2.4 for STDX 3000/6000 models.

Guidelines for Using Upload PRAM

Before you use the Upload PRAM function, review the following points:

- If you configure IP parameters using the Command Line Interface (CLI) instead of using CascadeView/UX, you need to upload the switch configuration to the NMS database.
- 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, the NMS database already contains a switch with that name.
- If you need to add SMDS Group Addresses and Group Address members, upload PRAM information from the CP before you upload PRAM from an I/O module.
  - The CP upload creates the Group Address in the CP.
  - The I/O module upload adds the members into the Group Address.
- If you need to add SMDS Address Screens and Address Screen members, upload PRAM information from the CP before you upload PRAM from an I/O module.

- The CP upload creates the Address Screen in the CP.
- The I/O module upload adds the members into the Address Screen.

• 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. Figure B-1 shows the relationships (parent/child) that exist between CascadeView/UX objects in the database.

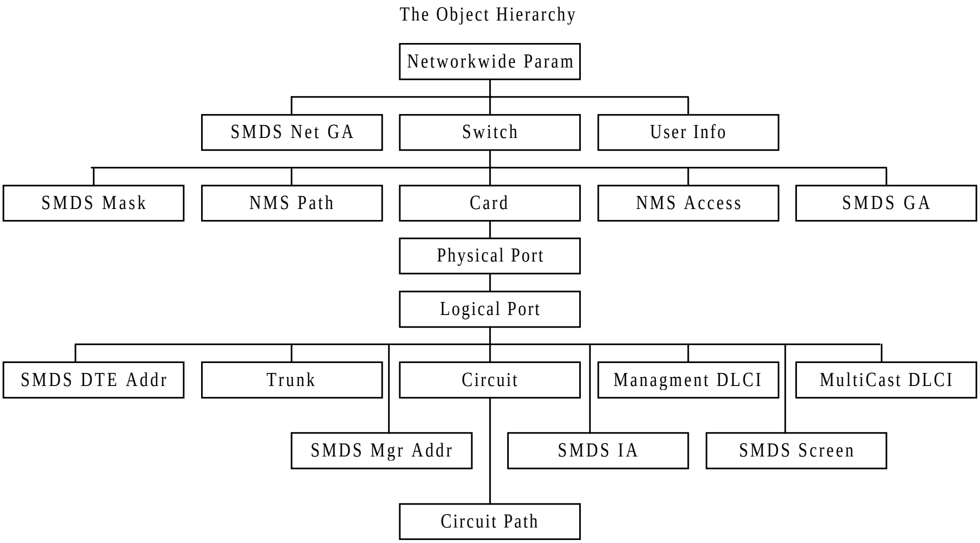

Figure B-1. CascadeView/UX Object Hierarchy

Supported Upload PRAM IP Objects

> *If you configure IP parameters using the Command Line Interface (CLI) instead of using CascadeView/UX, you need to upload the switch configuration to the NMS database.*

Upload PRAM currently supports the following IP objects:

IP Interface

- OSPF
- OSPF Neighbor
- OSPF Interface

BGP

- BGP Params
- BGP Peers
- BGP Networks
- BGP As
- BGP Aggregates

RIP

Route Filters

- Network Filters
- Network Access Lists
- Tag Filters
- Tag Access Lists
- As Path Filters

- As Access Lists
- Community Filters
- Route Policies

Uploading a Switch Configuration File

To upload the switch configuration file stored in PRAM, use the following steps:

1. On the network map, select the switch object.
2. From the Administer menu, select Cascade Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears, displaying the back panel of the selected switch.
3. Select either the I/O module or the CP module and choose the PRAM command. The CascadeView PRAM Sync dialog box appears.

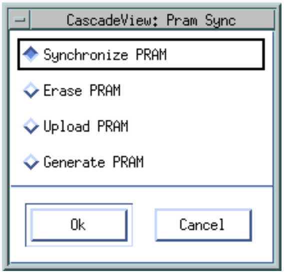

Figure B-2. Pram Sync Dialog Box

4. Select Upload PRAM and choose OK. The Card PRAM Upload and NMS Synchronization dialog box appears as shown in Figure B-3.

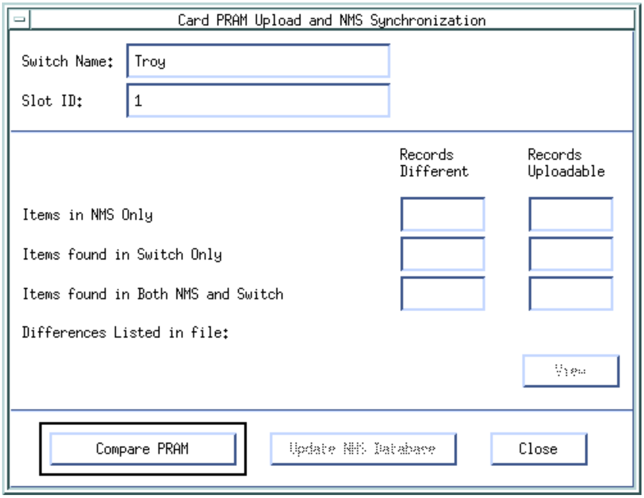

Figure B-3. Card PRAM Upload and NMS Synchronization Dialog Box

5. Choose Compare PRAM.

6. The dialog box displays information about the number of inconsistencies between the PRAM configuration file and the NMS database. If the field displays a zero, there are no differences between the PRAM and NMS configurations.

   An *item* can be a single physical port or logical port definition. This dialog box displays the following fields:

   Items in the NMS only — The item exists in the NMS database, but not in the switch PRAM. This situation occurs when you make configuration changes to an unmanaged switch.

   Items found in Switch only — The item exists in switch PRAM, but not in the NMS database. This situation occurs when you configure a switch using a third-party network management station or use the MIB to change configuration information.

Items found in both NMS and Switch — This item exists in both places, but there are discrepancies in the configuration. This can happen if you modified the configuration directly from the console. For example, if you used console commands to change the admin status of a logical port, the logical port definition in switch PRAM indicates that the logical port is Down; the NMS database records indicate the logical port as Up. These discrepancies can also occur if a PRAM synchronization or SET fails.

The name and location of the file that stores the inconsistencies appears on the dialog box.

7. Choose View to compare the files. (See the example in Figure B-4).

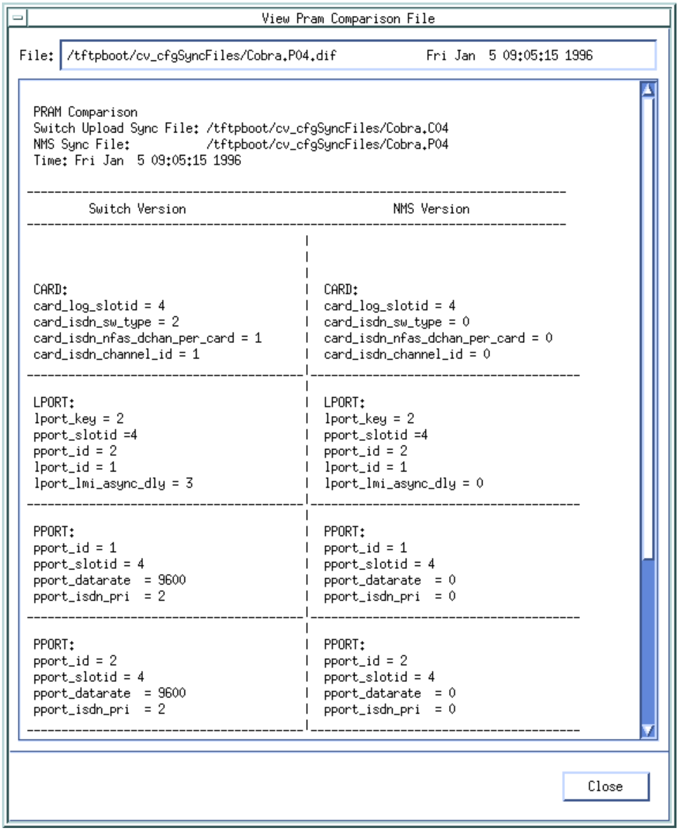

Figure B-4. View PRAM Comparison File Dialog Box

8. Choose Close to return to the Card PRAM Upload and NMS Synchronization dialog box.

9. To synchronize the information between switch PRAM and the NMS database, you can:

   - Choose the Update NMS Database command to use the configuration stored in switch PRAM.

   - To use the configuration stored in the NMS database, choose Close. Use the Synchronize PRAM command to update PRAM (refer to Refer to page 12-8 in the *Network Configuration Guide for B-STDX/STDX* for details).

10. Repeat Step 3 through Step 9 for each I/O module to complete the configuration upload process.

> *If an error occurs during the upload process, a message dialog appears. After closing this dialog box, you can choose Update NMS Database to continue the upload process for the remaining physical port and logical port definition.*
>
> *If there are problems with the PRAM configuration file, refer to Refer to page 12-2 in the Network Configuration Guide for B-STDX/STDX for details for instructions to download the configuration file stored in the NMS database.*

Index