NavisCore Frame Relay Configuration Guide

Ascend Communications, Inc.

Product Code: 80071
Revision 00
September 1998

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

Chapter 6 Configuring PVCs

Chapter 7 Configuring Management Paths

Chapter 8 Configuring Virtual Private Networks

Chapter 9 Configuring Fault-Tolerant PVCs

Chapter 10 Configuring Switched Virtual Circuit (SVC) Parameters

Chapter 11 Closed User Groups

Chapter 12 Port Security Screening

List of Figures

List of Tables

About This Guide

The *NavisCore Frame Relay Configuration Guide* provides detailed instructions for using NavisCore to configure Frame Relay services on an Ascend switch network. Specifically, this guide describes how to configure logical ports, trunks, permanent virtual circuits (PVCs), and switched virtual circuits (SVCs) to support Frame Relay services on either a B-STDX or CBX switch. This guide also explains how to configure a variety of features that enhance the Frame Relay service platform, including virtual private networks, closed user groups, and port security screening.

This guide also describes all the Frame Relay features supported in the following releases:

- NavisCore, Release 4.0
- B-STDX switch software, Release 6.0
- CBX switch software, Release 3.0

What You Need to Know

As a reader of this guide, you should know UNIX operating system commands and be familiar with HP OpenView. System administrators should be familiar with relational database software to properly maintain Sybase (the database used by NavisCore).

Before you read this guide, read the software release notice that accompanies the software. This guide assumes that you have installed the Ascend switch hardware, using one of the following guides:

- *B-STDX 8000/9000 Hardware Installation Guide*
- *CBX 500 Hardware Installation Guide*

Before you configure Frame Relay services, see the *NavisCore Physical Interface Configuration Guide* to configure processor and I/O modules and physical ports.

Reading Path

This section describes all of the documents that support the NavisCore Network Management Station (NMS) and switch software. The documents are grouped as follows:

- NMS Documentation
- Switch Software Documentation

NMS Documentation

Read the following documents to install and operate NavisCore Release 4.0.

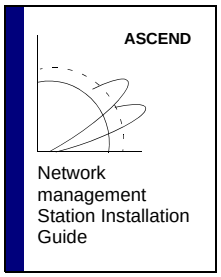

This guide describes prerequisite tasks, hardware and software requirements, and instructions for installing Solaris, HP OpenView, and naviscore on the NMS.

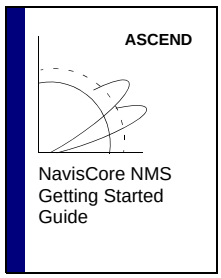

This guide describes how to configure and manage NavisCore, network maps, and Ascend switches.

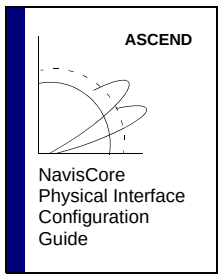

This guide describes how to configure processor and I/O modules on Ascend switches.

Switch Software Documentation

Read the following documents to configure switch software for B-STDX Release 6.0, CBX Release 3.0, and GX Release 1.0.

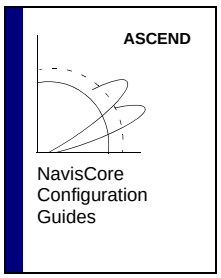

These guides describe how to configure WAN services on the STDX, B-STDX, CBX, and GX switch platforms:

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

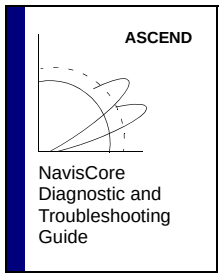

This guide describes how to diagnose and troubleshoot your NavisCore switch network.

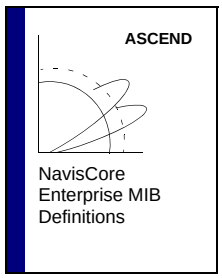

This document gives a brief overview of SNMP and describes the NavisCore Enterprise MIB definitions.

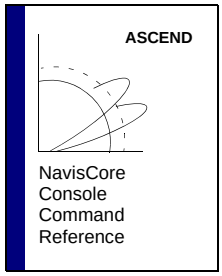

This reference lists and describes the NavisCore switch console commands.

How to Use This Guide

Before you read this guide, read the Software Release Notice (SRN) that accompanies the software. The following table provides a brief outline of this guide.

Read	To Learn About
Chapter 1	How the information in this guide is organized.
Chapter 2	Concepts you need to understand before you configure Frame Relay logical ports.
Chapter 3	Configuring Frame Relay logical ports on a B-STDX or CBX switch.
Chapter 4	Configuring Frame Relay trunks.
Chapter 5	Concepts you need to understand before you configure permanent virtual circuits (PVCs) in your network.
Chapter 6	Configuring Frame Relay PVCs, including Point-to-Point Protocol. This chapter also explains how to configure multicast data link connection identifiers (DLCIs).
Chapter 7	Configuring management PVC and management DLCI connection paths between the NMS or Internet Protocol (IP) host that you use to access the switch network.
Chapter 8	Configuring your Frame Relay services to provide virtual private networks (VPNs).
Chapter 9	Configuring fault tolerant (resilient UNI) PVC services to provide backup services should a logical port endpoint fail.
Chapter 10	Configuring Frame Relay switched virtual circuits (SVCs).
Chapter 11	Closed user groups (CUGs) that enable you to divide all network users into logically linked groups of users.
Chapter 12	Using the Port Security Screening feature to create screens that allow/disallow incoming and outgoing calls.
Appendix A	Reliable Scalable Circuit error messages and corrective actions.
Appendix B	Abbreviations and acronyms used in this guide.

What's New in This Release?

Table 1 lists the new product features that are supported in this release.

Table 1. NavisCore Release 04.00.00.00 Feature

Feature	Description	Described in
Priority Frame	Provides the following ATM-like Quality of Service (QoS) classes for frame-based B-STDX cards: • Real Time Variable Frame Rate (VFR-RT) – Provides committed bandwidth, low delay, low delay variation, and low frame loss service. VFR-RT service is used for special delay-sensitive applications, such as SNA and voice, which require low delay between end points. • Non-Real Time Variable Frame Rate (VFR-NRT) – Guarantees the loss ratio, and provides committed bandwidth, higher delay, and low frame loss. This service enables LAN-to-LAN and business class Internet/intranet access services. This service also offers *configurable* congestion control support. • Unspecified Frame Rate (UFR) – Provides no guarantees except for throughput. UFR provides a best-effort service that uses any remaining bandwidth. This service enables e-mail, file transfer, and residential Internet access services.	Chapter 2 and Chapter 3
Consolidated Link Layer Management (CLLM)	Transmits congestion threshold messages to the user device on the network for PVCs only. The CLLM message lists the DLCIs that correspond to the congested Frame Relay bearer connections. You can enable or disable CLLM on any Frame Relay UNI or NNI port.	Chapter 2
CBX 500 6-port DS3 Frame module *(revision 4 or greater)*	Supports high-speed Frame Relay access at DS3 rates. It provides capability to support native Frame Relay service on the CBX.	Chapter 2 and Chapter 3
B-STDX 1-port Channelized DS3-1-0 module *(requires switch software release 4.4)*	Provides 28 fractional T1 connections for frame-based traffic, and supports up to 488 DS0s per module. It allows flexibility with DTE/DCE and NNI Frame Relay interfaces, and with non-Frame Relay services via direct FRAD and Point-to-Point Protocol (PPP) according to RFC 1490.	Chapter 2 and Chapter 3

Table 1. NavisCore Release 04.00.00.00 Feature (Continued)

Feature	Description	Described in
Multilink Frame Relay (MLFR) Trunks	Provides a method of aggregating available bandwidth on a set of Frame Relay logical links on the same I/O module between a pair of B-STDX switches. The aggregated links, collectively referred to as the Multilink Frame Relay Bundle, can be used as a single logical link. Because it is a software solution, MLFR trunks provide cost-effective high speed trunking without using additional hardware.	Chapter 3
Switched Virtual Circuits (SVCs)	Provides as-needed logical connections across the Frame Relay network to any other node in the network. To support SVC services, each user endpoint is assigned a unique address that identifies the endpoint and enables the network to route the call. You can use E.164 or X.121 address formats to configure your network for SVCs.	Chapter 3 and Chapter 10
Multiple Open Shortest Path First (OSPF) Area Support	Allows large networks to move to a hierarchical architecture. This architecture can improve the performance of route look-ups and reduce routing table size.	Chapter 4 and Chapter 9
List PVCs that traverse trunk	The Set or Show All Trunks dialog box now lists the total number of PVCs that traverse the selected trunk, and provides logical port descriptions for each PVC endpoint.	Table 4-1 on page 4-7
Least OSPF delay metric routing and Admin Cost	Enables you to manage PVC routing using either OSPF Admin Cost or end-to-end delay metrics. These attributes are specified from the Add PVC dialog box. If you enable admin cost metrics, the PVC is routed over a path whose total administrative cost does not exceed the specified value. The NMS calculates the Admin Cost for a path by using the sum of the Admin Cost of each trunk in the path. If you enable end-to-end delay, the PVC is routed over a path whose total end-to-end delay does not exceed the specified value. The NMS calculates the total end-to-end delay for a path by using the sum of the end-to-end delays for each trunk in the path.	Chapter 6
Management PVC (MPVC)	Provides access to the switching network's management plane (which is IP-based). MPVCs offer an efficient, high performance data path capable of transferring large amounts of management data, such as NavisXtend Accounting or Statistics Server files.	Chapter 6

Table 1. NavisCore Release 04.00.00.00 Feature (Continued)

Feature	Description	Described in
Management DLCI Loop-back	Enables you to include PVC configuration information in the NMS initialization script file. This script file contains all the SNMP set requests necessary to replicate the entire switch configuration. Once you download this file to the switch, this PVC can be used to establish NMS-to-switch connectivity. This option is especially useful in some management DLCI configurations.	Table 6-1 on page 6-4
Reliable Scalable Circuit	Improves reliability when provisioning PVCs and is set to On by default. The NMS now verifies that the card state is up for each of the endpoints of a PVC before setting the cards in the switches. If the card status of either endpoint is not up, the system displays an error message indicating where the failure occurred. An abort option is provided from the error message to allow you to cancel the operation and prevent a card out-of-sync condition.	Chapter 6 and Appendix A
Network ID Addressing (for SVCs)	Identifies an inter-exchange carrier (IXC). This feature allows you to associate a network-to-network connection with a particular IXC and enables users to subscribe to a particular IXC and override this selection on a call-by-call basis.	Chapter 10
Closed User Groups (CUGs)	Divides all SVC network users into logically linked groups of users.	Chapter 11
Port Security Screening	Prevents your network from being compromised by unauthorized SVC access.	Chapter 12

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*]**
Boldface	User input in text.	Type **cd install** and...
Menu => Option	Select an option from the menu.	NavisCore => Logon
Italics	Book titles, new terms, and emphasized text.	*Network Management Station Installation Guide*
Boxes around text	Notes, warnings, cautions.	See examples below.

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

Related Documents

This section lists the related Ascend documentation that you may find helpful to read.

- *NavisCore Reading Roadmap* (Product Code: 80069)
- *B-STDX 8000/9000 Hardware Installation Guide* (Product Code: 80005)
- *CBX 500 Hardware Installation Guide* (Product Code: 80011)
- *Network Management Station Installation Guide* (Product Code: 80014)
- *NavisCore NMS Getting Started Guide* (Product Code: 80070)
- *NavisCore Physical Interface Configuration Guide* (Product Code: 80080)
- *NavisCore ATM Configuration Guide* (Product Code: 80072)
- *NavisCore Diagnostic and Troubleshooting Guide* (Product Code: 80074)
- *NavisCore Console Command Reference* (Product Code: 80075)
- *NavisXtend Accounting Server Administrator's Guide* (Product Code: 80046)

Customer Comments

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

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

Customer Support

To obtain patch software, release notes, or support, access the Ascend FTP Server or contact the Technical Assistance Center (TAC) at:

- 1-800-DIAL-WAN or 1-978-952-7299 (U.S. and Canada)
- 0-800-96-2229 (U.K.)
- 1-978-952-7299 (all other areas)

1

Overview

This chapter gives an overview of the information described in this guide. Some chapters provide information about Frame Relay network basics such as logical ports, trunks, and PVCs; other chapters explain how to configure optional features such as virtual private networks (VPNs) and closed user groups (CUGs). For more information about how various Frame Relay options can improve your network services, see the *Networking Services Technology Overview* or consult your Ascend representative.

Logical Ports

The following chapters describe Frame Relay logical ports:

- Chapter 2 provides an overview of Frame Relay logical port types and features. Read this chapter to learn about congestion control, Consolidated Link Layer Management (CLLM) notification, and Priority Frame Quality of Service (QoS).
- Chapter 3 describes how to configure Frame Relay logical ports on an Ascend switch. This chapter describes how to configure the following types of logical ports: UNI DCE/DTE, NNI, OPTimum PVC, FRAD, Direct Trunk, Point-to-Point Protocol (PPP) according to RFC 1490, and Multilink Frame Relay (MLFR) ML Member.

Trunks

Chapter 4 provides an overview of trunks and describes how to configure backup trunks and add the trunk-line connection. You can configure the following types of Frame Relay trunks:

- Frame Relay direct line trunks
- Frame Relay OPTimum PVC trunks
- MLFR direct line trunks

For information about these trunk types, review the trunk logical port descriptions in Chapter 2 and Chapter 3.

PVCs

The following chapters describe permanent virtual circuits (PVCs):

- Chapter 5 provides an overview of PVC features such as circuit routing priority, rate enforcement, and administrative tasks.
- Chapter 6 describes how to configure PVCs, specifically standard Frame Relay point-to-point PVCs. This chapter also describes how to configure fault tolerant PVCs and multicast DLCIs.
- Chapter 7 describes how to configure a management path between the Network Management Station (NMS) or Internet Protocol (IP) host to access the switch network. Use this chapter to configure both management PVCs and management DLCIs.

Virtual Private Networks

Chapter 8 describes a virtual private network (VPN), which is an *optional* software feature that enables network providers to dedicate resources for those customers who require guaranteed performance, reliability, and privacy. Use the instructions in this chapter to configure VPN services.

Fault Tolerant PVCs

Chapter 9 describes an optional logical port feature called fault tolerant PVC (sometimes referred to as resilient UNI). A fault tolerant PVC configuration enables a UNI DCE or DTE logical port to serve as a backup for any number of active UNI ports. If a primary port fails or if you need to take a primary port off-line for maintenance, you activate the backup port.

SVCs

Chapter 10 provides an overview of switched virtual circuit (SVC) concepts such as address formats, node prefixes, and network ID addressing. This chapter also describes how to configure SVC node and port prefixes and port addresses for each address format.

Closed User Groups

Chapter 11 describes closed user groups (CUGs). You can use CUGs to divide all SVC network users into logically linked groups of users. Members of the same CUG have particular calling privileges that members of different CUGs may not have. CUGs form one level of security between network users, allowing only those users who are members of the CUG to set up calls to each other.

Port Security Screening

Chapter 12 describes Port Security Screening, which is a mechanism you can use to ensure that your network cannot be compromised by unauthorized SVC access. You do this by creating screens that allow or disallow incoming and outgoing SVCs.

2

Frame Relay Services

This chapter describes Ascend Frame Relay networking services as defined by their logical port types. Frame Relay is the first packet-mode interface to Integrated Services Digital Network (ISDN) networks. Frame Relay offers the following advantages:

- Accommodates bursty LAN traffic
- Provides reliability without error-correction overhead
- Relies on upper-layer protocols at the workstation level
- Runs at very high speeds

For more information about Ascend's implementation of Frame Relay, see the *Networking Services Technology Overview*. To configure Frame Relay logical ports, see Chapter 3, "Configuring Frame Relay Logical Ports."

About Frame Relay Logical Ports

A single physical port may contain multiple logical port configurations. The logical port differs from the physical port configuration in that the physical port parameters specify only *clocking* and *clock speeds*. A logical port definition specifies how each channel is to communicate with the attached equipment. For example, a logical port configured as a Frame Relay *User Network Interface-Data Communication Equipment* (UNI-DCE) indicates that the port acts as the network for link management purposes. This UNI port may be physically set to provide clocking (defined as *Data Communications Equipment* or DCE) or no clocking (defined as *Data Terminal Equipment* or DTE).

The Set All Logical Ports in PPort dialog box, described on page 3-4, enables you to add, modify, or delete logical port configurations for a specified physical port. The logical port configuration defines which link management protocol is used, the amount of bandwidth allocated, and the individual link timer parameters.

Logical Port Types

Table 2-1 lists and describes the different types of Frame Relay logical ports.

Table 2-1. Frame Relay Logical Port Types

Logical Port Type	Physical Port Configuration	Description
Frame Relay Switch (UNI-DCE)	Frame Relay UNI-DCE	Performs link management and expects a Frame Relay DTE device to be attached. Frame Relay DTE devices refer to those user devices that perform the LMI/DTE and F or B protocols such as routers, bridges, cluster controllers, and front-end processors, or packetized voice and video.
Relay Feeder (UNI-DTE)	Frame Relay UNI-DTE	Specified for link management. Select this option to connect to a Frame Relay DCE (network switch) where the Ascend switch acts as the DTE. You can also use this logical port type as the link between two Ascend switches when configuring a Frame Relay OPTimum trunk on the same physical port.
Frame Relay NNI	Network-to-Network Interface	Functions according to the Frame Relay Forum NNI Specification. NNI enables two different switches or networks to connect together using a standard protocol. The NNI port performs both the DTE and DCE *Link Management Interface* (LMI) function. You can also use this port as the link between two Ascend switches when configuring a Frame Relay OPTimum trunk on the same physical port.
Frame Relay OPTimum PVC Trunk	Switch-to-switch Ascend trunk through a Frame Relay *public data network* (PDN)	Known as Open Packet Trunking (OPTimum trunk). You must first configure either a UNI-DTE feeder or a Frame Relay NNI logical port on the same physical port to enable link management between the two connections.
Encapsulation FRAD (Frame Relay Assembler Disassembler)	Frame Relay encapsulation/ de-encapsulation for HDLC/SDLC-based protocols	Encapsulates traffic entering the network and de-encapsulates it upon exiting the network. This configuration enables you to establish a single circuit between any FRAD port and another non-trunk port. The incoming HDLC/SDLC frames must have a start and end flag (hexadecimal '7E') and a 16-bit cyclic redundancy check (CRC 16). The remainder of the frame is transparent to the Ascend switch.
Direct Line Trunk	Trunk connection to another Ascend switch	Performs Frame Relay functions when the trunk connection carries traffic destined for other switches in the network using Ascend's trunk protocol.

Table 2-1. Frame Relay Logical Port Types (Continued)

Logical Port Type	Physical Port Configuration	Description
PPP	Point-to-Point Protocol according to RFC 1490	Enables a PPP DTE device to communicate with another DTE device configured for Frame Relay and encapsulating multi-protocols, according to the RFC 1490 Specification. This configuration enables you to establish a single circuit between the two devices. The switch performs the PPP Link Control Protocol (LCP) and Network Control Protocol (NCP) and translates PPP encapsulation into the RFC 1490 encapsulation.
ML Member	Multilink Frame Relay	Aggregates available bandwidth on a set of Frame Relay logical links between two networking devices. The aggregated links, collectively referred to as the Multilink Frame Relay Unit (MFRU), can be thought of as a single logical link. The MFRU provides a single logical link between the router and the Frame Relay switch.

Using Fault-Tolerant PVCs

You can configure Frame Relay UNI DCE, UNI DTE, and Network-to-Network Interface (NNI) logical ports for backup service by implementing a fault tolerant PVC configuration. A fault tolerant PVC configuration enables a logical port to serve as a backup for any number of active NNI and/or UNI ports. If the primary port fails, you can activate the backup port through NavisCore. See Chapter 9, "Configuring Fault-Tolerant PVCs," for more information.

Using Frame Relay OPTimum Trunks

A Frame Relay OPTimum trunk creates a switch-to-switch Ascend trunk through a public data network (PDN) into another Ascend Frame Relay network. This configuration maintains the Ascend header. The Ascend OPTimum trunk feature allows private enterprise networks to purchase lower-cost, public-carrier services as the trunk between two Ascend switches instead of using a more expensive leased line. See "Defining Frame Relay OPTimum PVC Trunk Logical Ports" on page 3-34 for configuration information.

About Congestion Control

Congestion control enables you to configure threshold values for each logical port. The congestion control parameters determine how the switch responds to frames and enable you to configure discard thresholds for red and amber frames.

As data travels through the network and is queued for transmit, the switch checks each transmit queue's state for congestion and monitors the behavior of each PVC. The switch marks each PVC as either "good-behaved" or "bad-behaved," based on the configured congestion commitment.

Congestion States and the Switch

When congestion occurs on a link, the switch sets the forward explicit congestion notification (FECNF) bit on packets traveling in the direction of the congestion and the backward explicit congestion notification (BECN) bit on packets traveling in the opposite direction of the congestion.

Closed-Loop Congestion Control and Congestion States

Closed-loop congestion control reduces the rate of excess data into the network during congested periods. The reduction in excess data is in relation to ill-behaved connections and is proportional to the PVC's configured Excess Burst Size (*Be*) value. Using OSPF, the trunk's congestion state is communicated to all switches in the network. The ingress switch uses this information to reduce the flow of excess data into the network. You can enable or disable the closed-loop congestion control feature for each logical port. The default is "Off" (closed-loop congestion control disabled).

There are three congestion states:

- Mild
- Severe
- Absolute

Table 2-2 shows how the congested and ingress switch reacts to congestion at each threshold state.

Table 2-2. Congested and Ingress Switch Behavior

Congestion State	Ingress Switch Be Reduction	Congested Switch Discards...	FECN/BECN Marking
Light-mild	Percent reduction for mild (Pm%) of Excess Burst Size (Be) of “bad” PVCs	“bad” red frames	“bad” PVCs
Heavy-mild	Pm% of Be of all PVCs	all red frames	“bad” PVCs
Light-severe	Percent reduction for severe (Ps%) of Be of “bad” PVCs Pm% of Be of other PVCs	all red and “bad” amber frames	all PVCs
Heavy-severe	Ps% of Be of all PVCs	all red and amber frames	all PVCs
Light-absolute	100% of Be of “bad” PVCs Ps% of Be of other PVCs	all red, amber, and “bad” frames	all PVCs
Heavy-absolute	100% of Be of all PVCs	all red, amber and green frames	all PVCs

Link State Updates

In switches that contain trunks, the OSPF agent in the switch monitors the trunk’s congestion state every *N* seconds. If one or more trunks become congested, OSPF sends a link state update (LSU) to all other switches in the network. When the switches receive the LSU, OSPF updates its routing table with the new congestion state. Similarly, if a trunk moves out of a congested state and remains non-congested for *Nc* seconds, OSPF sends a LSU to all switches in the network.

You can configure *N* (check interval) and *Nc* (clear delay) time intervals (see Table 2-3 on page 2-6). The default is 1 second and 3 seconds, respectively.

Congestion Parameters

Table 2-3 lists the congestion parameters you can configure for each logical port.

Table 2-3. Congestion Parameters

Parameter	Description	Default
Check Interval (N)	Congestion state check interval	1 second
Clear Delay (Nc)	Congestion state clear delay	3 seconds
Fb	“Bad” PVC factor	30
amber Pm (%)	Be reduction percentage level mild	50%
Amber Ps (%)	Be reduction percentage level severe	75%

Threshold Parameters

You can configure the mild, severe, and absolute congestion threshold parameters for each logical port. You change the existing (default) values for the threshold parameters by modifying the logical port congestion control attributes. When you configure the congestion thresholds, you set the threshold values incrementally as shown in Figure 2-1:

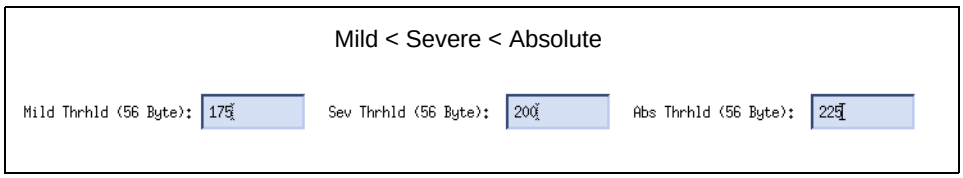

Figure 2-1. Congestion Threshold Example

In the example shown in Figure 2-1, if the congestion level is set at 175 for a mild threshold, the severe threshold must be greater than the mild threshold (200), and the absolute threshold must be greater than the severe threshold (225).

Logical Port Congestion Thresholds

Logical port maximum and default congestion threshold values vary depending on the type of service class configured for the logical port. You can configure congestion thresholds for both mono-class and Priority Frame QoS multi-class (VFR-NRT) services. See "Priority Frame Attributes" on page 3-20 for information about configuring mono- and multi-class services. See Table 2-11 on page 2-14 for descriptions of QoS classes of service.

Table 2-4 shows the maximum mono-class service thresholds for each type of card.

Table 2-4. Maximum Mono-Class Service Thresholds per Card Type

Card Type	56-Byte Buffers	Bytes
8-Port UIO	5450	305200
10-Port DSX	4668	261408
4-Port Channelized T1/T1 PRI	225[1]	12600
4-Port Channelized E1/E1 PRI	174[1]	9744
4-Port Unchannelized T1	5408	302848
4-Port Unchannelized E1	5408	302848
2-Port HSSI	23632	1323392
1-Port ATM UNI	60799	3404744
1-Port Channelized DS3	1922	107632
1-port Channelized DS3-1-0[2] (1-3 DS0s per logical port)	600	33600
1-port Channelized DS3-1-0[2] (4-24 DS0s per logical port)	1922	107632
6-port DS3 for CBX 500	9325 x 2	1036400

[1] For channelized T1/T1 Primary Rate Interface (PRI) and channelized E1/E1 PRI cards, if **n** DS0s are assigned per logical port, the maximum value allowed on the number of buffers is n x 225 (T1 card) and n x 174 (E1 card).

[2] The 1-port Channelized DS3-1-0 is supported *only* on B-STDX switches that are running Release 4.4 switch software. The DS3-1-0 supports mono-class services, only.

Do not exceed the maximum threshold value for each card type. The absolute congestion threshold cannot be greater than the maximum value allowed for each logical port.

Table 2-5 shows the maximum multi-class (VFR-NRT) service threshold values you can configure for each type of card.

Table 2-5. Maximum Multi-Class Service Thresholds per Card Type

Card Type	56-Byte Buffers	Bytes
8-Port UIO	2800 (if port speed is < or = 2048 Kbps)	156800
	5600 (if port speed is > 2048 and < or = 4096 Kbps)	313600
	11200 (if port speed is >4096 and < or = 8192 Kbps)	627200
10-Port DSX	2080	116480
4-Port Channelized T1/T1 PRI	225[1]	12600
4-Port Channelized E1/E1 PRI	180[2]	10080
4-Port Unchannelized T1	1600	89600
4-Port Unchannelized E1	1600	89600
2-Port HSSI	22400	1254400
1-Port ATM UNI	54504	3052224
1-Port Channelized DS3	2069	115864
6-port DS3 for CBX 500	9325	522200

[1]For Channelized T1 PRI cards, if the number of DS0s assigned per logical port is greater than or equal to 4 and less than or equal to 10, the maximum value allowed on the number of buffers is (225 * **n** DS0s –406)/2. If the number of DS0s assigned is greater than 10, the maximum value allowed is (225 * **n** DS0s –534)/2.

2 For Channelized E1 PRI cards, if the number of DS0s assigned per logical port is greater than or equal to 4 and less than or equal to 15, the maximum value allowed on the number of buffers is (180 * **n** DS0s –342)/2. If the number of DS0s assigned is greater than 15, the maximum value allowed is (180 * **n** DS0s –534)/2.

Table 2-6 lists the default values for a 4-port channelized T1/T1 PRI card configured for mono-class service. The threshold default values vary depending on the number of DS0s you assign to each logical port. For example, if you assign each DS0 to one channel on a 4-port channelized E1 card, you can assign a maximum of 225 (56-byte) buffers to each logical port.

Table 2-6. 4-Port Channelized T1/T1 PRI Default Mono-Class Thresholds

Congestion Level	1 DS0/Channel		2 DS0s/Channels		>2 DS0s/Channels	
	56-Byte Buffers	Bytes	56-Byte Buffers	Bytes	56-Byte Buffers	Bytes
Mild	175	9800	225	12600	225	12600
Severe	200	11200	294	16464	294	16464
Absolute	225	12600	450	25200	588	32928

Table 2-7 lists the default values for a 4-port channelized T1/T1 PRI card configured for multi-class (VFR-NRT) service.

Table 2-7. 4-Port Channelized T1/T1 PRI Default Multi-Class Thresholds

Congestion Level	4 DS0/Channels		5 DS0/Channels		6 DS0/Channels	
	56-Byte Buffers	Bytes	56-Byte Buffers	Bytes	56-Byte Buffers	Bytes
Mild	175	9800	225	12600	225	12600
Severe	200	11200	294	16464	294	16464
Absolute	247	13832	359	20104	472	26432
	7 DS0/Channels		> or = 8 DS0/Channels			
Mild	225	1260	225	12600		
Severe	294	16464	294	16464		
Absolute	584	13832	588	32928		

Table 2-8 lists the default values for a 4-port channelized E1/E1 PRI card configured for mono-class service. The threshold default values vary depending on the number of DS0s you assign to each logical port. For example, if you assign each DS0 to one channel on a 4-port channelized E1 card, you can assign a maximum of 174 (56-byte) buffers to each logical port.

Table 2-8. 4-Port Channelized E1/E1 PRI Default Mono-Class Thresholds

Congestion Level	1 DS0/Channel		2 DS0s/Channels		3 DS0s/Channels		>3 DS0s/ Channels	
	56-Byte Buffers	Bytes	56-Byte Buffers	Bytes	56-Byte Buffers	Bytes	56-Byte Buffers	Bytes
Mild	150	8400	225	12600	225	12600	225	12600
Severe	165	9240	294	16464	294	16464	294	16464
Absolute	174	9744	340	19040	520	29120	588	32928

Table Table 2-9 lists the default values for a 4-port channelized E1/E1 PRI card configured for multi-class (VFR-NRT) service.

Table 2-9. 4-Port Channelized E1/E1 PRI Default Multi-Class Thresholds

Congestion Level	4 DS0/Channels		5 DS0/Channels		6 DS0/Channels	
	56-Byte Buffers	Bytes	56-Byte Buffers	Bytes	56-Byte Buffers	Bytes
Mild	150	8400	165	9240	225	12600
Severe	165	9240	225	10584	294	16464
Absolute	189	10584	279	15624	369	29664
	7 DS0/Channels		8 DS0/Channels		> or = 9 DS0/Channels	
Mild	225	12600	225	12600	225	12600
Severe	294	16464	294	16464	294	16464
Absolute	459	25704	549	30744	588	32928

Table 2-10 lists the default values for the ATM, HSSI, UIO, 10-port DSX, channelized DS3 and DS3-1-0, unchannelized T1/E1, and CBX DS3 cards.

Table 2-10. ATM/HSSI, UIO/DSX/T1/E1, DS3/DS310, CBX DS3 Defaults

Card Type	Congestion Level	56-Byte Buffers	Bytes
ATM and HSSI (Mono- and Multi-Class Modes)	Mild	4268	239008
	Severe	8535	477960
	Absolute	17070	955920
UIO, 10-port DSX, unchannelized T1/E1 (Mono- and Multi-Class Modes)	Mild	225	12600
	Severe	294	16464
	Absolute	588	32928
Channelized DS3 (Mono- and Multi-Class Modes)	Mild	480	26880
	Severe	961	53816
	Absolute	1922	107632
Channelized DS3-1-0 (1-3 DS0s per logical port) (Mono-Class Mode)	Mild	300	16800
	Severe	450	25200
	Absolute	600	33600
Channelized DS3-1-0 (4-24 DS0s per logical port) (Mono-Class Mode)	Mild	480	26880
	Severe	961	53816
	Absolute	1922	107632
CBX DS3 Frame (Mono-Class Mode)	Mild	2000 x 2	224000
	Severe	4000 x 2	448000
	Absolute	8000 x 2	896000
CBX DS3 Frame (Multi-Class Mode)	Mild	2000	112000
	Severe	4000	224000
	Absolute	8000	448000

For CBX DS3 cards, all threshold values are by the number of 56-byte buffers.

CLLM Congestion Notification

Consolidated Link Layer Management (CLLM) congestion notification occurs with increases in network traffic load. Network congestion occurs when the traffic attempting to pass is greater than the available bandwidth. When a Frame Relay network reaches its congestion point, frames are discarded until congestion is alleviated.

The following types of congestion control are used to manage Frame Relay data transport:

Implicit Congestion — Involves certain events available in the data link layer to detect the frame loss.

Explicit Congestion — Involves the following types of notification:

Forward Explicit Congestion Notification (FECN) / Backward Explicit Congestion Notification (BECN) – Flow control is built into the Frame Relay address in the form of FECN and BECN bits.

Consolidated Link Layer Management (CLLM) – One DLCI address (1007) is reserved exclusively for transmitting congestion notification.

About CLLM

You can enable or disable CLLM on any Frame Relay UNI or NNI port. The CLLM mechanism applies to PVCs only. The switch reserves DLCI address 1007 exclusively for transmitting congestion notification messages to the user device. The CLLM message:

- Is sent periodically to the customer premise equipment (CPE) or network access device until congestion is alleviated.
- Notifies users of congestion activity outside the conventional framing structure.
- Contains a list of DLCIs that correspond to the congested Frame Relay bearer connections.
- Supports up to a maximum of 127 DLCIs. You can configure the time duration between each consecutive message.

If you are already using DLCI 1007, you must delete the PVC and assign a new DLCI number.

CLLM Threshold States

The configurable parameters, CLLMThresholdNone and CLLMThresholdMild, determine the virtual circuit (VC) congestion threshold type and congestion state. These parameters represent a percentage of BECN frames received since the last CLLM message. The following guidelines determine the VC congestion threshold:

Not congested — The percentage of BECN frames received on any VC on the logical port does not exceed the configured CLLM ThresholdNone.

Mild congested state — The percentage of BECN frames received on any VC on the logical port exceeds the configured CLLM ThresholdNone but does not exceed the configured CLLMThresholdMild.

Absolute congested state — The percentage of BECN frames received on any VC on the logical port exceeds the configured CLLM ThresholdMild.

See "Congestion Control Attributes (VFR-NRT only)" on page 3-12 for information about setting CLLM attributes.

CLLM Messages

Based on the congestion threshold state, there are two types of CLLM messages:

Absolute CLLM — Contains a list of all VCs that are in absolute congested state.

Mild CLLM — Contains a list of all VCs that are in mild congested state.

For example, in a network having two VCs on a Frame Relay UNI port with one VC in absolute congested state and the other VC in mild congested state, the absolute congested state is reported in absolute CLLM frame and the other VC is reported in mild CLLM frame.

Priority Frame QoS

The Priority Frame Quality of Service (QoS) feature enables you to configure "ATM-like service classes" on Frame Relay logical ports.

You must set the logical port service class type to "multi-class" to enable Priority Frame QoS. By default, all service classes are not selected. You can use the Set QoS Parameters option to set QoS parameters. See "Setting QoS Parameters" on page 3-22 for information about setting the multi-class service type and QoS parameters.

When you set the QoS service class, you can use the values listed in Table 2-13, or modify these settings.

For channelized T1/E1 cards, you can configure QoS Class of Service on ports with 4 or more DS0s only.

Table 2-11 briefly describes each class of service.

Table 2-11. QoS Class of Service Descriptions

Field	Description
Variable Frame Rate (VFR) Real Time	Used for special delay-sensitive applications, such as packet voice, which require low delay between endpoints.
Variable Frame Rate Non-Real Time (VFR-NRT)	Handles transfer of data streams with a committed information rate over a pre-established connection. This service provides low data loss but no delay guarantee. This service class also offers *configurable* congestion control support.
Unspecified Frame Rate (UFR)	Primarily used for LAN traffic. The CPE should compensate for any delay or lost traffic.

Using a T1/E1 Card

If you are configuring QoS Class of Service on a T1/E1 card, use the guidelines described in Table 2-12.

Table 2-12. T1/E1 I/O Module QoS Class of Service Guidelines

Number of DS0s	Number of allowed QoS Class of Service Combinations
1-3	1 one-class with VFR-NRT characteristics
4 or more	All valid traffic class combinations

Table 2-13 describes the QoS values for Frame Relay logical ports.

Table 2-13. Default QoS Values for Frame Relay Logical Ports

Service Type	Bandwidth Allocation	Routing Metric	Oversubscription Factor
VFR-RT	Dynamic	Admin Cost	100%
VFR-NRT	Dynamic	Admin Cost	100%
UFR	Dynamic	Admin Cost	100%

Administrative Tasks

This section describes how to:

- Use templates to define a new logical port
- Delete circuits
- Delete trunks
- Delete management or multicast DLCIs
- Delete Frame Relay logical ports

Using Templates

If you defined a logical port configuration and saved it as a template (see *Is Template* field on page 3-11), you can define a new logical port using the same parameters.

To define a logical port from a template:

1. Choose the Add Using Template command on the Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-3).
2. Do one of the following:
   - Choose Last Template to use the last template you defined for this switch.
   - Choose Template List to display a list of templates defined for this map. Select a template and choose OK.

You can define logical port templates for creating bulk (multiple) logical ports on the channelized DS3-1-0 module. For information about the Bulk LPort feature, see the *NavisCore Physical Interface Configuration Guide*.

Deleting Frame Relay Logical Ports

If “None” is not displayed in the loopback status field on the Set All Logical Ports dialog box, do not attempt to delete this logical port. See the *NavisCore Diagnostic and Troubleshooting Guide* for information about loopback testing.

Before you can delete a Frame Relay logical port, verify that the following conditions are met:

- No circuit uses this logical port as an endpoint.
- No trunk is defined that uses this logical port as an endpoint.
- No management DLCI or multicast DLCI exists on this port.
- This logical port is not defined as the feeder (FR UNI DTE/NNI) for an existing OPTimum PVC trunk logical port.
- If the MLFR trunk bundle logical port is an endpoint of a trunk, you can delete all but one logical port binding.
- No network ID exists on this logical port.

If any of the conditions mentioned above exist and use the logical port you want to delete, you must first delete them in the following order:

- Circuits
- Trunks
- Management or multicast DLCIs
- Logical port

Deleting Circuits

To delete a circuit:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Circuits. The Set All Circuits On Map dialog box appears.
2. To view the list of circuits, select the Search by Name field and press Return. If necessary, select each circuit and review each logical port endpoint.
3. Select the circuit to delete.
4. Choose Delete.
5. Choose Close to return to the network map.

Deleting Trunks

To delete a trunk:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Trunks. The Set All Trunks dialog box appears. If necessary, select each trunk and review each logical port endpoint.
2. Select the trunk to delete.
3. Choose Delete.
4. Choose Close to return to the network map.

Deleting Management or Multicast DLCIs

To delete management or multicast DLCIs:

1. From the Administer menu, select one of the following:
   - Ascend Parameters ⇒ Set All Multicast DLCIs. The Set All Multicast DLCIs dialog box appears, listing the defined Multicast groups in the network configuration.
   - Ascend Parameters ⇒ Set All Management DLCIs. The Set All Management DLCIs dialog box appears, listing the Management DLCIs already configured.
2. Select the DLCI to delete.
3. Choose Delete.
4. Choose Close to return to the network map.

Deleting the Logical Port

To delete the logical port:

1. Select the switch on which to delete a logical port.
2. Log in using either a provisioning or operator password.
3. From the Administer menu, select Ascend Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
4. Select the physical port. The Set Physical Port Attributes dialog box appears.
5. Choose Logical Port. The Set All Logical Ports in PPort dialog box appears.
6. Select the logical port to delete from the logical port list.

Make sure this logical port is not the UNI DTE or NNI logical port used as the feeder for a Frame Relay OPTimum trunk. You first need to take the OPTimum trunk out of service, or first define another feeder, before you can delete this logical port.

If the MLFR trunk bundle logical port is an endpoint of a trunk, you can delete all but one logical port binding. The system displays a warning message if deleting a logical port binding will cause the trunk endpoints to have a different number of logical ports.

7. Choose Delete. Make sure the Loopback field displays “NONE”.
8. Choose Delete.
9. Choose Close.

3

Configuring Frame Relay Logical Ports

This chapter provides instructions for configuring Frame Relay logical ports on a B-STDX or CBX switch. See Chapter 2, “Frame Relay Services,” for an overview of Ascend’s Frame Relay logical port services.

Accessing Frame Relay Logical Port Functions

To access logical port functions in NavisCore:

1. Log in to NavisCore using either a provisioning or operator password.
2. Select the switch to which you want to add a logical port.
3. From the Administer menu, select Ascend Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
4. Select the physical port you want to configure and press the right mouse button to display a popup menu. Select Logical Port. The Set All Logical Ports in PPort dialog box appears as shown in Figure 3-1 on page 3-3.

   For Channelized DS3 and DS3-1-0 I/O Modules only:

   a. From the Switch Back Panel dialog box, select the physical port you want to configure and press the right mouse button to display a popup menu. Select Physical Port. The Set Physical Port Attributes dialog box appears.

   b. Double-click the channel (button) that you want to configure. The Set Channel Attributes dialog box appears.

   c. Choose Logical Port.

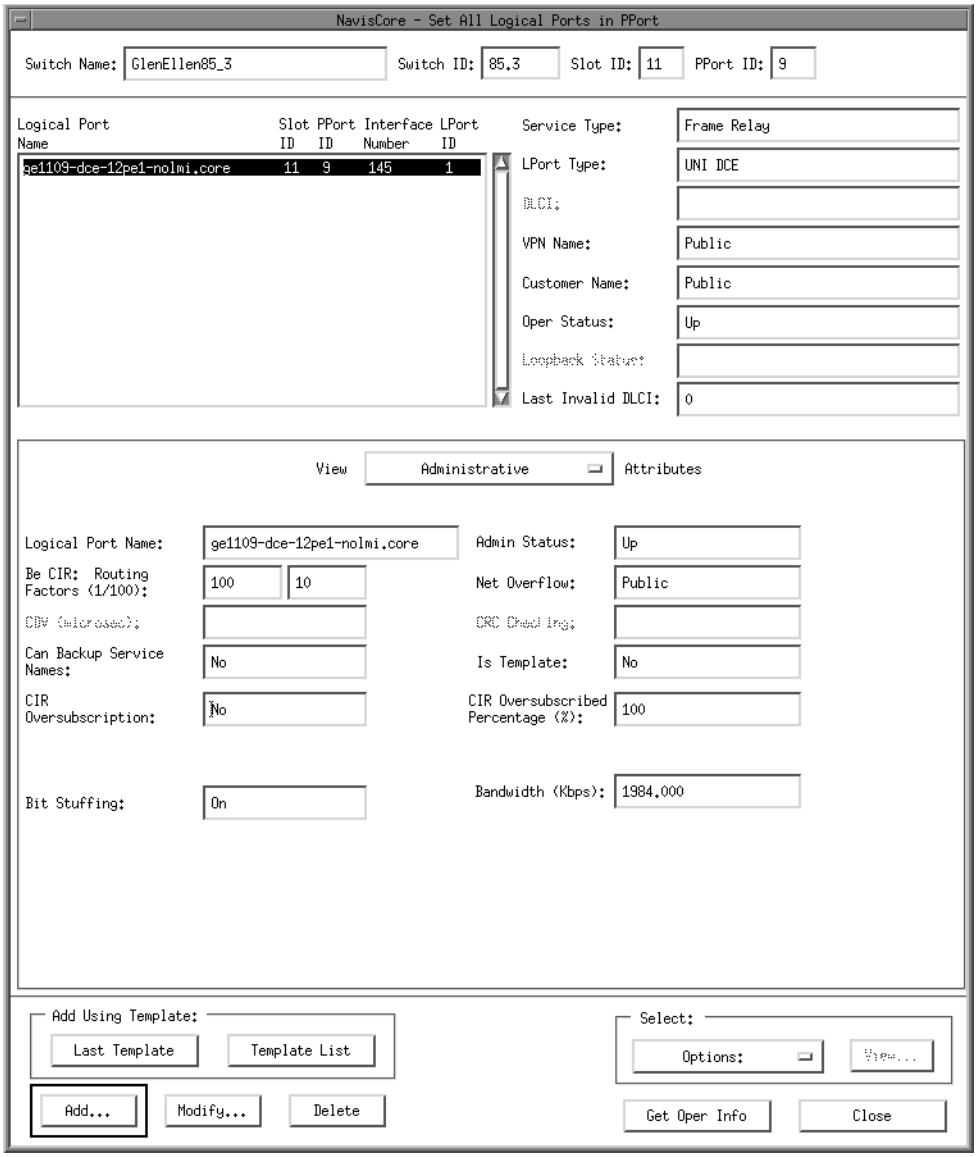

Figure 3-1. Set All Logical Ports in PPort Dialog Box

The following section describes the Set all Logical Ports in PPort dialog box fields and commands.

To begin adding a logical port, proceed to page 3-6.

About the Set All Logical Ports in PPort Dialog Box

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

Table 3-1 describes the Set All Logical Ports in PPort fields and commands.

Table 3-1. Set All Logical Ports in PPort Fields and Commands

Fields and Commands	Function
Service Type	Displays Frame Relay.
LPort Type	Displays the logical port type: either UNI DCE, UNI DTE, NNI, OPTimum Trunk, or Others.
DLCI	Displays the data link connection identifier (DLCI) assigned to this logical port. For more information, see page 3-34.
VPN Name	Displays the VPN name to which this logical port belongs.
Customer Name	Displays the name of the customer to which this logical port is dedicated.
Oper Status	Indicates whether this port is operationally Up, Down, or Unknown. Unknown indicates that the NMS is unable to contact the switch to retrieve status.
Loopback Status	Indicates whether loopback testing is enabled on this logical port. The default is None (no testing).
Last Invalid DLCI	Displays the last invalid DLCI that the switch detected. If this field displays a value, either the switch or the Customer Premise Equipment (CPE) was not configured properly. Check this value if you have a DLCI that is not receiving traffic.
View Attributes (option menu)	Displays the appropriate attributes configured for the selected option. See one of the following sections for more information: • "Defining Frame Relay UNI DCE/DTE or NNI Logical Ports" on page 3-8 • "Defining Frame Relay OPTimum PVC Trunk Logical Ports" on page 3-34 • "Defining Encapsulation FRAD, Direct Line Trunk, and PPP Logical Ports" on page 3-36

Table 3-1. Set All Logical Ports in PPort Fields and Commands (Continued)

Fields and Commands	Function
Add Modify Delete	Enables you to add a new logical port or *Modify* or *Delete* an existing logical port configuration. For information about deleting logical ports see "Deleting Frame Relay Logical Ports" on page 2-16.
Get Oper Info	Displays a status message in the *Oper Status* field for the selected logical port.
Last Template/ Template List	If you have already defined a logical port configuration and saved it as a template, you can use this option to define a new logical port using the same parameters. See the "Using Templates" on page 2-15 for more information.
Use the Select: Options button to view logical port options. Once you select an option from this list, choose View to access the information.  You can select any of the following options:	
IP Parameters	Access the Set IP Parameters dialog box for configuring IP logical port parameters. See the *NavisCore IP Navigator Configuration Guide* for more information.
Statistics	Displays the summary statistics for the selected logical port.
Diagnostics	Accesses diagnostic tests for the selected logical port.
VPN/Customer Info	Assigns a VPN and customer name to the selected logical port. See Chapter 8, "Configuring Virtual Private Networks," for more information.
QoS Parameters	Displays the quality of service parameters (including bandwidth and routing metrics) for the selected logical port. See page 3-22 for more information.
Accounting (*Optional*)	Accesses the NavisXtend Accounting server functions for a logical port.
Screen Assignments	Displays the SVC port security screen assignments for the selected logical port. See Chapter 12, "Port Security Screening," for more information.

Adding a Frame Relay Logical Port

To add a Frame Relay logical port:

1. On the Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-3), choose Add. The Add Logical Port Type dialog box (Figure 3-2) appears.

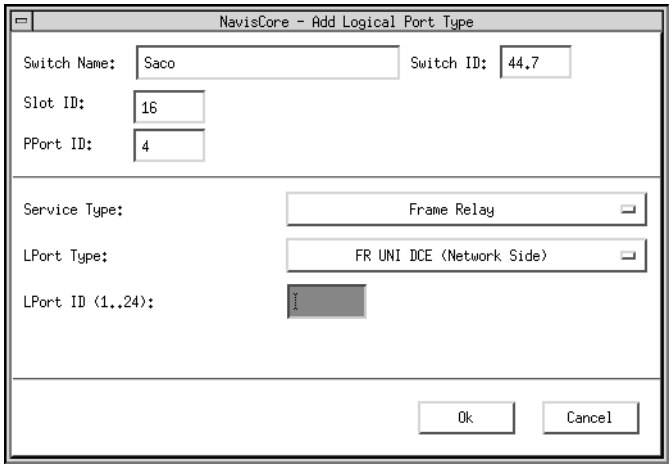

Figure 3-2. Add Logical Port Type Dialog Box

2. Define the specific logical port configuration. See Table 3-2 for more information on the different types of logical ports.

Table 3-2. Frame Relay Logical Port Configurations

Service Type	Logical Port Type	See...
Frame Relay	FR UNI-DCE	"Defining Frame Relay UNI DCE/DTE or NNI Logical Ports" on page 3-8
	FR UNI-DTE	
	FR UNI-NNI	
	Frame Relay OPTimum Trunk	"Defining Frame Relay OPTimum PVC Trunk Logical Ports" on page 3-34
Others	Direct Line Trunk, Encapsulation FRAD, and Point to Point Protocol	"Defining Encapsulation FRAD, Direct Line Trunk, and PPP Logical Ports" on page 3-36
	ML Member	"Defining ML Member Logical Ports" on page 3-44

3. Choose OK. The Add Logical Port dialog box appears. The sample dialog box in Figure 3-3 shows a Frame Relay UNI DCE logical port.

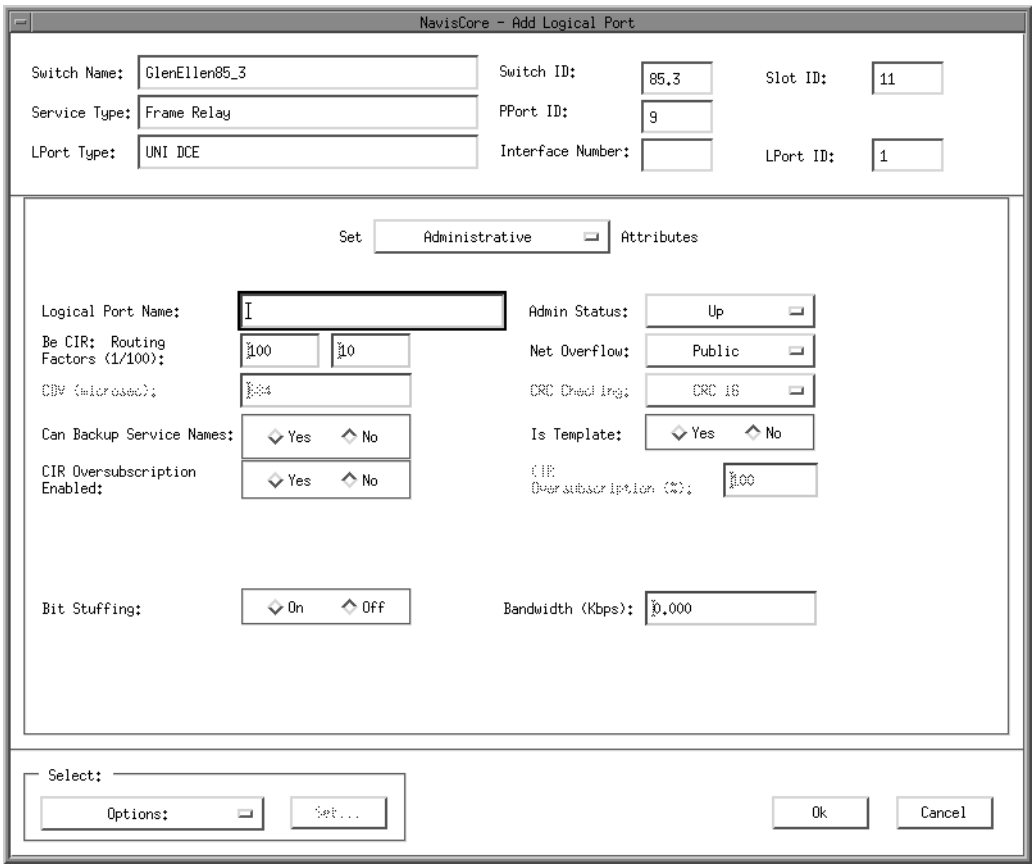

Figure 3-3. Add Logical Port Dialog Box (UNI DCE Logical Port)

The following section describes how to define a Frame Relay UNI DCE/DTE or NNI logical port. Later sections, which explain how to configure different types of logical ports, may refer to some of the procedures described in the following section.

See Table 3-2 on page 3-6 for information about other types of Frame Relay logical ports.

Defining Frame Relay UNI DCE/DTE or NNI Logical Ports

To define a Frame Relay UNI DCE, UNI DTE, or NNI logical port:

1. In the Add Logical Port dialog box (Figure 3-3 on page 3-7), complete the fields described in Table 3-3.

Table 3-3. Add Logical Port (UNI-DCE) Fields

Field	Action/Description
Service Type	Select Frame Relay.
LPort Type	Select either FR UNI DCE, FR UNI DTE, or FR NNI.
LPort ID	For a channelized T1 module, enter a number between 1 and 24 (or 1 and 30 for channelized E1). For all other modules, the LPort ID is a read-only field that automatically defaults to 1.

2. Choose OK. The Add Logical Port dialog box reappears, as shown in Figure 3-3.

When you define a logical port, you can set attributes and other optional parameters from the Set Attributes and Select Options menus. The next section, "Setting Logical Port Attributes," describes how to set attributes for a new logical port. See "Selecting Additional Logical Port Options" on page 3-21 for information about Quality of Service (QoS) and other optional settings.

Setting Logical Port Attributes

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

Administrative — Sets the admin status, net overflow, and bandwidth parameters.

Congestion Control — Sets the threshold parameters (mild, severe, and absolute) that determine how the switch responds to congestion in the network.

Link Management — Sets the link management protocol used in the network and the LMI update delay and error thresholds.

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

Priority Frame — Sets the logical port service class and transmit schedule mode.

Frame Relay SVC— Sets the Q.922 signaling parameters, CAC, and QoS classes.

SVC Parameters — Sets the calling party parameters and CUG status.

SVC Priorities — Assigns bandwidth and bumping priorities to SVCs based on ingress QoS class; assigns forward and reverse circuit discard priorities to SVCs that originate on a specific logical port.

This section explains how to set all attributes except SVC attributes. For information about setting SVC attributes, see "Configuring Logical Ports for Use With SVCs" on page 3-25.

Administrative Attributes

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [Administrative] Attributes and complete the fields described in Table 3-4.

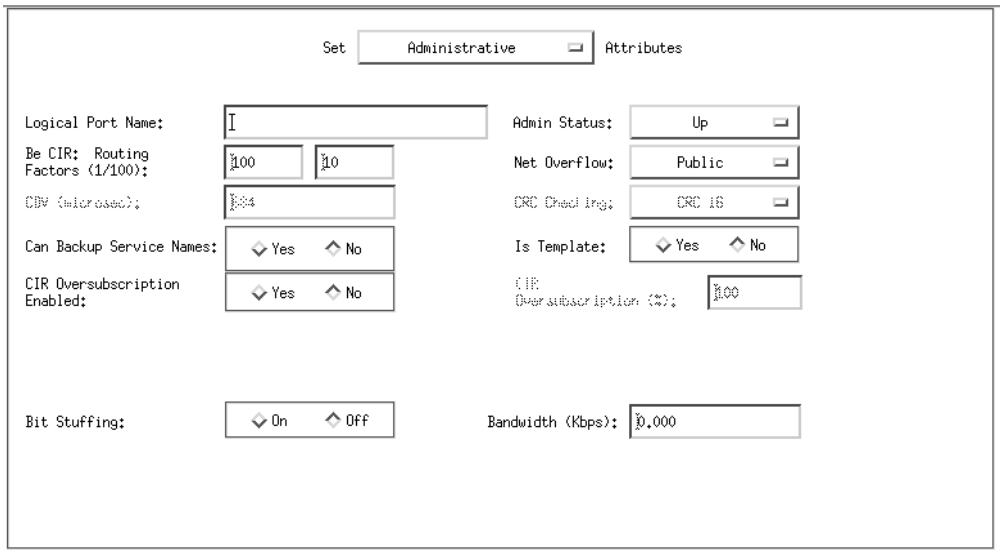

Figure 3-4. Set Administrative Attributes

Table 3-4. Set Administrative Attributes Fields

Field	Action/Description
Logical Port Name	Enter an alphanumeric logical port name (up to 32 characters in length) to assign this port.
Be CIR Routing	Enter a value between 0-100 percent. This value represents the UNI bandwidth percentage on all configured zero CIR circuits. The default is 100 percent.
Factors (1/100s)	Enter a value between 0-100 percent. This value represents the routing factor percentage on all rate enforcement circuits. The default is 10 percent.
Can Backup Service Names	*(Fault-tolerant PVC only)* Select Yes to configure a logical port for backup service. For more information, see Chapter 9, "Configuring Fault-Tolerant PVCs."
CIR Oversubscription Enabled	Select Yes to enable bandwidth subscription. The default is No.
Admin Status	Set the Admin Status as follows: *Up (default)* – Activates the port. *Down* – Saves the configuration in the database without activating the port or takes the port off-line to run diagnostics. ***Note:*** *When only one logical port exists on a physical port, and you set the admin status for the logical port down, the physical port is also considered "down." If more than one logical port exists on a physical port, and you set the admin status for each of these logical ports to down, the physical port is also considered down.*
Net Overflow	Determines how SVC traffic originating from this logical port is managed during trunk overflow or failure conditions. This feature is used with Virtual Private Networks. To assign this logical port to a specific VPN and customer, see Chapter 8. Select one of the following options: *Public (default)* – SVCs originating from this port are routed over dedicated VPN trunks. However, in the event of failure, the customer's traffic is allowed to run over common trunks (shared by a variety of different customers). *Restricted* – SVCs originating from this port can only use dedicated VPN trunks. A customer using this mode must purchase redundancy trunks to be used in the event of outages or other trunk failures.

Table 3-4. Set Administrative Attributes Fields (Continued)

Field	Action/Description
CRC Checking (HSSI and CBX DS3 modules, only)	Set this value to match the number of error checking bits used by the CPE connected to this port. Performs a cyclic redundancy check (CRC) on incoming data. Data is checked in either 4K (CRC 16) or 8K (CRC 32) frames.
Is Template	(*Optional*) Saves these settings as a template to configure another logical port with similar options. To create a template, choose Yes in the *Is Template* field. See "Using Templates" on page 2-15 for more information.
CIR Over-subscription (%)	If you enabled CIR oversubscription, enter the percentage of oversubscription. The default is 100%.
Channels allocated for a Logical Port are marked by their IDs (Channelized T1 or E1 modules, only)	If you are configuring a channelized T1 or E1 module, specify the DS0 (for T1) or TS0 (for E1) channel(s) assigned to the logical port. The logical port ID number appears in the box (channel) you select. To deselect DS0 channels, click on the channel to remove the X. You can select/deselect channels by using the following Channel Allocation editing buttons: -- To deselect all channels ++To select all channels - To deselect a specific channel + To select a specific channel ***Note**: The logical port bandwidth either increments or decrements depending on the number of channels you select or deselect. You can configure other logical ports with different attributes, to other DS0/TS0 channels on this same physical port.*
Bit Stuffing	Select the bandwidth that matches the bandwidth capability of the customer premise equipment (CPE) connected to this logical port. Enables bit stuffing on T1/E1/DSX-1 ports. Bit stuffing affects the available bandwidth of each DS0/TS0 channel on this port. *On* – Provides 56 Kbps of bandwidth. *Off* – Provides 64 Kbps of bandwidth.

Table 3-4. Set Administrative Attributes Fields (Continued)

Field	Action/Description
Bandwidth (Kbps)	Enter the amount of bandwidth you want to configure for this logical port. The default is the amount of bandwidth remaining from the physical clock rate, less any logical ports already configured. To define a trunk logical port on this same physical port, decrease the amount of bandwidth on this logical port to ensure sufficient remaining bandwidth. For example: *Physical port clock speed*: 1536 Kbps *Logical port UNI-DTE/NNI Feeder Bandwidth*: 56 Kbps *Logical port Frame Relay Trunk Bandwidth*: 1480 Kbps The example configuration allocates a PDN trunk with 1480 Kbps bandwidth between two Ascend switches, each attached to a PDN network.

Congestion Control Attributes (VFR-NRT only)

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [Congestion Control] Attributes and complete the fields described in Table 3-5 on page 3-13.

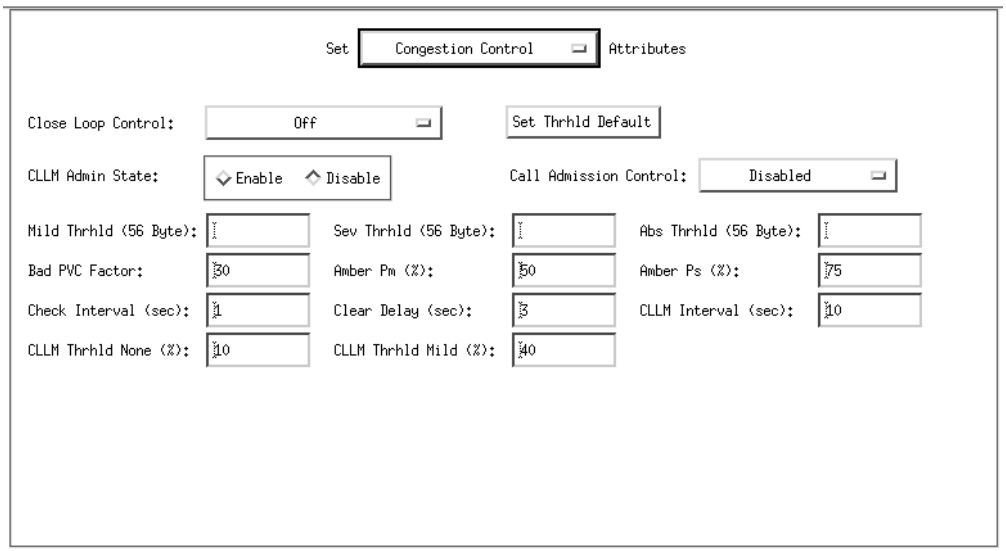

Figure 3-5. Set Congestion Control Attributes

Do not exceed the maximum threshold value for each card type (see Table 2-4 on page 2-7 for more information). The absolute congestion threshold cannot be greater than the maximum value allowed for each logical port.

For channelized T1/T1 PRI and Channelized E1/E1 PRI cards, if **n** DS0s are assigned per logical port, the maximum value allowed on the number of buffers is n x 225 (T1) and n x 174 (E1).

Table 3-5. Set Congestion Control Attributes Fields

Field	Action/Description
Close Loop Control	Set the congestion control parameters. This field enables/disables OSPF closed-loop congestion control for each logical port. For more information see "Closed-Loop Congestion Control and Congestion States" on page 2-4. Options include: *Off (default)* – Disables closed-loop congestion. *OSPF-based* – Enables closed-loop congestion.
CLLM Admin State	Set the admin state to enable or disable CLLM notification on this logical port. *Enable* - Enables CLLM notification. *Disable* (*default*) - Disables CLLM notification.
Set Thrhld Default	Sets the Mild, Severe, and Absolute threshold settings to the default settings described in Table 2-6 on page 2-9 through Table 2-10 on page 2-11.
Call Admission Control	When enabled, the port rejects a circuit creation request if there is not enough available bandwidth on that logical port. When disabled (*default*), the port attempts to create a circuit even if there is not enough available bandwidth on that logical port. ***Note:*** *If you disable Call Admission Control on a UNI logical port, you are effectively disabling Ascend's Call Master Connection Admission Control (CAC) function on that logical port.*

Table 3-5. Set Congestion Control Attributes Fields (Continued)

Field	Action/Description
Mild Thrshld (56 Byte) Severe Thrshld (56 Byte) Absolute Thrshld (56 Byte)	Accept the defaults or enter values for the mild, severe, and absolute threshold fields as defined in Table 2-6 on page 2-9 through Table 2-10 on page 2-11. ***Note**: Do not exceed the maximum threshold value for each card type (see Table 2-4 on page 2-7 for more information). The absolute congestion threshold cannot be greater than the maximum value allowed for each logical port.* ***Note**: If you are setting threshold parameters on a T1/E1 card, the default values will not appear until you set the bit stuffing and bandwidth allocation. See Table 2-4 on page 2-7 for more information.* ***Note**: For channelized T1/T1 PRI and Channelized E1/E1 PRI cards, if n DS0s are assigned per logical port, the maximum value allowed on the number of buffers is n x 225 (T1) and n x 174 (E1).*
Bad PVC Factor	Enter a value between 0-32. Determines the threshold for “bad” PVC detection. The following example shows the relationship between the “bad” PVC factor and threshold. $\text{Threshold} = \frac{Bc+(Be/2)}{2^{(32-F_b)}}$ The default is 30. ***Note:** If you select simple as the rate enforcement scheme this feature is disabled.*
Amber Pm (%)	Controls the reduction percentage of Be when mild congestion occurs. Enter a Pm% value. The default is 50%.
Amber Ps (%)	Enter a Ps% value. This value controls the reduction percentage of Be when severe congestion occurs. The default is 75%.
Check Interval (sec)	Enter and interval. This determines the number of seconds in which the switch monitors the trunk’s congestion on the port. The default is 1 second.
Clear Delay (sec)	Enter a value. This determines the number of seconds in which the switch monitors the trunk’s non-congestion state. The default is 3 seconds.

Table 3-5. Set Congestion Control Attributes Fields (Continued)

Field	Action/Description
CLLM Interval (sec)	The time duration (in seconds) between two consecutive CLLM messages sent on the logical port. The CLLM message is sent as long as at least one VC on this logical port remains in a congested state. The default value is 10 seconds. Enter a value between 5 and 30 seconds.
CLLM Thrhld None (%)	Displays the threshold percentage value (between 1-100) of BECN frames received on any VC on this port. The default value is 10. For more information, see "CLLM Threshold States" on page 2-13.
CLLM Thrhld Mild (%)	Displays the threshold percentage value (between 1-100) of BECN frames received on any VC on this port. The default value is 40. The value for the Mild threshold must be equal to or greater than the value for the None threshold. For more information, see "CLLM Threshold States" on page 2-13.

Link Management Attributes

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [LinkMgmt] Attributes and complete the fields described in Table 3-6.

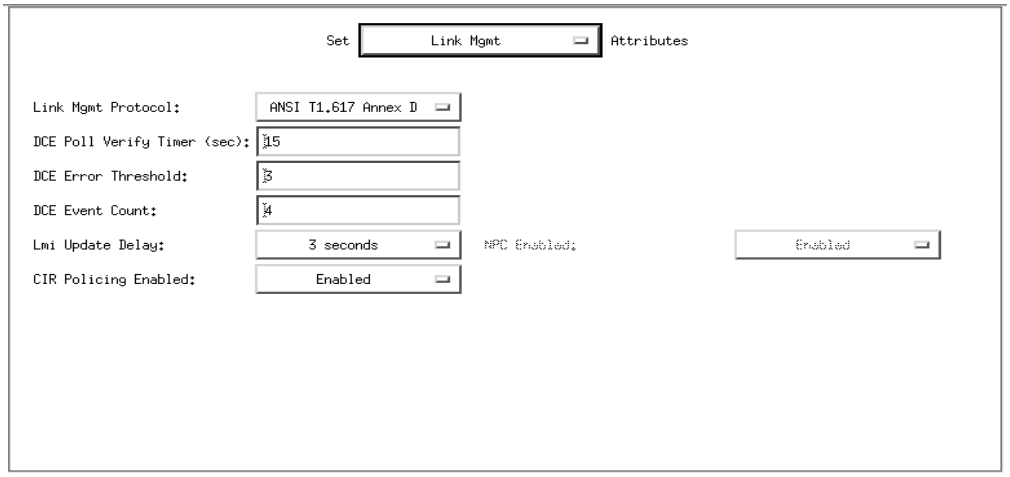

Figure 3-6. Set Link Mgmt Attributes

Table 3-6. Set Link Mgmt Attributes Fields

Field	Action/Description
Link Mgmt Protocol	Select the link management protocol that represents the type of Frame Relay implementation used in your network. Options include: *ANSI T1.617 Annex D (default)* – The network uses DLCI 0 for link management. *LMI Rev1* – The network uses DLCI 1023 for link management. *CCITT Q.933 Annex A* – For international standard (European) use only. The network uses DLCI 0 for link management. *Auto Detect* – Use this option only if the attached CPE provides the link management protocol. This logical port can then automatically detect which protocol is in use. *Disabled* – Use this option only if the attached CPE does not support link management or if you need to disable link management for troubleshooting purposes.
DCE Poll Verify Timer (sec) or DTE Poll Verify Timer (sec)	Set the poll verify timer (in seconds). This field specifies the value of the T392 timer, which sets the length of time the network waits between status inquiry messages. If the network does not receive a status inquiry message within the specified number of seconds, the network records an error. The default value is 15 seconds. ***Note**: The attached CPE must be set to a value that is less than the DCE (DTE) Poll Verify Timer.* Increase this value if the DTE (DCE) device has a poll frequency that is greater than or equal to the DCE (DTE) Poll Verify Timer. Decrease this value if the DTE's (DCE's) poll frequency is less than or equal to one-half that of the DCE (DTE) poll verify timer.
DCE Error Threshold or DTE Error Threshold	Specify an error threshold. This parameter is used with the DCE (DTE) Events Count (N393) parameter. The Local Management protocol monitors the specified number of events for the DCE (DTE) Event Count. If the number of events found in error exceeds the specified DCE (DTE) Error Threshold, the link is declared inactive. The default value is 3.

Table 3-6. Set Link Mgmt Attributes Fields (Continued)

Field	Action/Description
DCE Event Count or DTE Event Count	Specify the number of events in a sliding window of events monitored by the network. An event is the receipt of a valid or invalid status inquiry message, or the expiration of the T392 timer. For example, use the default DCE (DTE) Error Threshold value of 3 and the default DCE (DTE) Event Count value of 4. If three (N392) of the last four (N393) events are found in error, the link is declared inactive. The link remains inactive until the network receives four consecutive error-free events. ***Note**: The DCE (DTE) Error Threshold and the DCE (DTE) Event Count work together. The lower you set these values, the more sensitive the logical port is to LMI poll errors. To make the logical port less sensitive to errors, increase these values.*
LMI Update Delay	Set a timer from 1 to 9 seconds to enable asynchronous LMI updates. The default is 3 seconds. When you set this timer, the switch sends a signal (known as an *event*) to notify other network equipment (CPE) when a circuit on this logical port goes up or down. The specified time interval creates a buffer. If the circuit recovers within this period of time, no event is issued. • If you choose *No Updates*, the switch does not send a signal to the CPE. • If you choose *No Delay*, the switch sends an update immediately to the CPE. For example, if the network takes a significant amount of time to recover from trunk outages, increase the LMI update delay. This delay minimizes network downtime visibility to end users.
CIR Policing Enabled (UNI and NNI LPorts)	Enables or disables frame CIR policing. The default is Enabled. *Enabled* – When a circuit exceeds the established committed information rate (CIR), the Discard Eligible (DE) bit in the Frame Relay header is set *on* for incoming frames that exceed the CIR. *Disabled* – The DE bit is not changed for incoming frames. ***Note:** Whenever the network is congested, frames with the DE bit set* on *are discarded first.*
NPC Enabled (NNI, IISP, PNNI LPorts only)	Enables or disables the Network Parameter Control function, which enables you to communicate with other networks. *Enabled* – Frames that do not conform to the traffic parameters are dropped or tagged as they come in to the port. *Disabled* – All traffic, including non-conforming traffic, passes in through the port.

Trap Control Attributes

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [Trap Control] Attributes and complete the fields described in Table 3-7.

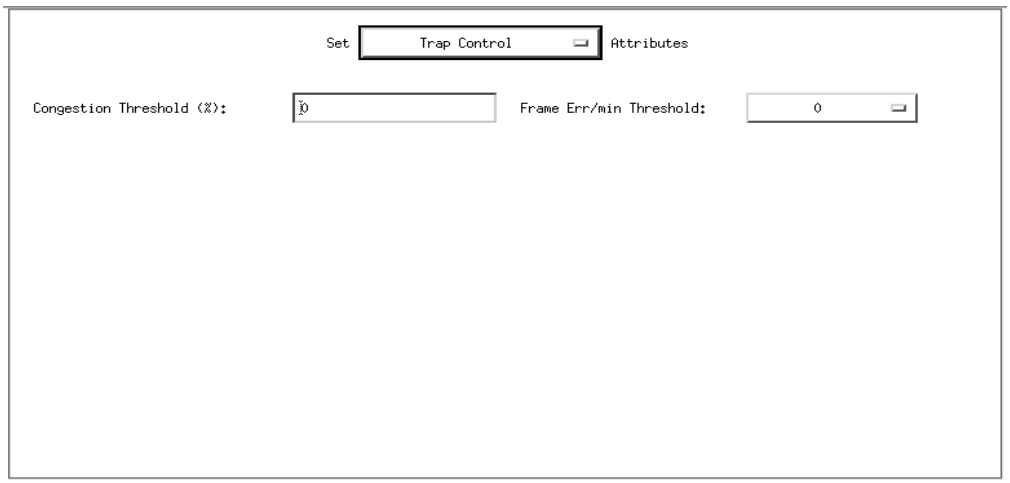

Figure 3-7. Set Trap Control Attributes

Table 3-7. Set Trap Control Attributes Fields

Field	Action/Description
Congestion Threshold (%)	Enter a value between 0 and 100 to indicate the threshold percentage for generating and sending traps to the NMS for this logical port. A congestion trap is generated and sent to the NMS if the rate of congestion over a one-minute period exceeds the percentage value you enter. Adjust the entered value according to how sensitive this port needs to be to network congestion. Options include: *Low* – Generates a trap at the first sign of congestion. *High* – Generates traps for serious network congestion. *Zero (default)* – Disables congestion threshold. If you enter zero, no traps are generated for this logical port.

Table 3-7. Set Trap Control Attributes Fields (Continued)

Field	Action/Description
Frame Err/Min Threshold	Enter a value from 0 to 16384 to configure the threshold of frame errors on this logical port. If the number of frame errors received in one minute exceeds the specified number, a trap is sent to the NMS. Adjust this value according to how sensitive this port needs to be to frame errors. Options include: *Low* – Port is sensitive to frame errors. *High* – Generates traps when a significant number of frame errors occurs within a one-minute period. *Zero (default)* – Disables this feature, which prevents traps from being generated for this logical port.
SMDS PDU Violation Threshold (0-255) (*Frame Relay OPTimum and Direct Trunks only*)	Specify the number of PDU violations that can occur before a trap is sent to the NMS. The software increments a counter every time an SMDS PDU violation takes place on a logical port. The software polls these counters every 60 seconds. If a particular counter exceeds the specified SMDS PDU violation threshold for the logical port, it generates a trap corresponding to that particular violation. The default is 10 PDU violations. Options include: *Low* – Sensitive to SMDS PDU violations. *High* – Issue traps only when there is a significant number of SMDS PDU violations.
SMDS PDU Violation Traps (*Frame Relay OPTimum and Direct Trunks only*)	Enable or disable this field. An SMDS PDU violation can be either an SIP 3 SMDS address failure or an invalid DXI2 frame header. These errors signify that incoming frames are bad, indicating problems with the CPE configuration. Options include: *Disable* (default) – Turns off traps. *Enable* – Issues traps for PDU violations.

Priority Frame Attributes

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [Priority Frame] Attributes and complete the fields described in Table 3-8.

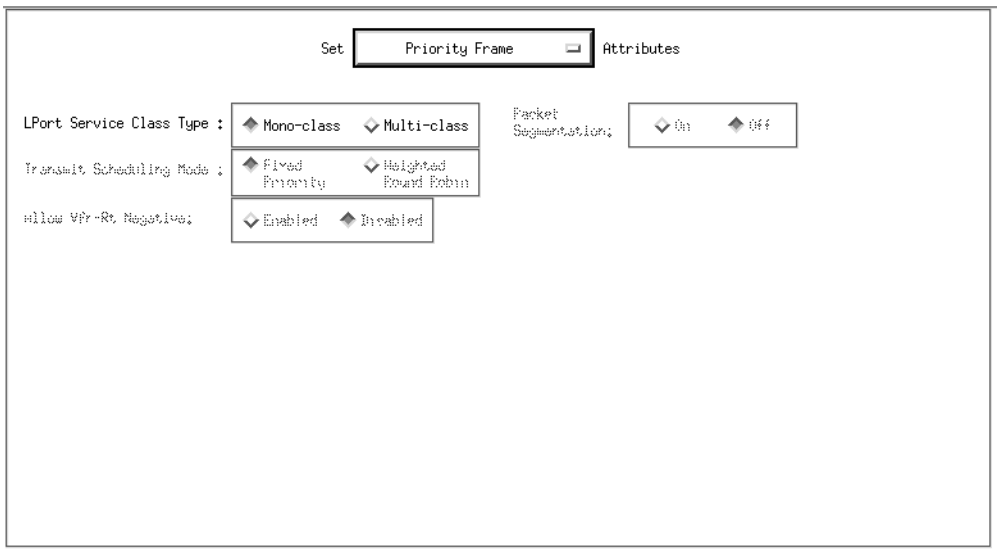

Figure 3-8. Set Priority Frame Attributes

Table 3-8. Set Priority Frame Attributes Fields

Field	Action/Description
LPort Service Class Type	Select the service class type for this LPort. Options include: *Mono-class* – If you select mono class, all circuits are transmitted using VFR-NRT characteristics. *Multi-class* – If you select multi-class, Frame Relay QoS is enabled and all QoS classes are supported. The multi-class setting default transmit scheduling mode is “Fixed Priority.”
Transmit Scheduling Mode	The transmit scheduling mode is available when you select the Multi-Class LPort Service Class. This mode determines the transmission scheduling method to schedule transmission among the three service class types (VFR-RT, VFR-NRT, and UFR). Select the transmit scheduling mode. Options include: *Fixed Priority* – Empties the VFR-RT, VFR-NRT, and UFR queues in a fixed order. Fixed is the default setting when you select the multi-class LPort service type. *Weighted Round Robin* – Empties the VFR-RT and VFR-NRT queues in a weighted order and the UFR queue last.

Table 3-8. Set Priority Frame Attributes Fields (Continued)

Field	Action/Description
Allow VFR-RT Negative (trunk logical port type, only)	The Allow VFR-RT Negative field is available when you select the Multi-Class LPort Service Class. If you choose enable, the trunk can be oversubscribed. This option is useful when a trunk has failed, and PVCs must be rerouted to a new trunk. In this event, trunk bandwidth can become negative and delay commitments are not guaranteed, but PVCs stay up. With this option disabled (*default*), VFR-RT PVCs from the failed trunk may not reroute and remain down; however, existing trunk bandwidth and service remain stable.

When you finish configuring the Priority Frame Attributes for this logical port, continue with the instructions on page 3-25 if you plan to configure SVC addresses for this logical port. Otherwise, use the following section to set additional logical port options.

Selecting Additional Logical Port Options

To select additional options for this new logical port:

1. From the Add Logical Port dialog box (Figure 3-3 on page 3-7), use the Select: Options: menu to review additional options. Choose Set to configure this information.

The Options button displays the commands described in Table 3-9. To invoke a command, select the option and choose Set.

Table 3-9. Add Logical Port Option Menu Commands

Option	Description
QoS Parameters	This option is available when you select the Multi-Class LPort Service Class Type (Figure 3-8 on page 3-20). To review QoS parameters and, if necessary, modify these defaults, see the next section, "Setting QoS Parameters."
Accounting (*Optional)*	Enables you to configure NavisXtend Accounting server parameters.
Screen Assignments (*Optional - Frame Relay SVCs*)	To configure screen assignments for port security screening, see "Configuring Port Security Screening" on page 12-7.

Setting QoS Parameters

The section describes how to set the Quality of Service (QoS) parameters for a logical port. These parameters enable you to specify the bandwidth and routing metrics (if applicable) for the various traffic service classes.

To set the QoS Parameters:

1. From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [Priority Frame] Attributes, and then select Multi-class LPort Service Class Type. Complete the fields described in Table 3-8 on page 3-20.
2. From the Add Logical Port Option menu commands (Table 3-9 on page 3-21), select QoS Parameters and choose Set. The Set Logical Port QoS Parameters dialog box (Figure 3-9) appears.

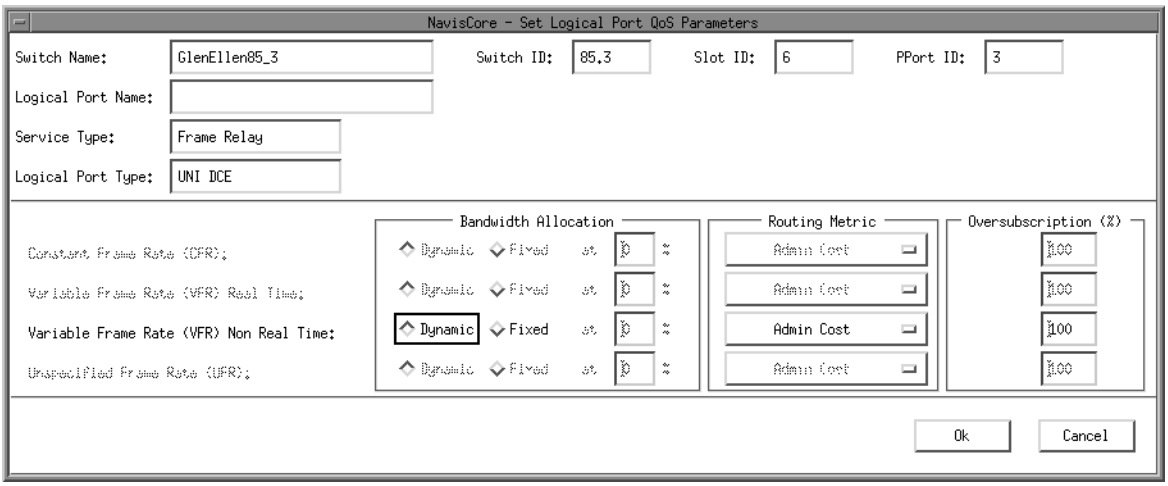

Figure 3-9. Set Logical Port QoS Parameters Dialog Box

3. Complete the required fields described in Table 3-10 for each service class.

Table 3-10. Set QoS Parameters Fields

Field	Action/Description
Bandwidth Allocation	Set the bandwidth allocation for each service class. Options include: *Dynamic* – Enables the bandwidth allocation to change dynamically according to bandwidth demands. Dynamic bandwidth allocation pools the remaining bandwidth for this logical port. This includes bandwidth that has not already been allocated to a specific queue or assigned to a connection. *Fixed* – Specifies the percentage of bandwidth you want to reserve for that service class. If all four service classes are set to Fixed, ensure that all four values add up to 100% so that you do not waste bandwidth. If you set the VFR service class bandwidth to "Fixed," you are specifying the maximum bandwidth to reserve for this type of traffic. If the network requests a circuit that exceeds the fixed value, the circuit cannot be created. If you set the UFR service class to "Fixed," you are guaranteeing that amount of service, at a minimum, for the UFR queue, provided the VFR queues are not oversubscribed. No bandwidth is actually allocated for UFR connections, so the port admits more connections into the UFR queue than it can service. ***Note**: If you have service classes set to Dynamic, any remaining bandwidth percentage is allocated to those service classes as needed. For example, if UFR is Fixed at 55%, and the two VFR classes are set to Dynamic, the remaining 45% of bandwidth will be dynamically allocated between the two VBR service classes.*
Routing Metric	Select one of the following *Routing Metrics* for each class of service. Routing metrics apply only if the port is configured as UNI DCE or UNI DTE logical port. Options include: *End-to-End Delay* – Measures the static delay of the logical port, which consists of both propagation and transmission delay. It is measured when the port initially comes up. It does not include queuing delays, and therefore does not account for port congestion. *Admin Cost* –Measures the Administrative Cost associated with the logical port.

Table 3-10. Set QoS Parameters Fields (Continued)

Field	Action/Description
Approx. Oversubscription (%)	(*Optional*) Specify the Oversubscription Factor percentage for each class of service (except CFR, which is set to 100% and cannot be modified). This value must be between 100% and 1000%. In general, you can leave these values set to 100%, since Ascend's Call Master Connection Admission Control (CAC) algorithm ensures that you can pack circuits on a port without losing data or quality of service. If, however, after monitoring your network, you determine that users of a particular service class are reserving more bandwidth than they are actually using, you can adjust the oversubscription values to suit your needs. By doing so, however, you may adversely impact the quality of service for this and lower-priority service classes.

Completing the Logical Port Configuration

When you finish selecting the additional options for this new logical port:

1. Choose OK from the Add Logical Port dialog box (Figure 3-3 on page 3-7). The Set All Logical Ports in PPort dialog box reappears (Figure 3-1 on page 3-3).
2. (*Optional*) To configure this logical port for a specific VPN and customer, see "Configuring a Logical Port for VPN" on page 8-7.
3. Choose Close to return to the Set Physical Port attributes dialog box. Then choose Cancel to return to the Switch Back Panel dialog box.

Configuring Logical Ports for Use With SVCs

If you plan to use SVCs in your network, you must configure three additional Set Attributes functions:

- The Set Frame Relay SVC option allows you to enable the Q.922 signaling feature. The fields on this dialog box also enable you to configure forward and reverse QoS class. To configure these attributes, continue with the next section, "Frame Relay SVC Attributes."

- The Set SVC Parameters option enables you to define various SVC screening and handling parameters for each logical port on the switch. Continue with the following section to configure these attributes. See page 3-27 to configure these attributes.

- The Set SVC Priorities option enables you to assign bandwidth and bumping priority to SVCs based on ingress QoS class. The network routes SVCs originating from this logical port according to the SVC ingress QoS class you select. The Set SVC Priorities option also enables you to assign forward and reverse circuit discard priorities to SVCs that originate on a specific logical port. See page 3-33 to configure these attributes.

For more information about SVCs, see Chapter 10, "Configuring Switched Virtual Circuit (SVC) Parameters," Chapter 11, "Closed User Groups,", and Chapter 12, "Port Security Screening."

Frame Relay SVC Attributes

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select [Frame Relay SVC] Attributes and complete the fields shown in Figure 3-10.

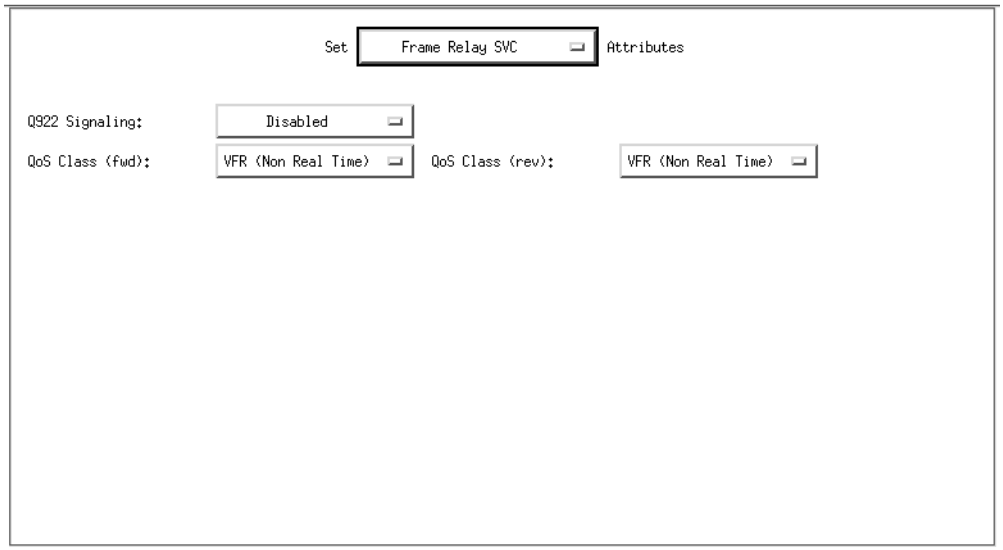

Figure 3-10. Set Frame Relay SVC Attributes

Table 3-11 defines how to configure the SVC attributes.

Table 3-11. Set Frame Relay SVC Attributes Fields

Field	Action/Description
Q.922 Signaling	Q.922 signaling must be set to *Enabled* for Frame Relay SVCs to function on the switch. The default value for this field is Disabled. For more information about Frame Relay SVCs and Q.922 Signaling, see Chapter 10, "Configuring Switched Virtual Circuit (SVC) Parameters."
QoS Class (fwd)	Select a QoS class for this logical port. For more information see Table 2-11 on page 2-14. Options include: VFR (Non Real Time) (default) VFR (Real Time) UFR
QoS Class (rev)	Select a QoS class for this logical port. For more information see Table 2-11 on page 2-14. Options include: VFR (Non Real Time) (default) VFR (Real Time) UFR

SVC Parameters Attributes

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [SVC Parameters] and complete the fields shown in Figure 3-11.

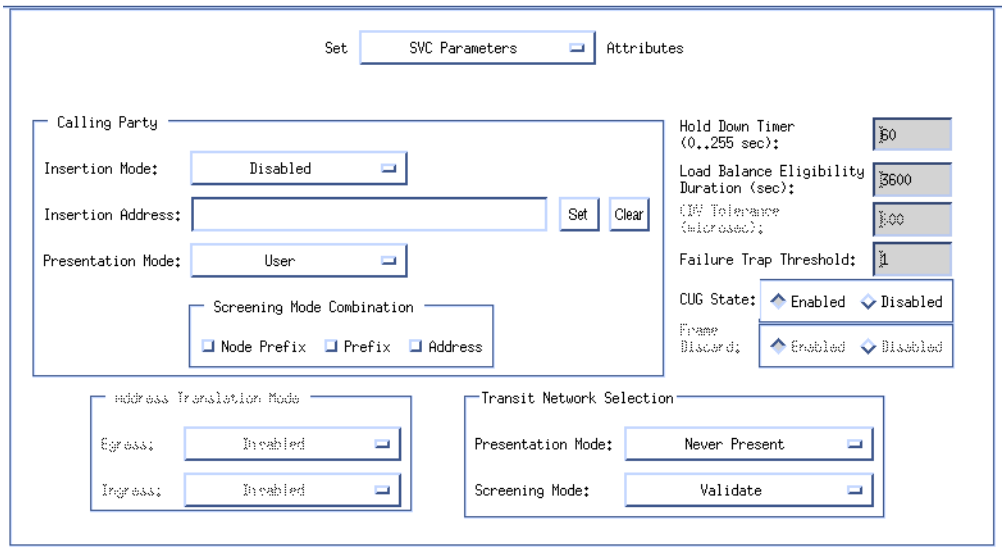

Figure 3-11. Set SVC Parameters Attributes

The following sections describe how to define these attributes.

Defining Calling Party Parameters

The following parameters configure the logical port for various address and screening options:

Insertion Mode and Insertion Address — Specifies how the logical port handles SVC requests.

Presentation Mode — Specifies whether or not to include the calling party address on outgoing SVCs.

Screening Mode Combination — Determines whether or not to process an ingress call at this logical port.

Insertion Address

1. Select one of the following Insertion Mode options:

For calling party screening to occur, set this field to Disable or Insert. If you select Replace, calling party screening is effectively disabled because the Calling Party Insertion Address is always considered valid. Also, if you select Insert, calling party screening occurs only when the caller signals the calling party address; if the caller does not signal the calling party address, the Calling Party Insertion Address, which is always considered valid, is used.

Option	Description
Disabled	The logical port does not insert or replace the calling party address. If you set the Insertion Mode field to Disable, skip to "Presentation Mode" on page 3-30.
Insert	If the logical port receives an SVC request that does not have a calling party information element, it inserts the address that is specified in the Calling Party Insertion Address field.
Replace	When the logical port receives an SVC request: • If there is no calling party address, it inserts the calling party address specified in the Calling Party Insertion Address field. • If there is a calling party address, it overwrites the existing calling party information element with the address specified in the Calling Party Insertion Address field.

2. Choose the Set command to the right of the Insertion Address field. The Set Insertion Address dialog box (Figure 3-12) appears.

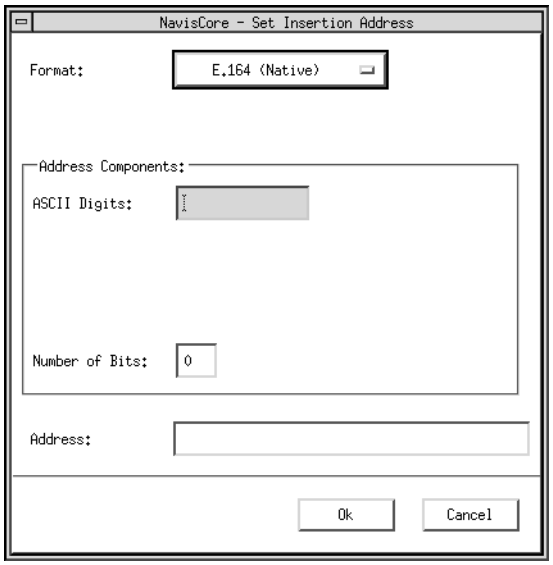

Figure 3-12. Set Insertion Address Dialog Box

The calling party insertion address is not used to route calls to this port. To use the calling party insertion address to route calls to this port, configure the address (or a prefix corresponding to the address) on this port. For more information, see "Configuring Port Addresses" on page 10-14.

3. Select the appropriate SVC Port Address Format. See Chapter 10, "Configuring Switched Virtual Circuit (SVC) Parameters," for information about Native E.164 and X.121 addresses. Then proceed to the following section to define the calling party presentation mode.

Presentation Mode

Select one of the following Presentation Mode options:

Option	Description
User	Include the calling party address based on the Presentation Indicator in the SETUP message of the user's SVC request.
Always	Always include the calling party address on outgoing calls, regardless of the Presentation Indicator in the SETUP message of the user's SVC request.
Never	Never include the calling party address on outgoing calls, regardless of the Presentation Indicator in the SETUP message of the user's SVC request.

Screening Mode Combination

Select one or more of the Screening Mode options. If you select more than one item, the ingress call is processed if it meets one or more of the selected criteria (for example, if you select both Node Prefix and Address, the calling party address must match either a valid node prefix or a valid port address).

If you enable screening at any level, and the calling party has no calling party address, the SVC fails unless you set the Calling Party Insertion Mode to Insert or Replace, and configure a Calling Party Insertion Address.

Select one of the following Screening Modes:

Option	Description
Node Prefix	Screens the calling party against all of the configured node prefixes. If a match is found, the screen is successful.
Prefix	Screens the calling party against all of the configured port prefixes. If a match is found, the screen is successful.
Address	Screens the calling party against all of the configured port addresses. If a match is found, the screen is successful.

Defining Transit Network Selection

Configure the Transit Network Selection options:

Option	Description
Presentation Mode	Select the egress presentation mode for the selected logical port. Options include: *Never Present* (default) – Never signal TNS in egress SVC requests. *Present Signaled TNS Only* – Signal TNS in egress SVC requests only if TNS was signaled by the user in the ingress SVC request. *Signaled or Source Default* – Signal TNS in egress SVC requests if TNS was signaled by the user in the ingress SVC request or a source default network ID was provisioned at the ingress user's logical port. ***Note:*** *Network IDs that do not match the adjacent network ID (see the Adjacent Network field in Table 10-8 on page 10-19) are processed according to the configured presentation mode; however, a network ID that matches the adjacent network ID will never be signaled in egress calls (as if presentation mode were Never).*
Screening Mode	Select the screening mode for the selected logical port. Options include: *Ignore* – Ignore the signaled TNS. *Accept* – Always accept the signaled TNS. *Validate* (default) – Screens the signaled TNS and ignores it if there is no match.

Defining Additional SVC Configuration Options

Set [SVC Parameters] Attributes (Figure 3-11 on page 3-27) provides additional SVC options, which you can configure using Table 3-12.

Although you can modify these fields, Ascend recommends you use the default parameters.

Table 3-12. Additional SVC Configuration Options

Field	Action/Description
Hold Down Timer	Enter the number of seconds to wait before the network initiates call clearing when a trunk has gone down. If you enter 0, the network clears the SVC immediately upon detection of a trunk outage.
Load Balance Eligibility Duration	Enter the number of seconds an SVC must be established before a call is eligible for load balance rerouting. The default is 3600 seconds. This feature is useful for those SVCs that are long term, and may encounter a forced reroute due to trunk failure.
Failure Trap Threshold	Enter the threshold crossing alarm value for SVC failure traps. The switch generates a trap if the internal SVC failure counter crosses this threshold during the current 15 minute time period. The internal counter is reset every 15 minutes. The default value of 1 means that if one SVC failure occurs on a logical port, a trap is issued and no additional traps are issued until the next 15-minute period expires. If you change the threshold value to 100, it means that to trigger a trap, 100 SVC failures must occur in a 15-minute window. If you enter 0, the switch never generates a failure trap.
CUG State	Select enable to allow CUG processing for this logical port.

SVC Priorities

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [SVC Priorities] Attributes and complete the fields described in Table 3-13. When you finish, continue with the instructions in "Selecting Additional Logical Port Options" on page 3-21.

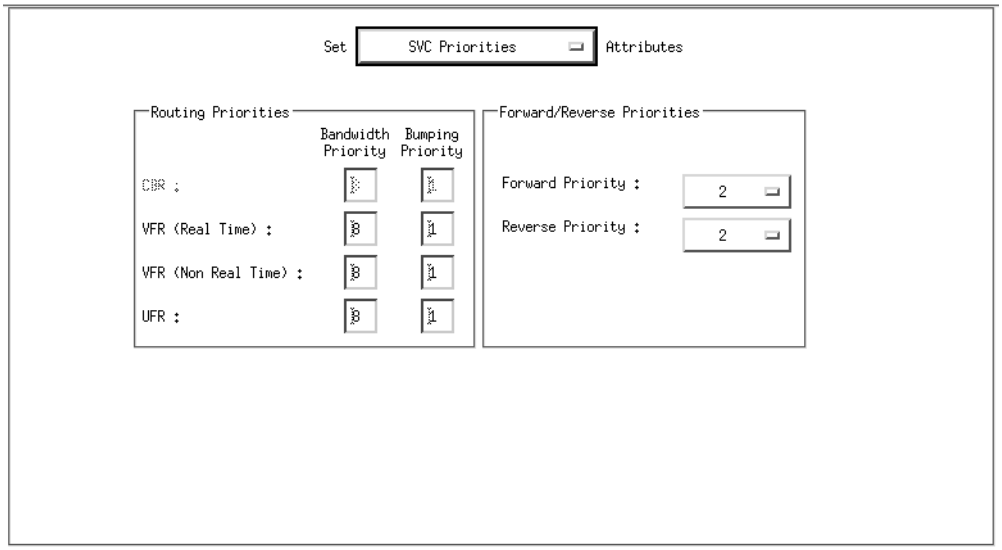

Figure 3-13. Set SVC Priorities

Table 3-13. SVC Priorities

Field	Action/Description
Bandwidth Priority	For each of the QoS queues, specify a value from 0 through 15 where 8 is the default and 0 indicates the highest routing priority.
Bumping Priority	For each of the QoS queues, specify a value from 0 through 7 where 1 is the default and 0 indicates the highest routing priority. (An active circuit with a bumping priority of 0 will not be bumped.)
Forward Priority	Specify a value from 1 through 3 where 2 is the default. The value sets the circuit discard priority in the ingress direction for SVCs that originate at this logical port.
Reverse Priority	Specify a value from 1 through 3 where 2 is the default. The value sets the circuit discard priority in the egress direction for SVCs that originate at this logical port.

Defining Frame Relay OPTimum PVC Trunk Logical Ports

To configure a Frame Relay OPTimum trunk, you must first configure either a UNI-DTE feeder or a Frame Relay NNI logical port on the same physical port.

> You cannot define a trunk logical port on a channelized T1/E1 module.

Use the following sequence to configure an OPTimum trunk:

1. Configure the physical port you want to use for the OPTimum trunk (see the *NavisCore Physical Interface Configuration Guide*).
2. Configure one of the following logical ports on this physical port:
   - Frame Relay DTE, or
   - NNI (page 3-8)

     Assign this logical port a minimum amount of bandwidth.
3. Follow the instructions in this section to configure a Frame Relay OPTimum trunk logical port. You can assign the remaining bandwidth to this logical port.

About DLCI Numbers

A data link connection identifier (DLCI) number is a 10-bit address that identifies PVCs. This DLCI number corresponds to the DLCI number the Frame Relay trunk uses to access the PDN. The PDN recognizes this as a normal PVC carrying user traffic.

Depending on your link management type, use the guidelines in Table 3-14 to define DLCI numbers.

Table 3-14. DLCI Number Guidelines

DLCI Number Range	Description
0-15	Reserved
16-991	Available for all link management types
16-1007	Available for LMI Rev 1 only
1008-1023	Reserved

Defining the OPTimum PVC Trunk

To define a logical port as a Frame Relay OPTimum PVC trunk:

1. Complete the Add Logical Port Type dialog box (Figure 3-2 on page 3-6) fields described in Table 3-15.

Table 3-15. Add Logical Port (OPTimum PVC Trunk) Fields

Field	Action/Description
Service Type	Select Frame Relay.
LPort Type	Select OPTimum PVC Trunk.
DLCI Number	Enter a data link connection identifier (DLCI) number that corresponds to the DLCI number the Frame Relay trunk uses to access the PDN. The PDN recognizes this as a normal PVC carrying user traffic.

2. Choose OK. The Add Logical Port dialog box appears and displays the Set Attributes option menu and fields shown in Figure 3-3 on page 3-7.
3. Complete the administrative attributes described in Table 3-4 on page 3-10.
4. Complete the congestion control attributes described in Table 3-5 on page 3-13.

Set the congestion control attributes on the feeder logical port only. You cannot define the threshold attributes on the OPTimum trunk logical port.

5. Complete the link management attributes described in Table 3-6 on page 3-16.
6. Complete the trap control attributes described in Table 3-7 on page 3-18.
7. Complete the priority frame attributes described in Table 3-8 on page 3-20.
8. Complete the QoS parameters described in Table 3-10 on page 3-23.
9. When you finish, proceed to "Completing the Logical Port Configuration" on page 3-24.

Defining Encapsulation FRAD, Direct Line Trunk, and PPP Logical Ports

This section describes how to define the following types of logical ports:

Encapsulation FRAD Services — Configure a logical port to perform Frame Relay encapsulation/de-encapsulation for the HDLC/SDLC-based protocol.

Direct Line Trunk Services — Configure the logical port for a trunk connection to another Ascend switch.

Point-to-Point-Protocol— Configure the logical port to enable a configured Point-to-Point Protocol (PPP) DTE device to communicate with another DTE device configured for Frame Relay and encapsulating multiprotocols, according to RFC 1490. This configuration enables you to define a single circuit between the two devices. PPP supports all IOP Type A cards, including the following modules:

- E1-PRI (ISDN)
- T1-PRI (ISDN)
- Unchannelized T1/E1
- Unchannelized 12-port E1
- Channelized T1/E1
- 4-port DSX-1
- 10-port DSX-1
- Channelized DS3
- Channelized DS3-1-0
- Universal IO (V.35, X.21)
- 2-port HSSI

To define encapsulated FRAD, direct line trunk, and PPP services:

1. Complete the Add Logical Port dialog box fields described in Table 3-16 .

Table 3-16. Add Logical Port (Other) Fields

Field	Action/Description
Service Type	Select Others.
LPort Type	Select a logical port type from the list.
LPort ID	For a channelized T1 module, enter a number between 1 and 24. For a channelized E1 module, enter a number between 1 and 30. For all other modules, the Logical Port ID is a read-only field that automatically defaults to one (1).

2. Choose OK. The Add Logical Port dialog box appears and displays the Set Attributes option menu and fields shown in Figure 3-3 on page 3-7.
3. Complete the administrative attributes described in Table 3-4 on page 3-10.
4. Complete the congestion control attributes described in Table 3-5 on page 3-13.
5. Complete the trap control attributes described in Table 3-7 on page 3-18.
6. Complete the priority frame attributes described in Table 3-8 on page 3-20.
7. When you finish, proceed to "Selecting Additional Logical Port Options" on page 3-21.

If you are defining PPP logical ports, proceed to the next section, "Completing the PPP Logical Port Configuration."

Completing the PPP Logical Port Configuration

When you configure PPP logical ports, you can define authentication attributes and other options for these ports.

Defining Authentication Attributes

Console authentication is a domain security feature that is handled by the Remote Access Dial-In User Service (RADIUS) protocol. Before you can define authentication attributes for PPP ports, you must add the authentication domain and configure the RADIUS server parameters. For more information about adding an authentication domain, see the *NavisCore NMS Getting Started Guide*.

To define authentication attributes for the two PPP ports:

1. Access the Set All Logical Ports in PPort dialog box as described in "Accessing Frame Relay Logical Port Functions" on page 3-2.
2. Choose Modify. The Modify Logical Port dialog box (Figure 3-14) appears.
3. Select Authentication from the Set Attributes option menu. The system displays the authentication attributes shown in Figure 3-14.

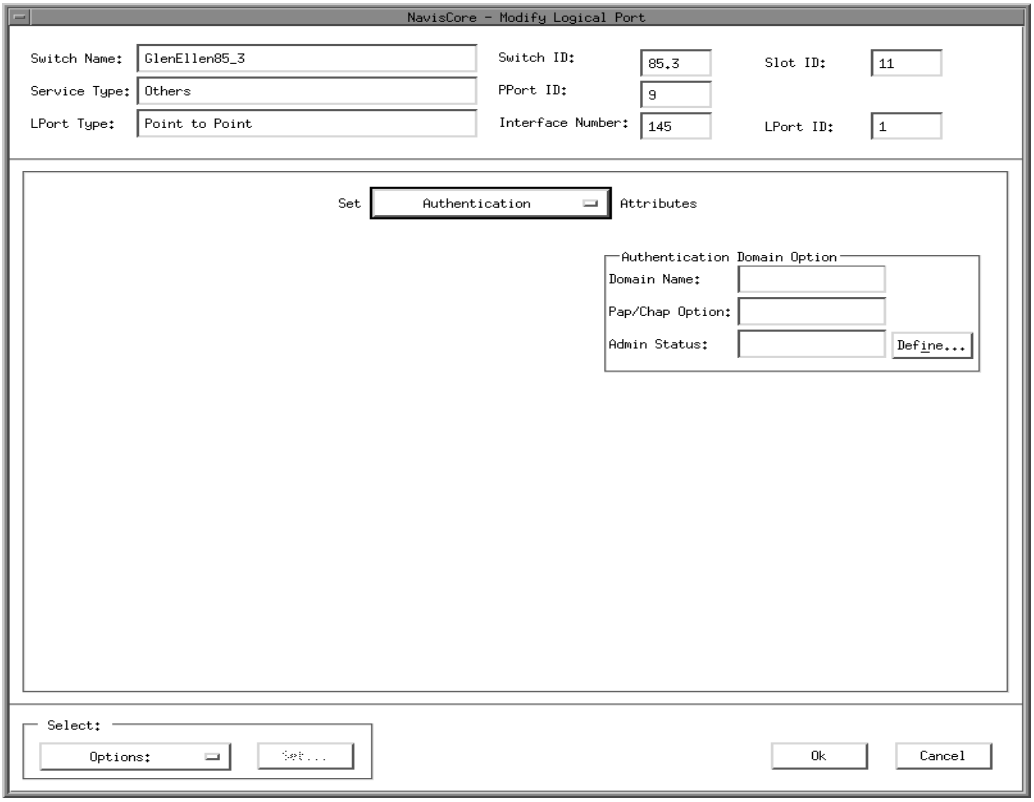

Figure 3-14. Modify Logical Port Dialog Box

4. Choose Define. The Set Authentication Info dialog box (Figure 3-15) appears.

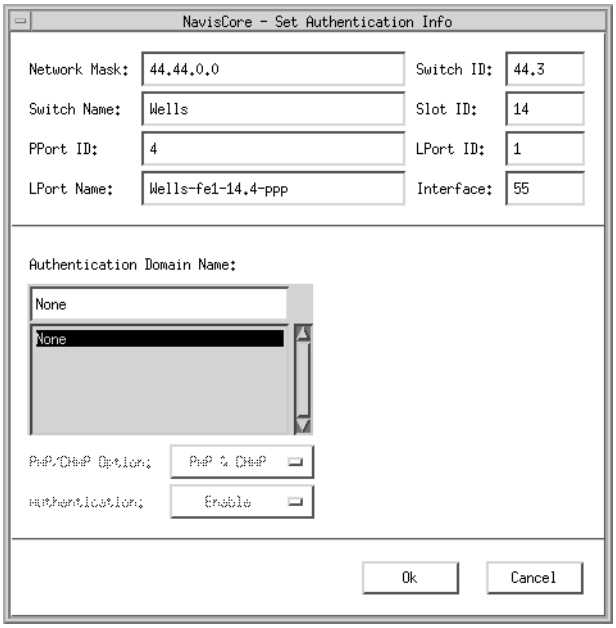

Figure 3-15. Set Authentication Information Dialog Box

5. Complete the Set Authentication Information dialog box fields described in Table 3-17.

Table 3-17. Set Authentication Attributes Fields

Field	Action/Description
Authentication Domain Name	Select the Authentication Domain. Each switch that has access to a RADIUS server has an Authentication Domain name. See the *NavisCore NMS Getting Started Guide* for more information about RADIUS.
PAP/CHAP Option	Select PAP only, CHAP only, or PAP & CHAP to establish Password Authentication Protocol, Challenge Handshake Authentication Protocol, or a combination of PAP and CHAP.
Authentication	Select Enable to enable the port to authenticate the connection to the RADIUS server. See the *NavisCore NMS Getting Started Guide* for more information about RADIUS.

6. Choose OK to accept the Authentication attributes. The Modify Logical Port dialog box reappears (Figure 3-14 on page 3-38).
7. Choose OK to return to the Set All Logical Ports in PPort dialog box.
8. Choose Close to exit the dialog box.

Defining the PPP Options

To define the PPP options:

1. Access the Set All Logical Ports in PPort dialog box as described in "Accessing Frame Relay Logical Port Functions" on page 3-2.
2. Choose Modify. The Modify Logical Port Type dialog box appears (Figure 3-14 on page 3-38).
3. Select Service Type: Others and LPort Type: Point-to-Point Protocol, and then choose OK. The Modify Logical Port dialog box appears (Figure 3-14 on page 3-38).
4. Select Options: PPP Option, and then choose Set. The system displays the PPP Options shown in Figure 3-16.

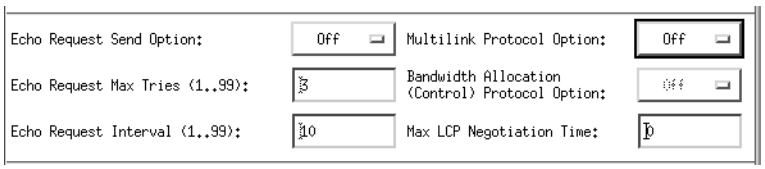

Figure 3-16. PPP Options

5. Complete the Set PPP Options dialog box fields described in Table 3-18.

Table 3-18. PPP Option Fields

Field	Action/Description
Echo Request Send Option	Select On to send keep-alive packets to the remote user.
Echo Request Max Retries (1..99)	Enter a number from 1 to 99 that represents the maximum number of keep-alive packets sent to the remote user.
Echo Request Interval (1..99)	Enter a number from 1 to 99 that represents the time interval between each keep-alive packet.
Multilink Protocol Option	This option, which enables Multilink PPP, is not supported in this release.
Bandwidth Allocation (Control) Protocol Option	Set this value to On if the associated router supports the Bandwidth Allocation Control Protocol (BACP). (See the *BACP/BAPP Internet Draft* for a detailed description of these protocols.) Set this value to Off if the associated router does not support BACP.
Max LCP Negotiation Time	Enter a number that represents the maximum time interval for the Link Control Protocol (LCP) to negotiate the exchange of packets.

6. Choose OK to return to the Modify Logical Port dialog box.

7. Choose OK to return to the Set All Logical Ports in PPort dialog box.

8. Choose Close to exit the dialog box.

Defining Multilink Frame Relay (MLFR) Trunks

Defining Multilink Frame Relay (MLFR) trunks requires creation of ML Member logical ports, which are then bound to the MLFR trunk bundle logical port.

About MLFR

MLFR is a method of aggregating available bandwidth on a set of Frame Relay logical links between two networking devices. The aggregated links, collectively referred to as the Multilink Frame Relay Unit (MFRU), can be thought of as a single logical link. As shown in Figure 3-17, the MFRU provides a single logical link (with 4*T1 bandwidth) between the router and the Frame Relay switch.

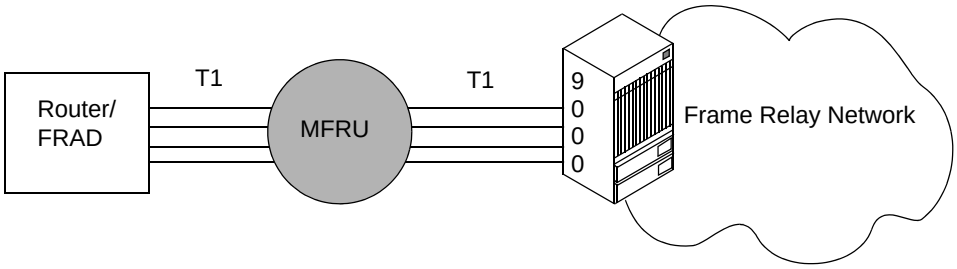

Figure 3-17. Multilink Frame Relay Unit (MFRU)

MLFR is implemented through the encapsulation of Frame Relay packets within a Multipoint-like frame. User and control packets are encapsulated enabling several logical links to be combined. PVC traffic is automatically distributed across the multiple links. MLFR provides a cost-effective, high-speed service without the need for additional hardware.

MLFR is supported on the following IOPs:

- 1-Port Channelized DS3
- 2-Port HSSI
- 4-Port channelized T1/E1
- 4-Port DSX
- 4-Port PRI E1/T1
- 8-Port UIO
- 10-Port DSX-1
- Unchannelized 12-Port E1

ML Member Logical Ports and MLFR Trunk Bundle Logical Ports

ML Member logical ports inherit their configuration from the logical port to which they are bound, therefore you only need to configure the administrative attributes described in Table 3-19 on page 3-46. The ML Member logical port can be bound to only one MLFR trunk bundle logical port, and the trunk bundle logical port must be on the same card.

A multilink trunk bundle logical port is created at the card level (not the physical port level). A maximum of 32 ML Member logical ports can be bound to a MLFR trunk bundle logical port. You should create the MLFR trunk with each MLFR trunk bundle endpoint containing the same number of bound MLFR logical ports and aggregate bandwidth (the NMS does not enforce this condition).

MLFR Logical Port Configuration Process

Figure 3-18 illustrates the process for defining MLFR over trunks. The next two sections describe these steps in more detail.

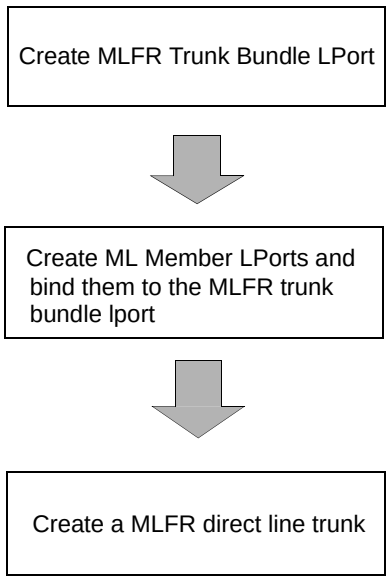

Figure 3-18. MLFR Over Trunks Configuration Process

Defining MLFR Trunk Bundle Logical Ports

Complete the following steps to define a MLFR trunk bundle logical port:

1. Select the switch to which you want to add a logical port.
2. From the Administer menu, select Ascend Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
3. Select the card you want to configure and press the right mouse button to display a popup menu. Select Card. The Set Card Attributes dialog box appears.
4. From the Set Card Attributes dialog box, choose MLFR Logical Ports. The Set All MLFR Trunk Bundle Logical Ports on Card dialog box appears.
5. Choose Add. The Add Logical Port dialog box appears.
6. Complete the administrative attributes described in Table 3-4 on page 3-10. The bandwidth field for this logical port type is read-only.
7. Complete the congestion control attributes described in Table 3-5 on page 3-13.
8. Complete the trap control attributes described in Table 3-7 on page 3-18.
9. Complete the priority frame attributes described in Table 3-8 on page 3-20.
10. Complete the QoS Service Classes attributes described in Table 3-10 on page 3-23.
11. When you finish, proceed to "Completing the Logical Port Configuration" on page 3-24.

Defining ML Member Logical Ports

Complete the following steps to define a ML Member logical port.

1. Select the switch to which you want to add a logical port.
2. From the Administer menu, select Ascend Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
3. Select the physical port you want to configure and press the right mouse button to display a popup menu. Select Logical Port. The Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-3) appears.
4. Choose Add. The Add Logical Port Type dialog box (Figure 3-2 on page 3-6) appears.
5. Select Others as the Service Type and choose ML Member as the LPort type.

6. Choose OK. The Add Logical Port dialog box appears (Figure 3-19 on page 3-45). ML Member logical ports inherit their configuration from the MLFR trunk bundle logical port to which it is bound, therefore you only need to configure the administrative attributes. The dialog box displays the MLFR trunk bundle that indicates the MLFR trunk bundle logical port to which it is bound.

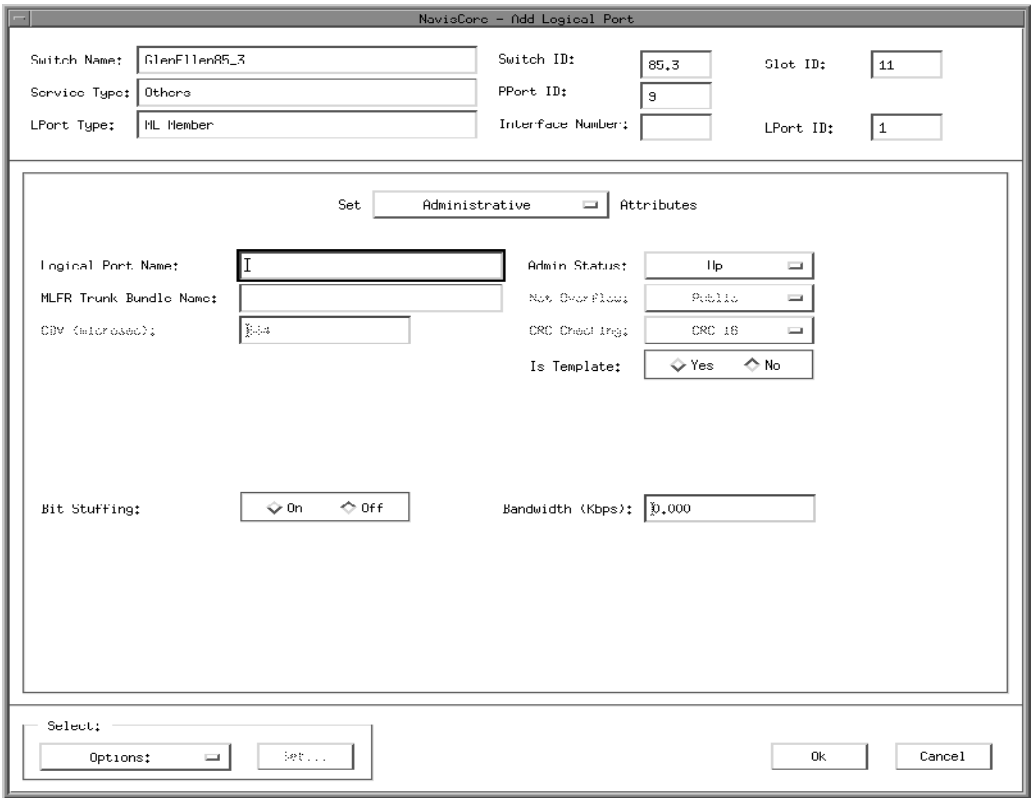

Figure 3-19. Add Logical Port Dialog Box

7. Complete the required administrative attributes fields described in Table 3-19.

Table 3-19. Set Administrative Attributes Fields

Field	Action/Description
Logical Port Name	Enter an alphanumeric logical port name (up to 32 characters in length) to assign this port.
Admin Status	Set the Admin Status. to *down* to save the configuration in the database without activating the port or to take the port off-line to run diagnostics. *Up (default)* – Activates the port. *Down* – Saves the configuration in the database without activating the port or takes the port off-line to run diagnostics.
Is Template	(*Optional*) Save these settings as a template to use again to quickly configure a logical port with the same options. To create a template, choose Yes in the *Is Template* field. See "Using Templates" on page 2-15 for more information.
Bit Stuffing	Select the bandwidth that matches the bandwidth capability of the customer premise equipment (CPE) connected to this logical port. Enables bit stuffing on T1/E1/DSX-1 ports. Bit stuffing effects the available bandwidth of each DS0/TS0 channel on this port. *On* – Provides 56 Kbps of bandwidth. *Off* – Provides 64 Kbps of bandwidth.
Bandwidth (Kbps)	Enter the amount of bandwidth you want to configure for this logical port. The default is the amount of bandwidth remaining from the physical clock rate, less any logical ports already configured. To define a trunk logical port on this same physical port, decrease the amount of bandwidth on this logical port to ensure sufficient remaining bandwidth. For example: *Physical port clock speed*: 1536 Kbps *Logical port UNI-DTE/NNI Feeder Bandwidth*: 56 Kbps *Logical port Frame Relay Trunk Bandwidth*: 1480 Kbps The example configuration allocates a PDN trunk with 1480 Kbps bandwidth between two Ascend switches, each attached to a PDN network.

8. Choose OK to save the parameters and exit the dialog box.

Binding and Unbinding ML Members to MLFR Bundle Logical Ports

To bind or unbind ML Member logical ports to MLFR bundle logical ports:

1. Select the switch to which you want to modify the MLFR bundle logical port.
2. From the Administer menu, select Ascend Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
3. Select the card on which this port resides and press the right mouse button to display a popup menu. Select Card. The Set Card Attributes dialog box appears.
4. From the Set Card Attributes dialog box, choose Configure MLFR Bundles. The Configure MLFR Trunk Bundle LPorts dialog box (Figure 3-20) appears .

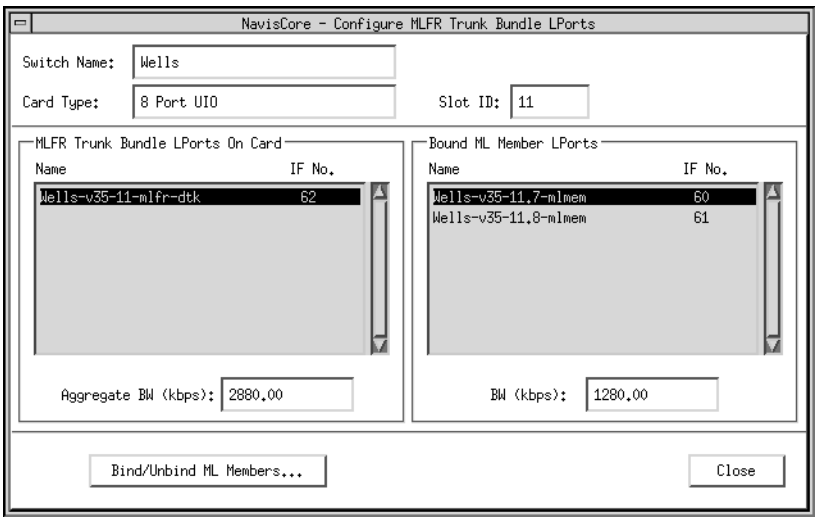

Figure 3-20. Configure MLFR Bundle Dialog Box

This dialog box displays existing MLFR trunk bundle logical ports on this card and ML Member logical ports bound to the selected MLFR trunk bundle logical port. Table 3-20 describes the dialog box fields.

Table 3-20. Configure MLFR Trunk Bundle Logical Ports Fields

Field	Description
Switch Name	Displays the name of the B-STDX switch on which the MLFR trunk bundle logical port resides.
Card Type	Displays the type of card on which the MLFR trunk bundle logical port resides.
Slot ID	Displays the ID of the slot in which the MLFR trunk bundle logical port resides.
MLFR Trunk Bundle LPorts on Card	Lists all existing MLFR trunk bundle logical ports currently defined on the card.
Aggregate BW (Kbps)	Displays the selected MLFR trunk bundle logical port's total aggregate bandwidth.
Bound ML Member LPorts	Lists all ML Member logical ports bound to the selected MLFR trunk bundle logical port.
BW (Kbps)	Displays the selected ML Member logical port bandwidth.

5. Select the desired MLFR trunk bundle logical port from the list box on the left.
6. Choose Bind/Unbind ML Members. The Create a MLFR Trunk Bundle dialog box appears as shown in Figure 3-21 on page 3-49. This dialog box enables you to bind and unbind ML Member logical ports to the selected MLFR trunk bundle logical port.

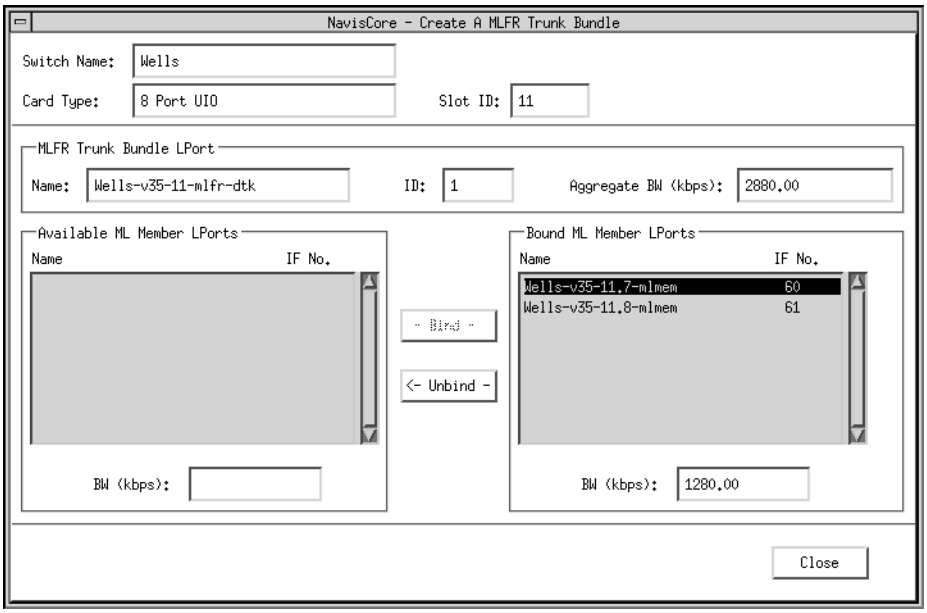

Figure 3-21. Create a MLFR Trunk Bundle

Table 3-21 describes the dialog box fields.

Table 3-21. Create a MLFR Trunk Bundle Fields

Field	Description
Switch Name	Displays the name of the B-STDX switch on which the MLFR trunk bundle logical port resides.
Card Type	Displays the type of card on which the MLFR trunk bundle logical port resides.
Slot ID	Displays the ID of the slot in which the MLFR trunk bundle logical port resides.
MLFR Trunk Bundle LPort	Displays the selected MLFR trunk bundle logical port name, ID, and aggregate bandwidth.
Available ML Member LPorts	Lists all available ML Member logical ports on the card that can be bound to the MLFR trunk bundle.
BW (Kbps)	Displays the selected ML Member logical port's available bandwidth.
Bound ML Member LPorts	Lists all bound ML Member logical ports bound to the selected MLFR trunk bundle logical port.
BW (Kbps)	Displays the selected ML Member logical port's bound bandwidth.

7. Do one of the following tasks:
   - To bind an ML Member logical port to the bundle:

     Select an ML Member logical port from the Available ML Member LPort list on the left and choose Bind. The selected logical port is removed from the available list and added to the bound list on the right. The system updates the Aggregate BW (Kbps) field to include the bound logical port's bandwidth.
   - To unbind a bound ML Member logical port from the bundle:

     Select a bound ML Member logical port from the Bound ML Member LPort list on the right and choose Unbind. The selected logical port is removed from the bound list and added to the available list on the left. The system updates the Aggregate BW (Kbps) field to include the deletion of the unbound logical port's bandwidth from the bundle.
8. Choose Close to exit the dialog box.

To delete a MLFR logical port binding from a MLFR trunk bundle logical port, see "Deleting Frame Relay Logical Ports" on page 2-16.

Create a MLFR direct link trunk between cards using MLFR trunk bundle logical ports as endpoints. See "Adding a Trunk" on page 4-11 for more information.

4

Configuring Trunks

This chapter describes how to configure a trunk in an Ascend switch network. A trunk is the communications circuit between two switches. The trunk enables two Ascend switches to pass data to each other and exchange internal control messages.

About Trunks

The Trunk oversubscription Factor and the OSPF Trunk administrative cost parameters enable you to better manage trunk traffic. The oversubscription factor enables you to configure more circuits to a trunk than can be supported at one time (over subscribe). Oversubscription assumes that due to the bursty nature of network traffic, not all circuits on the trunk are operating at the committed information rate (CIR) at the same time. Therefore, trunk bandwidth should remain sufficient.

The trunk administrative cost enables you to assign a cost value for the trunk. When multiple trunks are available, a circuit will use the trunk with the lowest administrative cost.

Trunk Oversubscription Factor

The trunk oversubscription factor percentage enables you to optimize the aggregate committed information rate (CIR) allowed over the trunk. The oversubscription factor represents the V value for this trunk. The bandwidth on trunks is reserved at runtime based on the CIR value of the PVCs that traverse that trunk.

The routing for PVCs is determined by either an OSPF algorithm or by the network administrator (if you manually define the circuit path). Each time a PVC attempts to come up, OSPF reserves bandwidth equal to the CIR of the PVC on the trunk with the shortest path. The amount of reserved bandwidth is deducted from the available virtual bandwidth pool. The formula used to determine virtual bandwidth is only used for allocating the initial path for the PVC. The system periodically reviews each PVC to optimize network resources according to the reroute tuning parameters (see the *NavisCore NMS Getting Started Guide*).

OSPF uses the following two formulas to determine the available virtual bandwidth value:

Formula 1

This formula determines the initial value of the available virtual bandwidth:

Initial Value = 0.95 (configured bandwidth) x V(%)

Note: V = trunk oversubscription factor

Formula 2

Available Virtual Bandwidth = Initial Value – (Sum of PVC CIR)

It is important to note that the available virtual bandwidth can become negative in extreme situations. If a number of trunks fail, PVC rerouting may cause the available virtual bandwidth value to become negative. Existing PVCs can be rerouted over a negative virtual bandwidth trunk. However, new PVCs cannot traverse trunks that have a negative virtual bandwidth.

If you configure the trunk oversubscription factor at a higher percentage, you increase the available virtual bandwidth (more PVC CIR) over the trunk. An oversubscription value of 200% effectively doubles the available virtual bandwidth. Ascend switches reserve 5% bandwidth for network management, routing updates, and other management traffic.

If all network traffic attempts to use the network resources at the same time (for example, during multiple file-transfer sessions over the same trunk), the overhead will degrade network performance.

OSPF Trunk Administrative Cost

OSPF trunk administrative cost is a function of OSPF that gives you more control over the specific path a virtual circuit will take through the network. Through OSPF, a circuit can choose the shorter hop path (most direct route across network), regardless of the available bandwidth.

OSPF trunk administrative cost only works in networks where all switches are running switch software, Release 4.1 or higher. If some switches are running an earlier release, OSPF only selects the path with the greatest amount of available bandwidth. This is not necessarily the most direct route (minimum number of hops) through the network.

When you first define a circuit, the circuit looks for a path that has enough virtual bandwidth available to handle its committed information rate (CIR). If the circuit finds more than one path with the available bandwidth, the circuit chooses the path with the lowest administrative cost. If there is more than one path with the same administrative cost, the circuit chooses the path that has the most available bandwidth.

Circuits are automatically rerouted around a trunk or switch failure. If the circuit cannot find a path with sufficient bandwidth, it chooses the path with the lowest administrative cost, even if this trunk has a negative bandwidth value. (The negative bandwidth indicates that the trunk is oversubscribed). Circuits use a path with a negative bandwidth only when a trunk fails.

Configuring Minimum-Hop Paths

If you use the default administrative cost value of 100, OSPF selects minimum-hop paths that respect the circuit's Quality of Services values. You can also use the following guidelines to configure this value:

- To minimize end-to-end delay, configure an administrative cost that is proportional to the propagation delay of the trunk. Set the cost of each trunk to the length of the trunk's physical media (in miles or kilometers).
- Set the administrative cost relative to the speed of the physical port. For example, a single T1 trunk hop may be equal to four HSSI trunk hops. You would set the HSSI trunk's cost to 25 and the T1 cost to 100. Keep in mind that since OSPF routing considers available bandwidth, administrative cost is not necessarily a function of bandwidth.

Link Trunk Protocol

Using Link Trunk Protocol (LTP), switches communicate by exchanging keep-alive (KA) control frames. Switches send KA requests at regular time intervals (one per second). After a switch receives a KA request, it returns a KA reply. A completed transaction consists of a KA request and a KA reply. The request and reply frame formats are identical.

Trunk Delay

Figure 4-1 illustrates the process of keep-alive frames used to measure trunk delay. When Switch A sends a KA request to Switch B, a time stamp is put into the KA request frame. When Switch B receives the KA request, it sends a KA reply to Switch A. Switch A receives the KA reply and calculates the round-trip delay from Switch A to Switch B.

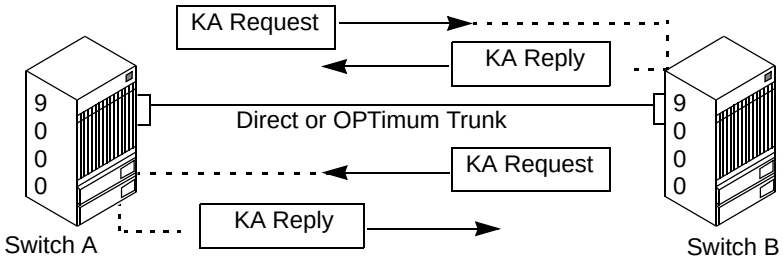

Figure 4-1. Trunk Delay - OSPF Metric and Keep-Alive Messaging

Keep-Alive Threshold

The Keep Alive Threshold field in the Set All Trunks dialog box (Figure 4-2 on page 4-6) represents the number of retries that the trunk protocol attempts before bringing the trunk down. The retry interval is represented in seconds. You can set the keep-alive threshold value between 3 and 255 seconds. The default is 5 seconds.

Trunk Backup

The Ascend switch supports a trunk backup option. Trunk backup can be automatic or manual, and it enables you to set up one or more backup trunks to replace a primary trunk. If an Ascend switch trunk line fails or requires maintenance, you can reroute PVCs from the primary trunk to the backup trunk. You can define primary and backup trunks on any I/O module.

You define a backup trunk in the Add Trunk dialog box (Figure 4-4 on page 4-13). A backup trunk can have a total bandwidth that is less than that of the primary trunk. To avoid congestion, you can configure multiple backup trunks to back up a single primary trunk. The Ascend switch allows you to define up to eight backup trunks for a single primary trunk.

Once you configure the primary and backup trunk(s), you can configure the primary trunk to automatically back up upon failure. If a trunk line requires maintenance, you can manually initiate and terminate a trunk backup.

Defining a Trunk

When you define a trunk, you must perform three steps:

Step 1. Configure a trunk logical port type. See one of the following sections:

"Defining Frame Relay OPTimum PVC Trunk Logical Ports" on page 3-34.

"Defining Encapsulation FRAD, Direct Line Trunk, and PPP Logical Ports" on page 3-36 (describes Direct Line Trunk).

Step 2. Define a trunk configuration between the two switches. See "Adding a Trunk" on page 4-11.

Step 3. Create the map line connection that corresponds to the trunk configuration. See "Creating a Trunk-Line Connection" on page 4-19.

Accessing Trunk Functions

The Set All Trunks function specifies the two endpoints for an Ascend-to-Ascend switch trunk. When you configure a trunk, you select endpoints that use the same type of logical port (such as Direct Line Trunk) and the same bandwidth.

To access the Set All Trunks dialog box, from the Administer menu, select Ascend Parameters ⇒ Set All Trunks. When the Set All Trunks dialog box appears, position the cursor in the Search by Name field and press Return to display the list of defined trunk names, as shown in Figure 4-2.

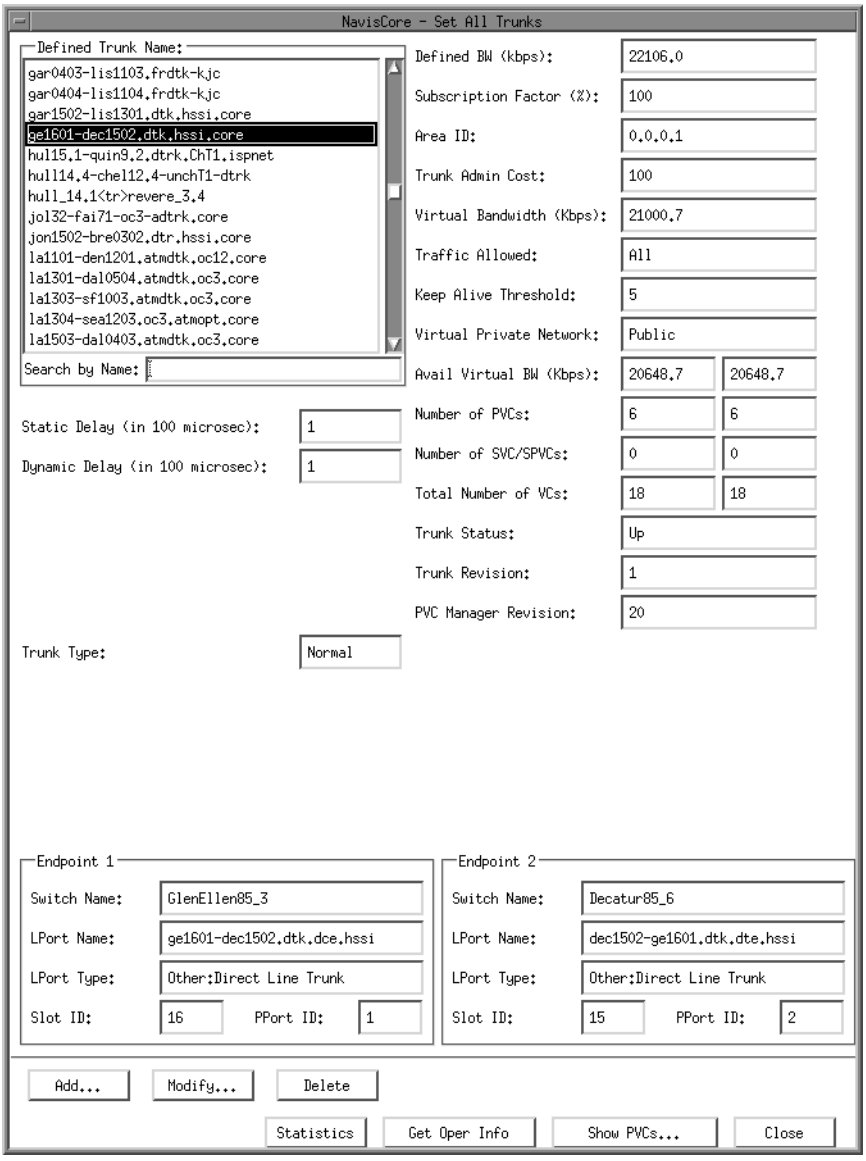

Figure 4-2. Set All Trunks Dialog Box

To learn more about the Set All Trunks dialog box fields and commands, continue with the following section. To begin defining trunks, proceed to "Adding a Trunk" on page 4-11.

About the Set All Trunks Dialog Box

The Set All Trunks dialog box (shown in Figure 4-2) displays information about the trunk you select from the Defined Trunk Names list. It also provides several commands that enable you to access additional trunk functions.

Table 4-1 describes the Set All Trunks dialog box fields and commands.

Table 4-1. Set All Trunks Dialog Box Fields and Commands

Field/Command	Action/Description
Defined Trunk Name	Displays the names of the trunks configured for the current network map.
Defined BW (Kbps) (*non-MLFR trunks*) Aggregate BW (Kbps) (*MLFR Direct Trunks only)*	Displays the amount of bandwidth, in Kbps, for the selected trunk line. Displays the maximum of the aggregate bandwidths of the endpoint bundles.
Subscription Factor (%)	Displays the percentage used to calculate the available virtual bandwidth for the selected trunk.
Area ID	Enter the area ID (x.x.x.x) for the area in which you want to locate this OSPF interface. The range of available values is from 0.0.0.0 to 255.255.255.255. Area 0.0.0.0 is the network backbone area. Area 0.0.0.1 is Area 1. Areas are collections of networks, hosts, and routers used for IP routing. The area ID identifies the area. If a trunk is in Area 1 and the OSPF Backwards Compatibility option (which is set through IP services) is set to Yes, external routes are not advertised across that link. ***Note:*** *Area 1 is reserved for Ascend switches. For a detailed description of OSPF areas, and how to use IP Services to configure multiple OSPF areas, see the* NavisCore IP Navigator Configuration Guide.
Trunk Admin Cost	Displays a value that defines the cost of using this trunk for a virtual circuit when a virtual circuit is being dynamically created on the switch.

Table 4-1. Set All Trunks Dialog Box Fields and Commands (Continued)

Field/Command	Action/Description
Virtual Bandwidth (Kbps)	Displays the amount of virtual bandwidth in Kbps. The value *.95* is used because *.05%* of the bandwidth is reserved for network management, routing updates, and other management traffic.
Traffic Allowed	Displays the type of management traffic allowed on this trunk.
Keep Alive Threshold	Displays the number of seconds that the trunk protocol will exchange keep-alive (KA) control frames without getting a response from the remote node. ***Note**: Service is disrupted if you change this value after the trunk is online.*
Virtual Private Network	Displays the virtual private network name. (Displays *Public* if the trunk is not dedicated to a specific VPN.) For more information about VPNs, see Chapter 8, "Configuring Virtual Private Networks."
Avail Virtual BW (Kbps)	Displays the amount of bandwidth, in Kbps, available for circuit configuration and allotment on the selected trunk.
Number of PVCs	Displays the total number of PVCs traversing the trunk logical port endpoint.
Number of SVC/SPVCs	The total number of SVCs and SPVCs traversing the trunk logical port endpoint.
Total number of VCs	The total number of PVCs, SVCs, SPVCs, MPTs, and any other type of VC traversing the trunk.

Table 4-1. Set All Trunks Dialog Box Fields and Commands (Continued)

Field/Command	Action/Description
Trunk Status	Displays the current status of the selected trunk. Options include: *Unknown* – The NMS cannot communicate with one or both switch endpoints that make up this trunk. *Down* – The switches cannot establish a communication link. *Attempt* – A switch is attempting to contact another switch but has not yet received a response. *Init* – A one-way communication exists between the two switches. *Two-way* – A bi-directional communication exists between the two switches. *Exchange Start* – The two switches are exchanging network topology. *Exchange* – The two switches are exchanging network topology. *Loading* – The two switches are requesting the most recent link state information. *Up* – The trunk is up and operational between the two switches.
Trunk Revision	Displays the revision of link trunk protocol software at each endpoint.
PVC Manager Revision	Displays the PVC manager software revision.
Static Delay (in 100 microsec)	Represents the measured one-way delay in units of 100 microseconds. This measurement is taken when the trunk initializes and it is only updated when the trunk changes state from down to up. The static delay value is used in conjunction with the end-to-end delay routing metric to enable you to route circuits over trunks with the lowest end-to-end delay.
Dynamic Delay (in 100 microsec)	Represents the measured one-way delay in units of 100 microseconds. This measurement is made continually on operational trunks. Under most conditions, the dynamic delay value will match the static delay value. However, if some characteristics of the underlying transmission media for the trunk change such that the dynamic delay changes, this value may differ from the static delay.

Table 4-1. Set All Trunks Dialog Box Fields and Commands (Continued)

Field/Command	Action/Description
Trunk Type	Displays the trunk type. *Normal* – Indicates a common trunk. *Primary* – Indicates that the trunk has a backup for fault tolerance. *Backup* – Indicates that it is the backup trunk (when failure occurs on the primary trunk).
Add/Modify/Delete	If you have already configured some trunk-line connections, the dialog box displays the names. The *Add*, *Modify* or *Delete* commands enable you to add, modify or delete trunk configurations.
Statistics	Displays the summary statistics for the selected trunk configuration. For more information about summary statistics, see the *NavisCore Diagnostic and Troubleshooting Guide*.
Get Oper Info	Displays a brief status for the selected trunk connection and a status message appears in the Oper Status field.
Show PVCs	Displays a dialog box that contains a list of the PVCs that traverse the selected trunk. This dialog box also provides logical port descriptions for each PVC endpoint. ***Note:*** *This function only works when both switches at either end of the trunk are running one of the following minimum switch software versions (or greater):* B-STDX switch software 06.00.xx.xx CBX 500 switch software 03.00.xx.xx

Adding a Trunk

To add a trunk:

1. Access the Set All Trunks dialog box, shown in Figure 4-2. (See "Accessing Trunk Functions" on page 4-6.)
2. Choose Add. The Select Logical Ports dialog box (Figure 4-3) appears.

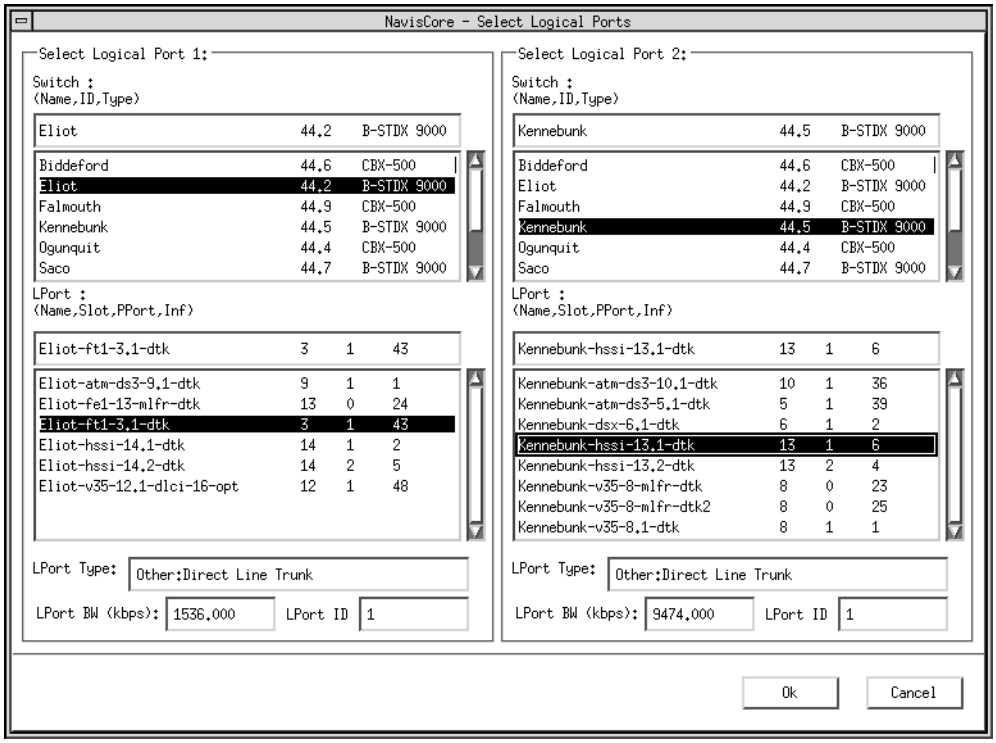

Figure 4-3. Select Logical Ports Dialog Box

3. From the Select Logical Ports window, select the switch and Lport for logical port 1 and logical port 2. Table 4-2 describes each of the Select Logical Ports fields.

Table 4-2. Select Logical Ports Fields

Field	Action/Description
Switch (Name, ID, Type)	Select a switch for each endpoint. The dialog box displays the parameters for the selected switch.
LPort (Name, Slot, PPort, Inf)	Select the same trunk logical port type for each endpoint. Choose from the following logical port types depending on the type of logical port service: • Frame Relay OPTimum Trunk • Other:Direct Line Trunk • SMDS OPTimum Trunk (see the *NavisCore SMDS Configuration Guide*) • ATM: Direct Trunk (see the *NavisCore ATM Configuration Guide*) • ATM OPTimum Frame Trunk (see the *NavisCore ATM Configuration Guide*) • ATM OPTimum Cell Trunk (see the *NavisCore ATM Configuration Guide*) • MLFR Direct Line Trunk. (This field also displays the ifnum, physical port number, and I/O slot (number) in which the module resides.) ***Note:*** *For non-MLFR trunks, review the LPort Bandwidth field for each endpoint to make sure the bandwidth is identical.*
LPort Type	Displays the configured logical port type.
LPort BW (kbps)	Displays the bandwidth configured for the logical port. For non-MLFR trunks, this must be the same for both endpoints.
LPort ID	Displays the logical port number.

4. Choose OK. The Add Trunk dialog box (Figure 4-4) appears, displaying the parameters for both switches in the trunk configuration.

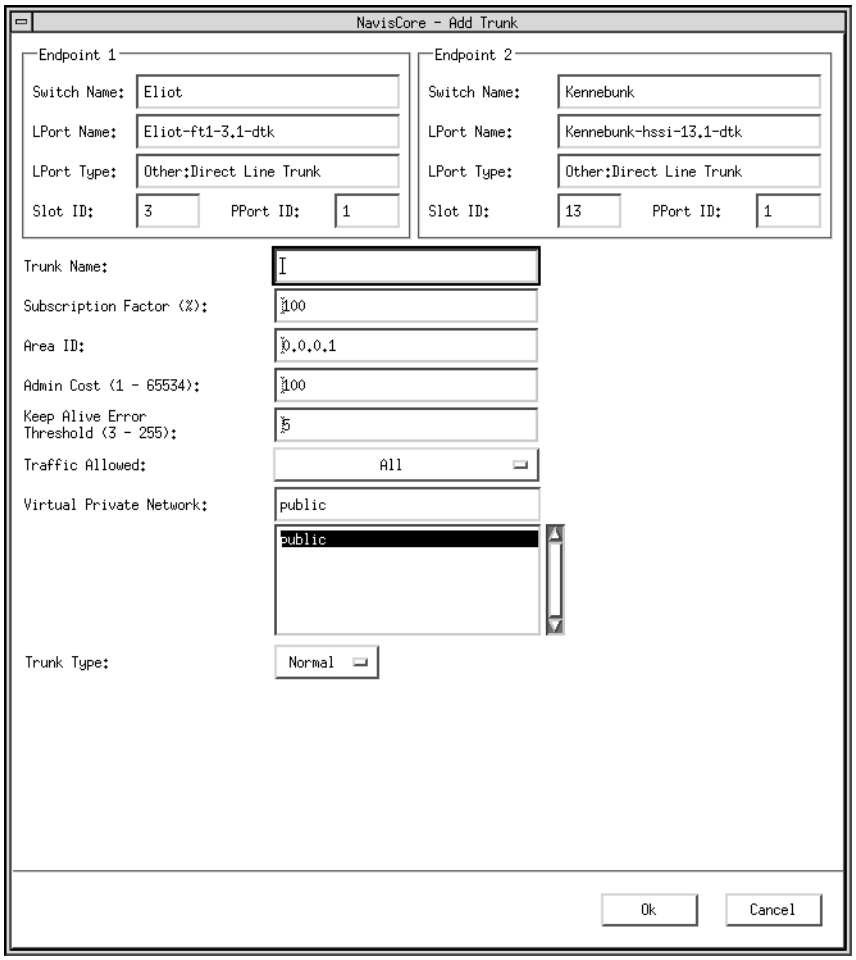

Figure 4-4. Add Trunk Dialog Box

5. Complete the Add Trunk fields described in Table 4-3.

Table 4-3. Add Trunk Fields

Field	Action/Description
Trunk Name	Enter a unique alphanumeric name to identify the trunk. You use this same name when you create the trunk connection (page 4-19).
Subscription Factor (%)	The trunk oversubscription factor percentage enables you to optimize the aggregate CIR you can configure on the trunk, by allowing you to over subscribe the trunk. The oversubscription factor represents the V value for this trunk. The bandwidth on a trunk is reserved at runtime, based on the configured CIR value of the PVCs that traverse that trunk. For example, you can set this factor to 200% to produce a virtual bandwidth that is two times greater than the defined bandwidth. For a detailed explanation of this parameter, see page 4-2. ***Note**: You can not over subscribe an ATM Direct Trunk.*
Area ID	Enter the area ID (x.x.x.x) for the area in which you want to locate this OSPF interface. The range of available values is from 0.0.0.0 to 255.255.255.255. Area 0.0.0.0 is the network backbone area. Area 0.0.0.1 is Area 1. Areas are collections of networks, hosts, and routers used for IP routing. The area ID identifies the area. If a trunk is in Area 1 and the OSPF Backwards Compatibility option (which is set through IP services) is set to Yes, external routes are not advertised across that link. ***Note:** Area 1 is reserved for Ascend switches. For a detailed description of OSPF areas, and how to use IP Services to configure multiple OSPF areas, see the* NavisCore IP Navigator Configuration Guide.

Table 4-3. Add Trunk Fields (Continued)

Field	Action/Description
Admin Cost (1-65534)	Assign an admin cost value of 1 to 65534. The lower the admin cost of the path, the more likely OSPF will select it for circuit traffic. *The default admin cost value is 100.* For a detailed explanation of this parameter, see page 4-3. ***Note**: When you increase or decrease the administrative cost of a trunk, the reroute tuning parameters control the rate at which the switch adds or removes circuits from the trunk. See the* NavisCore NMS Getting Started Guide *for information about reroute tuning. You cannot use trunk admin cost to force a trunk down.*
Keep Alive Error Threshold (%)	Configure the keep-alive threshold for a value between 3 and 255 seconds. *The default is 5 seconds.* For a detailed explanation of this parameter, see page 4-4. ***Note**: If you are running different switch code versions in your network (for example, Version 4.1 and Version 4.2), you must accept the default value of 5 seconds.* ***Note**: Service is disrupted if you change this value after the trunk is online.*
Traffic Allowed	Specify one of the following options to designate the type of traffic allowed on this trunk: *All* – The trunk can carry network management traffic, user traffic, and OSPF address distribution. *Mgt Only* – The trunk can carry *only* network management traffic, such as SNMP communication between a switch and the NMS. *Mgt & User* – The trunk can carry network management traffic and user traffic. ***Note**: To calculate the most efficient route for network management traffic, OSPF uses Trunk Admin Cost. OSPF ignores trunk bandwidth when it selects the best path or a route for management traffic. Management traffic can use a negative bandwidth trunk.*
Virtual Private Network	Select a VPN name. The default is *Public.* For more information, see Chapter 8, "Configuring Virtual Private Networks."

Table 4-3. Add Trunk Fields (Continued)

Field	Action/Description
Trunk Type	Select one of the following: *Normal* — is a common trunk. *Primary* — indicates that the trunk has a backup for fault tolerance. *Backup* — indicates that it is a backup trunk (when failure occurs on the primary trunk). See "Using the Automatic Trunk Backup Feature" on page 4-18 to configure a backup trunk. ***Note**: This parameter is not supported on trunks between CBX and B-STDX switches.* If you are configuring an ATM Direct Trunk, set this parameter to Normal.

6. (*Optional*) If you selected *Primary* as the Trunk Type, the system displays the fields shown in Figure 4-5.

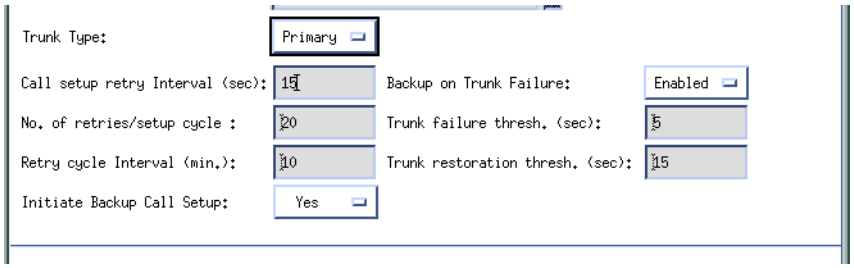

Figure 4-5. Add Trunk Dialog Box (Primary Trunk)

7. Complete the fields described in Table 4-4, or accept the default parameters.

Table 4-4. Add Primary Trunk Fields

Field	Action/Description
Call setup retry Interval (sec)	Specify the number of seconds between initiating a call. The default is 15 seconds.
No. of retries/setup cycle	Specify the number of retries per interval. The default is 20 retries.
Retry cycle Interval (min)	Specify a retry interval in minutes. The default is 10 minutes.
Initiate Backup Call Setup	Choose Yes (default) to initiate a backup call.
Backup on Trunk Failure	Enable (default) or disable trunk backup. If you enable trunk backup, the system automatically uses the backup trunk if the primary trunk fails. If you choose Disabled, the automatic trunk backup option is not used.
Trunk Failure thresh. (sec)	Specify the number of seconds (the default is 5). If you enabled trunk backup, this field specifies the number of seconds the system will wait before switching over to the backup trunk.
Trunk Restoration thresh. (sec)	Specify the number of seconds that the system will wait for the primary trunk to become functional before resuming use of the primary trunk. The default is 15 seconds. If the primary trunk is out of service and the backup trunk is in use, the system will not resume use of the primary trunk until it has been restored for the period of time you specify. The purpose of this field is to prevent a switch-over to a primary trunk that has only been temporarily restored.

8. When you complete the add trunk dialog box fields, choose OK.

9. Choose Close to return to the network map.

The next step is to create a trunk-line connection. Proceed to "Creating a Trunk-Line Connection" on page 4-19.

Using the Automatic Trunk Backup Feature

To use the automatic trunk backup function:

1. Access the Add trunk dialog box (Figure 4-4 on page 4-13).
2. Define a trunk that has a Trunk Type of *Primary.*
3. Specify all of the primary trunk field values shown in Table 4-3 on page 4-14 . Specify a value of *Yes* in the Initiate Backup Call Setup field on the Add Trunk dialog box.
4. Specify a value of *Enabled* in the Backup on Trunk Failure field on the Add Trunk dialog box.
5. Define from one to eight trunks that have a Trunk Type of *Backup*.
6. For each trunk with a Trunk Type of Backup, in the *Primary Trunk of the backup* field, select the name of the primary trunk specified in Step 2 .

Process for Switching Over to a Backup Trunk

In the event of trunk failure, the system uses the following process to automatically switch over to a defined backup trunk if you have used the steps in the previous procedure to enable Automatic Trunk Backup.

1. The system switches over to the backup trunk after the trunk is out of service for the amount of time specified for the primary trunk in the Trunk Failure Threshold field (Table 4-4).
2. The system resumes use of the primary trunk after it is in service for the period of time specified in the Trunk Restoration Threshold field (Table 4-4).

Defining the Manual Trunk Backup Feature

You can override the values for automatic trunk backup by using the manual trunk backup feature. To do this, use the Start Trunk Backup and Stop Trunk Backup options on the Modify Trunk dialog box.

Creating a Trunk-Line Connection

You must define the trunk configuration between two switches (page 4-11) before you create the trunk-line connection on the network map. The Add Connection function enables you to draw a line to connect the two switches on the network map.

To add a trunk line connection:

1. From the Edit menu, select Add Connection. The Add Connection dialog box appears as shown in Figure 4-6.

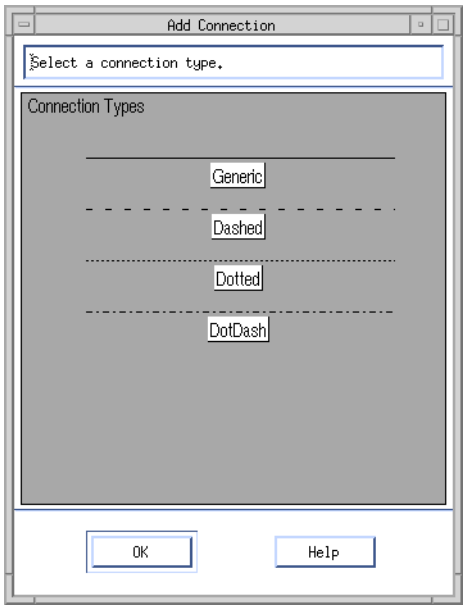

Figure 4-6. Add Connection Dialog Box

2. Select a Connection Type from the palette.
3. To create a trunk-line connection between the two Ascend switches on the network map, click on the first switch object (source symbol) and then the second switch object (destination symbol).
4. The Add Object dialog box appears as shown in Figure 4-7.

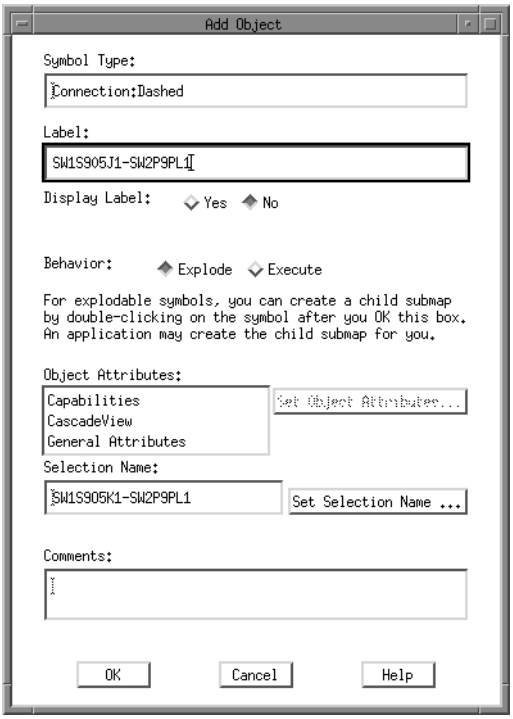

Figure 4-7. Add Object Dialog Box

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

Table 4-5. Add Object Fields

Field	Action/Description
Symbol Type	Displays the type of connection you are adding to the map.
Label	Enter the trunk name you specified on the Add Trunk dialog box (page 4-13).
Display Label	Select *Yes* to have the label name appear beneath the trunk line object on the network map. Select *No* if you do not want the label name displayed.
Behavior	Select *Explode* to create the basic NavisCore network configuration. See the *HP OpenView User's Guide* for more information about the Execute function.
Object Attributes	Select *NavisCore*. Then choose *Set Object Attributes*. The Add Object – Set Attributes dialog box appears as shown in Figure 4-8.

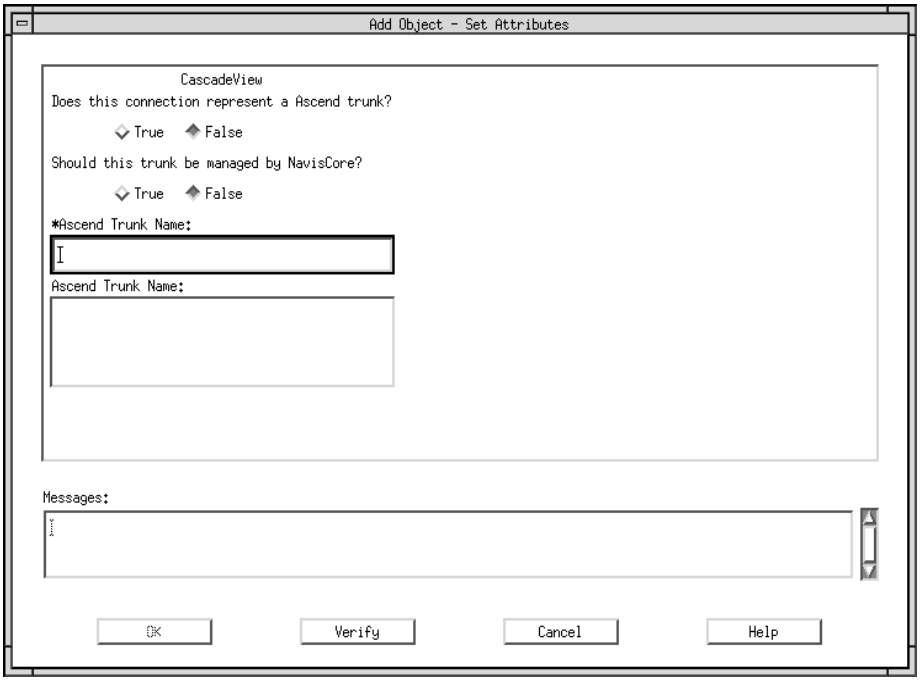

Figure 4-8. Add Object - Set Attributes Dialog Box

6. Complete the required dialog box fields described in Table 4-6.

Table 4-6. Add Object - Set Attributes Fields

Field	Action/Description
Does this connection represent an Ascend Trunk?	Select True.
Should this trunk be managed by NavisCore?	Select True.
Ascend Trunk Name	Enter the name you assigned to the trunk. This should be the same name you entered for the label in the Add Object dialog box on page 4-20.

7. Choose Verify to confirm your selections.

8. Choose OK to return to the Add Object dialog box.

9. Choose OK to return to the network map. The trunk line appears between the two switches on the network map.

Displaying Multiple Trunks Between Switches

If you configure more than one trunk between two switches, these trunks appear as a solid line between the switches.

To display all trunk connections between two switches:

1. Double-click the left mouse button on the solid line between the switches. A trunk submap window appears, as shown in Figure 4-9.

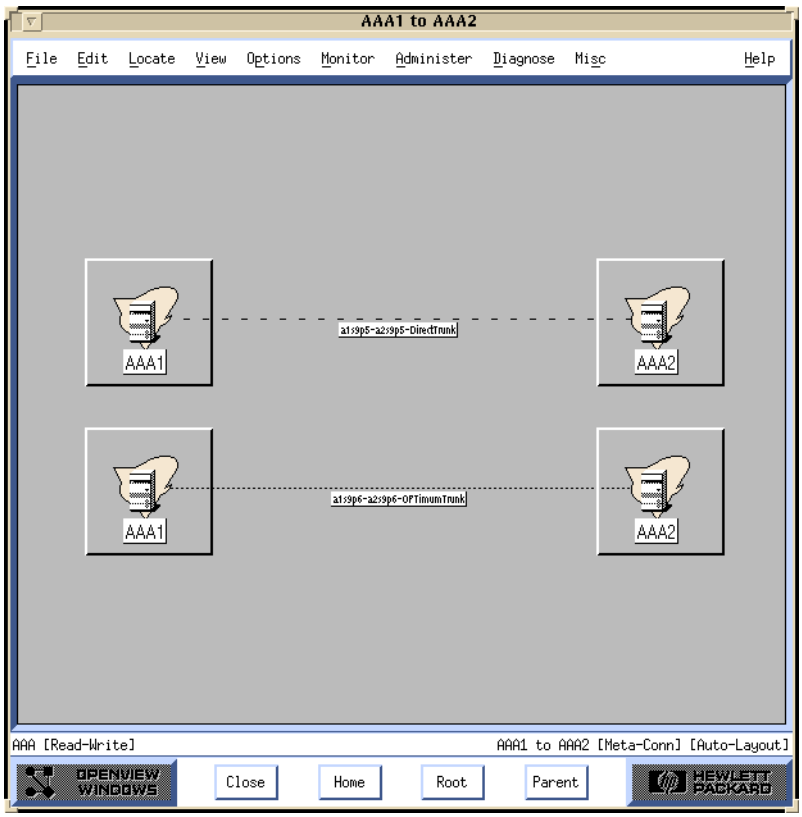

Figure 4-9. Displaying Multiple Trunks-Trunk Submap Dialog Box

2. Choose Close to return to the network map.

Trunk Coloring

All associated trunks are polled for status according to the trunk poll timer. The trunk lines on the network map change color. These colors indicate trunk status according to the polled status and the traps received by the Ascend Event Log. Table 4-7 documents the color scheme that identifies the status of a trunk connection on the network map.

Table 4-7. Trunk Color Status Indicators

Color	Status
Black	Either the line connection has not been defined as a trunk or the environment variable $XUSERFILESEARCHPATH does not point to **/opt/CascadeView/app-defaults**. [1]
Red	Trunk is down.
Blue	Trunk status is Unknown or Unmanaged.
Yellow	More than half the trunk connections are down.
Green	Trunk connection is Up.
Orange	Only one trunk connection, out of many connections, is Up.
Cyan	More than half the trunk connections are Up.

[1] If the Trunk graphic is black, set the following environment variable in .profile:
$ XUSERFILESEARCHPATH =/opt/CascadeView/app-defaults/%N
$ export XUSERFILESEARCHPATH
For more information about operational states and status, select Display Legend from the Help menu.

If you define more than one trunk connection between the same two switches, HP OpenView combines the status to display an orange, yellow, or cyan trunk line. To display a view of the individual connections, double-click on the trunk line and see Table 4-7 to interpret trunk-color status.

When you finish defining your trunk configuration(s) and trunk line connection(s), proceed to Chapter 6, “Configuring PVCs,” to complete the network configuration.

5

About Permanent Virtual Circuits (PVCs)

A Permanent Virtual Circuit (PVC) defines an end-to-end connection between two logical ports within the Ascend network. You can configure PVCs after you configure the switches, physical ports, logical ports, and trunks. The Set All Circuits function enables you to add, modify, or delete circuit configurations.

Circuit Routing Priority

A PVC routing priority enables you to specify the *bandwidth priority* and *bumping priority*, or level of importance, of each circuit in the network. The lower the number, the higher the priority. For example, on the Set All PVCs on Map dialog box, you specify these values as follows:

Bandwidth priority — A value from 0-3, where 0 is the highest priority (default).

Bumping priority — A value from 0-7, where 0 is the highest priority (default).

If you do not override the defaults, all circuits are defined at the highest priority (0, 0), which means all circuits in the network have the same routing priority. However, if you prioritize circuits in your network, the switch assigns circuits with the highest priority to the lowest-cost paths through the network. These high-priority circuits are guaranteed full bandwidth wherever possible. Circuit prioritization occurs at the cost of the lower-priority circuits.

Priority Routing and Path Cost

By assigning specific bumping and bandwidth priorities to circuits, you can guarantee that the needs of high-priority PVCs are met first. In addition, you can also accommodate PVCs where the path cost is not important. By assigning a routing priority to specific circuits, you can guarantee that when a link fails or congestion conditions occur, the higher priority PVCs are given preference in the network over PVCs with a lower priority.

Priority Routing Example

There are two paths (Path 1 and Path 2) between a pair of nodes (A and B). The cost of Path 1 is 100, while the cost of Path 2 is 200. The multiple circuits that are defined with a priority routing of 2,0 within the network use all of the bandwidth on the Path 1 link. Without priority routing, additional virtual circuits are forced to use Path 2, which could involve higher delays and more hops.

With priority routing, you can define additional circuits between A and B with a priority of 0,0. The switch running the priority-routing software can detect that Path 1 is entirely populated by the circuits with the 2,0 priority. The switch then forces enough 2,0 priority PVCs from Path 1 to ensure that every trunk in Path 1 has enough bandwidth to satisfy the Quality of Service (QoS) of the highest-priority (0,0) VCs. As a result, some 2,0 priority PVCs are forced to Path 2.

Special Condition

PVCs that have a low bandwidth priority are always at risk in a mixed priority network. For this reason, circuits with a low bandwidth priority can be left without a valid route. The occurrence of this condition depends on the network topology and the amount of available link bandwidth. You should consider this special condition when configuring your network.

Routing Priority Rules

The switch uses the following rules when implementing routing priority for circuit provisioning, trunk-failure recovery, and load-balance rerouting of circuits.

Circuit Provisioning

- When provisioned, a higher-priority PVC selects a path when the circuit is provisioned that is cost effective and that satisfies QoS.
- All PVCs with a lower priority are ignored.
- As the connection is established, higher-priority PVCs force lower-priority PVCs from their selected path until there is enough available link bandwidth to accommodate the higher-priority PVC QoS.
- NavisCore selects lower-priority PVCs that are forced from their path in the following order:
  - Bumping priority, where lowest bumping priority VCs are selected first.
  - The order of the Equivalent Bandwidth (EBW or CIR Frame Relay) within a group of VCs of the same bumping priority, where VCs with a higher EBW are selected first.
  - In the order of the channel identifier within a group of VCs that have the same bumping priority and the same EBW.

Trunk-Failure Recovery

PVCs always attempt to reroute themselves when a trunk goes down. The Ascend switch software allows a trunk to reach negative bandwidth for PVCs recovering from trunk failure if there is no other available path with positive bandwidth. Priority routing modifies these rules as follows:

- A higher bandwidth priority PVC selects an optimal path in response to trunk failure without taking into account the bandwidth consumed by lower-bandwidth priority VCs. The lower-priority VCs may be forced to use paths that are not optimal (as defined in the provisioning rules).
- Lower-bandwidth priority PVCs are not allowed to cross trunks where there is at least one higher-priority VC and the bandwidth is negative, with the exception of PVCs configured with 0 bumping priority. Bumping priority 0 PVCs are allowed to push a trunk to negative bandwidth and rely on reroute balancing to correct the negative bandwidth at a future time.
- Higher-priority PVCs may push a trunk to negative bandwidth if there are no more lower-priority PVCs to force off the trunk. In this case, all of the lower-priority PVCs (excluding 0 bumping priority PVCs) are forced off the trunk. PVCs configured with 0 bumping priority are given special permission to share the negative bandwidth trunk with higher-priority PVCs until the reroute balancing corrects this condition at a future time.

Balance Rerouting

Balance rerouting is a switch function that periodically tests the efficiency of each PVC route. A PVC that was rerouted due to trunk failure may not be on the most optimal path at any given time or may be traversing a negative bandwidth trunk. Balance rerouting corrects these conditions by rerouting the PVC to a new path.

Priority routing modifies the switch balance rerouting functions so that a PVC with a higher-bandwidth priority is given an optimal path, and the bandwidth used by the lower-priority PVCs is not considered by the switch. For this reason, the PVCs with the lower priority may be forced onto a path that is not optimal. See “Circuit Provisioning” on page 5-3 for details about path selection.

Interoperability with Previous Releases

To use circuit-routing priority in your network, the following interoperability restrictions apply:

- All switch software must be at least Release 4.1 or higher for B-STDX switches and Release 2.4 on STDX switches.
- On a trunk, if either end resides on a 4.1 B-STDX switch or a 2.4 STDX switch, the trunk treats all PVCs equally (assumes all have a 0,0 priority).
- On a circuit, if either end belongs to a 4.1 B-STDX switch or a 2.4 STDX switch, the circuit is automatically assigned a 0,0 priority. The NMS does not support any routing priority other than 0,0 on switches running Release 4.1 or lower.

Rate Enforcement

Rate enforcement prevents network congestion and allocates network resources to ensure the commitment of service contracts. Rate enforcement measures the actual traffic flow across a connection and compares it to the configured traffic flow parameters for that connection. Traffic outside the acceptable committed information rate (CIR) is tagged and discarded if congestion develops.

Rate enforcement is implemented on a per-DLCI basis on all circuits on ingress switches. When the switch receives data over time interval *Tc* (Tc=Bc/CIR), it classifies the frame as follows:

- Under the committed burst size (Bc)
- Over the committed burst size but under the excess burst size (Be)
- Over the excess burst rate

Color designators (green, amber, and red) identify packets travelling through the network. Congested nodes use the designators to determine which frames to discard first in various congested states or congestion conditions. Table 5-1 describes the designators (traffic colors) and discard policy.

Table 5-1. Rate Enforcement and Discard Policy

Traffic Color	Description	Discard Eligible (De)
Green	Accumulated number of bits received up to any time during the current time interval, excluding the current frame, less than Bc.	No
Amber	Accumulated number of bits received up to any time during the current time interval, excluding the current frame, greater than Bc but less than Be.	Frame is eligible for discard if it passes through a congested node.
Red	Accumulated number of bits received up to any time during the current time interval, excluding the current frame, greater than Be.	All red frames are discarded.

Graceful Discard

The *graceful discard* feature enables you to control network behavior and user traffic. You can set the graceful discard parameters as follows:

On — The switch allows some red frames to be transmitted. This maximizes network usage, but may overload the network.

Off — This option avoids potential congestion. This allows strict control of user traffic, but may waste network resources.

When graceful discard is set to On, you can configure the red-frame percent. The red-frame percent is used to limit the number of red frames the network is responsible to deliver. The red-frame percent (pr) is determined as follows.

$$Pr = \frac{\text{Allowed red frame bits}}{\text{Bc + Be + allowed red frame bits}}$$

For more information about the rate-enforcement discard process, see the *Networking Services Technology Overview Guide*.

Rate Enforcement Schemes

You can configure the rate enforcement scheme. This option provides additional flexibility, increased rate enforcement accuracy, and improved switch performance. You set the rate enforcement scheme in the Add PVC dialog box under the Traffic Type attributes (Table 6-5 on page 6-14).

Table 5-2 compares the accuracy and switch performance of the Jump and Simple rate enforcement schemes. Number 1 specifies the more accurate scheme and better switch performance, while 2 specifies a less-accurate scheme and slightly degraded switch performance.

Table 5-2. Rate Enforcement Schemes

Scheme	Rate Enforcement Accuracy	Switch Performance
Jump	1	2
Simple	2	1

About DLCI Numbers

A data link connection identifier (DLCI) number is a 10-bit address that identifies PVCs. DLCIs identify the logical endpoints of a virtual circuit and have local significance only.

Depending on your link management type, use the guidelines in Table 5-3 to define DLCI numbers.

Table 5-3. DLCI Number Guidelines

DLCI Number Range	Description
0-15	Reserved
16-991	Available for all link management types
992-1007	Available for LMI Rev1 only
1008-1023	Reserved

Administrative Tasks

This section describes how to:

- Move circuit endpoints
- Use templates to define a new circuit
- Delete circuits

Moving Circuits

The Move Circuit function enables you to move circuit endpoints defined for one logical port (the source) to another logical port (the destination). If you are upgrading a switch and do not want to lose PVC connections, you can use this function to move circuits to another switch.

This function has the following restrictions:

- You cannot move circuits you previously defined as part of a fault-tolerant PVC configuration (defined with a service name or designated as a backup).
- You cannot move a circuit that is currently in use.
- You cannot move a circuit if you receive an error that indicates there is a problem acquiring a lock for the circuit and all associated logical ports.
- You cannot move a circuit that has a manually defined circuit path.
- You cannot define more than one circuit for a Frame Relay Assembler/Disassembler (FRAD) logical port.
- The DLCI must be unique to the destination logical port.
- You cannot move a circuit if the source logical port type is not a valid type for the destination port. For example, you cannot move a Frame Relay or SMDS logical port type to an ATM DS3 module.
- The move circuit function will fail if moving a circuit exceeds the maximum number of circuits allowed for the destination logical port.
- You cannot move a circuit that is a member of a multicast DLCI configuration.

To move a circuit:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Circuits ⇒ Move Circuit Endpoint. The Select Source & Destination LPorts dialog box appears as shown in Figure 5-1.

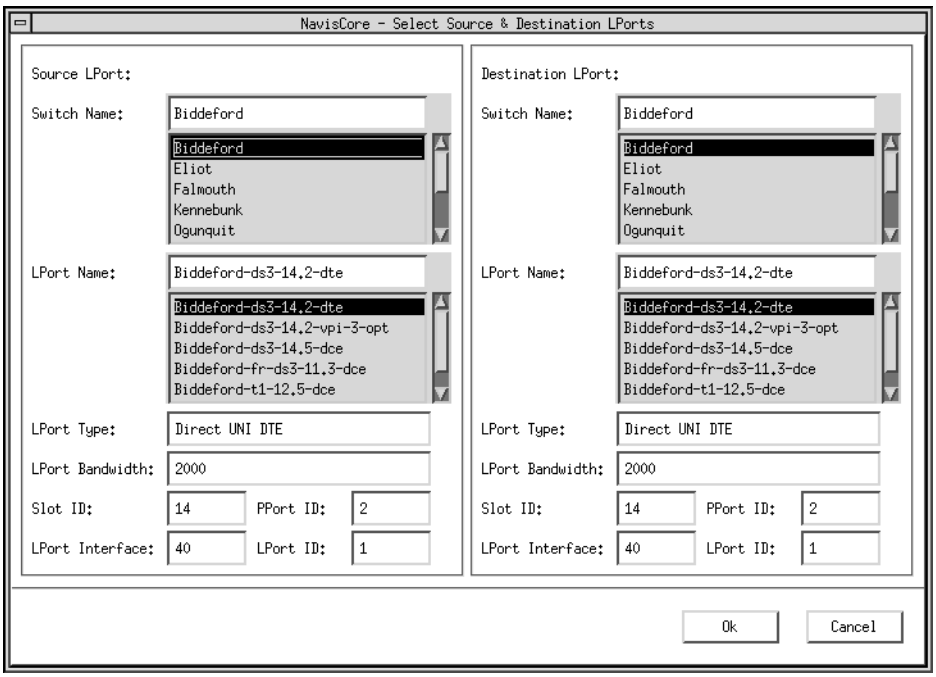

Figure 5-1. Select Source and Destination LPorts Dialog Box

2. Complete the following steps for the source logical port.
   a. Select the switch (name) that contains the circuit you want to move. A list of logical ports defined for this switch appears in the LPort Name field.
   b. Select the logical port (name) on which the circuit is defined. The fields below this list box display information about this port.
3. Complete the following steps for the Destination LPort.
   a. Select the switch (name) to which you want to move the circuit. A list of logical ports defined for this switch appears in the LPort Name field.
   b. Select the logical port (name). The fields below this list box display information about this port.

4. Choose OK. The Move Circuit Endpoint dialog box appears as shown in Figure 5-2. This dialog box displays the circuits that have the source logical port as an endpoint.

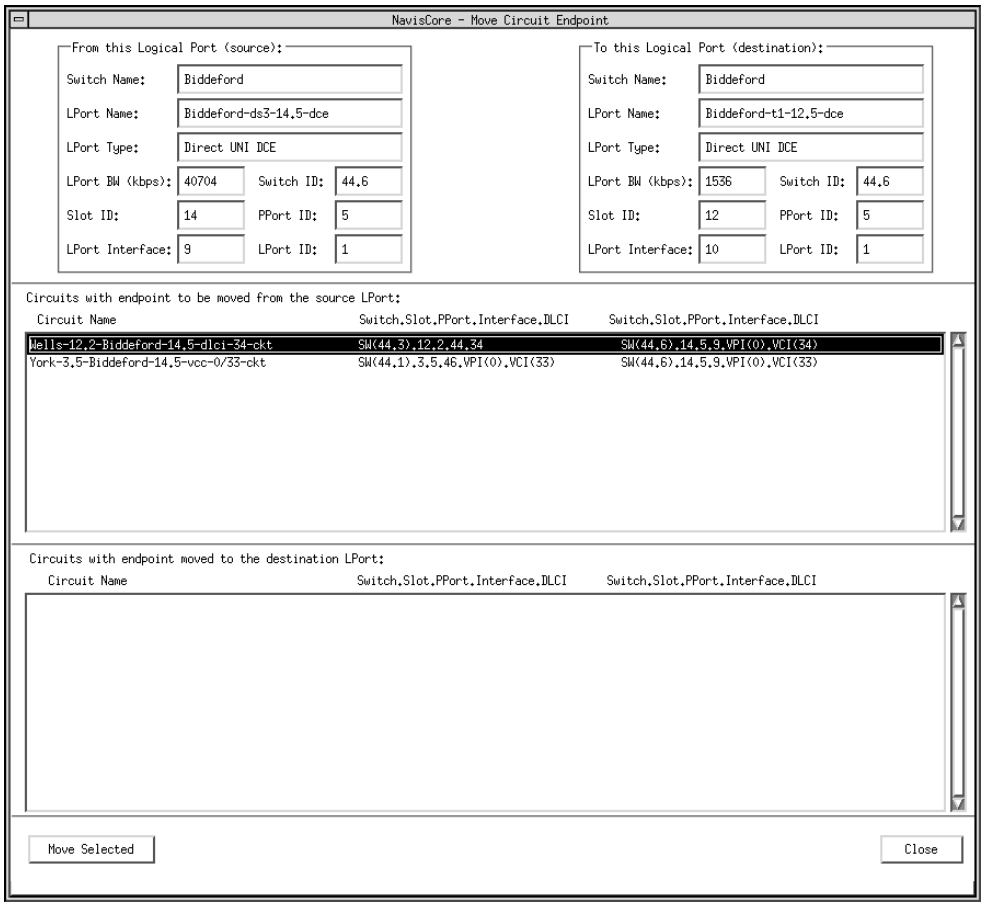

Figure 5-2. Move Circuit Endpoint Dialog Box

5. From the *Circuits with endpoint to be moved from the source LPort* list box, select the circuit you want to move.

6. Do one of the following:

   - Choose Move Selected. The selected circuit appears in the *Circuits with endpoint moved to the destination LPort* list box.
   - Choose Move All if you are moving circuits with endpoints on a 10-port DSX or channelized DS3 or DS3-1-0 I/O module. You do not have to highlight all circuits. The Move All command enables you to move all circuits at the same time. This command is only enabled if selected endpoints are on a 10-Port DSX or channelized DS3 or DS3-1-0 I/O module.

   This process takes a few minutes, depending on the number of circuits.

7. Repeat Step 4 and Step 6 for each circuit you want to move.
8. If some circuits were not moved in this process, check the restrictions on page 5-8.
9. When you finish, choose Close.

Using Templates

If you defined a circuit configuration and saved it as a template (see *Is Template* field in Table 6-4 on page 6-10), you can define a new circuit using the same parameters.

To define a circuit from a template:

1. Choose the Add Using Template command on the Set All PVCs on Map dialog box (Figure 6-1 on page 6-3).
2. Do one of the following:
   - Choose Last Template to use the last template you defined for this switch.
   - Choose Template List to display a list of templates defined for this map. Select a template and choose OK.

Deleting Circuits

To delete a circuit:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Circuits. The Set All Circuits On Map dialog box appears.
2. To view the list of circuits, select the Search by Name field and press Return. If necessary, select each circuit and review each logical-port endpoint.
3. Select the circuit you want to delete.
4. Choose Delete.

6

Configuring PVCs

This chapter describes how to configure the following types of Frame Relay Permanent Virtual Circuits (PVCs):

- Point-to-Point
- Multicast DLCI

See Chapter 7, “Configuring Management Paths,” for information about management DLCIs.

Reliable Scalable Circuit

The NavisCore Reliable Scalable Circuit feature (set to *On* by default) improves PVC configuration reliability. The NMS verifies that the card state for each PVC endpoint is up before sending the SNMP set command to the corresponding cards in each switch. If the card status of either endpoint is not up, the system displays an error message indicating where the failure occurred. The error message includes an abort option, which allows you to cancel the PVC configuration and prevent a card out-of-sync condition.

When enabled, the Reliable Scalable Circuit feature enables you to add, modify, or delete PVCs in the following scenarios:

- Both switches are unmanaged.
- Both switches are managed. Both cards (endpoints) have a status of *up*.
- One switch is unmanaged and one switch is managed. Both cards have a status of *up*.

For information about Reliable Scalable Circuit reported error types, see Appendix A.

Disabling the Reliable Scalable Circuit Feature

To disable this feature, edit the cascadeview.cfg file and remove the # signs from the following lines:

```
#CV_CARD_STATS=DISABLE
#EXPORT CV_CARD_STATS
```

Accessing Circuit Functions

The Set All PVCs on Map dialog box displays status information for the circuit you select from the Defined Circuit Name list. To access this dialog box, from the Administer menu, select Ascend Parameters ⇒ Set All Circuits ⇒ Point-to-Point. When the Set All PVCs on Map dialog box appears, press Return to display the list of defined circuit names, as shown in Figure 6-1.

Figure 6-1. Set All PVCs On Map Dialog Box

- To learn more about the Set All PVCs on Map dialog box fields and commands, continue with the following section.
- To begin defining a PVC, proceed to page 6-7.

About the Set All PVCs On Map Dialog Box

The Set All PVCs On Map dialog box displays information about the configured options for the selected circuit. To view a list of configured circuits, position the cursor in the Search by Name field and press Return. To use a wildcard search to find a specific circuit name, you can:

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

Table 6-1 describes the fields and commands on the Set All PVCs On Map dialog box.

Table 6-1. Set All PVCs On Map Status Indicators and Command Descriptions

Field/Command	Description
Defined Circuit Name	Displays a list of the circuits configured in the network. Use a wildcard search to find a specific circuit name. If applicable, this field also lists the configured circuit alias.
Logical Port	The dialog box displays the following logical port information for each circuit endpoint: *Switch Name* – Displays the name of the switch at each endpoint of the circuit. *LPort Name* – Displays the name of the logical port at each endpoint of the circuit. *LPort Type* – Displays the configured type of the selected logical port. *Slot ID* – Indicates the physical slot number where the I/O module containing the selected port is installed. *PPort ID* – Displays the number of the physical port for which the selected logical port is configured. *DLCI Number* – Displays the number that is used for the management DLCI. (See Chapter 7, "Configuring Management Paths," for information about management DLCIs.)
Fail Reason at endpoint 1 (2)	Displays the reason a selected circuit failed (if any) for a given endpoint. See the *NavisCore Diagnostic and Troubleshooting Guide* for a description of these fail reasons.
Defined Circuit Path	Displays the configured circuit path.

Table 6-1. Set All PVCs On Map Status Indicators and Command Descriptions (Continued)

Field/Command	Description
Actual Circuit Path	Displays the actual path that OSPF selected for this circuit to get to its destination.
Show Attributes (option menu)	Displays the appropriate attributes configured for the selected option. See one of the following sections for more information: "Administrative" attributes are defined later in this table. For more information, see Table 6-4 on page 6-10. "User Preference" attributes are defined on page 6-16. "Traffic Type" attributes are defined on page page 6-13.
Admin Status	Displays whether the selected circuit is online (Up) or offline (Down).
Oper Status	Displays the current operational status of the selected circuit. Messages include: *Active* – The circuit is operational through the network end-to-end. *Inactive* – The circuit is not operational. Check the Reason field for possible reasons. *Unknown* – The NMS cannot reach the higher-numbered node for status. (If the circuit is an intra-switch PVC, then the NMS cannot reach the highest-numbered LPort.) *Invalid* – The circuit definition is not found in the higher-numbered node. You may need to return to the Set Circuits dialog box and choose Apply to save the circuit definition. It may also be necessary to PRAM synch the host card.
VPN Name	Assigns the selected circuit to a specific VPN and customer name.
Customer Name	Displays the customer name for the selected PVC (if applicable).
Is Template	Displays Yes if you can use this circuit as a template to create other circuits using similar parameters.
Admin Cost Threshold	If you enable this option, the PVC will not be routed over a path whose total administrative cost exceeds the entered value. This means that if you enable this field and enter a value of 1000, the PVC will not be routed over a path whose total administrative cost exceeds 1000. The total administrative cost for a path is calculated by summing the administrative cost for each trunk in the path. The valid range of values for this field is 1 - 65534. (Do not enable this option if you use End-End Delay routing.)

Table 6-1. Set All PVCs On Map Status Indicators and Command Descriptions (Continued)

Field/Command	Description
Is Mgmt Dlci Loopback Ckt	Choose Yes to include this PVC configuration in the NMS initialization script file. This file contains all the SNMP set requests necessary to replicate the entire switch configuration. Once you download this file to the switch, this PVC can be used to establish NMS-to-switch connectivity. This option is especially useful in some management DLCI configurations. The default value is No. For more information about management DLCIs, see Chapter 7.
End-End Delay Thresh. (CTD µsec)	Displays the maximum acceptable delay for traffic using this PVC to traverse the network. For more information, see Table 6-4 on page 6-10.
Add/Modify/Delete	Enables you to add a new circuit or Modify or Delete an existing circuit. ***Note**: If "NONE" is not displayed in the PVC loopback status field, do not attempt to modify or delete the selected circuit. See the* NavisCore Diagnostic and Troubleshooting Guide *for more information about loopback testing.*
VPN/Customer	Assigns the selected circuit to a specific VPN and customer name.
Get Oper Info	Displays a status message in the *Oper Status* field for the selected circuit. For more information, see the *NavisCore Diagnostic and Troubleshooting Guide*.
Define Path	Manually defines a circuit path (page 6-18).
Statistics	Displays the summary statistics for the selected circuit. For more information about summary statistics, see the *NavisCore Diagnostic and Troubleshooting Guide*.
QoS	Displays the Quality of Service values for the selected circuit. For more information about QoS, see the *NavisCore Diagnostic and Troubleshooting Guide*. For a description of the QoS classes, see "Setting QoS Parameters" on page 3-22.
Add Using Template	If you have already defined a circuit configuration and saved it as a template, use this option to define a new circuit. • Choose Last Template to use the last template you defined for this switch. • Choose Template List to display a list of templates previously defined for this map.
Accounting	Accesses the NavisXtend accounting server functions for a PVC. For information about the accounting server functions, see the *NavisXtend Accounting Server Administrator's Guide*.
Close	Exits this dialog box and returns to the network map.

Adding a Circuit Connection

1. See "Accessing Circuit Functions" on page 6-3 for instructions on accessing the Set All PVCs on Map dialog box.
2. From the Set All PVCs On Map dialog box, choose Add. The Select End Logical Ports dialog box (Figure 6-2) appears.

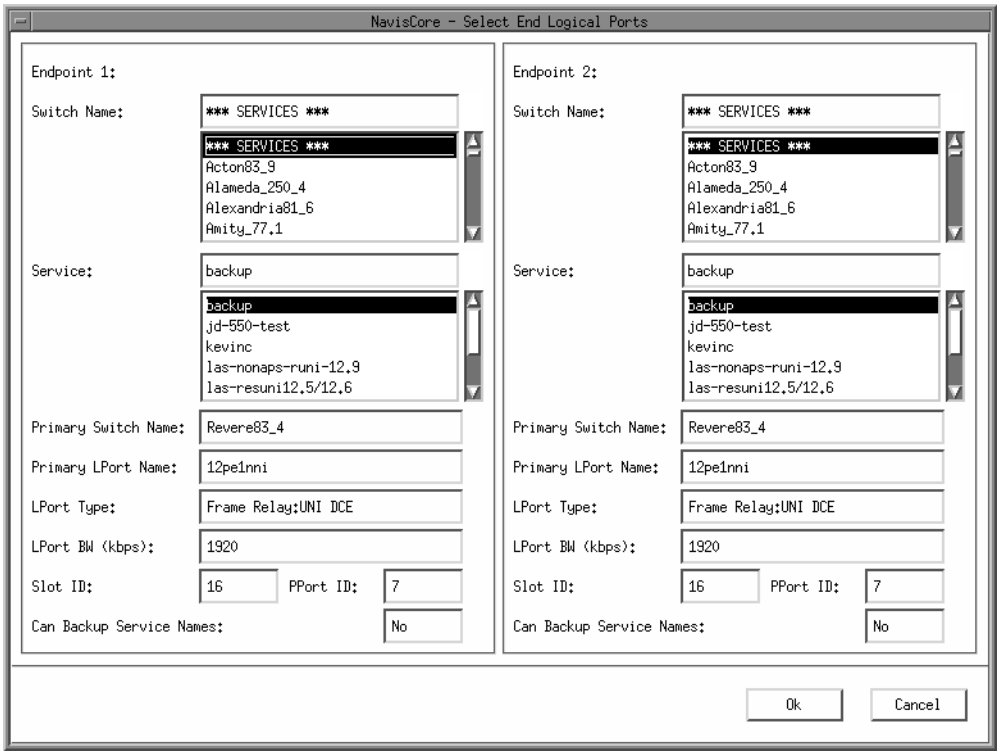

Figure 6-2. Select End Logical Ports Dialog Box

3. Configure Endpoint 1 and Endpoint 2 as follows:

For a fault-tolerant PVC Configuration

For more information about fault tolerant PVCs, see Chapter 9.

a. Choose *** SERVICES *** to configure a fault tolerant PVC.

b. Select a Service name from the list. You can configure a fault-tolerant PVC only for the following Frame Relay logical port types:

- UNI DCE
- UNI DTE
- UNI NNI

c. Continue with Step 4.

For a Standard Circuit Configuration

a. Select a switch name from the list.

b. Select an LPort name from the list of logical ports. Table 6-2 lists the standard logical port configurations. Configure Endpoint 1 and Endpoint 2 as shown in Table 6-2. Note that if you enable the VPN/Customer View function (page 8-8), only logical ports that belong to the VPN or customer you select appear in the list.

c. Continue with Step 4.

Table 6-2. Logical Port Endpoints for Circuits

Endpoint 1	Endpoint 2
FR UNI DCE/DTE, FR NNI	FR UNI DCE/DTE, FR NNI
FR UNI DCE/DTE, FR NNI	Encapsulated FRAD, PPP
Encapsulated FRAD	Encapsulated FRAD

4. The Select End Logical Ports dialog box displays information for both Endpoint 1 and Endpoint 2. Table 6-3 describes each field.

Table 6-3. 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 circuit configuration.
Slot ID	Displays the I/O slot (number) in which the module resides.
PPort ID	Displays the port number for the physical port.

5. Choose OK. The Add PVC dialog box (Figure 6-3) appears displaying the current parameters.

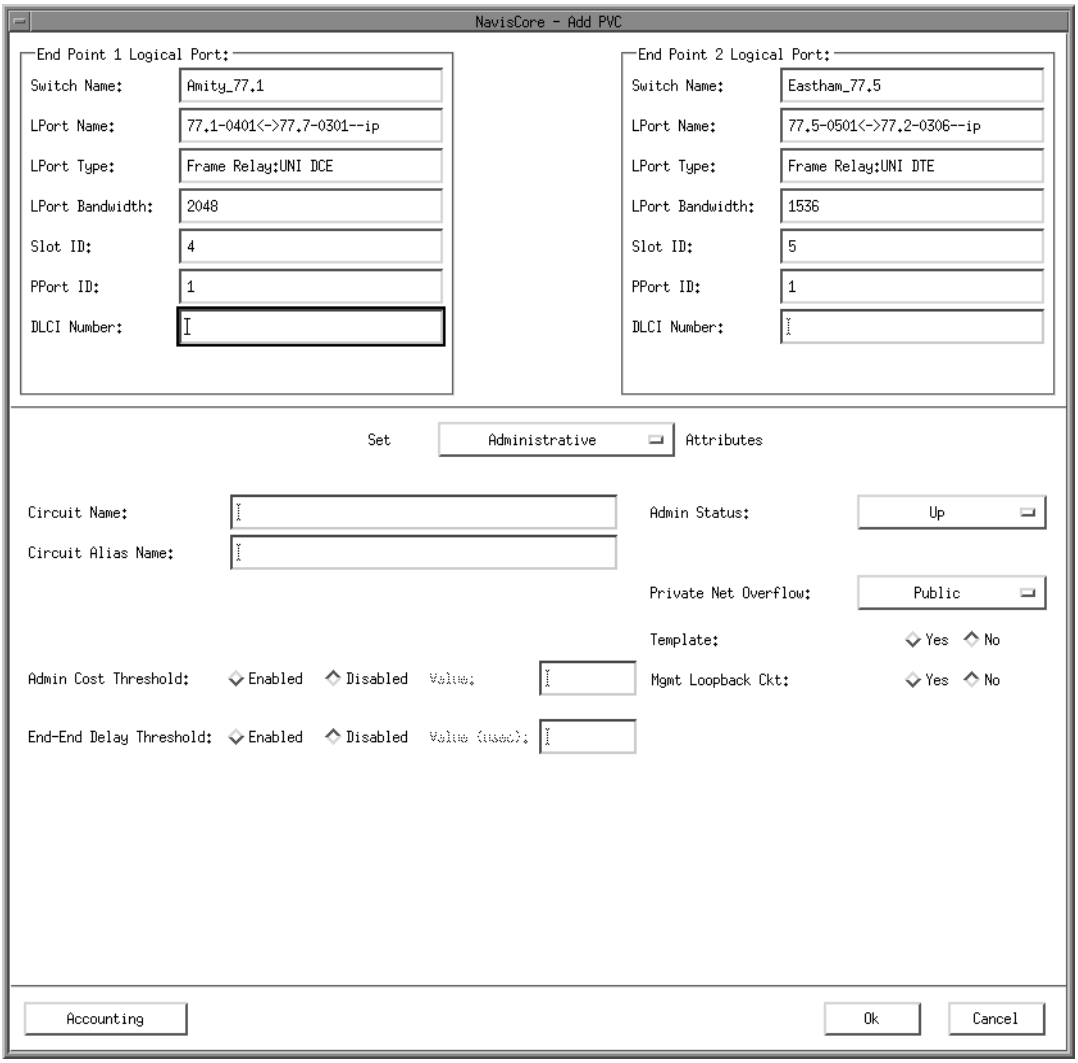

Figure 6-3. Add PVC Dialog Box (FR: UNI DCE)

Continue with the next section, "Defining Frame Relay Circuits," to configure these parameters.

Defining Frame Relay Circuits

When you configure a PVC, the Add PVC dialog box (Figure 6-3 on page 6-9) provides detailed parameters that you need to define for each endpoint. To define these parameters, you use the Set Attributes menu.

Setting Circuit Attributes

Use the Set Attributes menu on the Add PVC dialog box to configure the following information:

Administrative — Defines administrative information, such as circuit name, administrative status, and circuit type.

Traffic Type — Defines the traffic descriptor settings for forward and reverse traffic.

User Preference — Defines PVC features that deal with port congestion and traffic policing.

Continue with the following sections to define these parameters.

Administrative Attributes

From the Add PVC dialog box (Figure 6-3 on page 6-9), select Set [Administrative] Attributes and complete the fields described in Table 6-4.

Table 6-4. Set Administrative Attributes Fields

Field	Action/Description
DLCI Number	Enter a unique DLCI for this logical port. For more information, see "About DLCI Numbers" on page 5-7.
Circuit Name	Enter any unique, continuous, alphanumeric name to identify the circuit. Do not use parentheses and asterisks. This name must be unique to the entire map.
Circuit Alias Name	(*Optional*) The circuit alias is used by service providers to identify the circuit in a way that is meaningful to their customers. This option is often used in conjunction with the NavisXtend Report Generator. Enter any unique alphanumeric name to identify the circuit. Do not use parentheses and asterisks. This name must be unique to the entire map.

Table 6-4. Set Administrative Attributes Fields (Continued)

Field	Action/Description
Admin Status	Select either Up or Down to define whether the circuit is to be activated. *Up (default)* – Activates the circuit. *Down* – Takes the circuit off-line to run diagnostics such as PVC loopback.
Private Net Overflow	Determines how PVC traffic is managed during trunk overflow or failure conditions. This feature is used with Virtual Private Networks (VPNs). For more information about VPNs, see Chapter 8. Select one of the following options: *Public* – (*Default)* PVCs are routed over dedicated VPN trunks. However, in the event of failure, the customer's traffic is allowed to run over common trunks (shared by a variety of different customers). *Restricted* – PVCs can only use dedicated VPN trunks. A customer using this mode must purchase redundancy trunks to be used in the event of outages or other trunk failures.
Template	(*Optional*) Save these settings as a template to use again to configure another PVC with similar options. To create a template, choose Yes in the *Template* field. See "Using Templates" on page 5-11 for more information.
Mgmt Dlci Loopback Ckt	Choose Yes to include this PVC configuration in the NMS initialization script file. This file contains all the SNMP set requests necessary to replicate the entire switch configuration. Once you download this file to the switch, this PVC can be used to establish NMS-to-switch connectivity. This option is especially useful in some management DLCI configurations. The default value is No. (For more information about management DLCIs, see Chapter 7.)

Table 6-4. Set Administrative Attributes Fields (Continued)

Field	Action/Description
Admin Cost Threshold	If you enable this option, the PVC will not be routed over a path whose total administrative cost exceeds the entered value. This means that if you enable this field and enter a value of 1000, the PVC will not be routed over a path whose total administrative cost exceeds 1000. The total administrative cost for a path is calculated by summing the administrative cost for each trunk in the path. The valid range of values for this field is 1 - 65534. (Do not enable this option if you use End-End Delay routing.)
End-End Delay Threshold	If you enable this option, the PVC will not be routed over a path whose total end-to-end delay exceeds the entered value. This means that if you enable this field and enter a value of 500 µsec., the PVC will not be routed over a path whose total end-to-end delay exceeds 500 µsec. The total end-to-end delay for a path is calculated by summing end-to-end delay for each trunk in the path. The valid range for this field is 0 - 167777214 µsec. ***Note:*** *The value you enter should reflect your network topology. If a PVC will typically traverse high speed trunks, set the delay rate lower; increase the delay if the PVC must use low-speed trunks.*

Traffic Type Attributes

From the Add PVC dialog box (Figure 6-3 on page 6-9), select Set [Traffic Type] Attributes to specify Traffic Descriptor settings for forward and reverse traffic. Complete the dialog box fields described in Table 6-5.

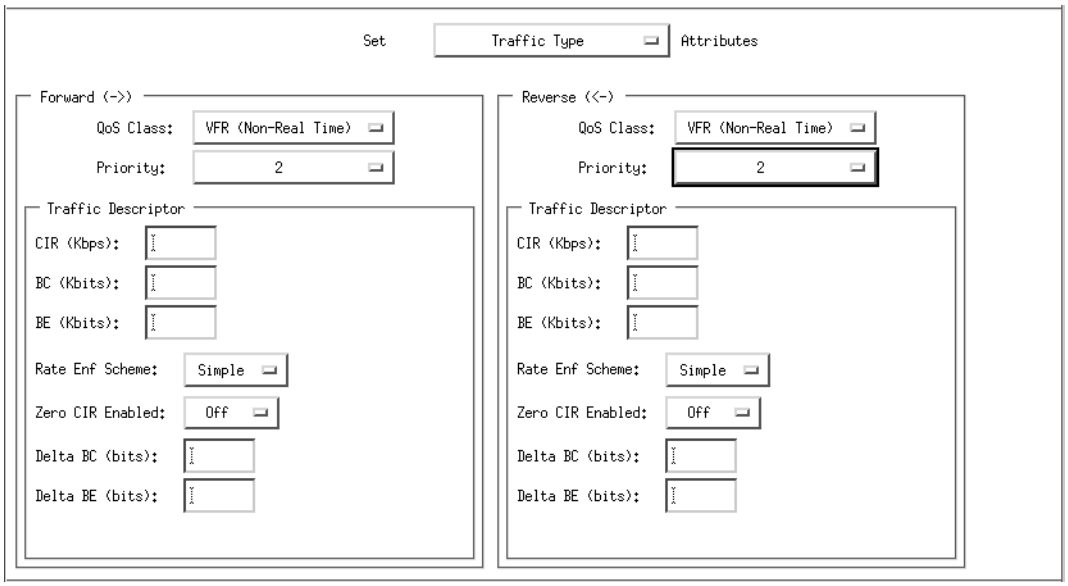

Figure 6-4. Set Traffic Type Attributes (FR: UNI DCE)

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 6-5. Set Traffic Type Attributes Fields

Field	Action/Description
QoS Class	Select one of the following Frame Relay Class of Service values: *VFR (Real-Time)* – Variable Frame Rate (VFR). Used for special delay-sensitive applications that require low delay variation between endpoints. *VFR (Non-Real Time)* – Handles transfer of long, bursty data streams over a pre-established connection. This service provides low data loss but no delay guarantee. Also used for short, bursty data, such as LAN traffic. CPE protocols adjust for any delay or loss incurred through the use of VFR-nrt. *UFR* – Unspecified Frame Rate (UFR). Used for LAN traffic, primarily. The CPE should compensate for any delay or frame loss.
Priority	Select 1, 2, or 3 to configure the priority of data being transmitted on this circuit. Circuit priority determines the data's forward priority. The highest priority is 1 (do not discard data); the lowest priority is 3 (discard data). The default priority for Frame Relay is 2. ***Note:*** *To configure the priority of transmitted data for a management PVC (MPVC) select 1, 2, or 3. The default priority is 1. (See Chapter 7 for more information about MPVCs.)*
CIR (Kbps)	Enter the Committed Information Rate (CIR) in Kbps at which the network transfers data under normal conditions. Normal conditions refer to a properly designed network with ample bandwidth and switch capacity. The rate is averaged over a minimum increment of the Committed Rate Measurement interval (Tc). The value on each PVC is asymmetric (you can set a different CIR in each direction), which provides more efficient use of bandwidth.
BC (Kbits) (Committed Burst Size)	Enter the maximum amount of data, in Kbits, that the network attempts to transfer 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.

Table 6-5. Set Traffic Type Attributes Fields (Continued)

Field	Action/Description
BE (Kbits) (Excess Burst Size)	Enter the maximum amount of uncommitted data, in Kbits, the network will attempt to deliver during a specified time interval, Tc. Tc is calculated as BC/CIR. The network treats this data as Discard Eligible (DE) data.
Rate Enf Scheme	Select *Simple (default)* or *Jump*. The configurable rate enforcement scheme provides additional flexibility, increased rate enforcement accuracy, and improved switch performance. See "Rate Enforcement" on page 5-5 for more information. ***Note**: Simple indicates time (Tc) as measured in periodic intervals. If you select the Simple scheme, the "bad" PVC detection feature is disabled.*
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 (1). *Off (default)* – Disables zero CIR. ***Note**: If you set Zero CIR Enabled to On, you can not set the CIR, Bc, and Be values.*
Delta BC (bits)	Set the number of Delta Bc bits for this circuit between 0 - 65528 (*default* 65528). 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). 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.

User Preference Attributes

From the Add PVC dialog box (Figure 6-3 on page 6-9), select Set [User Preference] Attributes and complete the fields described in Table 6-6.

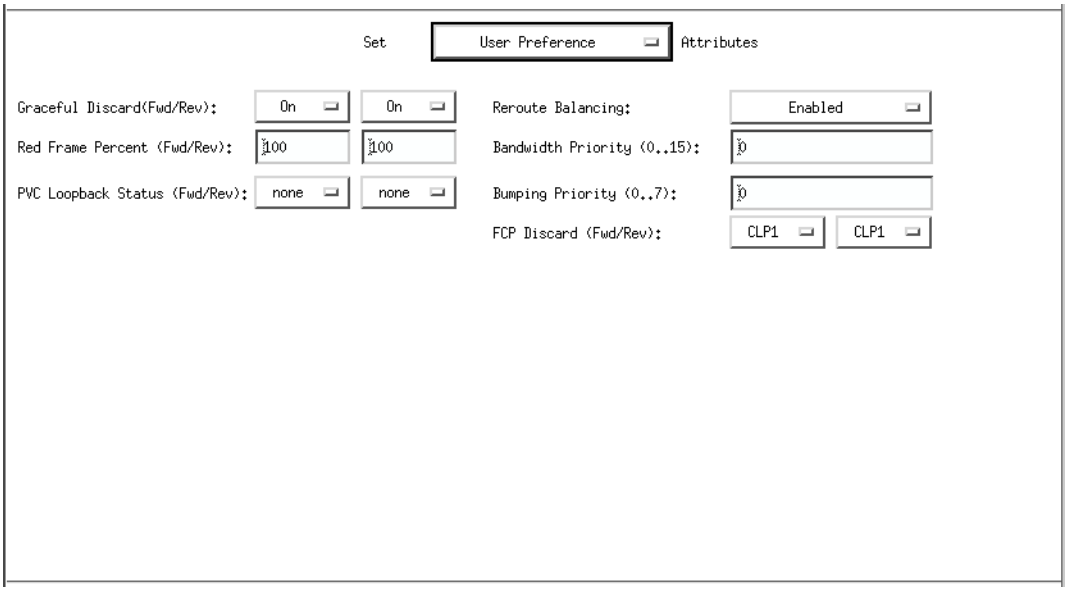

Figure 6-5. Set User Preference Attributes (FR: UNI DCE)

Table 6-6. 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. The Red Frame Percent limits the number of red frames the network is responsible for delivering.
PVC Loopback Status (Fwd/Rev)	Displays the current loopback state. If "None" is not displayed, do not attempt to modify or delete the selected circuit. See the *NavisCore Diagnostics and Troubleshooting Guide* for more information about loopback testing.

Table 6-6. Set User Preference Fields (Continued)

Field	Action/Description
Reroute Balancing	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 *NavisCore NMS Getting Started Guide*. If you set this option to *Disable*, this circuit does not use the reroute tuning parameters.
Bandwidth Priority	Set a value from 0 through 15 where 0 is the default and indicates the highest priority. See "Routing Priority Rules" on page 5-3 for more information.
Bumping Priority	Set a number from 0 through 7 where 0 is the default and indicates the highest priority. See "Routing Priority Rules" on page 5-3 for more information.
FCP Discard (Fwd/Rev)	This frame discard attribute applies only to a CBX 500 with an FCP; however, it is offered as a selection on non-CBX endpoints. Even though one or both endpoints may not be on a CBX with FCP, the PVC might traverse a CBX 500 FCP trunk. In this case, the provisioned attribute is used. Displays if a QoS class that supports FCP Discard is selected. One of the following options applies: *CLP1* – Selective CLP1 discard is provisioned for UBR, ABR, and VBR-nrt PVCs. If the current cell causes the queue for a PVC to exceed the discard thresholds, and the cell has CLP set to 1, the cell is discarded. Note that EPD is not performed in this case. *EPD* – Early Pack Discard. The ATM Flow-control processor can perform EPD for UBR, ABR, and VBR-nrt PVCs. If this option is selected, when a cell causes the queue for a PVC to exceed the discard thresholds, the VC enters the EPD state. The cells in the current packet of the VC are admitted to the queue. However, when the end of the current packet is detected, all of the cells in the next packet are discarded for that PVC.

Completing the Circuit Definition

Use the following steps to complete the circuit configuration:

1. (*Optional*) To configure NavisXtend Accounting Server parameters for this circuit, choose the Accounting command. For more information, see the *NavisXtend Accounting Server Administrator's Guide*.
2. 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 dialog box reappears.

If enabled, the Reliable Scalable Circuit feature verifies the card state of each PVC endpoint before sending the SNMP set command. If the card status at either endpoint is not up, the NMS displays an error message indicating where the failure occurred. If you receive such a message, see Appendix A for more information.

3. (*Optional*) To configure this PVC for a specific VPN and customer, see page 8-9.
4. To add more PVCs, repeat the steps in "Adding a Circuit Connection" on page 6-7.
5. When you finish, choose Close to return to the network map.

Manually Defining the Circuit Path

The Define Path function enables you to manually define a circuit path and bypass the OSPF algorithm to make circuit-routing decisions.

You cannot manually route a circuit that is configured with both endpoints in the same switch.

To manually define the circuit path:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Circuits ⇒ Point-to-Point. The Set All PVCs on Map dialog box (Figure 6-1 on page 6-3) appears.
2. Position the cursor in the Search by Name field and press Enter to view a list of configured circuits.
3. Select the circuit (name) for which you want to manually define the circuit path.

4. Choose Define Path. The Define Circuit Path dialog box appears as shown in Figure 6-6.

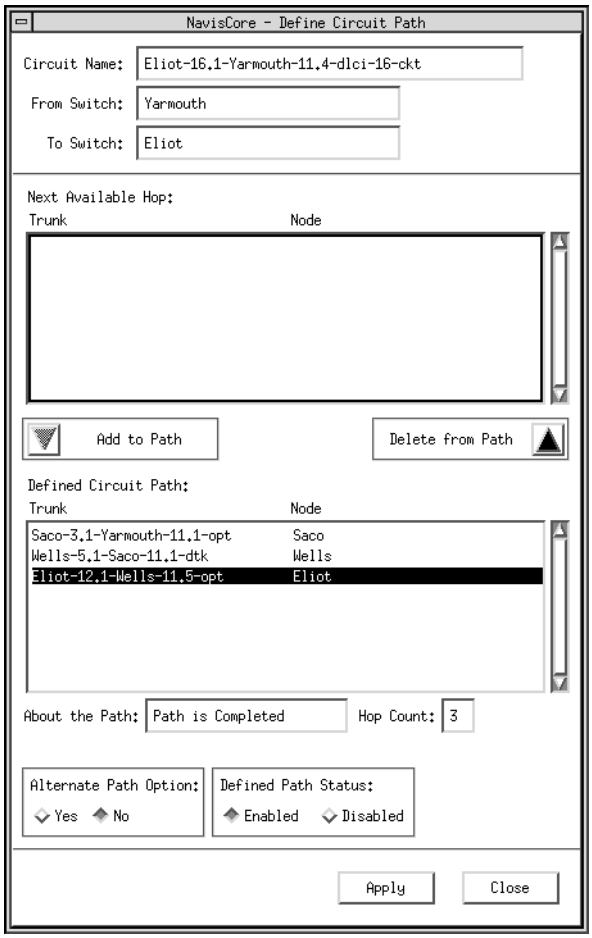

Figure 6-6. Define Circuit Path Dialog Box

5. Complete the Define Circuit Path dialog box fields described in Table 6-7.

Table 6-7. Define Circuit Path Fields

Field	Action/Description
Next Available Hop	Displays a list of the available hops (i.e., trunk-node pairs). Select the trunk-node pair to route the circuit through the network. When there are multiple trunks between two nodes, select [Any Trunk] to route the circuit bases on OSPF. Choose Add to Path. The specified trunk-node pair is added to the Define Circuit Path list, which displays all selected hops.
About the Path	Displays the path status. For example, after you define a path, this field displays "Path is Completed."
Hop Count	Displays the number of hops in the selected path.
Alternate Path Option	Select *Yes* or *No* to specify whether OSPF should route the circuit path if the manual route fails. *Yes* – Enables OSPF to route the circuit based on the best available path if the manually defined path fails. *No* – Prevents the circuit from being rerouted; the circuit remains down until the defined path is available.
Defined Path Status	Select *Enabled* or *Disabled* to define whether to use the defined path or to enable the network routing to specify the circuit path. *Enabled* – Routes the circuit based on the manually defined route. *Disabled* – Routes the circuit based on the network's OSPF algorithm.

6. Choose Apply and then choose Close. The specified values are sent to the switch(es), provided the switch(es) are currently communicating with the NMS.

Configuring Multicast DLCIs

The Set Multicast DLCIs function enables you to add, modify, and delete multicast DLCI configurations for a Frame Relay network. A multicast DLCI is a circuit configured as multiple groups of circuits on the same logical port. You can define up to 32 multicast groups per switch. You must first configure Frame Relay circuits to define the DLCIs. You then allocate these circuits as member DLCIs in the multicast configuration.

Ascend currently supports one-way multicast. A multicast DLCI enables the network to:

- Accept a frame on a single DLCI
- Replicate the frame
- Distribute the frame to multiple circuit destinations

This configuration requires you to enter a DLCI for a multicast group made up of several circuits. The DLCI represents the circuit endpoints. You must first configure the DLCIs, before you can allocate them as member DLCIs in the multicast group. See Table 6-4 on page 6-10 for information about defining a unique DLCI for a logical port.

Multicast DLCI Member Limits

The number of multicast members supported on a card is a function of the number of bytes available on the card and the frame size being transmitted, as follows:

```
member limit = bytes available / frame size
```

This formula determines the maximum number of multicast members supported at the egress card where the multicast DLCI actually resides. In general, the number of multicast members supported decreases as the frame size increases.

The number of bytes available depends on the card type:

- All IOPA cards (UIO, DSX, T1, E1, etc.) have a maximum of 32000 bytes available.
- All IOPB cards (HSSI, ATM, 12-port E1) have a maximum of 9500000 bytes.

The ATM CS and ATM IWU I/O modules do not support multicast DLCI.

Table 6-8 lists common frame sizes and the maximum number of multicast members supported on both card types.

Table 6-8. Multicast DLCI Member Limits

Frame Size (in bytes)	IOPA Card Maximum Multicast Members	IOPB Card Maximum Multicast Members
64	514	14936
128	257	7468
256	128	3734
512	64	1867
1024	32	933
2048	16	466
4096	8	233
8160	4	117

If you have already configured a multicast DLCI, the dialog box displays this information. You can use the Modify or Delete commands to modify or delete multicast DLCI configurations.

Adding a New Multicast DLCI

To configure multicast DLCIs:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Multicast DLCIs. The Set All Multicast DLCIs dialog box appears as shown in Figure 6-7.

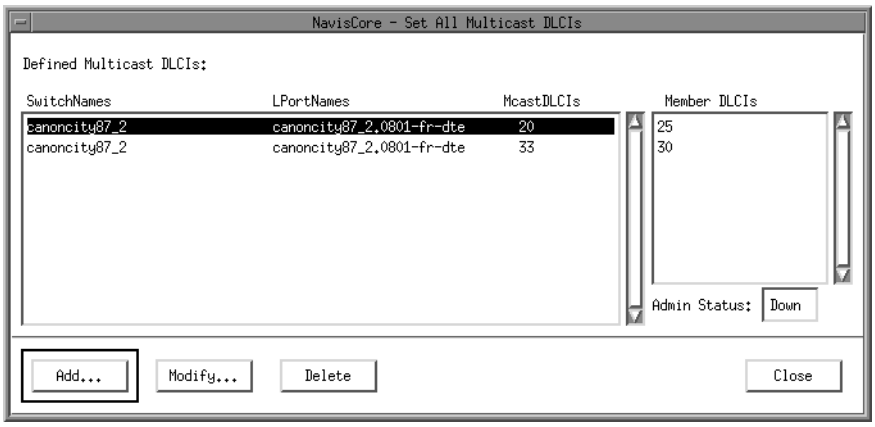

Figure 6-7. Set All Multicast DLCIs Dialog Box

2. Choose Add. The Select End Logical Port dialog box appears as shown in Figure 6-8.

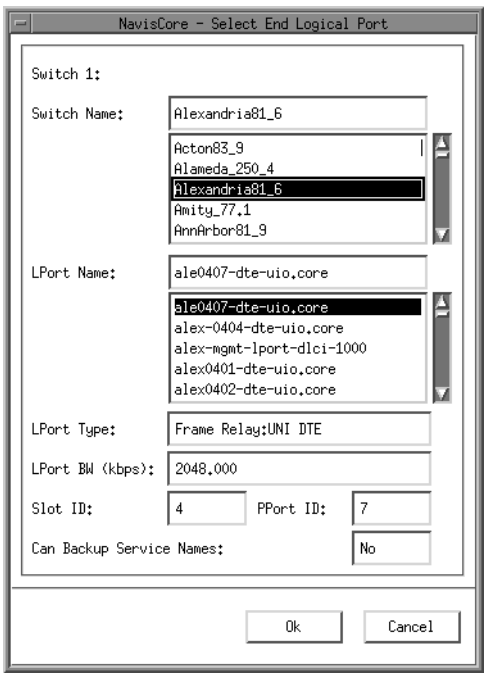

Figure 6-8. Select End Logical Port Dialog Box

3. Choose OK. The Add Multicast DLCI dialog box appears as shown in Figure 6-9.

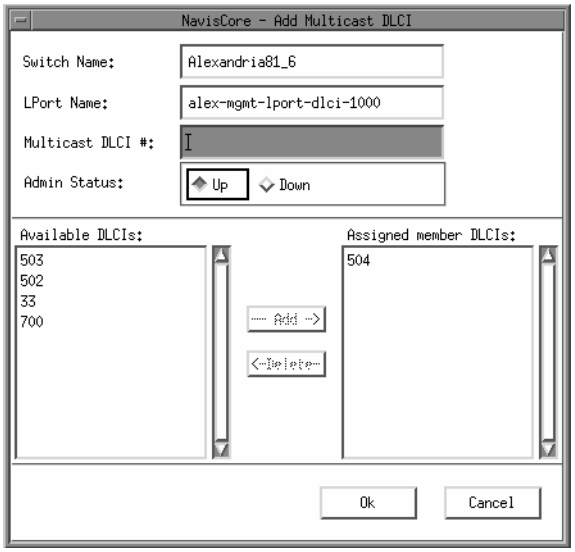

Figure 6-9. Add Multicast DLCI Dialog Box

4. Complete the required dialog box fields described in Table 6-9.

Table 6-9. Add Multicast DLCI Fields

Field	Action/Description
Switch Name	Displays the switch containing the multicast members.
LPort Name	Displays the name of the logical port to receive the multicast frames.
Multicast DLCI #	Enter a DLCI number to identify the multicast group. For more information, see "About DLCI Numbers" on page 5-7.
Admin Status	Select Down or Up to indicate whether to activate the multicast DLCI when the switch or port comes online.
Available DLCIs	Select the DLCIs (circuit endpoints) you want to allocate as members of the multicast group and choose Add. This DLCI now appears in the Assigned member DLCI list.
Assigned member DLCIs	Displays a list of the DLCIs you already selected for this multicast group. A multicast group must have at least one member. If you delete a circuit that is a member of a multicast group, the system automatically deletes it from the multicast group.

5. Choose OK to complete the configuration.

7

Configuring Management Paths

This chapter describes how to configure a management path between the network management station (NMS) and IP host that you use to access the switch network, either for configuration or Telnet purposes. The term *NMS* describes the workstation that is used to host NMS applications. You can use the same procedure to establish communication between the switch and any IP host (for example, the NavisXtend Accounting Server).

The connection between the NMS and the switch network is called the *NMS Path*. This connection sets up the link to send and receive management protocol requests and responses. To make this connection, you must know the IP address of the NMS. The NMS path configuration is node-specific and describes each NMS that attaches via the switch.

The management path options described in this chapter are available when the NMS or IP host connects to the switch via a router or Network Interface Card (NIC). You only need to define an NMS path for the switch that contains one of the following management connection elements:

- **Management PVC** — You can use this type of connection for all applications involving a switch and an attached NMS or IP host. Because the management PVC is an actual PVC between the UNI or NNI logical port (to which the NMS or IP host connects) and the remote switch processor module, the switch that connects the NMS or IP host is not burdened by the traffic traversing the management PVC.
- **Management DLCI** — You can use this type of connection when the NMS or IP host connects to a LAN through a router that provides the Frame Relay connection to the switch. (The switch does not need an Ethernet module for this type of NMS connection.) Network traffic is sent through the attached Frame Relay UNI-DCE connector as a PVC. This type of connection also enables you to move the NMS from one LAN to another with few reconfiguration requirements.

Using Management PVCs

A management PVC (MPVC) provides an access point to the switching network's management plane (which is IP-based). MPVCs offer an efficient, high-performance data path capable of transferring large amounts of management data, such as NavisXtend Accounting or Statistics Server files. This feature is available on B-STDX and CBX switch platforms.

MPVCs provide better performance than management DLCIs for transferring large amounts of data. Unlike DLCIs, MPVCs do not require that management traffic be processed by the background IP application at each switch on the path to the endpoint. For more information about DLCIs, see "Configuring Management DLCIs" on page 7-6.

MPVCs originate at the switch I/O interface: IOP for a B-STDX, and IOM for a CBX. They terminate at an internal logical port located on the switch processor module (CP or SP). MPVCs provide a data path that accesses internal network management functions. This enables you to use any physical port as a network management port.

The MPVC internal logical port is designated as MgmtLPort.SW[***switchname***]. It uses an interface number (ifnum) of 4093. To form the circuit, connect MgmtLPort.SW[***switchname***] endpoint to any Frame or ATM logical port for interworking MPVC. You can configure MPVCs across different switch platforms; for example, B-STDX UNI to CBX MPVC. Configure the remaining PVC attributes as you would for a standard PVC. Note that you can use the internal management port to terminate more than one MPVC.

MPVCs enable you to configure a management path to an autonomous system external (ASE). Once you define the management path, the IP process on the switch processor module can send (and receive) IP packets over the MPVC to (and from) the ASE. Note that IP packets are encapsulated within Frame Relay frames according to RFC 1490. The management path is described in the switch's arp cache and routing table.

Configuring a Management PVC

The following sections describe how to configure an NMS Path using a management PVC. As part of this process, you need to first configure an unused physical port for which you can then define, in this example, a Frame Relay UNI logical port.

Defining Physical Port Attributes

1. Select the switch for which you want to configure the Frame Relay UNI logical port endpoint.
2. Log in to NavisCore using either a provisioning or operator password.
3. From the Administer menu, select Ascend Parameters ⇒ Set Parameters. The switch back panel appears.
4. Select the physical port you want to configure and choose Attrs. to display the Set Physical Port Attributes dialog box.
5. Complete the dialog box fields. Refer to the *NavisCore Physical Interface Configuration Guide* if you need information about changing default values.
6. Choose Apply.

Defining a Frame Relay UNI Logical Port

1. From the Set Physical Port Attributes dialog box, choose Logical Port. The Set All Logical Ports in PPort dialog box appears (Figure 3-1 on page 3-3).
2. Choose Add to display the Add Logical Port dialog box (Figure 3-2 on page 3-6).
3. Select Frame Relay UNI DCE or DTE as the logical port type.
4. Choose OK. The Add Logical Port dialog box reappears (Figure 3-2 on page 3-6).
5. Use the instructions in Table 3-4 on page 3-10 to set the Administrative Attributes.
6. Choose OK to return to the Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-3).
7. Choose Close to return to the Set Physical Port Attributes dialog box. Choose Cancel to return to the Switch Back Panel dialog box.
8. Choose Close to return to the network map.

Defining the Management PVC Connection

1. From the Administer menu, select Ascend Parameters ⇒ Set All Circuits ⇒ Point-to-Point. The Set All PVCs on Map dialog box (Figure 6-1 on page 6-3) appears.
2. Choose Add. The Select End Logical Port dialog box (Figure 6-8 on page 6-23) appears.
3. Select the name of the switch where the management port (Endpoint 1) resides.
4. Select the logical port name "MgmtLPort.SW[***switchname***]" for Endpoint 1. The [***switchname***] should correspond to the name of the switch on which the management port endpoint resides. The LPort Type field should display Others:Multi Hop MPVC.
5. Select the name of the switch where Endpoint 2 resides.
6. Select the name of the logical port for Endpoint 2.
7. Choose OK. The Add PVC dialog box (Figure 6-3 on page 6-9) appears.
8. Enter a Circuit Name for the management PVC. You will select this name when you configure the NMS Path.
9. Use the instructions in Table 6-4 on page 6-10 to set the Administrative attributes.
10. Use the instructions in Table 6-5 on page 6-14 to set the Traffic attributes.
11. Use the instructions in Table 6-6 on page 6-16 to set the User Preference attributes.
12. *(Optional)* To configure NavisXtend Accounting Server parameters for this circuit, choose the Accounting command. For more information, see the *NavisXtend Accounting Server Administrator's Guide*.
13. Choose OK to define the circuit parameters. The Set All PVCs on Map dialog box reappears. Choose Close to return to the network map.

Defining the NMS Path

1. On the network map, select the switch to connect to the NMS.
2. From the Administer menu, select Ascend Parameters ⇒ Set All Management Paths.The Set All Management Paths dialog box (Figure 7-1) appears.

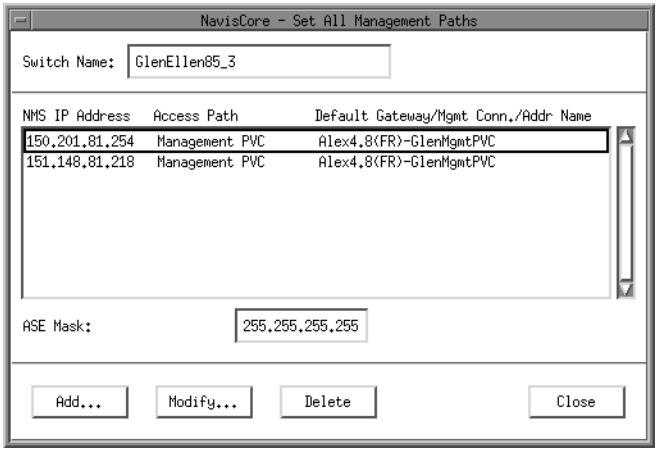

Figure 7-1. Set All Management Paths

3. Choose Add to display the Add Management Path dialog box, and select Management PVC as the Access Path. The dialog box now displays the Management PVC Name list box.

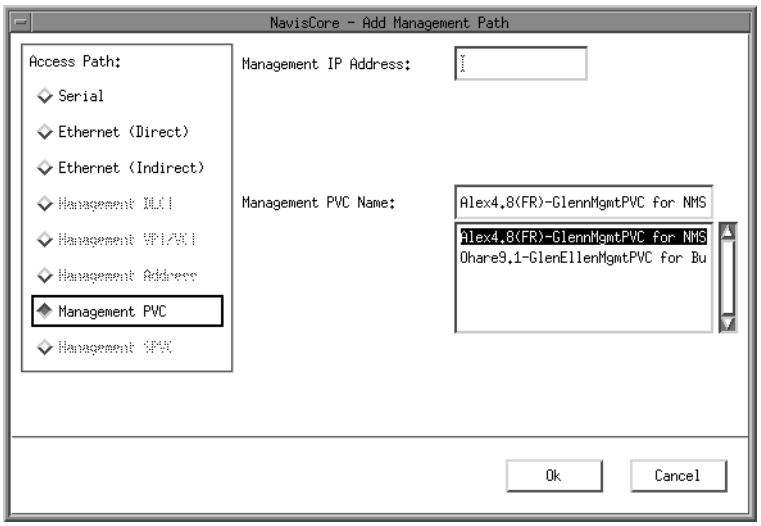

Figure 7-2. Add Management Path

4. Enter the NMS IP Address. This is the IP address of the SPARCstation to which this switch connects.

5. Select the Management PVC Name you entered in Step 8 on page 7-4.

6. Choose OK. Choose Close to return to the network map.

Using Management DLCIs

You use a management DLCI when the NMS connects to the gateway switch through a router, which provides the Frame Relay connection to the switch.

Configuring Management DLCIs

The following sections describe how to configure a management DLCI. This access method enables you to monitor the network without the use of an Ethernet module in the switch. It also provides the flexibility to move the NMS from one LAN to another with few reconfiguration requirements.

Adding a New Management DLCI

To configure a Management DLCI:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Management DLCIs. The Set All Management DLCIs dialog box appears listing the Management DLCIs already configured.

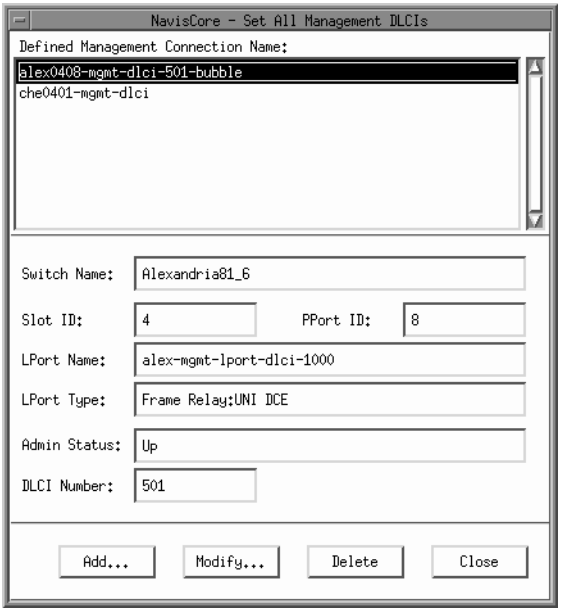

Figure 7-3. Set All Management DLCIs Dialog Box

If you have already configured a Management DLCI, the dialog box displays this information. You can use the Modify or Delete commands to modify or delete Management DLCI configurations.

2. Choose Add. The Select End Logical Port dialog box appears as shown in Figure 7-4.

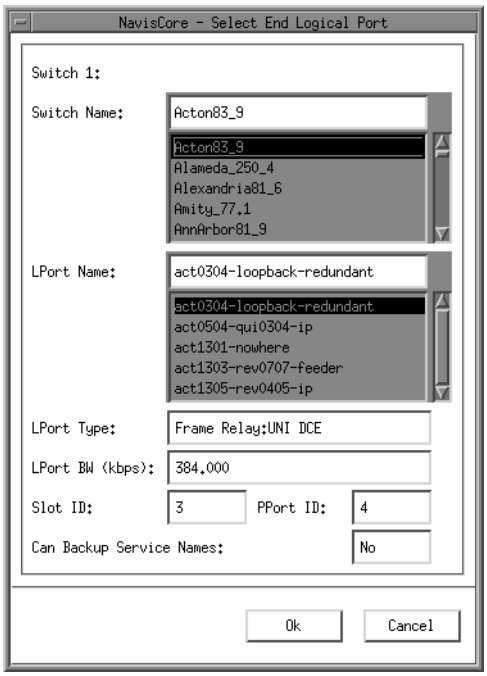

Figure 7-4. Select End Logical Port Dialog Box

3. Complete the required dialog box fields described in Table 7-1.

Table 7-1. Select End Logical Port Fields

Field	Action/Description
Switch Name	Select the name of the switch that connects to the router that serves as the Frame Relay interface for the Network Management DLCI.
LPort Name	Select the name of the logical port configured to access the router.
LPort Type	Displays the logical port type.
LPort BW (kbps)	Displays the logical port bandwidth.
Slot ID	Displays the I/O slot (number) in which the module resides.
PPort ID	Displays the port number for the port you are configuring.
Can Backup Service Names	*(Fault-tolerant PVC only)* Displays Yes or No to indicate whether this lport is configured for backup service. For more information, see Chapter 9, "Configuring Fault-Tolerant PVCs."

4. Choose OK. The Add Management DLCI dialog box appears as shown in Figure 7-5.

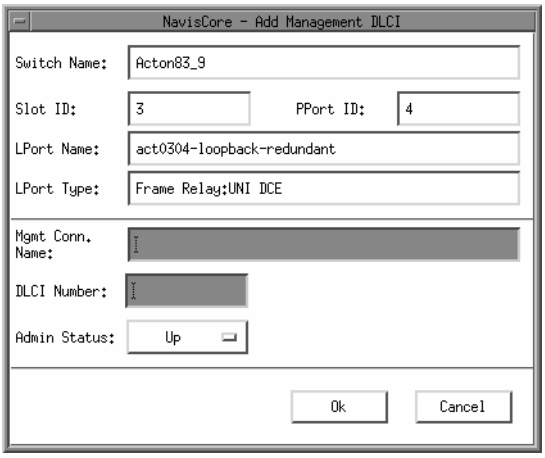

Figure 7-5. Add Management DLCI Dialog Box

5. Complete the dialog box fields described in Table 7-2.

Table 7-2. Add Management DLCI Fields

Field	Action/Description
Switch Name	Displays the name of the switch that connects to the router that serves as the Frame Relay interface for the Network Management DLCI.
Slot ID	Displays the I/O slot (number) in which the card resides.
PPort ID	Displays the port number for the physical port.
LPort Name	Displays the name of the logical port you configured for the router.
LPort Type	Displays the logical port type.
Mgmt Conn. Name	Enter a unique, continuous, alphanumeric name to identify the DLCI. Do not use hyphens, dashes, parentheses, and asterisks.
DLCI Number	Enter the number that is used for the Management DLCI. For more information, see "About DLCI Numbers" on page 5-7.
Admin Status	Select either Up or Down to define whether the DLCI is activated when the switch or port comes online.

6. Choose OK to complete the configuration.

Defining the NMS Path

1. On the network map, select the switch to connect to the NMS.
2. From the Administer menu, select Ascend Parameters ⇒ Set All Management Paths. The Set All Management Paths dialog box (Figure 7-1 on page 7-5) appears.
3. Choose Add. The Add Management Path dialog box (Figure 7-2 on page 7-5) appears.
4. Select Management DLCI as the Access Path.
5. Enter the NMS IP Address. This is the IP address of the SPARCstation to which this switch connects.
6. Select the Management DLCI Name (Management Conn. Name) you defined in Step 5 on page 7-3.
7. Choose OK. Choose Close to return to the network map.

Defining the Static Route

To complete the management DLCI configuration, you must enter a static route in the router and the NMS workstation to access the internal IP network.

8

Configuring Virtual Private Networks

Virtual Private Network (VPN) is an *optional* software feature that enables network providers to dedicate resources for those customers who require guaranteed performance, reliability, and privacy. This feature is sometimes called Application Specific Routes (ASR) or Customer Specific Routes (CSR).

A VPN enables you to provide dedicated bandwidth to the customer. When you configure a trunk, you can dedicate it to a specific VPN and, if desired, allow customers to monitor their own networks. However, switch control and configuration stays with you as the network provider.

About Virtual Private Networks

The VPN feature allows you to create multiple private networks from a single public network. After you create a VPN name and ID, you associate one or more customer names and IDs with the VPN. When all VPNs and customers are created in the database, you assign UNI/NNI logical ports to specific VPN/customer associations. In addition, you need to dedicate selected public network trunks to specific VPNs.

You must configure all PVCs that you create on UNI/NNI logical ports for selected VPN/customer associations. SVCs, however, inherit the VPN/customer associations of the host logical port.

When you configure the logical port or PVC, you also set the Net Overflow attribute. This attribute specifies whether PVCs or SVCs are restricted to trunks of their own VPN or can use public (shared) trunks during outages. Customers that operate in restrictive mode need to purchase redundant trunks. Figure Figure 8-1 provides a restrictive mode example.

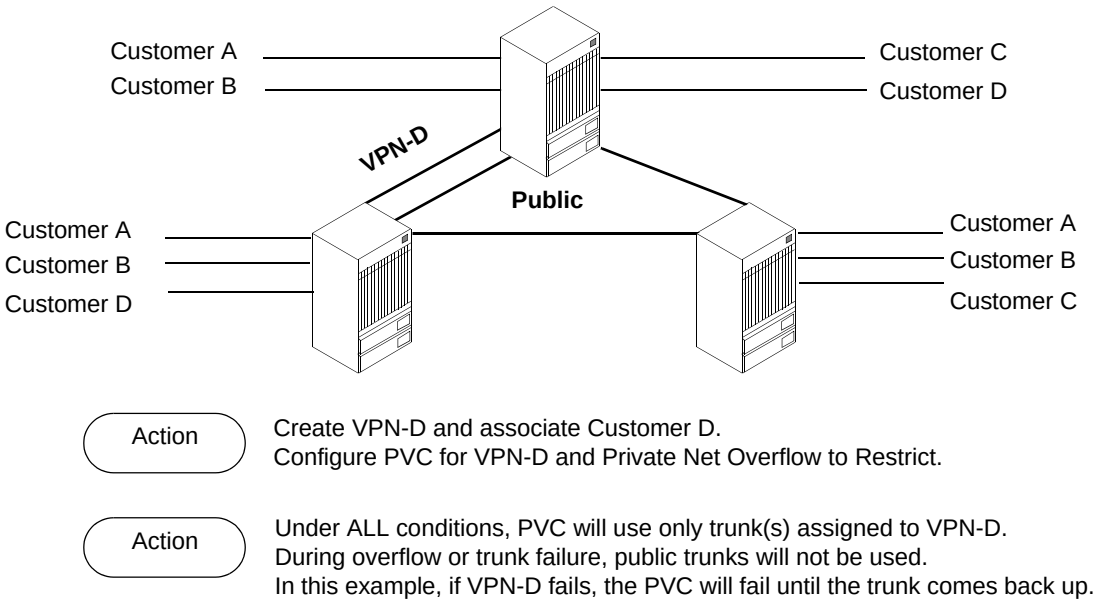

Figure 8-1. VPN Restrictive Mode Example

If you set the Net Overflow parameter to shared, a private network can also use public trunks as a backup. This is called inclusive mode (shown in Figure 8-2 on page 8-3). The identifier, VPN 0, is reserved to indicate the public part of the network. Trunks that have non-zero VPNs are reserved for data traffic matching that VPN, although they can also carry management traffic for the entire network.

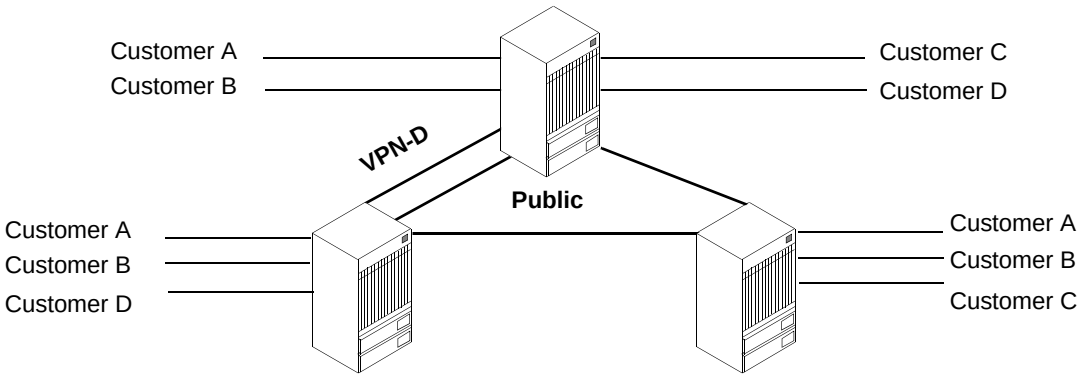

Figure 8-2. VPN Inclusive Mode Example

Configuring a Virtual Private Network

Use the following sequence to set up a VPN:

Step 1. Create the VPN (page 8-4).

Step 2. Add customers to a specific VPN (page 8-5).

Step 3. Dedicate a trunk to a specific VPN (refer to page 4-14).

Step 4. *For SVC traffic*, when you configure the UNI or NNI logical port, specify the net overflow attribute (page 3-10). Then, dedicate this logical port to a specific VPN and customer (page 8-7).

Step 5. *For PVC traffic*, specify the net overflow attribute for the circuit (page 3-10). Then, dedicate the circuit to a specific VPN and customer (page 8-9).

Creating a VPN

Use the following steps to create a VPN and add customers to this network:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Virtual Private Networks. The Set All Virtual Private Networks dialog box (Figure 8-3) appears.

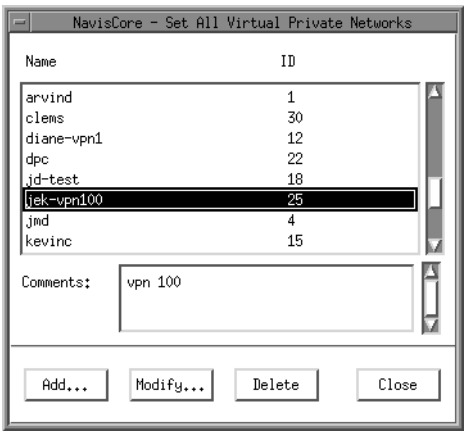

Figure 8-3. Set All Virtual Private Networks Dialog Box

2. Choose Add. The Add Virtual Private Network dialog box (Figure 8-4) appears.

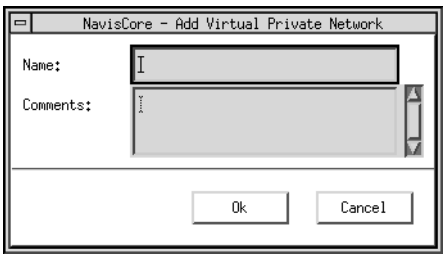

Figure 8-4. Add Virtual Private Network Dialog Box

3. Enter a name for this VPN and add any additional comments.
4. Choose Apply.
5. Choose Close to return to the network map.

Adding Customers to the VPN

To add customers to the VPN:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Customers. The Set All Customers dialog box (Figure 8-5) appears.

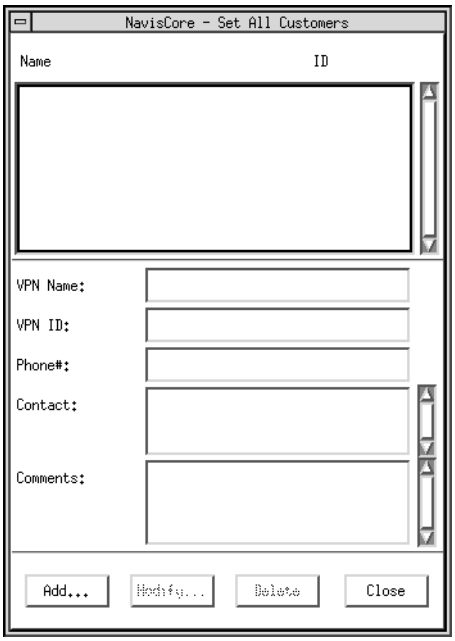

Figure 8-5. Set All Customers Dialog Box

2. Choose Add. The Add Customer dialog box (Figure 8-6) appears.

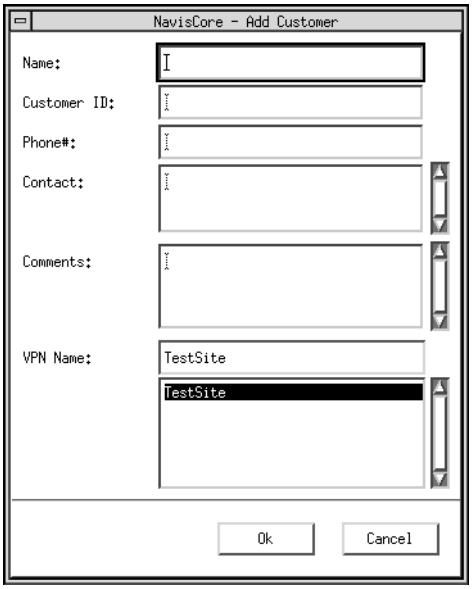

Figure 8-6. Add Customer Dialog Box

3. Enter a customer name.
4. Assign a value from 1 to 65535 for the customer ID.
5. (*Optional*) Enter the phone number, contact name, and any additional comments.
6. Select the VPN name to which this customer belongs.
7. Choose Apply.
8. Choose Close to return to the network map.

Configuring a Logical Port for VPN

To implement VPN for a network that contains SVCs, specify the net overflow attribute when you configure a UNI logical port (Table 3-4 on page 3-10). This parameter determines whether SVCs originating from this port are restricted to trunks of their own VPN, or whether SVCs can use public (shared) trunks during overflow conditions.

Once you configure a logical port, use the following steps to dedicate it to a VPN:

1. See page 3-3 to access the Set All Logical Ports in PPort dialog box.
2. From the list of logical port names, select the one you need to dedicate to a VPN.
3. Using the Select:Options command, select VPN/Customer Info and choose Set. The Select Customer and VPN dialog box (Figure 8-7) appears.

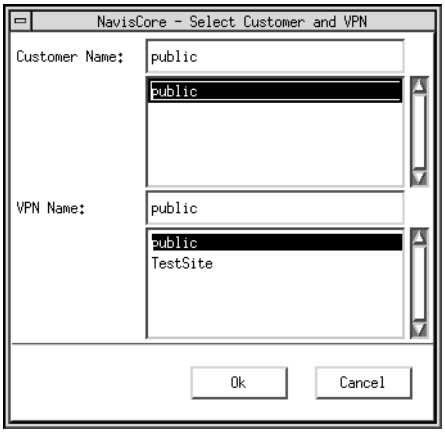

Figure 8-7. Select Customer and VPN Dialog Box

4. Select the customer and VPN name.
5. Choose OK.
6. Choose Close to exit.

Using the VPN/Customer View Feature

When you need to create PVCs for a specific VPN or customer, use the Select VPN/Customer view feature. This feature allows you to enable a network map view for a specific VPN or customer. VPN/Customer View makes it easy to identify those logical ports that belong to the VPN for which you need to configure PVCs; with this feature enabled, the Select End Logical Ports dialog box (page 6-7) only displays the logical ports that belong to the VPN or customer you select.

As you configure logical ports, use the instructions in "Configuring a Logical Port for VPN" on page 8-7 to assign the port to a VPN or customer.

To give a customer the ability to monitor network resources without the ability to provision, edit either the .cshrc or the .profile file for an NMS user and add the following lines:

```
OVwRegDir=/opt/CascadeView/registration
export OVwRegDir
```

These lines disable the Administer menu and all its provisioning functions; the NMS user only sees the Monitor menu functions.

To use VPN/Customer view:

1. From the Administer menu, select Ascend Object:Select Customer/VPN. The Select Customer/Virtual Private Network View dialog box (Figure 8-8) appears.

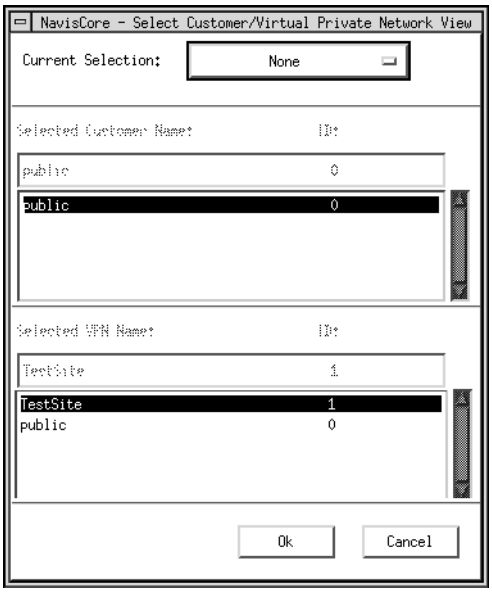

Figure 8-8. Select Customer/Virtual Private Network Dialog Box

2. Use the Current Selection command to select either Customer or VPN.

   Use None (default) to disable VPN/Customer view. (With the VPN/Customer view disabled, you can configure PVCs using logical port endpoints that belong to any VPN or customer.)

3. Depending on the option you select, review either the Selected Customer Name or Selected VPN Name list.
4. Select the Customer or VPN name.
5. Choose OK.

Configuring a PVC for VPN

When you configure a PVC for VPN, first specify the private net overflow attribute (Table 6-4 on page 6-10). This parameters determines whether this PVC is restricted to trunks of its own VPN, or can use public (shared) trunks during overflow conditions.

After you configure a PVC, use the following steps to dedicate it to a VPN:

1. Refer to page 6-3 to access the Set All PVCs on Map dialog box.
2. From the list of PVC names, select the one you need to dedicate to this VPN.
3. Choose VPN/Customer. The Select Customer and VPN dialog box appears (see Figure 8-7 on page 8-7).
4. Select the customer and VPN name.
5. Choose OK.
6. Choose Close to exit.

9

Configuring Fault-Tolerant PVCs

A fault-tolerant PVC configuration enables UNI DCE and DTE logical ports to serve as a backup for any number of active UNI ports. If a primary port fails or if you need to take a primary port off-line, you activate the backup port.

Use the following sequence to configure fault-tolerant PVCs:

Step 1. Define a UNI-DCE or UNI-DTE logical port as described in Chapter 3, "Configuring Frame Relay Logical Ports." To designate a backup port, choose Yes for the option "Can Backup Service Names" (page 3-10).

Step 2. Specify a service name for the primary port.

Step 3. Configure circuits to use a service name as the endpoint.

Step 4. Activate the backup port (refer to page 9-5).

Ascend recommends that you avoid configuring SVCs on a logical port that is also designated as a backup port in a fault-tolerant PVC configuration.

Creating a Backup Port

To create a backup port, first define a UNI DCE or DTE logical port and select Yes for the option, “Can Backup Service Names” (refer to Table 3-4 on page 3-10). When a backup port is not in use, the port is idle and does not use network resources.

Creating a Primary Port

To create a primary port, you assign a service name to a UNI logical port. (Do not choose a port that you already configured for backup.) When you configure the circuit, choose this service same as the endpoint, instead of a switch and logical port combination. When you activate the backup port, all PVCs on the failed primary port are rerouted, preserving VPI/VCIs in the process.

Ascend’s fault-tolerant PVC feature is transparent to the end user, meaning that you do not have to configure the CPE to accommodate the new functionality. Therefore, end users can benefit from this feature through the public Ascend-based ATM network, or by combining their private Ascend switches with services provided by their public carrier.

Creating Service Names

The *service name binding* is a name you define to identify the primary port. A circuit recognizes its service endpoint by this name, instead of the logical port name.

To create the service name bindings:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Service Name Bindings. The Set All Service Name Bindings dialog box (Figure 9-1) appears, displaying any service names you have already configured.

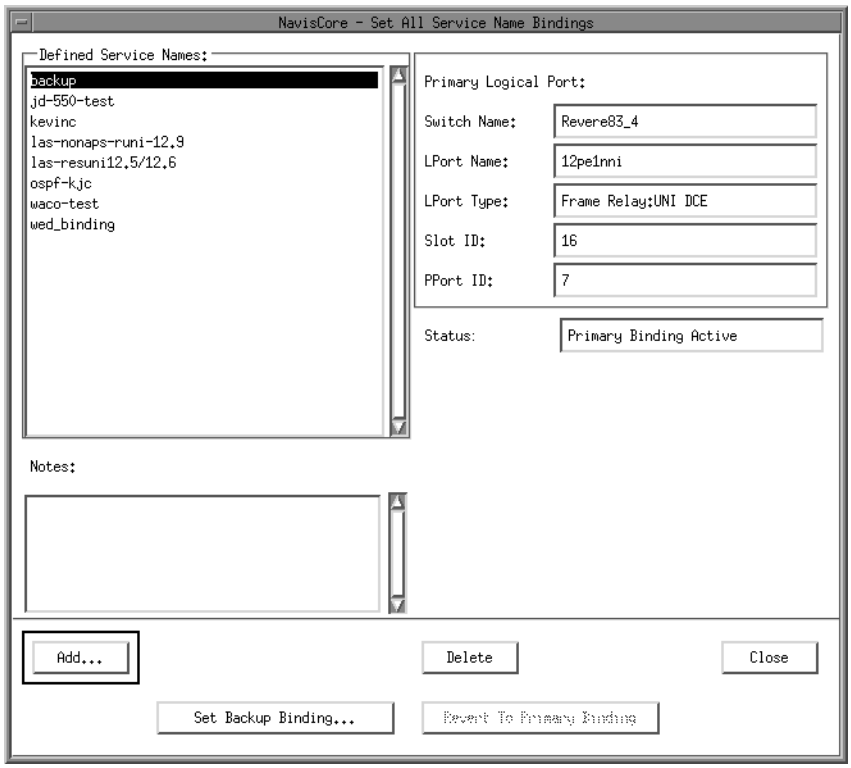

Figure 9-1. Set All Service Name Bindings Dialog Box

You can access the following functions from the Set All Service Name Bindings dialog box:

- To return the primary logical port to service, select the Service Name and choose Revert to Primary Binding.
- To delete a service name, select the service name and choose Delete.
- To modify a backup service binding, choose Set Backup Binding.

2. Choose Add. The Select End Logical Port dialog box (Figure 9-2) appears.

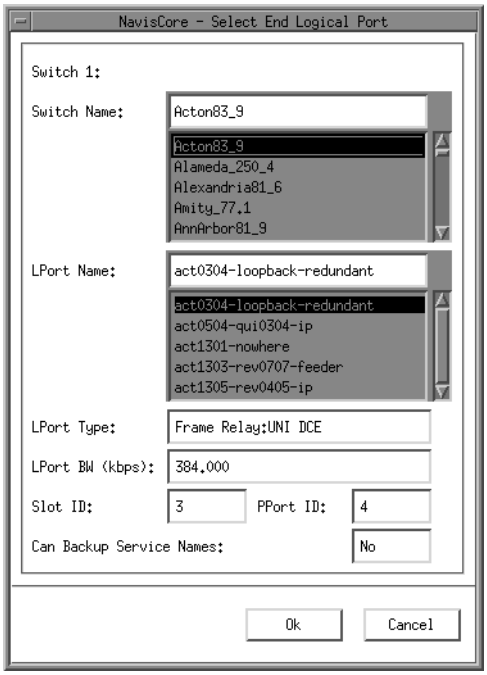

Figure 9-2. Select End Logical Port Dialog Box

3. Select the switch name and the primary logical port name.

Make sure that the Can Backup Service Names field displays No. You cannot configure a Service Name for a logical port you designated as a backup.

4. Choose OK. The Add Service Name Binding dialog box (Figure 9-3) appears.

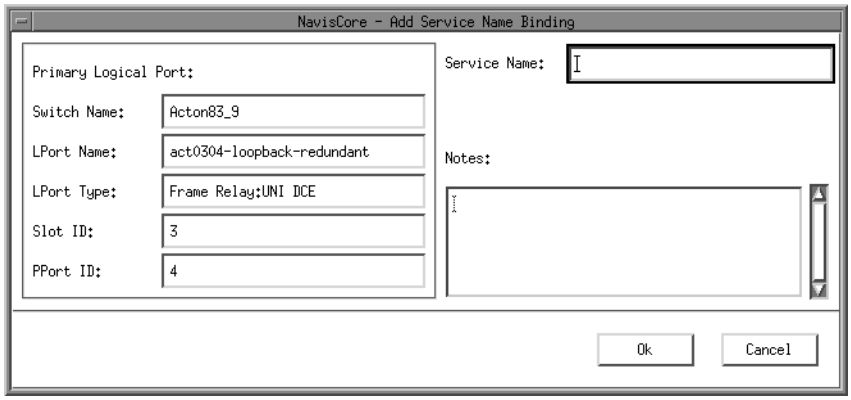

Figure 9-3. Add Service Name Binding Dialog Box

5. Type a service name (up to 32 characters). Optionally, you can enter a brief comment or description of the service in the Notes box.
6. Choose OK.
7. Continue with the instructions in Chapter 6, "Configuring PVCs," to configure the circuits for fault-tolerant PVCs.

To reroute a PVC if it fails, refer to "Activating a Backup Binding Port."

Activating a Backup Binding Port

If a primary port fails, you reassign the service name of the primary port to the backup port. Since circuits use the service name as the endpoint, all circuits configured for the primary port are rerouted to the backup port.

To enable the backup binding:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Service Name Bindings. The Set All Service Name Bindings dialog box (Figure 9-1 on page 9-3) appears.

2. Choose Set Backup Binding. The Select End Logical Port dialog box (Figure 9-4) appears.

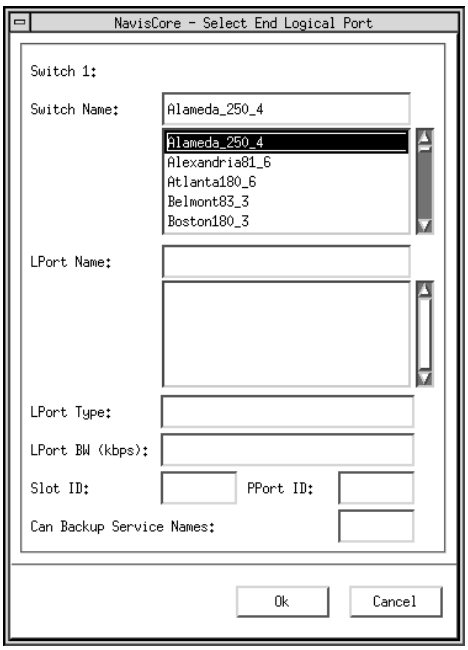

Figure 9-4. Select End Logical Port Dialog Box

3. Select the Switch Name for the backup service name binding you want to use.
4. The LPort Name field displays a list of logical ports configured for this service. Select an LPort Name that has the *same* logical port type as the port you need to back up.

Make sure that the Can Backup Service Names field displays Yes. This indicates you can use this logical port as a backup.

5. Choose OK. The Set/Modify Backup Service Name Binding dialog box (Figure 9-5) appears, displaying the Service Name that corresponds to the switch and logical port names you selected.

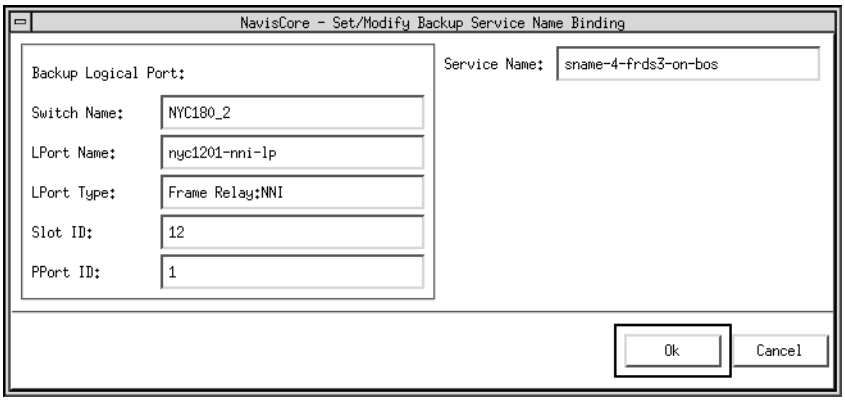

Figure 9-5. Set/Modify Backup Service Name Binding Dialog Box

6. Choose OK. The Set Service Name Bindings dialog box reappears (Figure 9-1 on page 9-3). The Status field should now display the message, Backup Binding Active.

10

Configuring Switched Virtual Circuit (SVC) Parameters

This chapter describes how to use switched virtual circuits (SVCs). With SVCs, connections are not predefined as they are for PVCs. Instead, end stations use a signaling protocol to indicate to the Frame Relay network the endpoint to which it should route the call (*called party*). To support SVC services, each user endpoint is assigned a unique address that identifies the endpoint and enables the network to route the call.

Address Formats

Before you can begin to configure your network for SVCs, you must decide which of the following address format types to use:

Native E.164 address format — E.164 addresses are phone numbers. This address format is simple and familiar; native E.164 addresses are a convenient choice for service providers using public Frame Relay or ATM networks (e.g., RBOCs) that already "own" E.164 address space.

Both Frame Relay and ATM support native E.164 address formats. However, Frame Relay does *not* support E.164 ATM End System Address (AESA) formats.

X.121 address format — X.121 addresses are an ITU-T standard used in X.25 networks. X.121 addresses are sometimes referred to as IDNs (International Data Numbers) and consist of 14 ASCII digits. Only number values between 0-9 are valid.

X.121 address formats are not configurable on ATM ports.

Designing an Address Format Plan

You use address formats to develop a network numbering plan. The SVC address formats you select must support the equipment and services your network needs to provide. Keep in mind that some CPEs may not support certain address formats. To avoid address conflicts, apply for globally-recognized address space in the formats you need to use.

Regardless of the address format you choose, the network numbering plan should satisfy the following goals:

- Intelligently assign network addresses
- Simplify network topology using a hierarchal organization
- Minimize the size of network routing tables
- Uniquely identify each endpoint
- Provide a high level of network scalability

About Route Determination

The node prefixes, port prefixes, and port addresses that are configured on network nodes are used to determine the route for a given SVC. The route is determined by a “best match” hierarchy, starting from the left-most digit of the called party address.

Keep in mind that you use node prefixes to summarize the common address parts of individual nodes. For example, if all addresses on a node contain the digits 15085551, you would define this as the node prefix. Using node prefixes to summarize – or *aggregate* – port prefixes and/or port addresses can result in more efficient routing determination. If more than one node has the same node prefix, this aggregation does not occur.

The following example shows three nodes configured with a combination of native E.164 node prefixes, port prefixes, and port addresses:

Table 10-1. E.164 Node Prefix, Port Prefix, and Port Address Example

	Node 1	**Node 2**	**Node 3**
Node Prefixes	508 6	None	508 603
Port Prefixes	508551 508552 508553	5085 508553 6035	508554 508555
Port Addresses	5085511111 5085511112 5085511113 5085555555 5085555556	None	None

The following table shows the node to which a call is routed for certain called-party addresses, and why the call is routed to that node:

Table 10-2. Routing by Called Party Address Example

Called Party Address	Node	Reason
5085511234	1	Port prefix 508551 on Node 1 is a longer match than port prefix 5085 on Node 2 and node prefix 508 on Node 3.
5085555555	1	This calling party address exactly matches a port address defined on Node 1. This is a longer match than port prefix 5085 on Node 2 and port prefix 508555 on Node 3.
5085555557	3	Port prefix 508555 on Node 3 is a longer match than port prefix 50855 on Node 2 and node prefix 508 on Node 1.
5085561111	2	Port prefix 5085 on Node 2 is a longer match than node prefix 508 on Node 1 and node prefix 508 on Node 3.
6175551111	1	Node prefix 6 on Node 1 is the only match.
6035551111	2	Port prefix 6035 on Node 2 is a longer match than node prefix 6 on Node 1 and node prefix 603 on Node 3.
6038558888	3	Node prefix 603 on Node 3 is a longer match than node prefix 6 on Node 1. There is no matching prefix or address on Node 2.
5085531111	1 or 2	Since the longest match occurs on both Nodes 1 and 2, the Admin Cost value assigned to port prefix 5085 on each node determines where the call is routed. The call is routed to the node with the lowest Admin Cost value for port prefix 5085.
5145551234	None	The call is not routed to any of these nodes because there are no matching node prefixes, port prefixes, or port addresses. If, however, you set up a default route on a port being used for network-to-network connections, all non-matching calls are routed to that port (refer to "Defining Default Routes for Network-to-Network Connections" on page 10-13).

Network ID Addressing

A network ID can be used to identify an inter-exchange carrier (IXC). You can configure network ID addressing on Frame Relay UNI and ATM UNI logical ports.

Depending on the administering authority, a network ID may be a 3-, 4-, or 8-digit Carrier Identification Code (CIC) or a 4-digit Data Network Identification Code (DNIC, X.121). Network ID addressing enables you to associate a network-to-network connection with a particular IXC using a route determination ID. It enables end-users to presubscribe to a particular IXC using a source default network ID, and override this selection on a call-by-call basis using a signaled Transit Network Selection (TNS). Signaled TNSs are screened by matching them against a list of presubscribed source validation network IDs. It is also possible to “ignore” the signaled TNS to allow routing based on the called party address instead of the TNS value; the signaled TNS is essentially stripped at ingress port.

An SVC is routed based on one of the following addresses provided at the ingress port (selected in listed order):

- Signaled TNS
- Signaled Called Party
- Configured Default TNS

Routing is performed based on *either* the signaled/provisioned TNS value or the signaled called party address (not both). Route determination network IDs *and* route determination port prefixes/addresses can be configured on a logical port at a network-to-network connection. A combination of source validation network IDs and route determination network IDs can coexist on the same port. You can provision network IDs on FRF.4, ATM UNI 3.x, 4.0, and IISP ports.

You can configure a maximum of 512 configurable addresses for a logical port (where configurable addresses equal the sum of all port addresses, prefixes, user parts, and network IDs). The maximum number of network IDs for a logical port equals 512 minus the sum of port addresses, prefixes, and user parts.

I/O Modules for SVC Frame Relay Service

The following modules support Frame Relay SVCs.

Table 10-3. Frame Relay SVC Modules

Low-Speed (IOPA)	High-Speed (IOPB)
UIO-8	HSSI
4-port channelized T1/E1	Channelized DS3
4-port unchannelized T1/E1	
DSX-10	

Configuring Node Prefixes

Node prefixes apply to all ports on the switch and are used for routing aggregation. You can configure multiple node prefixes on a switch; however, you do not need to configure any if you have port prefixes or port addresses defined on the node.

At the very least, a node prefix consists of at least one digit of the 1-15 digit native E.164 address. You can define the node prefix to be part of or all of the E.164 address. For example, for E.164 addresses that begin with 508555, you can configure the node prefix as 5 (at a minimum), 50, 508, 5085, etc. The level of required granularity depends on your network.

Node prefixes do not have to be unique to a particular node. For example, you can define node prefix 508 on multiple nodes. However, if more than one node has the same node prefix, routing aggregation does not occur.

When a node acts as an Area Border Router (that is, when the node interfaces to trunks assigned to different OSPF areas), the node prefix OSPF Area ID is used to unambiguously assign addresses configured on that node to a particular OSPF area.

Defining a Node Prefix

To define a node prefix:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Node Prefixes. The Set All Node Prefixes dialog box (Figure 10-1) appears.

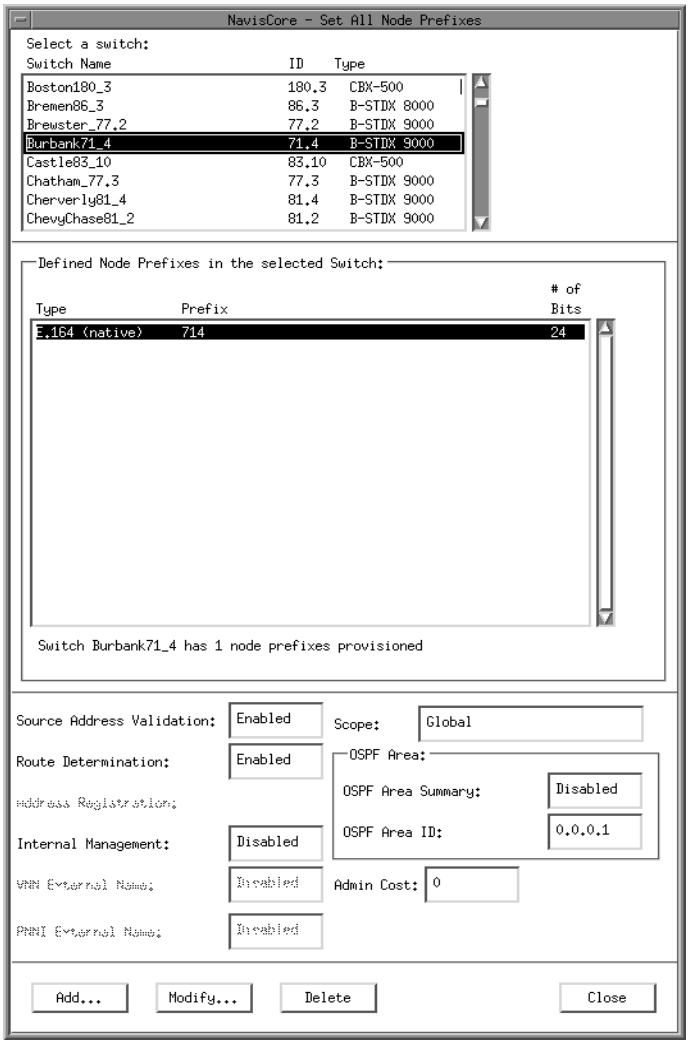

Figure 10-1. Set All Node Prefixes Dialog Box

In the Set All Node Prefixes dialog box, the top list box (*Select a switch*) displays all switches that are accessible from this NMS. The bottom list box (*Defined Node Prefixes...*) displays all the node prefixes configured on the selected switch.

2. From the top list box, select the switch for which you want to configure node prefixes.

3. Choose Add. The Add Node Previx dialog box (Figure 10-2) appears.

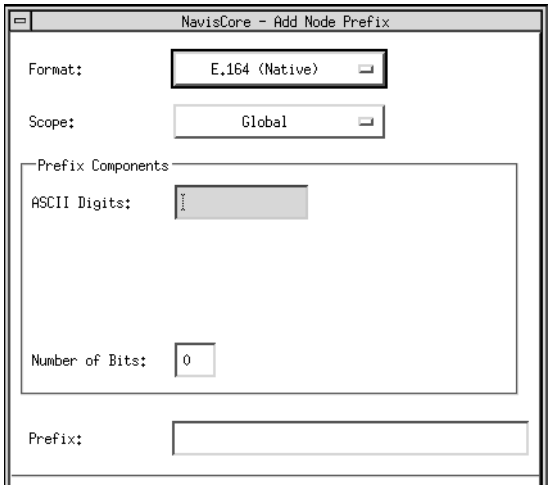

Figure 10-2. Add Node Prefix Dialog Box (E.164 Native Format)

4. Complete the Add Node Prefix dialog box fields described in Table 10-4.

Table 10-4. Add Node Prefix Fields

Field	Action/Description
Format	Select the address format. Valid options include: • E.164 (Native) (default) • X.121 For more information, see "Configuring Node Prefixes" on page 10-6.
Scope	Organizational Scope defines how far into the hierarchical PNNI domain the switch should advertise this prefix or address. ***Note:*** *This release supports Private Network-to-Network Interface (PNNI) on CBX 500 switches. PNNI is a standard designed by the ATM Forum. For more information about PNNI, refer to the* ATM Forum PNNI specification. *For information about Ascend ATM configuration and PNNI, refer to the* NavisCore ATM Configuration Guide.

Table 10-4. Add Node Prefix Fields (Continued)

Field	Action/Description
ASCII Digits	Enter the ASCII digits that represent the E.164 or X.121 address. For example, for E.164 addresses enter 5085552600 (a standard 10-digit U.S. phone number) or enter a partial number (such as 508). The value is converted to the ASCII hex values that represent each digit in the number. If you entered 5085552600, it converts to 35308383535353236303O. This value is also displayed in the Prefix column on the Set All Node Prefixes dialog box (Figure 10-1 on page 10-7).
Number of Bits	As you type the address, the value in the Number of Bits field changes to indicate the number of address bits that are checked during call screening and call routing.
Prefix	Displays the prefix that you specified in the ASCII Digits field.
Source Address Validation	Select enable to validate the calling party address against the node prefix associated with the UNI logical port that received the call setup message. If you disable this option, this node prefix is not used to validate calling party addresses.
Route Determination	If enabled, the OSPF protocol uses this node prefix for routing aggregation. If disabled, OSPF does not use it.
Internal Management	Select enable to configure the prefix that corresponds to the switch itself as an addressable entity. Select disable to disregard this feature.
OSPF Area Summary	Select enable if the node prefix summarizes an area border router. Then enter an OSPF Area ID. When a node acts as an Area Border Router (that is, when the node interfaces to trunks assigned to different OSPF areas), the node prefix OSPF Area ID is used to unambiguously assign addresses configured on that node to a particular OSPF area.
OSPF Area ID	If you enable OSPF Area Summary, enter the Area ID. The OSPF Area ID is used to assign addresses configured on the node to a particular OSPF area.
Admin Cost	Enter the administrative cost associated with this node prefix. When an SVC is being created, if more than one node in the network is found with the same node prefix, the call is routed to the node that has the lowest administrative cost associated with the node prefix.

5. Choose OK to return to the Set All Node Prefixes dialog box. The new entry appears in the *Defined Node Prefixes...* list box.
6. Choose Close to exit the dialog box.

Configuring Port Prefixes

The Set All Port Prefixes function enables you to define how calls are routed to the port. Port prefixes are also used for calling party screening.

To define a port prefix:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Port Prefixes. The Set All Port Prefixes dialog box (Figure 10-3) appears.

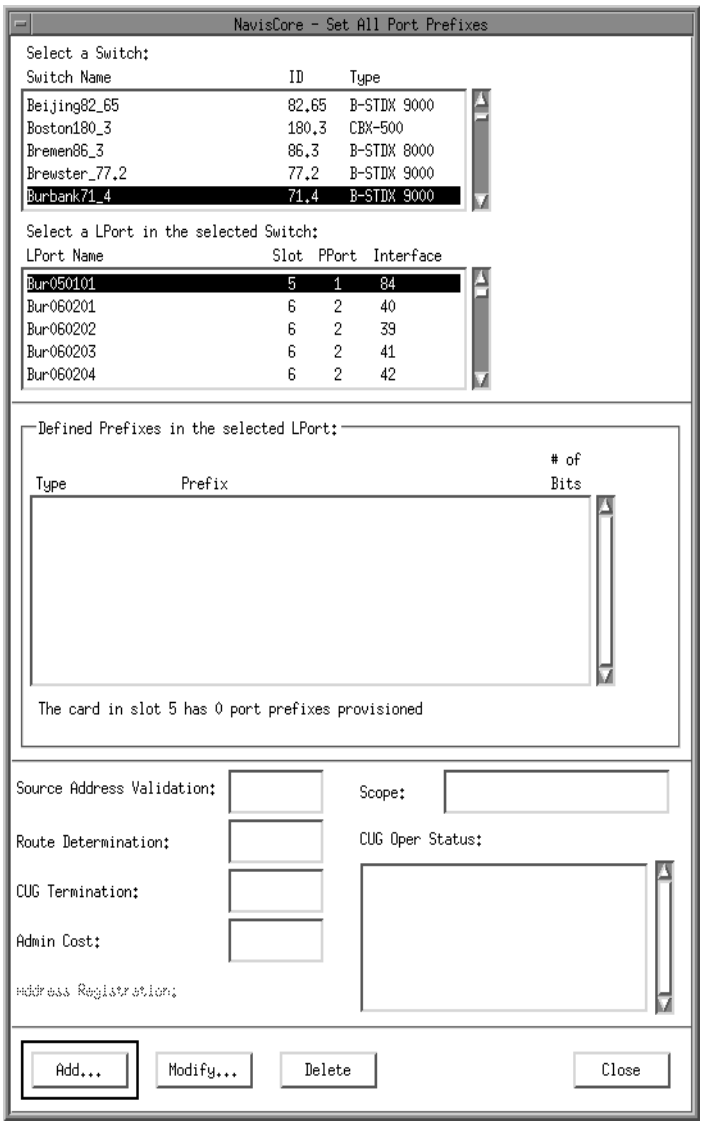

Figure 10-3. Set All Port Prefixes Dialog Box

- The *Select a Switch* list box displays all the switches that this NMS can access.
- The *Select a LPort...* list box displays the logical ports that are configured for the selected switch and their slot, physical port, and MIB interface numbers.
- The *Defined Prefixes...* list box displays the port prefixes that have already been defined on the selected logical port.

2. Select the switch for which to configure port prefixes.
3. Select the logical port for which to configure port prefixes.
4. Choose Add. The Add Prefix dialog box (Figure 10-4) appears.

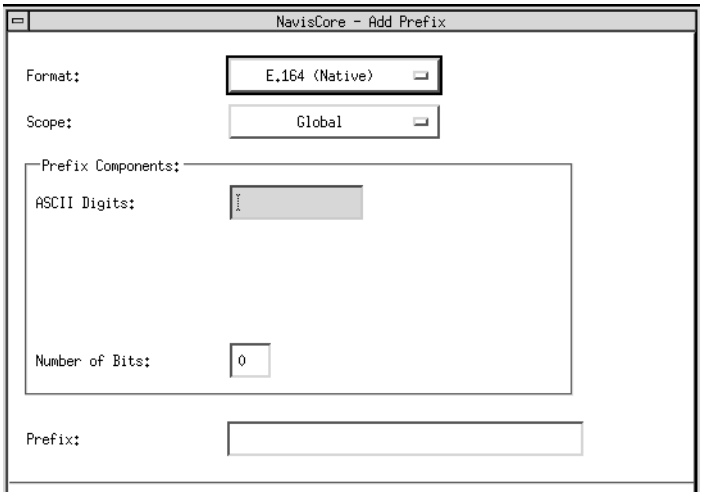

Figure 10-4. Add Prefix Dialog Box

5. Complete the Port Prefix fields as described in Table 10-5.

Table 10-5. Port Prefix Fields

Field	Action/Description
Format	Select the address format. Valid options include: • E.164 (Native) (default) • X.121
ASCII Digits	Enter all or part of the ASCII digits that represent the address. For example, enter 5085552600 (a standard 10-digit U.S. phone number) or enter a partial number (such as 508).
Number of Bits	As you type the address, the value in the Number of Bits field changes to indicate the number of address bits that are checked during call screening and call routing.
Prefix	The value that you enter in the ASCII digits field is converted to the ASCII hex values that represent each digit in the number. If you entered 508555260, it converts to 35308385353532363030. This value also appears in the Prefix column on the Set All Port Prefixes screen.
Source Address Validation	Select enable to validate the calling party address against the port prefix associated with the UNI/NNI port that received the call setup message. If you disable this option, this port prefix is not used to validate calling party addresses.
Route Determination	If enabled, the OSPF protocol uses this port prefix for route determination. If disabled, OSPF registration is not used.
CUG Termination	Select enable to use this prefix as part of a Closed User Group (CUG). Incoming and outgoing calls with a calling or called party address that matches this prefix are subject to CUG security checks. For more information about CUGs, see Chapter 11, "Closed User Groups."
Admin Cost	Enter the administrative cost associated with the port prefix. When an SVC is being created, if more than one port in the network is found with the same port prefix, the call is routed to the port in the network that has the lowest administrative cost associated with the port prefix.

6. Choose OK to return to the Set All Port Prefixes dialog box (Figure 10-3 on page 10-10). The new entry appears in the *Defined Prefixes...* list box.
7. Choose Close to exit the dialog box.

Defining Default Routes for Network-to-Network Connections

For ports being used as network-to-network connections, you can define a default route (which is automatically assigned 0x00 as its address with a length of 0 bits).

If the network receives a call and the called-party address does not match any port prefixes or addresses, it reroutes the call to the port on which the default route is defined. If more than one port has a default route defined, the administrative cost value is used to determine the port to which the call is routed.

When a default route is used, the switch provides partial protection from routing loops by preventing a call from being routed out the logical port on which it was received. It is important to note that, depending on the network topology, routing loops can still occur when multiple logical ports are provisioned with the default route.

You can define multiple default routes within a node or the network. The default route typically applies to network-to-network logical ports.

To define a default route:

1. From the Set All Port Prefixes dialog box (Figure 10-3 on page 10-10), select the switch and logical port on which to configure a default route. Choose Add. The Add Prefix dialog box shown in Figure 10-4 on page 10-11 appears.
2. In the Format field, select Default Route.
3. Enter the Administrative Cost for the default route on this logical port, then choose OK to return to the Set All Prefixes dialog box. The new Default Route entry appears in the *Defined Prefixes* list box.
4. Choose Close to exit the dialog box.

Configuring Port Addresses

To fully specify an address to be used for calling party screening, you can define SVC addresses on all the logical ports on that physical port. For native E.164 addresses, you enter the 1-15 digit E.164 address; for X.121 addresses, you enter the 1-14 digit X.121 address.

To define SVC port addresses:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Port Addresses. The Set All Port Addresses dialog box (Figure 10-5) appears.

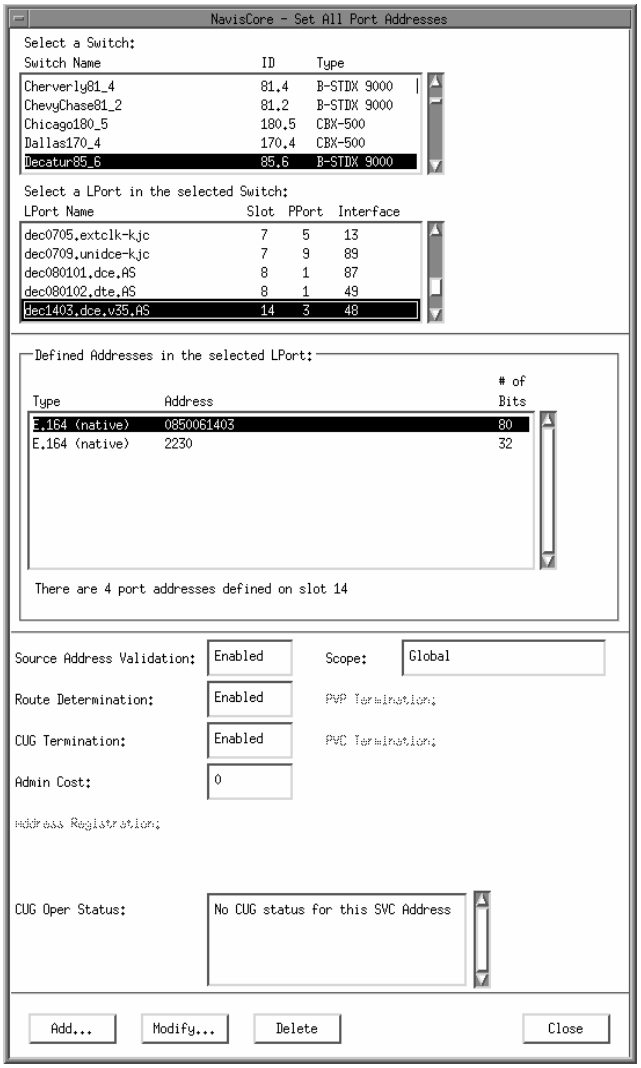

Figure 10-5. Set All Port Addresses Dialog Box

- The *Select a Switch* list box displays all switches this NMS can access.
- The *Select a LPort...* list box displays the logical ports that are configured for the selected switch and their slot, physical port, and MIB interface numbers.
- The *Defined Prefixes...* list box displays the port prefixes that have already been defined on the selected logical port.

2. Select the switch for which to configure SVC port addresses.
3. Select the logical port for which to configure SVC port addresses.
4. Choose Add. The Add Address dialog box appears.
5. Complete the port address fields described in Table 10-6.

Table 10-6. Port Address Fields

Field	Action/Description
Format	Select the address format. Valid options include: • E.164 (Native) (default) • X.121 For more information, see "Configuring Port Addresses" on page 10-14.
ASCII Digits	Enter all or part of the ASCII digits that represent the address. For example, enter 5085552600 (a standard 10-digit U.S. phone number) or enter a partial number (such as 508).
Number of Bits	As you type the address, the value in the Number of Bits field changes to indicate the number of address bits that are checked during call screening and call routing.
Address	The value that you enter in the ASCII digits field is converted to the ASCII hex values that represent each digit in the number. If you entered 508555260, it converts to 353038353535323630. This value also appears in the Prefix column on the Set All Port Prefixes screen.
Source Address Validation	Select enable to validate the calling party address against the port address associated with the UNI port that received the call setup message. If you disable this option, this address is not used to validate calling party addresses.
Route Determination	If enabled, the OSPF protocol uses this address for route determination.
CUG Termination	Select enable to use this address as part of a Closed User Group (CUG). Incoming and outgoing calls with a calling or called party address that match this address are subject to CUG security checks. For more information about CUGs, see Chapter 11, "Closed User Groups."

Table 10-6. Port Address Fields (Continued)

Field	Action/Description
Admin Cost	Enter the administrative cost associated with the port address. When an SVC is being created, if more than one port in the network is found with the same port address, the call is routed to the port in the network that has the lowest administrative cost associated with the port address.

6. Choose OK to return to the Set All Port Addresses dialog box (Figure 10-5 on page 10-14). The new entry appears in the *Defined Addresses in...* list box.
7. Choose Close to exit the dialog box.

Defining Network ID Parameters

This section describes how to add, delete, and modify network IDs. For more information about Network IDs, see page 10-6. See the following sections for configuration information:

- "Adding a Network ID" on page 10-17
- "Modifying a Network ID" on page 10-20
- "Modifying a Network ID" on page 10-20

Adding a Network ID

To add a network ID:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Port Network IDs. The Set All Port Network IDs dialog box (Figure 10-6) appears.

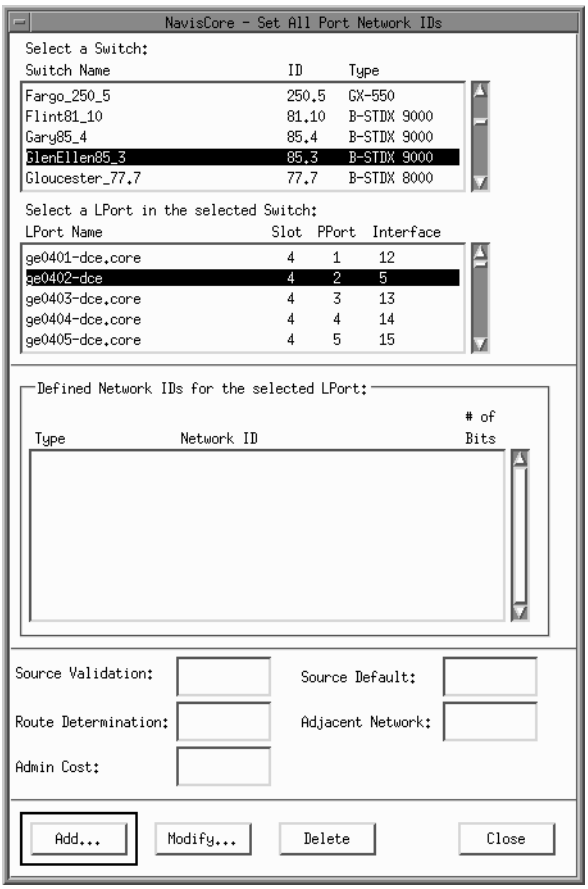

Figure 10-6. Set All Port Network IDs Dialog Box

Table 10-7 describes the dialog box fields.

Table 10-7. Set All Port Network IDs Fields

Field	Action/Description
Select a Switch	Displays the Switch Name, ID, and Type for all existing switches in the network.
Select an LPort in the selected Switch	Displays the LPort Name, Slot, PPort, and Interface for all ATM UNI and NNI, and Frame Relay UNI LPorts for the selected switch.
Defined Network IDs for the selected LPort	Displays the Type, Network ID, and # of Bits for all Network IDs created for the selected LPort.
Source Validation Route Determination Source Default Adjacent Network	Displays Enabled or Disabled for the selected Network ID.
Admin Cost	Displays the administrative cost for the selected Network ID.

2. To add a network ID, select a switch name from the Select a Switch list box and select a logical port for the selected switch in the Select an LPort in the selected Switch list box. Choose Add. The Add Network ID dialog box (Figure 10-7) appears.

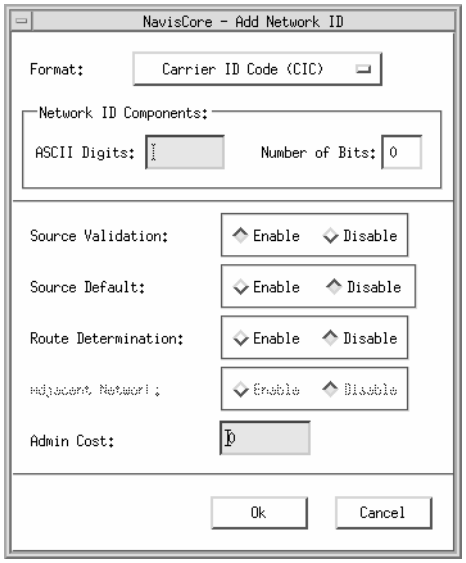

Figure 10-7. Add Network ID Dialog Box

3. Complete the Add Network ID fields described in Table 10-8.

Table 10-8. Add Network ID Fields

Field	Action/Description
Format	Select an ID format. Options include: *Carrier Identification Code (CIC)* *Data Network ID Code (DNIC)*
ASCII Digits	Enter a number between 0-9 for CIC or DNIC formats. CIC IDs are 1-8 digit values. DNIC IDs are 4 digit values.
Number of Bits	Displays the number of bits in the network ID.
Source Validation	Enable (*default*) or disable source validation for this network ID. When enabled, a signaled TNS may be screened against this network ID. If you enable this field, route determination is disabled and the adjacent network parameter becomes inactive.
Source Default	Enable or disable (*default*) source default for this network ID. Only one network ID on each port may have this attribute. When enabled, this network ID represents the preferred IXC for user calls originating on this logical port that do not signal a transit network selection.
Route Determination	Enable or disable (*default*) route determination for this network ID. If you enable this field, source validation is disabled and the source default parameter becomes inactive. If enabled, the OSPF protocol uses this network ID for route determination.
Adjacent Network	Enable or disable (*default*) adjacent network for this network ID. This information is used by billing. Only one network ID on each logical port may have this attribute. When enabled, this network ID is considered to be the adjacent network (as opposed to another network reachable through the actual adjacent network). This adjacent network ID will not be signaled from this logical port.

Table 10-8. Add Network ID Fields (Continued)

Field	Action/Description
Admin Cost	Enter an administrative cost between 0 - 65535 for this network ID. The default is 0.

4. Choose Ok to add the network ID and exit the dialog box.
5. Choose Cancel to exit the dialog box.

Modifying a Network ID

To modify an existing network ID:

1. From the Set All Port Network IDs dialog box (Figure 10-6 on page 10-17), select the desired network from the Defined Network IDs for the selected LPort list box.
2. Choose Modify to display the Modify Network ID dialog box. This dialog box contains the same fields that are described in Figure 10-7 on page 10-18.
3. Modify the network ID parameters. (You can not modify the format or the ASCII Digits fields.)
4. Choose OK to modify the network ID and exit the dialog box.
5. Choose Cancel to exit the dialog box.

Deleting a Network ID

To delete an existing network ID:

1. From the Set All Network IDs dialog box (Figure 10-6 on page 10-17), select the desired network from the Defined Network IDs for the selected Lport list box.
2. Choose Delete. The selected network is deleted from the list.
3. Choose Cancel to exit the dialog box.

11

Closed User Groups

A closed user group (CUG) is a division of all SVC network users into logically linked groups of users. Members of the same CUG have particular calling privileges that members of different CUGs may not have. CUGs form one level of security between users of a network, allowing only those users who are members of the CUG to set up calls to each other. Information about CUG membership and rules is available throughout the network.

A CUG is comprised of a set of rules called members. These rules represent SVC port addresses and prefixes for which you have enabled the CUG termination option (refer to Table 10-6 on page 10-15). You configure CUG member rules in either X.121 or E.164 address format. When you configure a member rule, you can replace some digits with the * or ? UNIX wild card characters. If a member rule does not contain a wild card character, it maps to a specific network user. If the member rule includes a wild card, then this member can potentially map to multiple network users.

Throughout this document, most address descriptions use the term “SVC address.” Unless otherwise noted, the term SVC address is used interchangeably with term “SVC prefix.”

About CUG Member Rules

CUG member rules correspond to SVC addresses. You can enter a rule as a UNIX-style expression. You can use the * as a wildcard to replace zero (0), one, or more digits, or the ? as a wild card to replace a single digit. You can only use the * once in a string. Keep in mind that an X.121 digit is 4 bits and an E1.64 digit is 8 bits.

The following examples show how you can use wild cards to represent multiple E.164 addresses.

Example	Description
1508952*	This CUG includes all numbers using area code 508 and exchange number 952.
1508952148?	This CUG includes all numbers using area code 508, exchange number 952, and an extension starting with 148 (i.e., 1480 – 1489).

When you define a CUG member, these addresses define the *member value* for the CUG member rule. Each CUG member rule is defined by an ASCII name, an address type (either E.164 or X.121), and the CUG member value (rule).

Defining Incoming and Outgoing Access

In addition to defining CUG member address values, you can also define the incoming and outgoing access attributes that complete the CUG member rule.

The *incoming access* (IA) attribute enables you to define how a CUG member handles calls coming from other CUGs or non-CUG users. A user mapping to a CUG member with incoming access enabled can receive calls coming from non-CUG users as well as calls coming from other CUGs. If you disable incoming access, the CUG member can only receive calls from other members of the same CUG.

The *outgoing access* attribute (OA) enables you to define how a CUG member handles calls to other CUGs and non-CUG users. A user mapping to a CUG member with outgoing access enabled can make calls to other CUGs and non-CUG users. If you disable outgoing access, the CUG member can only make calls to other members of the same CUG.

Member Rule Example

You define the following CUG member rule:

Member Rule Name	rule1
Member Value/Type	1508* (E.164)
Incoming Access	Y
Outgoing Access	N

This member rule applies to E.164 addresses beginning with digits 1508. Users that map to this rule can receive calls from members of their own CUG, members of other CUGs, and non-CUG users (incoming access is enabled), but they cannot make calls outside their own CUG.

Developing Closed User Groups

For each CUG you create, you can assign up to 128 different member rules; you can use an individual member rule in up to 16 different CUGs. In this way, a CUG is made up of all users that map to the addresses that these rules define. You can configure up to 1024 CUGs per switch.

When you create a CUG (“CUG A”), the attributes you configure for each CUG member rule (“Rule1”) that you associate with the CUG define how the CUG handles calls between members. For example, if you enable the *incoming calls barred* (ICB) attribute for Rule1, users that map to Rule1 cannot receive calls from other CUG A members. Conversely, disable ICB to allow users that map to Rule1 to receive calls from other CUG A members.

If you enable the *outgoing calls barred* (OCB) attribute for Rule1, users that map to Rule1 cannot make calls to other CUG A members. Conversely, disable OCB to allow users that map to Rule1 to make calls to other CUG A members.

Using CUGs in the Network

The following example illustrates how you can implement CUGs in your network.

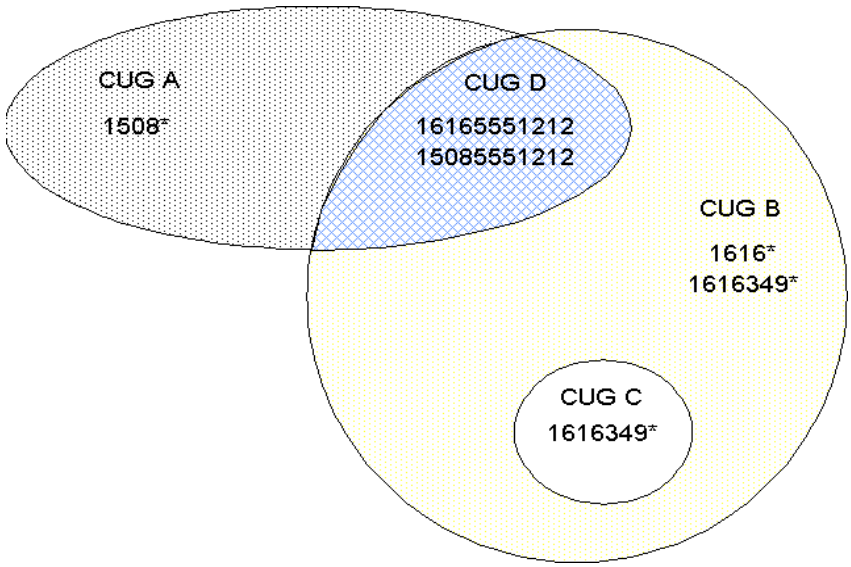

Figure 11-1. Implementing CUGs

The CUGs used in this example represent the following:

- CUG A: Business Unit A
- CUG B: Business Unit B
- CUG C: Independent entity within Unit B
- CUG D: Joint venture between Units A and B

For each of these CUGs, the following table defines the ICB and OCB attributes and member rules. Each member rule is made up of an expression that represents an E.164 address and an incoming access (IA) and outgoing access (OA) attribute.

Table 11-1. ICB/OCB Attributes and Member Rules

	ICB	OCB	Member Rules	IA	OA
CUG A	No	No	1508*	No	No
CUG B	No Yes	No Yes	1616* 1616349*	No No	Yes No
CUG C	No	No	1616349*	No	No
CUG D	No No	No No	16165551212 15085551212	No Yes	Yes No

Call Setup Examples

- A call is made from 15085551212 to 16165551212:
  - 15085551212 (IA enabled): Address belongs to CUG A and CUG D
  - 16165551212 (OA enabled): Address belongs to CUG B and CUG D

  Result: Call succeeds because both addresses belong to CUG D.

- A call is made from 16163498888 to 16165551212:
  - 1616349: Address belongs to CUG B (ICB, OCB enabled) and CUG C
  - 16165551212 (OA enabled): Address belongs to CUG B and CUG D

  Result: Although both addresses belong to CUG B, the call fails because the outgoing calls barred (OCB) attribute is enabled on CUG B for member 1616349*. Users mapping to matching rule 1616349* cannot make calls to other CUG B members.

- A call is made from 12035551212 to 15085551212:
  - The address 12035551212 does not belong to any CUG.
  - 15085551212 (IA enabled): Address belongs to CUG A and CUG D

  Result: Call succeeds because the incoming access (IA) attribute enabled is enabled for 15085551212. This member rule allows users mapped to 15085551212 to receive calls from non-CUG users.

Configured Addresses and CUG Membership

Using the CUG design depicted in Figure 11-1 on page 11-4, Table 11-2 illustrates how a single configured address can match multiple member rules, and can belong to more than one CUG.

Table 11-2. Configured Address and Corresponding CUG Membership

Address	OA	IA	CUG	ICB	OCB
15085551212	N	Y	A	N	N
			D	N	N
16165551212	Y	N	B	N	N
			D	N	N
15082178989	N	N	A	N	N
16161234567	Y	N	B	N	N
16163498888	Y	N	B	Y	Y
			C	N	N

Member rules that specify an address prefix only can simplify call routing since the logical port only needs to check the address prefix digits to route the call. However, CUG membership must be recalculated at call time if the port to which this address is routed contains other CUGs with member rules that begin with the digits 1616.

For example, if a CUG contains a member rule that uses a prefix format (i.e.,1616*) as well as other member rules that are more specific (1616349*), you are likely to encounter performance issues due to address ambiguity.

The more specific you make the CUG member rules, the more quickly CUG membership can be determined.

Configuring Closed User Groups

Use the following sequence to configure CUGs. Remember that each member rule should correspond to at least one SVC address.

Step 1. Create SVC addresses and enable CUG termination (refer to page 10-15).

Step 2. Define the CUG member rules that represent the member addresses and call access (page 11-7).

Step 3. Define the CUG names (page 11-10).

Step 4. Associate CUG members to specific CUGs. You can also modify call access attributes for a specific CUG (page 11-11).

Defining CUG Members

A CUG member is defined by a rule that matches one or more port addresses/prefixes and attributes that specify incoming and outgoing call access. Once you define these members, you can associate them with specific CUGs.

To define a CUG member:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All SVC CUG Members. The following dialog box appears (Figure 11-2 on page 11-8).

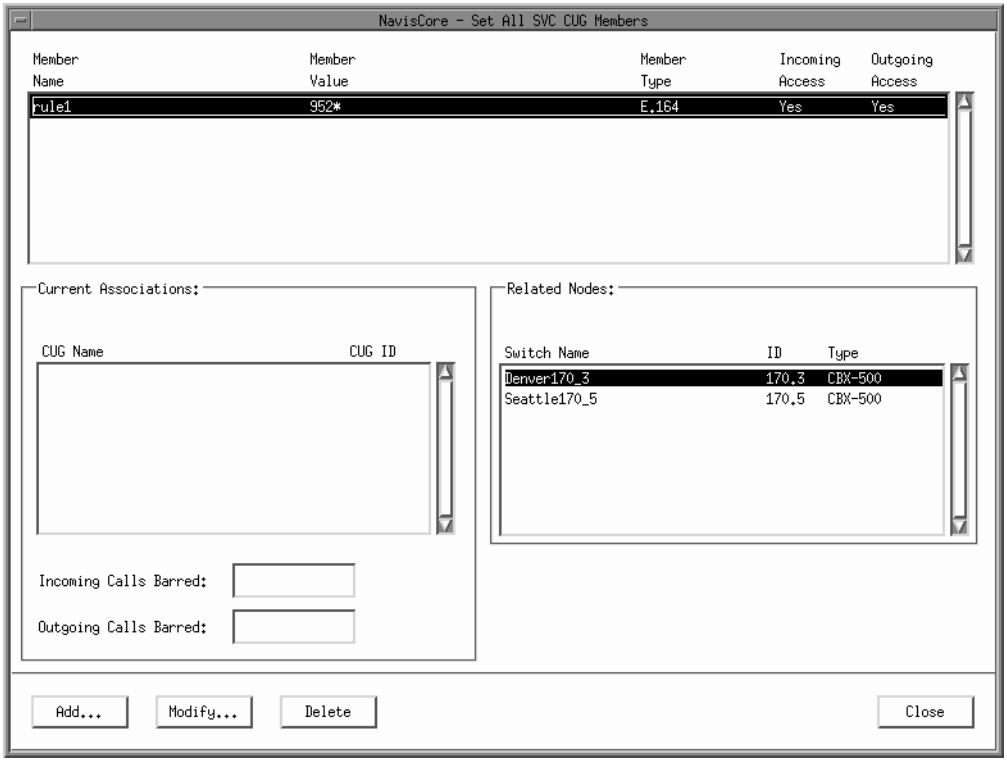

Figure 11-2. Set All SVC CUG Members Dialog Box

This dialog box provides a list of previously defined CUG members and their rules. The CUG Name list contains the name of the CUG(s) to which the selected member name is associated.

- To modify an existing member name, select a name from this list and choose Modify.
- To delete an existing member name, select a name from this list and choose Delete

2. Choose Add to define a CUG member. The Add SVC CUG Member dialog box (Figure 11-3) appears.

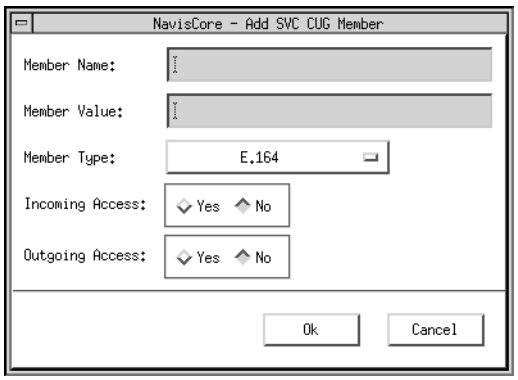

Figure 11-3. Add SVC CUG Member Dialog Box

3. Configure the member attributes described in Table 11-3.

Table 11-3. Add SVC CUG Member Fields

Field	Description
Member Name	Enter a name (up to 32 characters).
Member Value	Enter the CUG member rule using the guidelines on page 11-2. Do not enter more than 15 characters for an E.164 address or more than 14 characters for an X.121 address.
Member Type	Select X.121 or E.164.
Incoming Access	This attribute specifies how incoming calls from non-CUG users or users of a different CUG are handled. • Select Yes to accept calls from users that do not belong to the same CUG. • Select No (default) to reject calls from users that do not belong to the same CUG.
Outgoing Access	This attribute specifies how outgoing calls to non-CUG users or users of a different CUG are handled. • Select Yes to allow calls to users not belonging to the same CUG. • Select No (default) to block calls to users not belonging to the same CUG.

4. When you finish, choose OK.

5. You can use these fields to define additional members, or choose Cancel to exit this dialog box.

Defining a Closed User Group

Next, set up the CUGs for your network. This is a simple process of supplying a name for each CUG. NavisCore supports up to 1024 CUGs per switch.

To create a CUG:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All SVC CUGs. The Set All SVC CUGs dialog box Figure 11-4) appears.

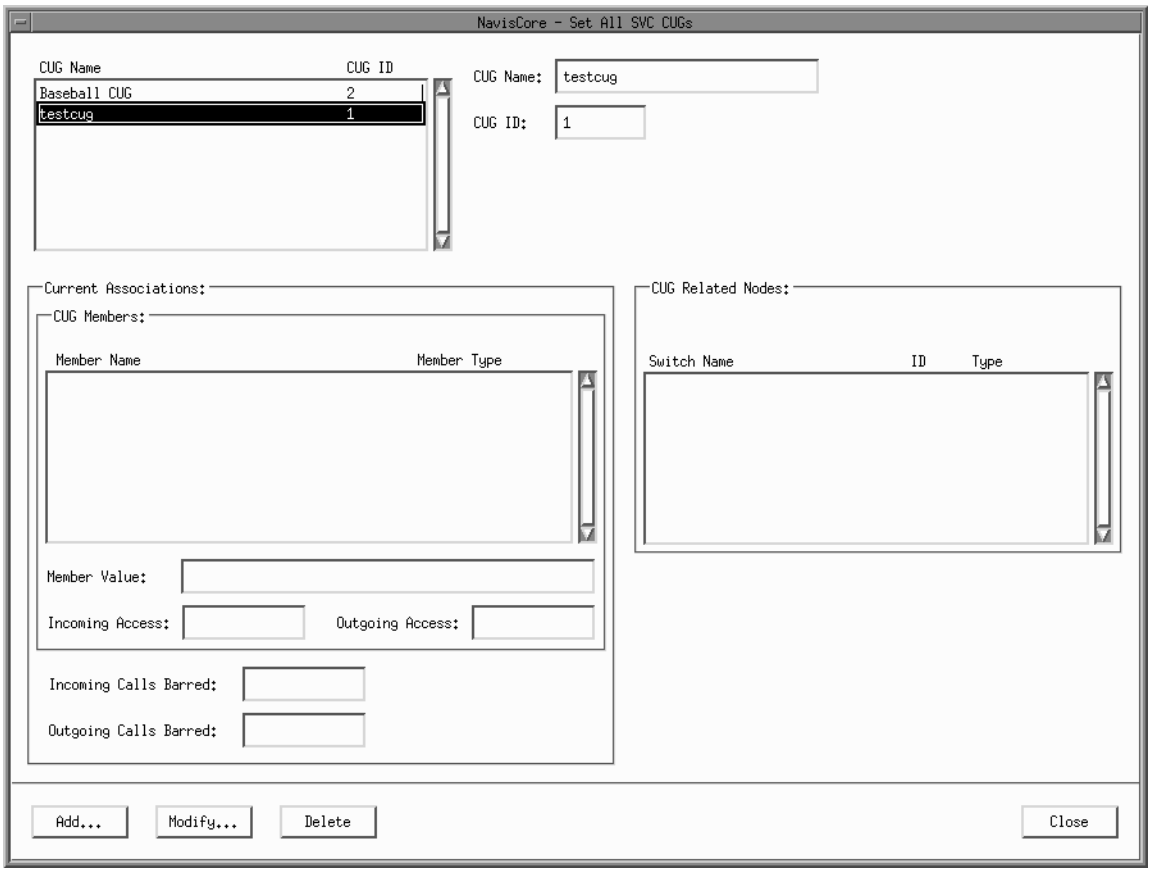

Figure 11-4. Set All SVC CUGs Dialog Box

2. This dialog box provides a list of previously configured CUG names as well as a listing of members for each CUG you select.
   - To modify an existing CUG, select a name from this list and choose Modify.
   - To delete an existing CUG, select a name from this list and choose Delete.
3. Choose Add to create a new CUG. The Add SVC CUG dialog box (Figure 11-5) appears.

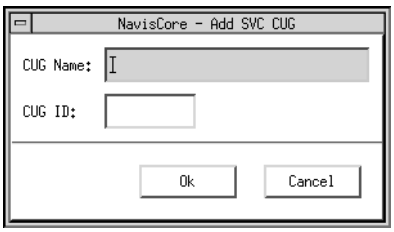

Figure 11-5. Add SVC CUG Dialog Box

4. Enter a CUG name (up to 32 characters). The NMS assigns a CUG ID.
5. Choose OK.

Assigning Member Rules to CUGs

To complete the CUG definition process, you need to assign member rules to each CUG. You can assign up to 128 members per CUG. You can assign each member to as many as 16 CUGs.

To assign members to a CUG,

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All SVC CUGs. The Set All SVC CUGs dialog box appears (Figure 11-4).
2. From the CUG Name list, select the CUG to which you want to add members.

3. Choose Modify. The Modify CUG dialog box (Figure 11-6) appears.

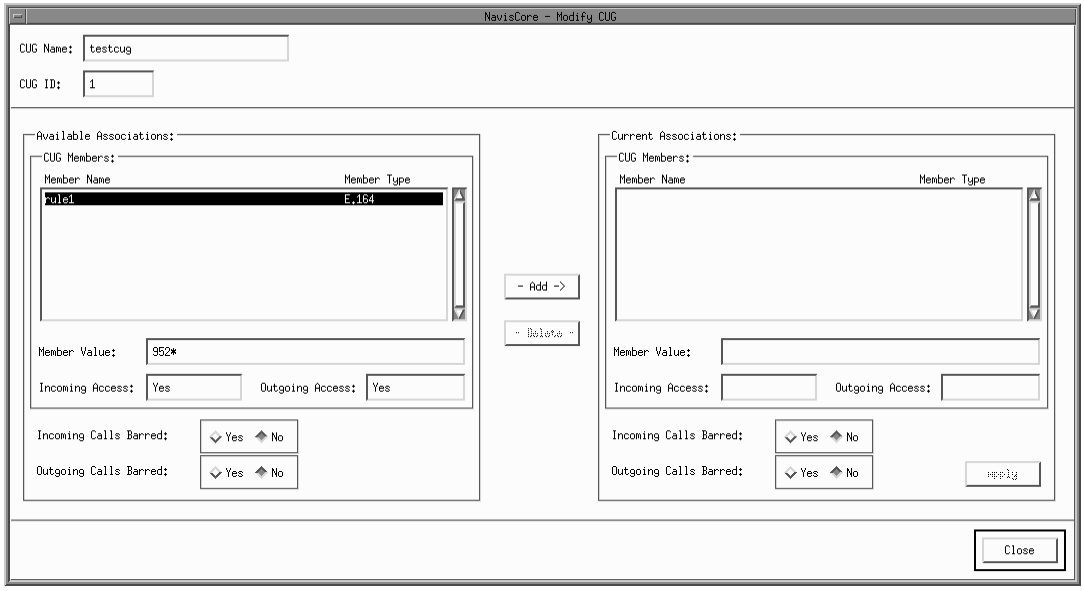

Figure 11-6. Modify CUG Dialog Box

This dialog box displays the current list of CUG member rules with Current Associations. It also provides a list of member names you can associate with this CUG.

4. From the list of Available Associations, select the member you want to associate with this CUG and specify Incoming Calls Barred and Outgoing Calls Barred:

   Incoming Calls Barred — Specifies how incoming calls from the same CUG are handled. Select Yes to reject calls from users of the same CUG. Select No (default) to allow calls from users of the same CUG.

   Outgoing Calls Barred — Specifies how outgoing calls to the same CUG are handled. Select Yes to block calls to users of the same CUG. Select No (default) to allow calls to users of the same CUG.

5. Choose Add. The member name appears in the Current Associations list. All SVC addresses and prefixes that match the member rule take on the attributes specified for this CUG.

6. To associate additional member names, repeat Step 4 and Step 5. When you finish, choose Close to exit this dialog box.

Modifying Call Access for CUG Members

Use the following steps to modify the incoming and outgoing call access for an existing CUG member.

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All SVC CUGs. The Set All SVC CUGs dialog box appears (Figure 11-4).
2. From the CUG Name list, select the CUG that contains this member.
3. Choose Modify. The Modify CUG dialog box appears (Figure 11-6 on page 11-12).
4. Select the CUG member name from the Current Associations list.
5. Use the instructions on page 11-9 to modify incoming and outgoing call access.
6. Choose Apply.
7. Choose Close to exit this dialog box.

12

Port Security Screening

The Port Security Screening feature ensures that your network cannot be compromised by unauthorized SVC access. You do this by creating screens that can allow/disallow incoming and outgoing SVCs. You configure each screen with the following information:

SVC direction — Screen either ingress (incoming) or egress (outgoing) SVCs.

Screen type — Pass or block SVCs according to the configured screen.

Address type — Any address type used in a public or private UNI. This includes E.164 and X.121 formats for calling and called party addresses, and Network Service Access Point (NSAP) AESA format for calling and called subaddresses.

Matching information — Address criteria that either allows or disallows the SVC.

Once you develop a set of screens, you can apply them to any UNI or NNI logical port in your network. You can use a maximum of 16 different screens per port. Using these screens, the port checks every SVC it receives and/or sends for the matching criteria specified in the screen(s). If the SVC meets the matching criteria specified in at least one of these screens, the port either passes or blocks that SVC according to the security screen design.

Implementing Port Security Screening

Although you can apply multiple security screens to a single logical port, the decision as to whether an SVC is passed or blocked is made based on the combined effects of the following:

- The default ingress/egress screen mode for the logical port.
- The security screens you assign to this logical port.
- The incoming/outgoing SVC address criteria defined in the security screen.

Default Screens

For each logical port, you configure default screen criteria that specifies the behavior of any SVC on this port. You can use security screens on both ingress user ports (which represent SVC originating endpoints) or egress user ports, which in turn represent SVC terminating endpoints. The default screens enable you to quickly override the security screens you assign to the logical port; use the default screens to either pass or block all incoming or outgoing SVCs.

Table 12-1 describes the default ingress and egress security screen options. These defaults represent the port screen activation parameters.

Table 12-1. Default Screens

Default	Value	Description
Ingress Screen Mode	All Screens	All ingress screens you apply to this port are used to determine whether an incoming SVC is passed or blocked.
	Default Screen (*default*)	Disables the ingress security screens applied to this port. Incoming SVCs are screened according to how you set the Default Ingress Screen.
Default Ingress Screen	Pass (*default*)	If you set the Ingress Screen Mode to Default Screen, all incoming SVCs to this port are passed; if it is set to All Screens, all incoming SVCs are passed, unless one of the ingress security screens assigned to this port blocks the SVC.
	Block	If you set the Ingress Screen Mode to Default Screen, all incoming SVCs to this port are blocked; if it is set to All Screens, all incoming SVCs are blocked unless one of the ingress security screens assigned to this port passes the SVC.

Table 12-1. Default Screens (Continued)

Default	Value	Description
Egress Screen Mode	All Screens	All egress screens you apply to this port are used to determine whether an outgoing SVC is passed or blocked.
	Default Screen (*default*)	Disables the egress security screens applied to this port. Outgoing SVCs are screened according to the Default Egress Screen.
Default Egress Screen	Pass (*default*)	If you set the Egress Screen Mode to Default Screen, all outgoing SVCs from this port are passed; if it is set to All Screens, all outgoing SVCs are passed, unless one of the egress security screens assigned to this port blocks the SVC.
	Block	If you set the Egress Screen Mode to Default Screen, all outgoing SVCs from this port are blocked; if it is set to All Screens, all outgoing SVCs are blocked, unless one of the egress security screens assigned to this port passes the SVC.

Security Screens

The security screens you assign to a logical port represent exceptions to the default screens. You can assign up to 16 security screens per logical port. Once you assign security screens to a port and set the ingress/egress screen mode to All Screens, the logical port uses these security screens to screen SVCs that match the criteria they specify.

You define a security screen based on two attributes: SVC direction and screen type. SVC direction defines the SVCs to which this screen applies, either ingress (incoming) or egress (outgoing). The screen type attribute determines whether or not the port passes or blocks these SVCs.

About Security Screen Addresses

To provide a more detailed level of SVC screening, you can specify either an E.164 or X.121-style address for calling or called addresses, or an NSAP AESA-style address for calling or called subaddresses. You can enter the entire address as a number, or enter a UNIX-style expression using wild cards. When you use a UNIX expression, a single screen can match multiple endpoint addresses. Use the ? wild card to replace a single digit or the * wild card to replace one or more digits. You can only use the * once in a string. See “Configuring Node Prefixes” on page 10-6 for more information about addressing.

The following examples show how you can use a UNIX expression to represent an E.164 North American address.

Example	Description
1508952*	This screen applies to all numbers using area code 508 and exchange number 952.
1508952148?	This screen applies to all numbers using area code 508, exchange number 952, and an extension starting with 148 (i.e., 1480 – 1489).
150895?*5?	This screen applies to all numbers using area code 508, with an exchange number value of 950 – 959. The number 5 must appear as one digit from the end of the address.

Table 12-2 describes some examples using the port security screens.

Table 12-2. Security Screens

SVC Direction	Screen Type	Calling Address	Calling Subaddress	Called Address	Called Subaddress	Description
Ingress	Pass	Ignore	Ignore	1800* Type: E.164	Ignore	Pass all incoming calls to 1800 numbers.
Ingress	Block	Ignore	Ignore	1800* Type: E.164	Ignore	Block all incoming calls to 1800 numbers.
Egress	Block	Ignore	Ignore	* Type: E.164	Ignore	Block all outgoing calls with E.164 called addresses.
Egress	Block	15089700705 Type: E.164	Ignore	1908870* Type: E.164	Ignore	Block all calls to called address 1908870* from calling address 15089700705.

Port Security Screening Sample Configuration

Once you assign security screens to a logical port, if you set the ingress and egress screen modes to All Screens (Figure 12-3 on page 12-12), the port checks incoming/outgoing SVCs for the matching criteria specified in each assigned screen. If an SVC meets the criteria specified in at least one screen, then the SVC is screened according to the action this screen recommends. The SVC is further checked for the matching criteria of this screen's default behavior. If it meets the matching criteria specified in at least one of these screens, then the SVC exhibits the default behavior (either pass or block).

Although you can apply multiple screens to a single port, the decision on whether the port should block or pass an SVC is made based on:

- The combined effect of the default screens specified for the logical port
- The security screens you assign to that port
- The matching address criteria defined in each screen (if applicable)

If you set the ingress/egress screen mode to Default Screens, the port does not check SVCs for the matching criteria specified in an assigned security screen. It takes the action (either pass or block) specified in the Default Screen.

The following example provides a logical port configuration that blocks all incoming SVCs, except incoming 1800 SVCs, with one exception. You want to block all incoming SVCs that contain the 234 exchange number.

Logical Port Configuration Examples

1. For the logical port, configure the following default screen:

Ingress Screen Mode:	All Screens
Default Ingress Screen:	Block

Setting the default ingress screen to *block* enables you to block all incoming SVCs on this port by default; setting the ingress screen mode to *all screens* enables the port to screen SVCs based on the ingress security screens you assign.

2. Create and assign two security screens.

– The following screen passes all incoming 1800 SVCs:

Screen Name:	pass_in_800
SVC Direction:	Ingress
Screen Type:	Pass
Calling Address:	Ignore
Calling Subaddress:	Ignore
Called Address:	Type: E.164 1800*
Called Subaddress:	Ignore

– The following screen blocks all SVCs from the 234 exchange:

Screen Name:	blk_234_exchg
SVC Direction:	Ingress
Screen Type:	Block
Calling Address:	Ignore
Calling Subaddress:	Ignore
Called Address:	Type: E.164 1???234*
Called Subaddress:	Ignore

Summary

As you begin to design port security screening features for your network, keep the following points in mind:

1. Configure the default screen for a logical port. This default mode determines whether or not to pass or block SVCs from certain addresses. The previous example blocks all incoming SVCs for the logical port. You can quickly revert back to the default mode if necessary.
2. Configure and assign the security screen exceptions. The previous example passes all incoming 1800 SVCs.
3. Configure and assign any exceptions to these screen. The previous example specifically blocks incoming SVCs from the 234 exchange; this includes incoming SVCs from 1800234*.

Configuring Port Security Screening

Use the following sequence to configure port security screening.

Step 1. Configure logical ports (see Chapter 3, “Configuring Frame Relay Logical Ports”).

Step 2. Configure SVCs (see Chapter 10, “Configuring Switched Virtual Circuit (SVC) Parameters”).

Step 3. Create a set of security screens (see page 12-8).

Step 4. Define the logical port security screening defaults. If necessary, assign the security screens that provide exceptions to these defaults (see page 12-11).

Creating Port Security Screen Definitions

To create a security screen:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Port Security Screens. The Set All Port Security Screens dialog box (Figure 12-1) appears.

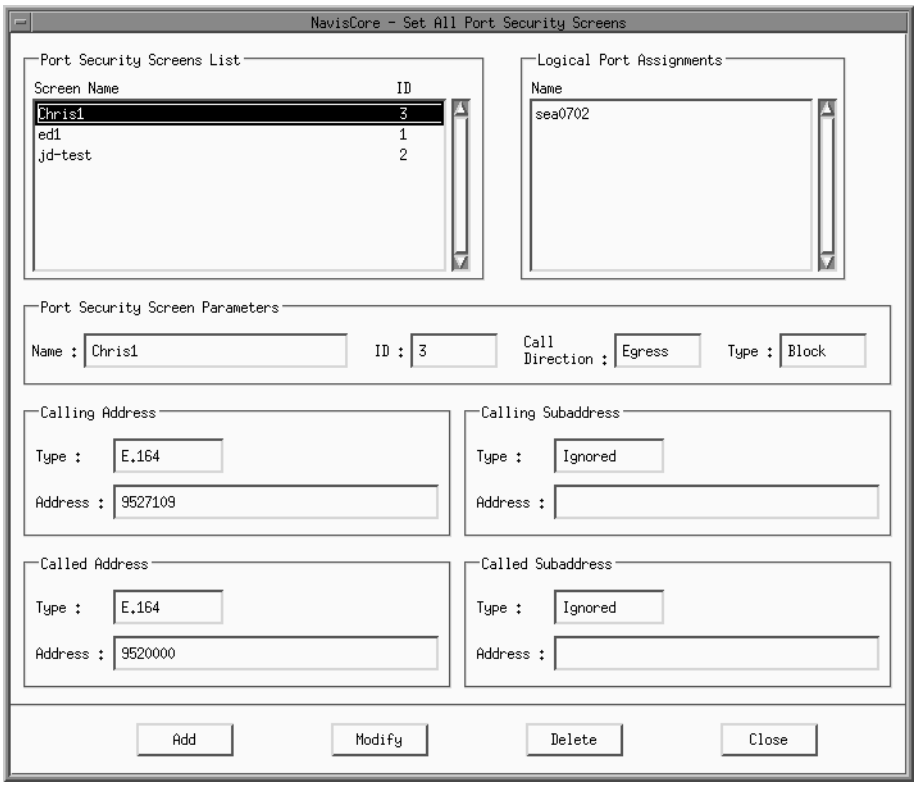

Figure 12-1. Set All Port Security Screens Dialog Box

This dialog box displays a list of previously configured security screens. It provides the configured parameters for each screen you select from the Port Security Screens List.

- To modify an existing screen, select a screen name and choose Modify.
- To delete an existing screen, select a screen name and choose Delete.

2. Choose Add to create a new screen. The Adding Port Security Screens dialog box (Figure 12-2) appears.

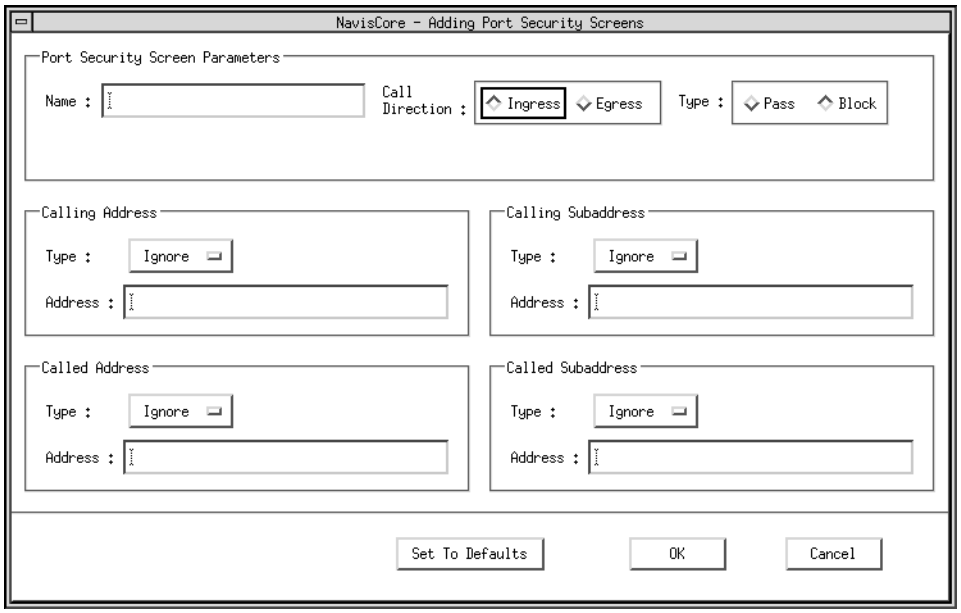

Figure 12-2. Adding Port Security Screens Dialog Box

3. Complete the dialog box fields as described in Table 12-3.

Table 12-3. Adding Port Security Screens Fields

Field	Action/Description
Name	Enter a name (up to 32 characters) for this security screen.
Call Direction	The screen you configure is only applied to these SVCs. *Ingress* – (Default) Screen incoming SVCs. *Egress* – Screen outgoing SVCs.
Type	Select the Type of screen. This determines the action this screen performs. *Block* – (Default) Blocks all SVCs that match the criteria. *Pass* – Passes all SVCs that match the criteria.
Calling Address	Configure the Calling Address: *Type* – Select the address type, either E.164 or X.121. Select Ignore (default) if the screen does not use this parameter. *Address* – Enter the address screen using the guidelines on page 12-3. Enter up to 15 characters for an E.164 address; enter up to 14 characters for an X.121 address.

Table 12-3. Adding Port Security Screens Fields (Continued)

Field	Action/Description
Calling Subaddress	Configure the Calling Subaddress. This parameter provides an optional level of screening. *Type* – Select AESA. Select Ignore (default) if the screen does not use this parameter. *Address* – Enter the address screen (up to 40 characters) using the guidelines on page 12-3.
Called Address	Configure the Called Address: *Type* – Select the address type, either X.121 or E.164. Select Ignore (default) if the screen does not use this parameter. *Address* – Enter the address screen using the guidelines on page 12-3. Enter up to 15 characters for an E.164 address; enter up to 14 characters for an X.121 address.
Called Subaddress	Configure the Called Subaddress. This parameter provides an optional level of screening. *Type* – Select AESA. Select Ignore (default) if the screen does not use this parameter. *Address* – Enter the address screen (up to 40 characters) using the guidelines on page 12-3.

4. Choose OK to create the new screen.
5. The Adding Port Security Screen dialog box (Figure 12-2 on page 12-9) is designed to allow you to create several screens in a single session. To create additional screens, repeat Step 3 and Step 4 on page 12-10. Choose *Set To Defaults* to retrieve the default values if necessary.
6. When you finish creating your screens, choose Cancel to exit this dialog box.

Assigning Security Screens to Logical Ports

Once you create the security screens, you must modify existing logical ports to assign these screens to the individual logical ports. The default security screens you configure for each logical port enable you to quickly pass or block incoming or outgoing SVCs, without having to remove or modify the screen you have applied. For information about reverting back to the default security screen, see "Activating Default Screens" on page 12-13.

You also have the option of assigning several different security screens to this port, but configuring them as "inactive." You can then activate them as necessary, at a later time. For more information, see "Activating and Deactivating Security Screens" on page 12-14.

To assign security screens to a port:

1. Use the instructions on page 3-2 to access the Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-3).
2. Select the logical port to which you will assign a screen and choose Modify.
3. Review the logical port type and ID and choose OK. The Modify Logical Port dialog box reappears.

4. From the Select:Options: menu, select Screen Assignments and choose Set. The Assigning and Activating Port Security Screens dialog box (Figure 12-3) appears.

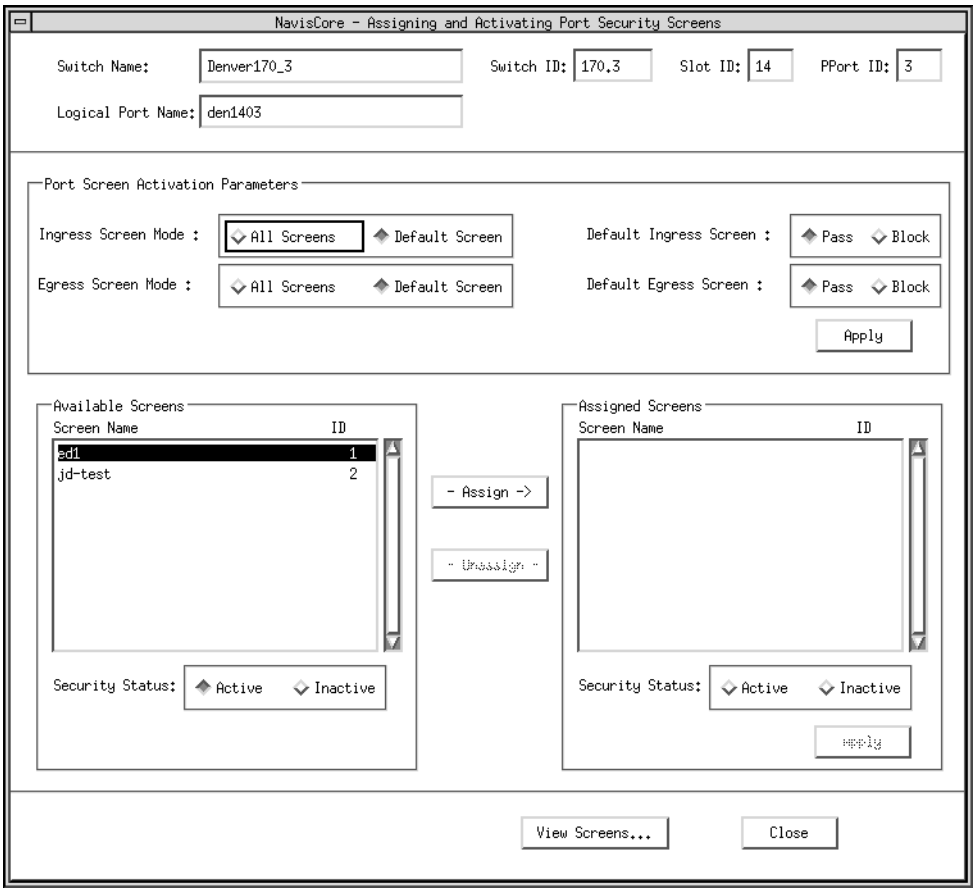

Figure 12-3. Assigning and Activating Port Security Screens

5. See Table 12-1 on page 12-2 to configure the Port Screen Activation Parameters to meet your network needs.
6. Choose Apply to set the Port Screen Activation Parameters.
7. The Available Screens list provides the list of security screens you can assign to this port. Select the name of the screen you want to assign.

8. The Security Status of the screen you select defaults to Active. Using the Active Security Status, after you choose Apply the logical port begins screening SVCs according to the rules of this screen.

   To assign a screen to this logical port without making it active immediately, select Inactive and choose Apply.

> You can choose View Screens to view the parameters configured for the screen you want to use.

9. Choose Assign to assign a screen to this logical port. The screen name appears in the Assigned Screens list.
10. The Assigning and Activating Port Security Screens dialog box (Figure 12-3 on page 12-12) is designed to allow you to assign several screens in a single session. To create additional screens, repeat Step 7 through Step 9. (You can assign up to 16 screens per logical port.)
11. When you finish creating your screens, choose Close to exit this dialog box.

Deleting Security Screen Assignments

Use the following steps to remove a security screen assignment for a logical port:

1. Use Step 1 through Step 4 starting on page 12-13 to access the Assigning and Activating Port Security Screens dialog box (Figure 12-3 on page 12-12).
2. Review the list of Assigned Screens and select the screen.
3. Choose Deassign. This screen should now appear in the Available Screens list.

Activating Default Screens

Use the following steps to activate the default screening parameter(s) to temporarily override assigned security screens.

1. Use the instructions on page 3-2 to access the Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-3).
2. Select the logical port for which you will activate the default screen(s) and choose Modify.
3. Review the logical port type and ID and choose OK. The Modify Logical Port dialog box reappears.
4. From the Select:Options: menu, select Screen Assignments and choose Set. The Assigning and Activating Port Security Screens dialog box appears (Figure 12-3 on page 12-12).

5. Review the information configured in the Port Screen Activation Parameters group box. See Table 12-1 on page 12-2 if you need information to modify these parameters.

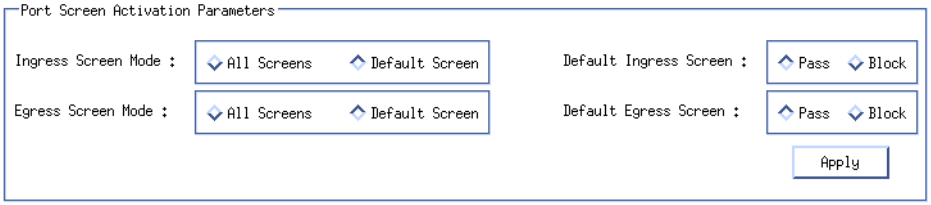

Figure 12-4. Port Screen Activation Parameters Group Box

6. Choose Apply to activate the default screen.
7. Choose Close to exit this dialog box.

Activating and Deactivating Security Screens

Use the following steps to activate or deactivate a security screen according to your network needs:

1. Use the instructions on page 3-2 to access the Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-3).
2. Select the logical port for which you will change the security status and choose Modify.
3. Review the logical port type and ID and choose OK. The Modify Logical Port dialog box reappears.
4. From the Select:Options: menu, select Screen Assignments and choose Set. The Assigning and Activating Port Security Screens dialog box appears (Figure 12-3 on page 12-12).
5. In the Assigned Screens list, select the screen and modify the Security Status as necessary (Active or Inactive).

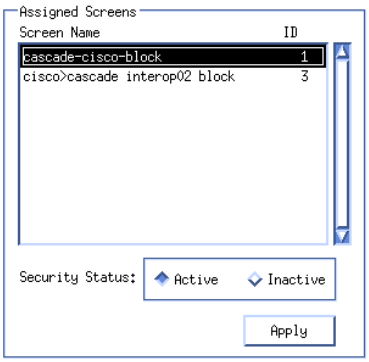

Figure 12-5. Assigned Screens List

You can choose the View Screens command to view the parameters configured for the screen you want to modify.

6. Choose Apply. The change takes effect immediately.
7. Repeat Step 5 and Step 6 to activate/deactivate additional screens.
8. Choose Close when you finish to exit this dialog box.

Viewing Screen Assignments

Use the following steps to view screen assignments for a specific logical port:

1. Use the instructions on page 3-2 to access the Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-3).
2. Select the logical port for which you want to view screen assignments.
3. Use the Select:Options menu to select Screen Assignments. Choose View. The Assignment of Port Security Screens dialog box (Figure 12-6) appears.

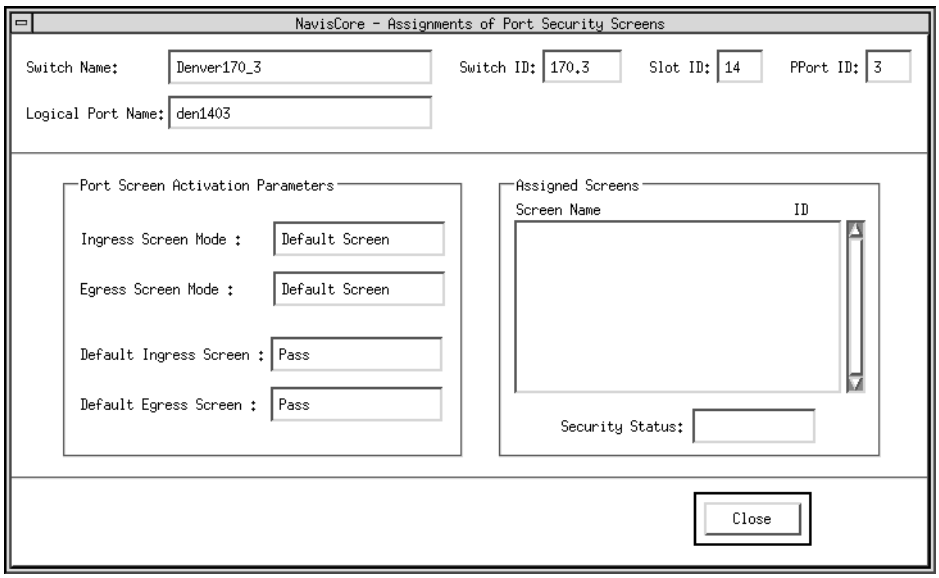

Figure 12-6. Assignments of Port Security Screens Dialog Box

- The *Port Screen Activation Parameters* fields provide the default security screen settings for this logical port. See Table 12-1 on page 12-2 for field descriptions.
- The *Assigned Screens* list provides each screen name assigned to this logical port.

4. Choose Close to exit this dialog box.

A

Reliable Scalable Circuit

SNMP errors can occur while attempting to add, modify, or delete circuits. These are reported to the user and, when possible, the circuit endpoint causing the error is identified. The options presented to the user in the case of an error (Abort, Retry, and Ignore) are sensitive to which endpoint caused the failure.

The tables in this appendix list the NMS SNMP set errors during the circuit Add, Modify, and Delete operations. This presentation is documented as a function of which endpoint experiences the SNMP set failure and the type of SNMP set failure (time-outs usually caused by switch reachability problems, and circuit not present conditions usually caused by disabled or missing endpoint cards). For each error combination of circuit operation, type of error, and endpoint failure, the effect on NMS database, state of both switches, out of sync status, effect of performing a PRAM sync and other special considerations is indicated.

In these tables, endpoint switches and cards are designated as 1st and 2nd, indicating the send order for the SNMP set commands. An SNMP set is sent to the 1st endpoint, and (if successful) it is then sent to the 2nd endpoint. Note that for Circuit Add and Modify operations, the 1st endpoint is the lower-numbered node. For Circuit Delete, the 1st endpoint is the higher-numbered node.

Several of the table descriptions list the “Nothing marked out of sync” after choosing Abort. This is only true if the configuration variable CV_PRAM_UPLOAD_ ABORT_ENABLE is set to 1 (the default). Any other variable setting results in both endpoint cards being placed out of sync when the indicated failure occurs.

Circuit Add Errors

The following table describes error messages and lists choice buttons for typical SNMP set failures encountered during attempts to add a circuit.

Table A-1. Errors Encountered During Circuit Add Procedure

Type of Failure	SNMP Set Failure Reason	Available Choices
1st switch unreachable (lower-numbered node)	The SNMP request timed out [1st endpoint identified]	**Abort** – Discontinue attempt to add circuit. NMS database, switches, and out-of-sync status unmodified. **Retry** – Attempt to add circuit again.
2nd switch unreachable (higher-number node)	The SNMP request times out [2nd endpoint identified]	**Abort** – Discontinue attempt to add circuit (NMS database unmodified, circuit dangling on 1st switch, nothing marked out-of-sync). PRAM sync of endpoint cards will remove traces of circuit from switches. **Ignore** – Discontinue attempt to add circuit, but add the circuit to the NMS database (circuit dangling on 1st switch, 2nd endpoint card marked out-of-sync). PRAM sync of endpoint cards will put circuit into switches. **Retry** – Attempt to add the circuit again. Dangling circuit on 1st switch will not interfere with the retry.
Circuit not present on 1st switch (lower-numbered node)	There is no such variable name in this MIB - possibly the card is down or not present [Specific endpoint not identified]	**Abort** – Discontinue attempt to add circuit (NMS database unmodified, nothing marked out-of-sync). PRAM sync of endpoint cards will remove traces of circuit from switches. **Retry** – Attempt to add the circuit again. Dangling circuit on 1st switch will not interfere with the Retry.
Circuit not present on 2nd switch (higher-numbered node)	There is no such variable name in this MIB - possibly the card is down or not present [Specific endpoint not identified]	**Abort** – Discontinue attempt to add circuit (NMS database unmodified, circuit dangling on 1st switch, nothing marked out-of-sync). PRAM sync of endpoint cards will remove traces of circuit. **Retry** – Attempt to add the circuit again. Dangling circuit on 1st switch will not interfere with the Retry.

Circuit Modify Errors

The following table describes error messages and lists choice buttons for typical SNMP set failures encountered during attempts to modify an existing circuit.

Table A-2. Errors Encountered During Circuit Modify Procedure

Type of Failure	SNMP Set Failure Reason	Available Choices
1st switch unreachable (lower-numbered node)	The SNMP request timed out [1st endpoint identified]	**Abort** – Discontinue attempt to modify circuit (NMS database, switches, and out-of-sync status unmodified). **Retry** – Attempt to modify circuit again.
2nd switch unreachable (higher-number node)	The SNMP request times out [2nd endpoint identified]	**Abort** – Discontinue attempt to modify circuit (NMS database unmodified, circuit dangling on 1st switch, nothing marked out-of-sync). PRAM sync of endpoint cards will remove circuit modification. **Ignore** – Discontinue attempt to modify circuit, but modify the circuit in the NMS database (circuit modify on 1st switch, 2nd endpoint card marked out-of-sync). PRAM sync of endpoint cards will modify circuit on both switches. **Retry** – Attempt to modify the circuit again. Dangling circuit modification on 1st switch will not interfere with the retry.
Circuit not present on 1st switch (lower-numbered node)	There is no such variable name in this MIB - possibly the card is down or not present [Specific endpoint not identified]	**Abort** – Discontinue attempt to modify circuit (NMS database unmodified). **Retry** – Attempt to modify the circuit again.
Circuit not present on 2nd switch (higher-numbered node)	There is no such variable name in this MIB - possibly the card is down or not present [Specific endpoint not identified]	**Abort** – Discontinue attempt to modify circuit (NMS database unmodified, circuit dangling on 1st switch, nothing marked out-of-sync). PRAM sync of endpoint cards will remove circuit modification. **Retry** – Attempt to modify the circuit again. Will begin with 1st switch, where dangling circuit modification will not interfere with the Retry.

Circuit Delete Errors

The following table describes error messages and lists choice buttons for typical SNMP set failures encountered during attempts to delete an existing circuit.

For a circuit delete, the SNMP set is first sent to the higher-numbered node (switch circuit endpoint), not the lower numbered node as is done with a circuit add or modify.

Table A-3. Errors Encountered During Circuit Delete Procedure

Type of Failure	SNMP Set Failure Reason	Available Choices
1st switch unreachable (higher-numbered node)	The SNMP request timed out [1st endpoint identified]	**Abort** – Discontinue attempt to delete circuit (NMS database, switches, and out-of-sync status unmodified). **Ignore** – Discontinue attempt to delete circuit, but delete the circuit from the NMS database (circuit not deleted on either switch, both endpoint cards marked out-of-sync). PRAM sync of endpoint cards will delete circuit on switches. **Retry** – Attempt to delete the circuit again.
2nd switch unreachable (lower-numbered node)	The SNMP request timed out [2nd endpoint identified]	**Abort** – Discontinue attempt to delete circuit, (NMS database unmodified, circuit deleted on 1st switch but left dangling on 2nd switch, nothing marked out-of-sync). PRAM sync of cards will restore the circuit on switches. **Ignore** – Discontinue attempt to delete circuit, but delete the circuit from the NMS database (circuit deleted on 1st switch but left dangling on 2nd switch, both endpoint cards marked out-of-sync). PRAM sync of endpoint cards will delete circuit on switches. **Retry** – Attempt to delete the circuit again, which now will not be able to succeed completely. ***Note:*** *Retry process starts with 1st switch, which has a deleted circuit that results in an error message. See the next table row for more information.*

Table A-3. Errors Encountered During Circuit Delete Procedure (Continued)

Type of Failure	SNMP Set Failure Reason	Available Choices
Circuit not present on 1st switch (higher-numbered node)	There is no such variable name in this MIB - possibly the card is down or not present [Specific endpoint not identified]	**Abort** – Discontinue attempt to delete circuit (NMS database, switches, and out-of-sync status unmodified). **Ignore** – Discontinue attempt to delete circuit, but delete the circuit from the NMS database (circuit not deleted on 1st switch [but it may not be there in the first place, which caused the error] or 2nd endpoint). Both circuit endpoint cards marked out-of-sync. PRAM sync cards delete circuits on switches. **Retry** – Attempt to delete the circuit again.
Circuit not present on 2nd switch (higher-numbered node)	There is no such variable name in this MIB - possibly the card is down or not present [Specific endpoint not identified]	**Abort** – Discontinue attempt to delete circuit (NMS database unmodified, circuit deleted from 1st switch, but left dangling on 2nd switch, nothing marked out-of-sync). PRAM sync of cards will restore the circuit on switches. **Ignore** – Discontinue attempt to delete circuit, but delete the circuit from the NMS database (circuit deleted on 1st switch, but is left dangling on the 2nd switch [but it may not be there in the first place, which caused the error]). 2nd endpoint card marked out-of-sync. PRAM sync of endpoint cards will delete circuits on switches. **Retry** – Attempt to delete the circuit again which will not be able to succeed completely.

B

Abbreviations and Acronyms

This appendix lists abbreviations for units of measure (in specifications) and for terms and acronyms used in Ascend documentation. Refer also to the glossary at the end of this guide, which provides definitions for many of these terms.

Abbreviations

The following table lists some of the abbreviations used in documentation and product specifications.

Table B-1. Abbreviations

Abbreviation	Meaning
Bc	Committed Burst Size
Be	Excess Burst Size
bit	binary digit
bpi	bits per inch
bps	bits per second
GB	gigabyte(s)
Gbps	gigabits per second
hex	hexadecimal
Hz	hertz (cycles per second)
ID	identification
i.e.	id est (that is)
in.	inch(es)
k	kilo (1,000)
KB	kilobyte(s)
Kbps	kilobits per second
kg	kilogram
kHz	kilohertz
MB	megabyte(s)
Mbps	million bits per second
MHz	megahertz
min	minute(s)
modem	modulator/demodulator
msec	millisecond

Table B-1. Abbreviations (Continued)

Abbreviation	Meaning
Pm	percent reduction for mild
Ps	percent reduction for severe
usec	microsecond (abbreviate with lowercase “u” for micro)
sec	second
Tc	time interval
vs.	versus
#	number; pound
x	by (multiplication)
>	greater than
<	less than
=	equal to

Acronyms

The following table lists some of the acronyms used in this guide.

Table B-2. Acronyms

Acronym	Description
AESA	ATM End System Address
ANSI	American National Standards Institute
ASCII	American Standard Code for Information Exchange
ASE	Autonomous System External
ASR	Application Specific Route
ATM	Asynchronous Transfer Mode
ATM UNI	ATM User Network Interface
BACP	Bandwidth Allocation Control Protocol
BECN	backward explicit congestion notification
BW	bandwidth
CAC	Connection Admission Control
CCITT	Consultative Committee for International Telegraph and Telephone
CHAP	Challenge Handshake Authentication Control
CFR	constant frame rate
CIC	Carrier Identification Code
CIR	committed information rate
CLLM	Consolidated Link Layer Management
CP	control processor
CPE	customer premise equipment
CRC	cyclic redundancy check
CSR	customer specific route
CSU	Channel Service Unit
CUG	closed user group
DE	discard eligibility

Table B-2. Acronyms (Continued)

Acronym	Description
DCE	data communications equipment
DLCI	data link connection identifier
DNIC	data network identification code
DS0	digital signal level 0 (64 kbps)
DS1	digital signal level 1 (1.544 Mbps)
DS3	digital signal level 3 (44.7326 Mbps)
DTE	data terminal equipment
EBW	equivalent bandwidth
EPD	early packet discard
FCP	Flow Control Processor
FECN	forward explicit congestion notification
FR	Frame Relay
FRAD	Frame Relay Assembler/Disassembler
FTP	File Transfer Protocol
HDLC	High-level Data Link Control
HSSI	High Speed Serial Interface
IA	incoming access
ICB	incoming calls barred
IDN	international data numbers
ILMI	Interim Local Management Interface
IOM	input/output module
IOP	input/output processor
IP	Internet Protocol
ISDN	Integrated Services Digital Network
ISO	International Standards Organization
ISP	internet service provider

Table B-2. Acronyms (Continued)

Acronym	Description
ITU	International Telecommunications Union
ITU-T	ITU Telecommunication Standardization Sector (formerly CCITT)
IXC	inter-exchange carrier
KA	keep alive
LAN	local area network
LCP	Link Control Protocol
LMI	Link Management Interface
LSU	link state update
LTP	Link Trunk Protocol
MBS	maximum burst size
MFRU	Multilink Frame Relay Unit
MIB	Management Information Base
ML Member	Multilink Member
MLFR	Multilink Frame Relay
MLFU	Multilink Frame Relay Unit
MPT	Multipoint-to-point tunnel
MPVC	management permanent virtual circuit
NAK	negative acknowledgment
NCP	Network Control Protocol
NIC	network interface card
NMS	Network Management Station
NNI	Network-to-Network Interface
NPC	network parameter control
NSAP	Network Service Access Point
OA	outgoing access
OC	optical carrier

Table B-2. Acronyms (Continued)

Acronym	Description
OCB	outgoing calls barred
OSPF	Open Shortest Path First
PAP	Password Authentication Protocol
PDN	public data network
PDU	protocol data unit
PNNI	Private Network-to-Network Interface
PPP	Point-to-Point Protocol
PRI	Primary Rate Interface
PVC	permanent virtual circuit
QoS	Quality of Service
RADIUS	remote authentication dial-in user service
RBOC	Regional Bell Operating Company
RFC	request for comments
SDLC	Synchronous Data Link Control
SMDS	Switched Multimegabit Data Service
SNMP	Simple Network Management Protocol
SP	switch processor
SPVC	soft permanent virtual circuit
SVC	switched virtual circuit
TAC	Technical Assistance Center
TNS	transit network selection
TS0	telecom signal level
UFR	unspecified frame rate
UIO	universal input/output
UNI	user-to-network interface
VC	virtual circuit

Table B-2. Acronyms (Continued)

Acronym	Description
VCI	virtual channel identifier
VFR-RT/NRT	variable frame rate-real time/non-real time
VPI	virtual path identifier
VPN	virtual private network
WAN	wide area network

Glossary

A

absolute congestion

In Frame Relay, a congested condition in the network that occurs when the queue length reaches a third threshold (64 buffers full), and there is no more room on the queue for any packets, regardless of the type of packet.

access rate

The data rate of the user access channel. The speed of the access channel determines how quickly (maximum rate) the end user may inject data into the network. See also *bandwidth*.

address

The logical location or identifier of a network node, terminal, pc, peripheral device, or location in memory where information is stored. See also *NavisCore*.

alternate path

An optional automatic feature of OSPF (Open Shortest Path First) that reroutes the PVC should a trunk fail within a manually defined path.

amber frames

Ascend's own class of packet frames used to identify packets as they travel through the Frame Relay network. The network forwards amber frames with the Discard Eligible bit set; therefore the packet is eligible for discard if it passes through a congested node.

American National Standards Institute

A private, non-governmental, non-profit organization, which develops US standards required for commerce.

American Standard Code for Information Interchange

A code representing characters in binary form.

ANSI

See *American National Standards Institute*.

ASCII

See *American Standard Code for Information Interchange*.

Asynchronous Transfer Mode

A method used for transmitting voice, video, and data over high-speed LAN and WAN networks.

ATM

See *Asynchronous Transfer Mode*.

B

backbone

The part of a network that carries the bulk of the network traffic, e.g. over Ethernet cabling, fiber-optic cabling.

Backward Explicit Congestion Notification

A bit in the Frame Relay header that indicates the frame has passed through a congested node from traffic traveling in the opposite direction.

bandwidth

The range of frequencies, expressed in Kilobits per second, that can pass over a given data transmission channel within a Frame Relay network. The greater the bandwidth, the more information that can be sent in a given amount of time.

Bc

See *Committed Burst Size*.

Be

See *Excess Burst*.

BECN

See *Backward Explicit Congestion Notification*.

broadband network

A type of network that allows for the transmitting of large amounts of information, including voice, data, and video over long distances using the same cable.

burst mode

A method of data transmission in which information is collected and then sent in a single high-speed transmission, rather than one character at a time.

C

CAC

See *Connection Admission Control*.

channel

Any connecting path that carries information from a sending device to a receiving device. May refer to a physical medium (e.g., coaxial cable) or a specific frequency within a larger channel.

Channel Service Unit

A device that functions as a certified safe electrical circuit, acting as a buffer between the customer's equipment and a public carrier's WAN.

CIR

See *Committed Information Rate*.

circuit

A communications channel or path between two devices.

circuit switching

A temporary communications connection that is established as needed between a sending node and a receiving node.

CLLM

See *Consolidated Link Layer Management*.

closed user group

A division of all SVC network users into logically linked groups of users. CUGs form one level of security between users of a network, allowing only those users who are members of the CUG to set up calls to each other.

Committed Burst Size

The maximum amount of data, in bits, that the network agrees to transfer information under normal conditions, during a time interval Tc. Committed Burst Size is defined for each PVC.

Committed Information Rate

The rate at which the network agrees to transfer information under normal conditions. The rate is averaged over a minimum increment of time, Tc. See also *bandwidth*.

Committed Rate Measurement Interval

The time interval during which the user is allowed to send only Bc committed amount of data and Be excess amount of data. In general, the duration of Tc is proportional to the burstiness of the traffic. Tc is computed from CIR and Bc as Tc=Bc/CIR.

communications protocol

A standard way of communicating between computers, or computers and terminals; also a hardware interface standard, such as RS-232C for communication between DTE and DCE devices.

congestion threshold

The point at which devices in the network are operating at their highest utilization. Congestion is handled by employing a congestion avoidance mechanism. See also *mild congestion*, *absolute congestion*, and *severe congestion*.

Connection Admission Control

The Cascade Connection Admission Control (CAC) algorithm performs connection admission control for all ATM service classes. The CAC enables you to control circuit creation on physical ports based on QoS objectives.

Consolidated Link Layer Management

A type of congestion control that reserves one DLCI address (1007) for transmitting congestion notification.

control processor

A module that makes up the hardware architecture of a B-STDX 9000 switch. A CP provides network and system management and routing functions in support of the real-time switching functions provided by the multiple, IO Processor modules (IOPs).

CP

See *control processor*.

CRC

See *Cyclic Redundancy Check*.

CRC error

A condition that occurs when the CRC in a frame does not agree with the CRC frame received from the network.

CSU

See *Channel Service Unit*.

Cyclic Redundancy Check

A calculation method used to check the accuracy of digital transmission over a communications link.

CUG

See *closed user group*.

D

Data communications equipment

Any device that connects a computer or terminal to a communications channel or public network.

Data Link Connection Identifier

A 10-bit address that identifies PVCs. See also *Local Management Interface* and *globally significant DLCI*.

data-link layer

The second of seven layers of the ISO/OSI model for computer-to-computer communications. This layer ensures data flow and timing from one node to another by synchronizing blocks of data and controlling the flow of data.

data packet

One unit of information transmitted as a discrete entity from one network node to another. In packet-switched networks, a data packet is a transmission unit of a fixed maximum length that contains a header, a set of data, and error control information.

data service unit

A device that connects DTE to digital communications lines. A DSU formats the data for transmission on the public carrier WAN, and ensures that the carrier's requirements for data formats are met.

data terminal equipment

Any device, such as a terminal or computer, that is connected to a communications device, channel, or public network.

data transfer rate

The speed at which data is transferred, usually measured in megabits per second (Mbps) or megabytes (MB) per second.

DCE

See *Data communications equipment*.

DE

See *Discard Eligible (DE)*.

define path

A function that allows a manual path to be defined for the PVC, thereby bypassing the OSPF (Open Shortest Path First) algorithm to make PVC routing decisions.

delay

In communications, a pause in activity, representing the time that a message must wait for transmission-related resources to become available.

destination address

The address portion of a packet or datagram that identifies the destination node.

Digital Signal (Digital Service)

A classification of digital circuits. The DS defines the level of common carrier digital transmission service. DS0 = 64 Kbps (Fractional T1), DS1 = 1.544 Mbps (T1), DS2 = 6.312 Mbps (T2), DS3 = 44.736 Mbps (T3), and DS4 =274-176 Mbps (T4).

direct Ethernet

A connection method used by the NMS to the network. The NMS communicates directly to the gateway switch through the Ethernet port on the NMS to the Ethernet port on the switch.

Discard Eligible (DE)

A bit in the Frame Relay header used to indicate that a frame is eligible for discard by a congested node to maintain the committed information rate (CIR).

DLCI

See *Data Link Connection Identifier*.

domain

A network community of users sharing the same database information.

DS

See *Digital Signal (Digital Service)*.

DSU

See *data service unit*.

DTE

See *data terminal equipment*.

dynamic routing

A routing technique that allows a message's route to change "en route" through the network.

E

E.164

A public network addressing standard utilizing up to a maximum of 15 digits. Frame Relay and ATM use E.164 addressing for public network addressing.

E1

The European counterpart to the North American T1 transmission speed. Adopted by the Conference of European Posts and Telecommunications Administrations, the E1 standard carries data at the rate of 2.048 Mbps.

egress

Frame Relay frames leaving a Frame Relay network toward the destination device. Contrast with *ingress*.

encapsulation

The wrapping of data in a particular protocol header. For example, Ethernet data is wrapped in a specific Ethernet header before being transmitted. Also, when bridging dissimilar networks, the entire frame from one network is simply placed in the header used by the data link layer protocol of the other network.

error rate

In communications, the ratio between the number of bits received incorrectly and the total number of bits in the transmission.

Ethernet

A popular LAN protocol and cabling scheme with a transfer rate of 10 Mbps.

Ethernet address

A 48-bit number physical address. Each Ethernet address is unique to a specific network card or PC on a LAN, which forms the basis of a network-addressing scheme.

Excess Burst

The maximum allowed amount of uncommitted data (in bits) in excess of Bc that the network attempts to deliver during time interval Tc. In general, this data (Be) is delivered with a lower probability than Bc.

fault-tolerant PVCs

A set of backup ports (Permanent Virtual Circuits) on the switch used to restore connections from a failed data center to the backup data center. When enabled, a fault-tolerant PVC automatically reroutes all affected circuits to the set of backup ports.

F

FECN

See *Forward Explicit Congestion Notification*.

File Transfer Protocol

A method of transferring information from one computer to another, either over a modem and telephone line, or over a network. FTP is a TCP/IP application utility.

Forward Explicit Congestion Notification

A bit in the Frame Relay header that indicates the frame has passed through a node that is experiencing congestion in the same direction in which the frame is traveling.

FRAD

See *Frame Relay Assembler/Disassembler*.

frame

In Frame Relay, a block of data that can be transmitted as a single unit.

frame check sequence

In a frame, a field that contains the standard 16-bit cyclic redundancy check used to detect errors in HDLC and LAPD frames. See also *Cyclic Redundancy Check*.

Frame Relay

A type of data transmission based on a packet-switching protocol, with transmission rates up to 2 Mbps. Frame Relay provides for bandwidth-on-demand.

Frame Relay Assembler/Disassembler

A function that enables a logical port to perform Frame Relay encapsulation/de-encapsulation for HDLC/SDLC-based protocols. The FRAD function encapsulates HDLC/SDLC traffic entering an Ascend Frame Relay network and de-encapsulates it upon exiting the network. This function is restricted to one point-to-point PVC.

Frame Relay Frame

A variable-length unit of data in Frame Relay format that is transmitted through a Frame Relay network as pure data. Contrast with *packet*.

Frame Relay Network

A telecommunications network based on Frame Relay technology. Data is multiplexed. Contrast with packet-switched network.

Frame Relay RFC1294 Multi-protocol Encapsulation

A specification allowing for a single circuit to be established between two devices.

FTP

See *File Transfer Protocol*.

G

globally significant DLCI

A feature of the Local (or Link) Management Interface (LMI) enhancement to Frame Relay that enables DLCIs to use the same connection-identification scheme across the network (global values) to specify individual end devices.

graceful discard

When enabled, this function turns red frames into best-effort frames. When disabled, this function discards frames.

green frames

Ascend's own class of packet frames used to identify packets as they travel through the network. Green frames are never discarded by the network except under extreme circumstances, such as node or link failure.

group addressing

The ability to send a single datagram/packet to multiple locations.

H

HDLC

See *High-level Data Link Control*.

header

The initial part of a data block, packet, or frame, which provides basic information about the handling of the rest of the block, packet or frame.

heartbeat polling process

An exchange of sequence numbers between the network and a user device to ensure that both are operational and communicating.

Hello protocol

Protocol used by OSPF systems for establishing and maintaining neighbor relationships.

High-level Data Link Control

An international protocol defined by ISO. In HDLC, messages are transmitted in variable-length units known as frames.

High-Speed Serial Interface

A high-speed interface (up to 52 Mbps full duplex) between a DTE and a DCE. The DCE provides the timing for the interface. HSSI can operate over a 50-ft (15m) shielded twisted-pair cable.

hop (count)

The number of links that must be “jumped” to get from a source node to a destination node.

host name

A unique name identifying a host system.

HP OpenView

The UNIX-based network management application used with NavisCore on an NMS to manage an Ascend switch network.

HSSI

See *High-Speed Serial Interface*.

I

ILMI

See *Interim Local Management Interface*.

indirect Ethernet

A LAN topology or an extended LAN where the NMS and the switch reside on different LANs and must use a router for access.

ingress

Frame Relay frames leaving an access device toward the Frame Relay network. Contrast with *egress*.

Input/Output Processor

A module in a switch that manages the lowest level of a node’s trunk or user interfaces. An IOP performs physical data link and multiplexing operations on external trunks and user links.

Interim Local Management Interface

Specifications developed by the ATM forum for incorporating network-management capabilities into the ATM UNI.

Integrated Services Digital Network

A CCITT standard for a worldwide digital communications network, intended to replace all current systems with a completely digital transmission system.

International Standards Organization

An international standards group based in Geneva, Switzerland that establishes global standards for communications and information exchange.

International Telecommunication Union Telecommunication Standard Sector

An advisory committee established under the United Nations to recommend worldwide standards for voice and data. One of the four main organizations of the International Telecommunications Union.

Internet Protocol

The TCP/IP session-layer protocol that regulates packet forwarding.

IOP

See *Input/Output Processor*.

IP

See *Internet Protocol*.

ISDN

See *Integrated Services Digital Network*.

ISO

See *International Standards Organization*

ITU-T

See *International Telecommunication Union Telecommunication Standard Sector*.

K

KA

See *keep-alives*.

Kbps

Kilobits per second.

keep-alives

A series of polling messages used in the Local (or Link) Management Interface (LMI) of a Frame Relay port to verify link integrity between devices.

L

LAN

See *Local Area Network*.

Link Management Interface

A set of enhancements to the basic Frame Relay specification. LMI dynamically notifies the user when a PVC is added or deleted. The LMI also monitors each connection to the network through a periodic heartbeat “keep alive” polling process.

Link Management Interface Rev 1

A synchronous polling scheme used for the link management of a Frame Relay channel where the user polls the network to obtain status information of the PVCs configured on the channel. LMI exchanges this information using DLCI 1023.

link-state routing protocol

A sophisticated method of determining the shortest paths through the network. See also *Open Shortest Path First*.

LMI

See *Link Management Interface*.

LMI Rev 1

See *Link Management Interface Rev 1*.

load balancing

A technique that distributes network traffic along parallel paths to maximize the available bandwidth while providing redundancy at the same time.

Local Area Network

Any physical network technology that connects a number of devices and operates at high speeds (10 Mbps through several gigabits per second) over short distances. Compare with *Wide Area Network*.

Local Management Interface

See *Link Management Interface*.

locally significant DLCI

In Frame Relay, an identifier or address that specifies a local router, PVC, SVC, or endpoint device. It is reusable at non-overlapping endpoints and allows for scalability. Compare with *globally significant DLCI*.

logical port

A configured circuit that defines protocol interaction.

loopback test

A diagnostic that directs signals back toward the transmitting source to test a communications path.

M

Management DLCI

A value that specifies a PVC or SVC from a LAN connected via a router to a Ascend switch over a Frame Relay network.

Management Information Base

The set of variables forming a database contained in a CMIP or SNMP-managed node on a network. Network management stations can fetch/store information from/to this database.

Management PVC

Provides access to the switching network's management plane, which is IP-based. MPVCs offer an efficient, high performance data path capable of transferring large amounts of management data, such as accounting or bulk statistics files.

Mbps

Megabits per second.

MIB

See *Management Information Base*.

mild congestion

In Frame Relay, the state of a link when the threshold (more than 16 buffers by default) is exceeded.

MLFR

See *Multilink Frame Relay*.

ML member

A type of logical port configuration that can be bound to an MLFR trunk bundle logical port. See also *Multilink Frame Relay*.

mono-class service

A logical port service class type for which all circuits are transmitted using VFR-nrt characteristics.

MPVC

See *Management PVC*.

multicast

A type of broadcast transmission that sends copies of the message to multiple stations, but not to all possible stations.

multicast DLCI

A circuit configured to send multiple groups of circuits on the same logical port.

multi-class service

A logical port service class type for which all QoS classes are supported. Multi-class service requires specification of a transmit scheduling mode.

Multilink Frame Relay

A method of aggregating available bandwidth on a set of Frame Relay logical links between two networking devices. MLFR requires creation of ML member logical ports, which are then bound to the MLFR trunk bundle logical port. MLFR combines the multiple logical links between two networking devices into a single greater logical connection.

multiplexing

A technique that transmits several signals over a single communications channel.

N

NavisCore

The UNIX-based graphical user interface used to configure and monitor an Ascend switch network.

network address

A network layer address refers to a logical, rather than a physical network device; also called protocol address.

Network Interface Card

A card, usually installed in a PC, that enables you to communicate with other users on a LAN; also called adapter.

Network Management Station

The device used to configure and manage the network.

network parameter control

The set of actions taken by the network to monitor and control traffic from the NNI. Its main purpose is to protect network resources from malicious as well as unintentional misbehavior, which can affect the QoS of previously established connections, by detecting violations of negotiated parameters and taking appropriate actions.

Network-to-Network Interface

The standard that defines the interface between ATM switches and between Frame Relay switches. In an SMDS network, an NNI is referred to as Inter-Switching System Interface (ISSI).

NIC

See *Network Interface Card*.

NMS

See *Network Management Station*.

NNI

See *Network-to-Network Interface*.

node

Any device such as a pc, terminal, workstation, etc., connected to a network and capable of communicating with other devices.

node number

A unique number that identifies a device on the network.

NPC

See *network parameter control*.

O

Open Shortest Path First

A routing protocol that takes into account network loading and bandwidth when routing information over the network. Incorporates least-cost routing, equal-cost routing, and load balancing.

OPTimum PVC trunk

A logical port configuration that optimizes interoperability in performance and throughput in networks where both ends are connected by Ascend switches.

OPTimum trunking

A software function that allows public data networks based on Frame Relay, SMDS, or ATM to be used as trunk connections between Ascend switches.

OSPF

See *Open Shortest Path First*.

out of frame

A T1 error condition where two or three framing bits of any five consecutive frames are in error.

P

packet

Any block of data sent over a network. Each packet contains sender, receiver, and error-control information in addition to the actual message; sometimes called payload or data bits.

packet assembler/disassembler

A device connected to a packet-switched network that converts a serial data stream from a character-oriented device (e.g., a bridge or router) into packets suitable for transmission. It also disassembles packets into character format for transmission to a character device.

packet-switched network

A network that consists of a series of interconnected circuits that route individual packets of data over one of several routes and services.

packet switching

Type of networking in which nodes share bandwidth with each other by intermittently sending logical information units (packets). In contrast, a circuit-switching network dedicates one circuit at a time to data transmission.

PAD

See *packet assembler/disassembler*.

parameter random access memory

The PRAM on a switch that contains the module's downloaded configuration file, which is stored in battery backup.

payload

The portion of a frame that contains the actual data.

PDN

See *Public Data Network*.

PDU

See *Protocol Data Unit*.

Permanent Virtual Circuit

A logical connection across a packet-switched network that is always in place and always available along a predetermined network path. See also *Virtual Circuit*.

Point-to-Point Protocol

A protocol that provides router-to-router and host-to-network connections.

PPP

See *Point-to-Point Protocol*.

PRAM

See *parameter random access memory*.

PRI

See *Primary Rate Interface*.

primary group

The main group to which associated users belong. The system identifies the primary group by the group field in the user account (stored in the /etc/password file) and by the group ID associated with a new file.

Primary Rate Interface

An ISDN interface to primary rate access, which consists of a single 64-Kbps D channel plus 23 (T1) or 30 (E1) B channels for voice or data.

Priority Frame

Provides ATM-like Quality of Service (QoS) for Frame Relay. Priority Frame enables setting QoS parameters for logical ports, which allows selection of bandwidth and routing metrics for the various traffic service classes.

protocol

A set of rules governing communication between two entities or systems to provide interoperability between services and vendors. Protocols operate at different layers of the network, e.g., data link, network, and session.

Protocol Data Unit

A unit of data consisting of control information and user data exchanged between peer layers.

Public Data Network

Any government-owned or controlled commercial packet-switched network, offering WAN services to data processing users.

PVC

See *Permanent Virtual Circuit*.

Q

QoS

See *Quality of Service*.

Quality of Service

A statistical report that specifies certain characteristics of network services, sessions, connections, or links. For example, a NavisCore statistics report describes the lost packets and round-trip delay measurements.

R

RADIUS

See *Remote Authentication Dial-In User Service*.

rate enforcement

A process used to measure the actual traffic flow across a given connection and compare it to the total admissible traffic flow for that connection. Traffic outside of the acceptable level can be tagged and discarded en route if congestion develops. ATM, Frame Relay, and other types of networks use rate enforcement.

Receive Data

A hardware signal, defined by the RS-232-C standard, that carries data from one device to another. Compare with *Transmit Data*.

red frames

In Frame Relay, a type of frame to be discarded. Color designators green, amber, and red identify packets as they travel through the network.

redundancy

The duplication of hardware or software within a network to ensure fault-tolerant or back-up operation.

Remote Authentication Dial-In User Service

A Distributed security system that uses an authentication server to solve the security problems associated with remote computing.

remote connection

A workstation-to-network connection made using a modem and telephone line or other WAN services equipment. Remote connections enable you to send and receive data over greater distances than you can with conventional cabling methods.

Request For Comment

A series of notes and documents available on-line that describe surveys, measurements, ideas, techniques, and observations, as well as proposed and accepted Internet protocol standards, such as Telnet and FTP. For example, RFC 1294 and RFC 1490 describe multiprotocol access over Frame Relay.

RFC

See *Request For Comment*.

route recovery

In Frame Relay, an OSPF routing function in the Ascend switch. When a tandem node or trunk is down, new shortest-path routes for those affected PVCs are recalculated immediately at the ingress nodes, due to fast convergence of the link-state updates. The PVCs are then rerouted to the new route. Recovery time is typically under four seconds. The network reports PVC rerouting as an event/alarm.

router

An intelligent LAN-connection device that routes packets to the correct LAN segment destination address(es). The extended LAN segments may or may not use the same protocols. Routers link LAN segments at the ISO/OSI network layer.

routing

The process of directing data from a source node to a destination node.

routing protocol

A protocol that maintains a list of accessible networks and calculates the lowest hop count from a particular location to a specific network.

S

severe congestion

In Frame Relay, a state or condition that occurs when the queue size is greater than a second predetermined threshold (32 buffers full). In this state, the continued forwarding of amber and red packets jeopardize the successful delivery of green packets.

shortest path routing

A routing algorithm that calculates the path distances to all network destinations. The shortest path is then determined by a cost assigned to each link. See also *OSPF*.

Simple Network Management Protocol

A standard network management protocol used to manage and monitor nodes and devices on a network.

SMDS

See *Switched Multimegabit Data Services*.

SNMP

See *Simple Network Management Protocol*.

SP

See *Switch Processor*.

static route

A route or path that is manually entered into the routing table. Static routes take precedence over routes or paths specified by dynamic routing protocols.

SVC

See *Switched Virtual Circuit*.

Switch Processor

A control module present in the CBX 500 switch that controls the switch and interacts with multiple Input/Output Processor (IOP) modules.

Switched Multimegabit Data Services

A high-speed WAN service based on the 802.6 standard for use over T1 or T3 circuits.

Switched Virtual Circuit

A logical connection across a packet-switched network providing as-needed connections to any other node in the network. See also *Virtual Circuit*.

synchronization

The timing of separate elements or events to occur simultaneously. In communications, hardware and software must be synchronized so that file transfers can occur.

synchronous transmission

A data transmission method that uses a clock signal to regulate data flow.

T

T1

A long-distance, point-to-point circuit that provides 24 channels at 64 Kbps each (for a total of 1.544 Mbps). See also *E1*.

T3

A long-distance, point-to-point circuit that provides up to 28 T1 channels. T3 can carry 672 channels of 64 Kbps (for a total of 44.736 Mbps).

Tc

See *Committed Rate Measurement Interval*.

TCP

See *Transmission Control Protocol*.

telnet

The Internet standard protocol for remote terminal-connection services.

throughput

The actual speed of the network.

time interval “T”

The time interval over which the number of bits used to average the number of bits transmitted, is averaged. To calculate **T,** use the following formula: Bc/CIR=T.

topology

The map or configuration design of a network. Physical topology refers to the location of hardware. Logical topology refers to the paths that messages take to get from one node to another.

traffic shaping

In Frame Relay, a set of rules that describes traffic flow. The sender has a mechanism to ensure that the transmission of its guaranteed packets behaves in a certain way. The network knows what kind of traffic to expect, and can monitor the behavior of the traffic.

Transmission Control Protocol

The Internet standard, transport-level protocol that provides the reliable, full duplex, stream service on which many application protocols depend.

Transmit Data

A hardware signal, defined by the RS-232-C standard, used by the DTE to transmit data to the DCE. Compare with *Receive Data*.

trap

An unsolicited message generated by an SNMP agent on a network device (e.g., switch) due to a predefined event occurring or alarm threshold being exceeded, which triggers an alarm at the NMS.

trunk

The communications circuit connecting a Frame Relay-compatible device to a Frame Relay switch.

trunk backup

A configuration setting specified by a network operator via the NMS. The network operator can initiate or terminate primary trunk backups at any time via the NMS. Trunk backups take over a connection should the primary trunk fail.

trunk failure

A condition (alarm) that occurs when the Ascend switch status indicates that a trunk is no longer available.

trunk restoration

A process that reroutes the PVCs carried on the backup trunk, and frees up the circuit on the backup trunk.

U

UFR

See *unspecified frame rate*.

UIO module

See *Universal Input/Output Module*.

UNI

See *User-to-Network Interface*.

UNI DCE

See *User Network Interface Data Communications Equipment*.

UNI DTE

See *User Network Interface Data Terminal Equipment*.

Universal Input/Output Module

In the Ascend switch, a module that has three 80-pin connectors and is used for redundancy, and also as an I/O module for X.21, RS449, V.35, EIA530, and EIA530A interfaces.

unspecified frame rate

ATM-like QoS class provided by Priority Frame for Frame Relay networks. UFR is used primarily for LAN traffic.

User Network Interface Data Communications Equipment

A device that performs the Frame Relay DCE functions for link management and expects a Frame Relay DTE device (e.g., Ascend switch) to be attached to it.

User Network Interface Data Terminal Equipment

A device that performs the Frame Relay DTE functions for link management. The user specifies this option on the NMS to connect to a Frame Relay DCE, where the Ascend switch acts as the DTE.

User-to-Network Interface

A standard defined by the ATM Forum for public and private ATM network access. UNI connects an ATM end system (such as a router) and an ATM switch, and is also used in Frame Relay. UNI is called SNI (Subscriber Network Interface) in SMDS.

V

variable frame rate

ATM-like QoS class provided by Priority Frame for Frame Relay networks. VFR is subdivided into a real time (rt) class and non-real time (nrt) class. VFR-rt is used for packaging special delay-sensitive applications, such as packet video, which require low delay variation between endpoints. VFR-nrt handles packaging for transfer of long, bursty data streams over a pre-established connection, or for short, bursty data, such as LAN traffic. VFR-nrt also provides congestion control support.

VC

See *Virtual Channel*; *Virtual Circuit*.

VCI

See *Virtual Circuit Identifier*;*Virtual Path Identifier*.

VFR

See *variable frame rate*.

virtual bandwidth

Channel capacity calculated to allow for oversubscription of channel usage.

Virtual Channel

A connection between two communicating ATM networks.

Virtual Circuit

A logical circuit set up to ensure reliable communication between two network devices. See also *PVC* and *SVC*.

Virtual Circuit Identifier

A 16-bit field in the ATM cell header that is used as an addressing identifier to route cell traffic.

Virtual Path

A group of VCs carried between two points that provides a way to bundle traffic headed in the same direction.

Virtual Path Identifier

An 8-bit field in the ATM cell header that is used as an addressing identifier to route cell traffic.

Virtual Private Network

A network that provides dedicated bandwidth and guaranteed performance, reliability, and privacy.

VP

See *Virtual Path*.

VPI

See *Virtual Path Identifier*.

VPN

See *Virtual Private Network*.

W

WAN

See *Wide Area Network*.

Wide Area Network

A network that usually consists of packet-switching nodes over a large geographical area.

X

X.121

An ITU-T addressing standard used in X.25 networks. X.121 addresses are sometimes referred to as IDNs (International Data Numbers). X.121 addresses consist of 14 ASCII digits. Only number values between 0-9 are valid.

Index

A

B

C

D

E

F

G

I

K

L

M

N

O

P

Q

R

S

T

U

V

X