NavisCore ATM Configuration Guide

Ascend Communications, Inc.

Product Code: 80072
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 4 Configuring ATM Logical Ports on the B-STDX

Chapter 5 Configuring Trunks

Chapter 7 Configuring Management Paths

Chapter 8 Configuring ATM Traffic Descriptors

Chapter 9 Configuring Virtual Private Networks

Chapter 10 Configuring Fault-Tolerant PVCs

Chapter 11 About SVCs

Chapter 12 Configuring SVC Parameters

Chapter 13 Configuring SVC Proxy Signaling

Chapter 14 Configuring SPVCs

Chapter 15 Closed User Groups

Chapter 16 Port Security Screening

Figures

List of Tables

About This Guide

The *Naviscore ATM Configuration Guide* provides detailed instructions for using NavisCore to set up and manage a network map and configure ATM services on an Ascend switch network. Specifically, this guide describes how to configure logical ports, trunks, PVCs, and SVCs to support ATM services on either a CBX 500, GX 550 or B-STDX switch. This guide also explains how to configure a variety of features that enhance the ATM service platform, including virtual private networks, closed user groups, and port security screening.

This guide supports the following NMS and switch software releases:

- NavisCore, Release 4.0
- B-STDX, Release 6.0
- CBX 500, Release 3.0
- GX 550, Release 1.0.

What You Need to Know

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

This guide assumes you have already installed the Ascend switch hardware, using one of the following guides:

- *STDX 6000 Hardware Installation Guide*
- *B-STDX 8000/9000 Hardware Installation Guide*
- *CBX 500 Hardware Installation Guide*
- *GX 550 Hardware Installation Guide*

You should have also installed the NMS software using the *Network Management Station Installation Guide*. Before you begin to configure ATM services, see the *NavisCore Physical Interface Configuration Guide* to configure processor and I/O cards and physical ports.

Reading Path

This section describes all of the documents that support the NavisCore 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. Be sure to review the NavisCore Customer Software Release Notice for any changes not included in these guides.

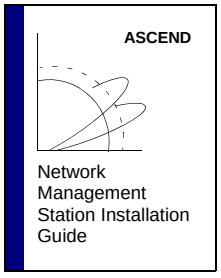

This guide describes prerequisite tasks, hardware and software requirements, and instructions for installing Solaris, HP OpenView, and NavisCore on the NMS.

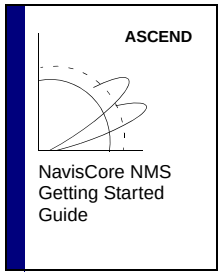

This guide describes how to configure and manage NavisCore, network maps, and Ascend switches.

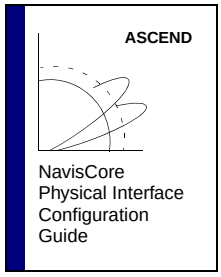

This guide describes how to configure processor and I/O modules on Ascend switches.

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. Be sure to review the appropriate switch software customer software release notice(s) for any changes.

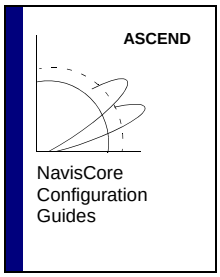

The following guides describe how to configure WAN services on the STDX, B-STDX, CBX, and GX switch platforms:

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

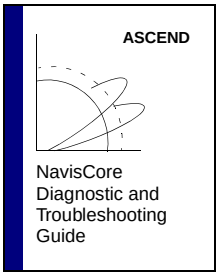

This guide describes how to diagnose and troubleshoot your NavisCore switch network.

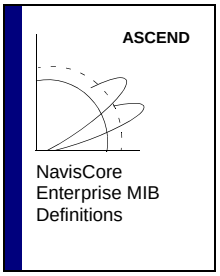

This document gives a brief overview of SNMP and describes the NavisCore Enterprise MIB definitions.

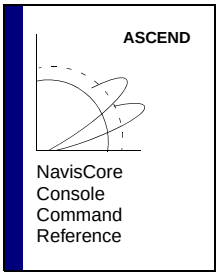

This reference lists and describes the NavisCore switch console commands.

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 ATM logical ports. These concepts include: virtual paths and channels, signaling, and Interim Local Management Interface (ILMI).
Chapter 3	Configuring ATM logical ports on a CBX 500 or GX 550.
Chapter 4	Configuring ATM logical ports on a B-STDX.
Chapter 5	Configuring ATM trunks and APS trunk backup.
Chapter 6	Configuring point-to-point and point-to-multipoint PVCs.
Chapter 7	Configuring NMS paths using either a management PVC or management VPI/VCI connection.
Chapter 8	Configuring traffic descriptors to manage Quality of Service (QoS) throughout your ATM network.
Chapter 9	Configuring your ATM services to provide Virtual Private Networks (VPNs).
Chapter 10	Configuring fault tolerant (resilient UNI) PVC services to provide backup services should a logical port endpoint fail. This chapter also describes APS Resilient UNI.
Chapter 11	ATM Switched Virtual Circuit (SVC) concepts you need to understand before you can configure SVCs. These include: address formats and registration, route determination, and address translation.
Chapter 12	Configuring ATM SVCs on a CBX 500 or GX 550.
Chapter 13	Configuring Proxy Signaling.
Chapter 14	Configuring soft PVCs (SPVCs) within the network using signaling.
Chapter 15	Closed User Groups (CUGs) that enable you to divide all network users into logically linked groups of users.
Chapter 16	Using Port Security Screening to create screens which allow/disallow incoming and outgoing calls.
Chapter 17	Configuring the ATM Private Network-to-Network Interface (PNNI) routing protocol in your Ascend network.

Read	To Learn About
Appendix A	Tuning the Ascend Call Master Connection Admission Control (CAC) to achieve a desired cell loss ratio objective across all physical ports in your network.
Appendix B	How each traffic descriptor combination affects the cell streams under different traffic conditions.
Appendix C	Allocating logical port bandwidth on CBX 500 shared SP threads.
Appendix D	Using the CBX 500's *optional* Flow Control Processor (FCP) functions.
Appendix E	Using priority routing to prioritize PVC traffic.
Appendix F	Using the Reliable Scalable Circuit function to troubleshoot PVC provisioning problems.

What's New in This Release

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

Table 1. NavisCore Release 04.00.02.00 Features

Feature	Description	Reference . . .
GX 550 switch hardware	This release of NavisCore supports the GX 550 switch hardware platform, which includes the following hardware components: • Node Processor (NP) card and NP Adapter (NPA) • Switch Fabric (SF) module • Timing Module (TM) • 4-Thread 16-Port Base Input/Output (BIO) module supporting a mix of 4-port OC3/STM1 and 1-port OC12/STM4 physical interface modules	This manual describes how to configure ATM services for the GX 550, as well as the CBX 500 and B-STDX.
VPI 0 for ATM and Frame Relay OPTimum trunks	It is now possible to configure an OPTimum trunk that uses a VPI of zero.	For CBX/GX, see Chapter 3. For B-STDX, see page 4-33
ITU Variant UNI Signaling (Q.2931/Q.2971)	This release supports Q.2931/Q.2971, the international variant for UNI signaling. ITU UNI signaling provides all UNI 4.0 signaling features with some exceptions, including: • No support for ATM Anycast • Provides additional E.164 numbering plans: private, national	"ATM Attributes" on page 3-11

Table 1. NavisCore Release 04.00.02.00 Features (Continued)

Feature	Description	Reference . . .
VP Shaping *(CBX and B-STDX)*	VP shaping provides the ability to shape OPTimum trunk and virtual UNI connections at a specified peak cell rate (PCR) while preserving QoS integrity. For example, VP shaping is useful when multiple OPTimum trunk logical ports exist on one physical port and each trunk must traverse a VP of a fixed PCR. This new feature ensures that the maximum rate of the OPTimum trunk traffic does not exceed the specified PCR. For a CBX 500, you can enable VP shaping only if the host IOM is equipped with certain revisions of the ATM flow control processor (FCP) module. You can configure the FCP to either enable flow control or VP shaping; it does not support both functions at the same time. See either the CBX or NavisCore software release notices for appropriate revision levels. For the B-STDX, this feature is only supported on certain revisions of the ATM IWU and CS modules. See either the B-STDX or NavisCore software release notices for appropriate revision levels.	For CBX 500, see page 2-14 and page 3-10 For B-STDX, see page 4-6
CBX 500 6-port DS3 Frame module	Supports high-speed Frame Relay access at DS3 rates. This module provides IP routing services that enable the CBX 500 to support IP switching. This module can also act as an IP server to enable switching of IP traffic that arrives at the CBX from any supported ATM UNI interface (T1/E1 through OC12/STM4).	Chapter 4
Priority Frame	This feature provides ATM-like quality of service (QoS) classes for frame-based B-STDX cards. By configuring these QoS classes, you can offer extended congestion control support.	"Priority Frame" on page 4-20
DE to CLP and FECN to EFCI mapping (*B-STDX CS/IWU cards*)	This feature allows you to specify how a trunk logical port on a CS/IWU card maps discard eligible (DE) bits to cell loss priority (CLP) bits and EFCI bits to FECN bits for data sent to and received from the ATM interface. The default is to map the corresponding Frame Relay and ATM bits in each direction. Using this default, when data is sent out the ATM trunk, the current switch setting of the DE bit is transferred to the ATM cell's CLP bit. The current switch setting of the FECN bit is transferred in the ATM cell header as the EFCI bit. ATM cells received from the ATM interface use the CLP bit set in the cell header for a frame as the DE bit associated with that frame. Similarly, the EFCI bit received in cells will be translated to the FECN bit that is carried with the frame within the switch.	"Discard/Congestion Mapping" on page 4-30

Table 1. NavisCore Release 04.00.02.00 Features (Continued)

Feature	Description	Reference . . .
Additional per trunk VC counts	The Set All Trunks dialog box now lists the total number of PVCs, SVCs, and VCs that traverse the selected trunk. The total number of VCs field represents the total number of PVCs, SVCs, SPVCs, MPTs, and any other VC traversing the trunk.	"Accessing Trunk Functions" on page 5-7
APS Trunk Backup (*GX and CBX*)	This feature automates the new Trunk Backup feature through the use of Automatic Protection Switching (APS). With the combination of these two features, it is possible to define and associate a backup trunk with a primary trunk. In the event a physical port failure occurs on the primary trunk, the APS function will automatically switch traffic to the backup trunk. See either the NavisCore or corresponding switch software release notice for OC3/STM-1 revision level requirements.	"About APS Trunk Backup" on page 5-4
Multiple 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.	page 5-14
Extended QoS Parameters	As part of Ascend's support of UNI 4.0 and TM 4.0, it is now possible to configure cell delay variation (CDV), cell loss ratio (CLR), and cell transfer delay (CTD) parameters for individual PVCs. This PVC capability is in addition to the SVC capability that is inherent in Ascend's support of UNI 4.0 and TM 4.0.	"Extended QoS Parameter Attributes" on page 6-19
Least OSPF delay metric routing and Admin Cost	This feature enables you to manage PVC routing using either OSPF admin cost or end-to-end delay metrics. 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.	"Administrative Attributes" on page 6-12

Table 1. NavisCore Release 04.00.02.00 Features (Continued)

Feature	Description	Reference . . .
Management DLCI Loopback	This feature 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.	"Administrative Attributes" on page 6-12
PVC Frame Discard (EPD/PPD support)	The Add PVC screen provides parameters to enable forward/reverse frame discard features on a per call basis. This feature offers congestion-control support using early packet discard (EPD) and partial packet discard (PPD) functions.	"Frame Discard Attributes" on page 6-18
ABR support for interworking PVCs and Frame Relay endpoints	The NMS now offers the ABR service class for Frame Relay endpoints and for ATM endpoints on interworking PVCs. This feature is used in cases where the PVC transits one or more CBX switches equipped with an FCP.	"User Preference Attributes" on page 6-30
B-STDX Management PVC	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.	"Using Management PVCs" on page 7-2
Configurable Control Circuits (*GX and CBX Only*)	This feature allows you to control the amount of bandwidth associated with various control traffic circuits. You may also select the QoS class of this control traffic. Control traffic circuits include those used by UNI signaling, ILMI, trunk signaling, PNNI RCC, and node-to-node management traffic.	"Configuring ATM Traffic Descriptors" on page 8-7

Table 1. NavisCore Release 04.00.02.00 Features (Continued)

Feature	Description	Reference . . .
APS Resilient UNI (*GX and CBX Only*)	This feature automates the existing Resilient UNI feature through the use of Automatic Protection Switching (APS). With the combination of these two features, it is possible to define and associate a backup logical port with a primary logical port. In the event a physical port failure occurs on the primary port, the APS function will automatically switch traffic to the backup port. See either the NavisCore or corresponding switch software release notice for revision level requirements.	"Configuring APS with Resilient UNI" on page 10-6
CBX/GX UNI 4.0 signaling support	This release offers UNI 4.0 signaling support. UNI 4.0 includes support for the following: • ATM Anycast addressing • Switched Virtual Path Service • Frame Discard support • Virtual UNI enhancements • Proxy Signaling	"Address Formats" on page 11-2 "About SPVCs" on page 14-2 "Frame Discard Attributes" on page 6-18 "Using Interim Local Management Interface (ILMI)" on page 2-3 "About Proxy Signaling" on page 13-2
PNNI	This release enables you to configure Private Network-to-Network Interface (PNNI) routing using ATM NNI logical ports. PNNI is a standard designed by the ATM Forum. This standard defines both an ATM routing protocol and an ATM signaling protocol. Ascend supports PNNI on the CBX 500 and GX 550.	Chapter 17
VC Priority Reroute	Priority rerouting allows a VC (PVC or SVC) with a higher configured priority to select a more optimal path that satisfies its Quality of Service (QoS) than VCs configured with lower priority. In selecting this optimal path, higher priority VCs may elect to force lower priority traffic away from the optimal path to sub-optimal paths. Priority rerouting guarantees that under any circumstances, such as link failure or congestion, the higher priority VCs will always be given preference on the network over the lower priority VCs.	See Appendix E, "Priority Routing"
Reliable Scalable Circuit	This feature improves reliability when provisioning PVCs and is set to On by default. The NMS verifies that the card state is up for each PVC endpoint before sending an SNMP Set command to the switches. If the card status of either endpoint is not up, the NMS displays an error message to indicate where the failure occurred.	Appendix F

Conventions

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

Convention	Indicates	Example
`Courier Bold`	User input on a separate line.	**`eject cdrom`**
[*bold italics*]	Variable parameters to enter.	**[*your IP address*]**
`Courier Regular`	Output from a program.	`Please wait...`
Boldface	User input in text.	Type **cd install** and ...
Menu => Option	Select an option from the menu.	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.

Cautions notify the reader to proceed carefully in order to avoid damaging equipment or losing data.

Warnings notify the reader to proceed carefully in order to avoid personal harm.

Related Documents

This section lists the related Ascend and third-party documentation that you may find helpful to read.

Ascend

- *NavisCore Reading Roadmap* (Product code: 80069)
- *B-STDX 8000/9000 Hardware Installation Guide* (Product code: 80005)
- *CBX 500 Hardware Installation Guide* (Product code: 80011)
- *GX 550 Hardware Installation Guide* (Product code: 80077)
- *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 Frame Relay Configuration Guide* (Product code: 80071)
- *NavisCore Diagnostic and Troubleshooting Guide* (Product code: 80074)
- *NavisCore Console Commands User's Guide* (Product code: 80075)

Third Party

- *Solaris 2.4 System Configuration and Installation Guide*
- *HP OpenView Windows User's Guide* (for HP 9000 Series and Sun SPARCstation)
- *SYBASE Command Reference Manual*
- *SYBASE System Administration Guide*

Customer Comments

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

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

Customer Support

To obtain release notes, technical tips, 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. It provides a suggested reading path to follow, depending on your network needs. Some chapters provide information on ATM 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. For more information on how various ATM options can improve your network services, see the *Networking Services Technology Overview* or consult your Ascend representative.

Logical Ports

The following chapters describe ATM logical ports:

- Chapter 2 provides an overview of ATM logical port types and features. Read this chapter if you are unfamiliar with basic ATM UNI concepts such as ILMI and signaling, or if you need more information on ATM VPI/VCI addresses. This chapter also describes the administrative tasks you perform for all logical ports.
- Chapter 3 describes how to configure ATM logical ports on a CBX 500 or GX 550 switch. This chapter includes information on configuring the logical port options you need if you plan to use SVCs in your network.
- Chapter 4 describes how to configure ATM logical ports on either the B-STDX or the CBX 500 6-port Frame DS3/E3 module. Note that since the B-STDX is not a true ATM switch, many of the parameters you need to configure for the various ATM logical port types are different from the CBX or GX; in addition, the B-STDX does not provide ATM features for signaling and SVCs. These same ATM exceptions exist for the CBX 6-port Frame DS3/E3 module.

ATM Trunks

Chapter 5 describes how to configure the following types of ATM trunks:

- ATM Direct Trunks
- ATM OPTimum (Cell) Trunks
- ATM OPTimum Frame Trunks (B-STDX only)

For information on each of these trunk types, review the trunk logical port descriptions in Chapter 2 and Chapter 4. Chapter 5 also provides instructions for using trunk backup.

PVCs

Chapter 6 describes how to configure point-to-point and point-to-multipoint PVCs. In addition, Chapter 7 explains how to configure optional Management VPI/VCI and Management PVC connections.

A Management VPI/VCI connection enables the NMS to connect to the gateway switch via an ATM router or ATM network interface card (NIC). A Management PVC provides access to the CBX 500 management port.

Network-wide Features

The following chapters explain how to configure features that you can use throughout your ATM network.

- Chapter 8 describes how the CBX 500 and GX 550 use traffic descriptors to define a service contract which guarantees that a specified amount of data is delivered. You configure a set of traffic descriptors that you can use when you define PVCs throughout your ATM network; this *configurable control circuit* feature enables you to ensure Quality of Service (QoS). Note that ATM services for a B-STDX do not use traffic descriptors.

- Chapter 9 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 10 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.

Using this feature, a logical port is given a service name. When you configure a PVC, select this service name as the logical port endpoint. If you activate the backup port, all PVCs on the failed primary port are automatically rerouted. Note that you should not configure SVCs on a logical port that is also designated as a backup port.

If you use resilient UNI features in conjunction with OC3/STM-1 or OC12/STM-4 automatic protection switching (APS) functions, you can configure a PVC to automatically revert to the backup port if the primary port fails.

SVCs

The CBX 500 and GX 550 offer switched virtual circuit (SVC) features. With SVCs, connections are not predefined as they are for PVCs. Instead, end stations use a signaling protocol to indicate to the ATM network the endpoint to which it routes the SVC request. To support SVC services, each user endpoint is assigned a unique address which identifies the endpoint and enables the network to route the SVC request.

The following chapters describe basic SVC concepts and configuration:

- Chapter 11 provides an overview of SVC concepts. Read this chapter if you are unfamiliar with SVC address formats and registration or need more information on route determination or address translation. This chapter also describes how to use network ID addressing.
- Chapter 12 describes how to configure SVC node and port prefixes and port addresses for each SVC address format. This chapter includes information on configuring network ID addressing.

The following sections describe *optional* SVC features you can use in your network to take advantage of ATM signaling functions.

SVC Proxy Signaling

Chapter 13 describes SVC proxy signaling. SVC proxy signaling is an *optional* CBX 500/GX 550 feature that enables a single signaling entity to signal on behalf of multiple endpoints. You can use proxy signaling to allow endsystems that do not understand ATM signaling to set up SVCs via a proxy signaling agent (PSA). The PSA performs all signaling functions on behalf of the endsystem, known as the proxy signaling client (PSC).

SPVCs

Chapter 14 describes SPVCs. The network uses signaling to establish a soft PVC (SPVC). The network management system provisions one end of the SPVC with the address identifying the egress interface from the network. Once the SPVC configuration is in place, the switch at one end of the SPVC initiates the signaling. This calling end is responsible for establishing, releasing, and re-establishing the SVC request.

Closed User Groups

Chapter 15 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 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.

Port Security Screening

Chapter 16 describes Port Security Screening. This feature 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 can allow/disallow incoming and outgoing SVCs.

PNNI

Chapter 17 describes how to configure the ATM PNNI routing protocol in your Ascend network. Private Network-to-Network Interface (PNNI) is a standard designed by the ATM Forum. This standard defines both an ATM routing protocol and an ATM signaling protocol. Ascend supports PNNI on both the CBX 500 and GX 550 switch platforms. For a detailed explanation of PNNI routing, see the *ATM Forum Technical Committee Private Network-Network Interface Specification Version 1.0* (af-pnni-0055.000), available from the ATM Forum's web site (http://www.atmforum.com).

2

About ATM Logical Ports

This chapter describes ATM concepts you need to understand before you can configure ATM services for an Ascend B-STDX, CBX 500, or GX 550 switch. Note that the B-STDX does not support all ATM features. For specific information about the B-STDX ATM implementation, see Chapter 4.

For details on configuring an ATM NNI logical port for PNNI routing, see Chapter 17.

ATM UNI Concepts

This chapter describes the following CBX 500 and GX 550 logical port types:

- UNI DCE and DTE
- NNI
- Direct (Cell) Trunk
- OPTimum Cell Trunk

For information about the logical port types you can configure on a B-STDX and the CBX 500 6-port Frame DS3/E3 module, see Chapter 4, "Configuring ATM Logical Ports on the B-STDX."

ATM UNI DCE and DTE

This section describes some of the concepts you need to know when defining ATM UNI DCE and ATM UNI DTE logical ports for a CBX 500 and GX 550 switch. You can configure a single ATM UNI logical port on a physical port to support the following standard protocol functions:

- ATM UNI 3.0, 3.1, and 4.0
- ITU UNI
- IISP 3.0 and 3.1

You use the ATM UNI DCE logical port type to communicate with most ATM customer premise equipment (CPE). An ATM UNI DCE logical port represents the "network side" equipment. This logical port supports all types of PVCs as well as SVCs. For SVC applications, the ATM UNI DCE logical port assumes the role of the network side of the UNI signaling interface.

You can also use the ATM UNI DCE as a feeder port for Ascend OPTimum trunks or Virtual UNIs. When used as a feeder port, you can still use the ATM UNI DCE logical port for PVC and SVC applications.

The ATM UNI DTE logical port type has the identical functionality of the ATM UNI DCE logical port with one exception. For SVC applications, the ATM UNI DTE assumes the role of the "user side" of the UNI signaling interface.

Using Interim Local Management Interface (ILMI)

ILMI is a management information base (MIB) that provides status and communication information to ATM UNI devices; this information includes status and statistics for virtual paths, connections, and address registration. Both the CBX 500 and GX 550, as well as the B-STDX, support the ILMI MIB.

If you want to use ILMI, make sure both endpoints of the UNI connection support this MIB. When you enable ILMI on an ATM UNI DCE logical port, the switch polls the attached device every five seconds. Five seconds is the *polling period*. If no response is received after four consecutive polls (*loss threshold)*, the switch considers the ILMI state to be down.

Ascend recommends that you enable ILMI support before you provision circuits. Under certain conditions, enabling ILMI after you provision circuits on a logical port may cause negative bandwidth with the associated QoS classes (including CBR). Note that if you enable ILMI on a logical port, and for some reason the ILMI state is down, the logical port does not go down.

Table 2-1 describes the differences between UNI DCE and DTE logical ports with ILMI enabled and disabled.

Table 2-1. Logical Ports and ILMI Settings

Port Type	Effect On	With ILMI Enabled	With ILMI Disabled
UNI DCE	Address Registration	• Send node prefixes • Send port prefixes • Accept addresses (qualified against configured prefixes)	None
	Remainder of ILMI MIB	Switch responds to get and get next commands sent by attached devices.	None
UNI DTE	Address Registration	Accept prefixes (and optionally qualify prefixes against configured prefixes).	None
	Remainder of ILMI MIB	Switch responds to get and get next commands sent by attached devices.	None

Using Logical Port Signaling

This section describes the default signaling tuning parameters for an ATM UNI logical port. (Note that ATM logical ports on a B-STDX or the CBX 500 6-port Frame DS3/E3 module do not support signaling.) In an ATM network, signaling is responsible for establishing and releasing SVCs. Signaling is used only on ingress and egress ports, including user-to-network, network-to-user, and network-to-network ports.

On ATM UNI DTE or ATM UNI DCE logical ports, if you change the default values and later change the UNI version for the port, the NMS prompts you to overwrite current settings with the default tuning parameters for the new UNI version. Ascend recommends that you set the logical port signaling options before you provision circuits. Under certain conditions, enabling signaling after you provision circuits on a logical port may cause negative bandwidth with the associated QoS classes (including CBR).

ILMI and Signaling Example

Under certain conditions, enabling ILMI and/or signaling after you provision circuits on a logical port may cause negative bandwidth for the associated QoS classes. For example, you create an ATM DS3 logical port with both ILMI and signaling disabled. You then create a full-bandwidth CBR circuit (PCR=96000 cps) on this logical port. If you later enable ILMI and/or signaling on the logical port, the bandwidth now appears to be negative. If you need to modify the logical port admin status or if you modify the circuit, the circuit will no longer come back up due to insufficient bandwidth.

Configurable Control Circuits

The CBX 500 and GX 550 configurable control circuit feature enables you to configure forward and reverse traffic descriptors for the following:

- ATM UNI ILMI and Signaling control channels
- ATM Direct and OPTimum trunk signaling and node-to-node management traffic

The switch software views a control circuit as a virtual circuit link (VCL) between the logical port and the internal switch processor. When you configure a control circuit, the switch creates a VCL between this port and the ATMizer. The logical port uses the forward traffic descriptor to police traffic flowing into the switch. It uses the backward traffic descriptor to determine the service category and equivalent bandwidth for the control circuit. The backward traffic descriptor is also used to calculate the effective bandwidth of the circuit to be used for bandwidth management on the logical port.

For more information about traffic descriptors, see Chapter 8, “Configuring ATM Traffic Descriptors.”

ATM OPTimum Cell Trunk

The CBX 500 and GX 550 ATM OPTimum Trunk logical port type provides trunk connectivity between two Ascend switches that are not directly connected. In this application, some other network elements are separating the two Ascend switches. These network elements usually consist of ATM switches in another network. The network provider who manages the other ATM switches provisions a virtual path connection (VPC) to carry the Ascend trunk traffic. This VPC supports the trunk and carries all the associated trunk protocol, management data, PVCs, and SVCs between the two Ascend switches.

Starting in this release, it is possible to have VPCs traverse OPTimum trunks. This capability is dependent on the logical port configuration as well as the configuration of the interfacing network (see "Configuring the Virtual Path ID" on page 2-5). Prior to this release, VPCs could not traverse OPTimum trunks.

Before you can configure an ATM OPTimum trunk logical port, you must first configure an ATM UNI or NNI logical port with a minimal amount of bandwidth; this logical port acts as the feeder port.

Configuring the Virtual Path ID

The Virtual Path ID (VPI) is the identifier used for all circuits routed over this OPTimum trunk. The range of valid VPI values depends upon the number of valid VPI bits you set for the ATM UNI feeder port (see Table 2-2 on page 2-9). Enter a number from 0-*nnnn* to identify the virtual path for the ATM logical port. *nnnn* is equal to 2^P-1, where *P* is the value specified in the Valid Bits in VPI field for the UNI feeder port that shares this physical port (see the example on page 2-8).

For example, if you entered 4 in the Valid Bits in VPI field for the UNI feeder port, you can have up to 15 virtual paths on this port (2^4-1=15); if you entered 8 in the Valid Bits in VPI field, you can have up to 255 virtual paths on this port (2^8-1=255). The highest value you can enter (and therefore, the greatest number of virtual paths you can configure on the port) depends on the value you entered in the Valid Bits in VPI field for the ATM UNI feeder port. The OPTimum trunk's VPI must be unique to the port.

The network that interfaces with the OPTimum trunk must be configured to accept circuits with this VPI and any of its valid VCIs. To accomplish this, in the interfacing network create a PVC using this VPI and define the PVC circuit type as VPC (see Table 6-2 on page 6-12).

Using VPI 0 for OPTimum Trunks

Configuring OPTimum trunk logical ports to use VPI 0 allows up to 4096 PVCs to be created on one trunk. The use of VPI 0 instead of VPI 1 doubles the number of PVCs supported on a single OPTimum trunk.

ATM Direct Trunk

The CBX 500 and GX 550 ATM Direct Trunk is used to provide trunk connectivity between two directly connected Ascend switches. The ATM Direct Trunk carries all types of PVC, SVC, and management data.

ATM NNI

The CBX 500 and GX 550 ATM Network-to-Network Interface (NNI) logical port type enables you to connect ATM-based public networks belonging to two different carriers. This logical port type implements the B-ISDN-Inter-Carrier Interface (B-ICI) protocol, which facilitates the multiplexing of services for inter-carrier (RBOC and IXC) delivery. You can use an ATM NNI logical port as a feeder port for Ascend OPTimum trunks and virtual UNIs.

ATM NNI logical ports also support the PNNI routing protocol. To configure PNNI routing in your Ascend network, see Chapter 17. For a detailed explanation of PNNI routing, see the *ATM Forum Technical Committee Private Network-Network Interface Specification Version 1.0* (af-pnni-0055.000), available from the ATM Forum's web site (http://www.atmforum.com).

Virtual UNI/NNI

A Virtual UNI/NNI forms an extension of the standard "direct" UNI DCE/DTE or NNI logical port types. In an ATM network you can use virtual UNI/NNI logical ports to enable VP tunneling or to connect to a VP multiplexer. VP tunneling allows you to connect two switches using signaling via a virtual path through the ATM network (network-to-network connection class). See the example in Figure 2-1 on page 2-7.

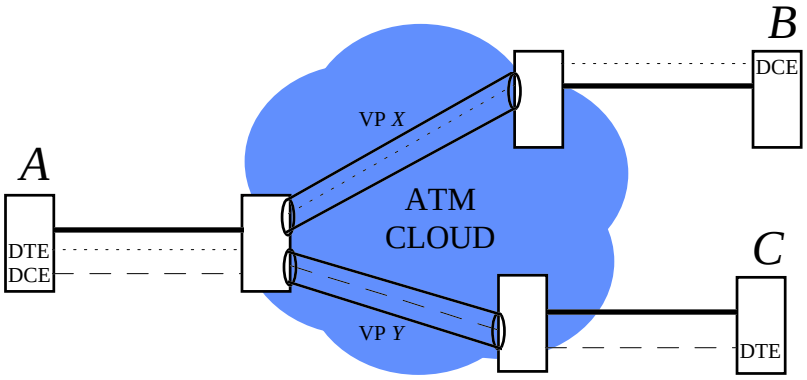

Figure 2-1. Two Virtual UNIs Through Central Network

VP multiplexing enables you to connect the GX 550/CBX 500 switch to a VP multiplexer using a direct UNI (or NNI) logical port on which you have configured several “virtual” UNI (or NNI) ports. The VPI address range you define for each virtual UNI/NNI port corresponds to a port on the VP multiplexer. The Virtual UNI/NNI terminates the signaling of an endsystem across a physical link. This method does not use VPCs and the configured logical port bandwidth can be used by any PVC on any VPI (network-to-endsytem connection class). See the example in Figure 2-2.

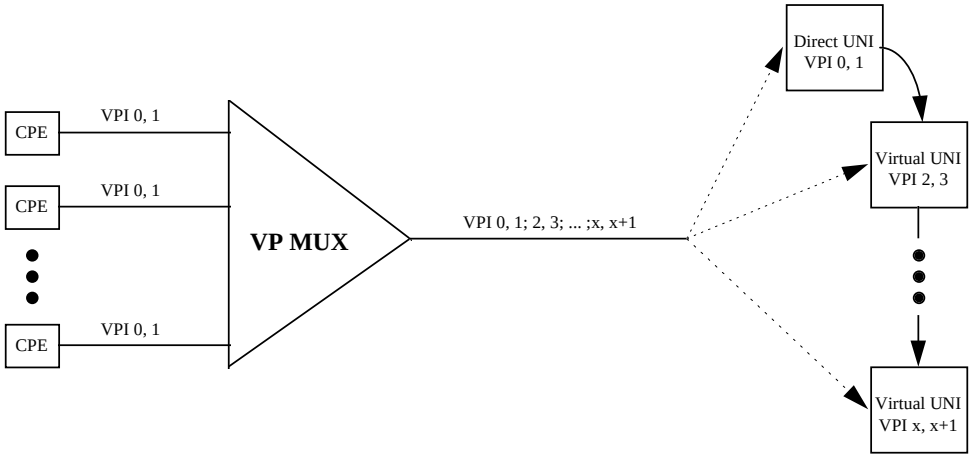

Figure 2-2. Virtual UNI with VP Multiplexer

Virtual Paths and Virtual Channels

To establish connections, ATM uses *virtual channels (VCs)* and *virtual paths (VPs)*. A virtual channel is a connection between two communicating ATM entities. It may consist of a group of several ATM links, CPE to central office switch, and switch-to-switch, or switch-to-user equipment. All communications proceed along this same VC, which preserves call sequence and provides a certain quality of service.

A virtual path is a group of VCs carried between two points. VPs provide a way to bundle traffic headed in the same direction.

Virtual path identifiers (VPIs) and *virtual channel identifiers (VCIs)* are hardware addressing identifiers (similar to Frame Relay's DLCI) that route cell traffic. The ATM cell header contains both a VPI and a VCI, which gives an ATM cell a unique VCI and associates it with a particular virtual path. Every ATM cell uses these VPI/VCI identifiers.

Switching equipment checks the VPI portion of the header to route traffic over certain trunks. It uses the VCI portion of the address to deliver the cell to an individual user within that destination. For more information about VPI/VCI hardware addressing, see the *Networking Services Technology Overview.*

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

Setting the Number of Valid Bits in the VPI/VCI

The Number of Valid Bits setting applies to the VPI and VCI range that you can use for VCCs (both PVCs and SVCs). The default values of VPI = 4 and VCI = 10 mean that you can use VCCs over the range of VPI = 0 – 15 (4 bits of VPI) and a VCI range of VCI = 32 – 1023 (10 bits of VCI). The values have no effect on VPCs, which you can provision anywhere over the VPI = 0 – 255 range; you can provision VPCs over the VPI = 0 – 4095 range if you use the NNI cell header format.

For the CBX 500, the valid range for the VPI field is 0 – 8 and the valid range for the VCI field is 6 – 14; for the GX 550, the valid range of the VPI field is 0 – 12 and the VCI field is 6 – 13. You may have to adjust these values in the following situations:

- In cases where the required VPI/VCI(s) of the attached devices are outside the range that the default values provide (VPI = 0 – 15 and VCI 32 – 1023).
- If you use this logical port as a feeder for OPTimum trunks or virtual UNIs, the VPI value limits the number of OPTimum trunks you can create on this physical port. The VCI value limits the number of circuits you can route over each OPTimum trunk.

  This OPTimum trunk/circuit trade-off is shown by the following formulas, where *P* represents the value in the Valid Bits in VPI field, and *C* represents the value in the Valid Bits in VCI field:

  Maximum virtual paths = $2^P - 1$
  Maximum virtual channels = $2^C - 32$
  $P+C \leq 14$

  For example, if you set the VPI value to 3 and the VCI value to 11, you can have up to 7 virtual paths on the port, and up to 2,016 virtual channels on each path.

Use Table 2-2 and Table 2-3 as a guide to set these values.

When you configure an OPTimum trunk or virtual UNI between two endpoints, the logical ports must match the VPI of the VPC that provides the connectivity between the two switches. The VPI range for the VPI/VCI valid bits setting for each endpoint must accommodate this VPI.

Table 2-2. Number of Valid Bits in VPI/VCI for CBX 500

[1] If Number of Valid VPI Bits =	Valid VPI Range Is	If Number of Valid VCI Bits =	[2] Valid VCI Range Is
0	0	0	Not Valid
1	0 - 1	1	Not Valid
2	0 - 3	2	Not Valid
3	0 - 7	3	Not Valid
4	0 - 15	4	Not Valid
5	0 - 31	5	Not Valid
6	0 - 63	6	32 - 63
7	0 - 127	7	32 - 127
8	0 - 255	8	32 - 255
Not Valid	–	9	32 - 511
Not Valid	–	10	32 - 1023
Not Valid	–	11	32 - 2047
Not Valid	–	12	32 - 4095
Not Valid	–	13	32 - 8191
Not Valid	–	14	32 - 16383

[1] Only 8 bits of the VPI are available on UNI type interfaces per ATM Forum standards.
[2] VCI 0 - 31 are reserved and cannot be used per ATM Forum standards.

Table 2-3. Number of Valid Bits in VPI/VCI for GX 550

[1] If Number of Valid VPI Bits =	Valid VPI Range Is	If Number of Valid VCI Bits =	[2] Valid VCI Range Is
0	0	0	Not Valid
1	0 - 1	1	Not Valid
2	0 - 3	2	Not Valid
3	0 - 7	3	Not Valid
4	0 - 15	4	Not Valid
5	0 - 31	5	Not Valid
6	0 - 63	6	32 - 63
7	0 - 127	7	32 - 127
8	0 - 255	8	32 - 255
9	0 – 511	9	32 - 511
10	0 – 1023	10	32 - 1023
11	0 – 2047	11	32 - 2047
12	0 – 4095	12	32 - 4095
Not Valid	–	13	32 - 8191

[1] Only 8 bits of the VPI are available on UNI type interfaces per ATM Forum standards.
[2] VCI 0 - 31 are reserved and cannot be used per ATM Forum standards.

Configuring VCC VPI Start and Stop for Virtual UNI/NNI

The CBX 500 and GX 550 provide a virtual UNI/NNI feature. The direct UNI/NNI provides the range of VCC VPI Start and Stop values from 0 to 15. The range of VPI start and stop values you define for the first virtual UNI/NNI must fall within this range; it cannot overlap with the range you define for subsequent virtual UNI/NNI ports.

For example:

Logical Port	VPI Start	VPI Stop
First Virtual UNI/NNI	2	5
Second Virtual UNI/NNI	6	10

For network-to-network virtual UNI/NNI configurations, the first VPI is used for VCCs only; additional VPIs are for VPCs only and provide “best effort” QoS. For network-to-endsystem virtual UNI/NNI configurations, there are no restrictions.

About Logical Port Bandwidth

The maximum amount of logical port bandwidth does not equal the physical port bandwidth due to the overhead associated with packaging ATM cells into the physical layer frames. This overhead is different for each physical media type as well as the different packaging methods. Table 2-4 provides a guide to mapping and converting physical layer bandwidth to logical port bandwidth.

Table 2-4. Physical and Logical Port Bandwidth Conversions

Physical Port Media Type	Physical Port Bandwidth (kbs)	Exact Logical Port Bandwidth (kbs)	Exact Logical Port Bandwidth (cps)	NMS Rounded Maximum Logical Port Bandwidth (kbs)	NMS Rounded Maximum Logical Port Bandwidth (cps)
OC-12/STM-4	622080	599040	1412830.19	599040	1412830
OC-3/STM-1	155520	149760	353207.55	149760	353207
ATM DS3 (with PLCP)	44736	40704	96000	40704	96000
ATM DS3 (with HCS direct mapping)	44736	44209.694	104268.15	44209	104266
ATM E3 (with HCS direct mapping)	34368	33920	80000	33920	80000
ATM E3 (with G.751 PLCP)	34368	30528	72000	30528	72000
T1	1544	1536	3622.64	1536	3622
E1	2048	1920	4528.3	1920	4528

In some cases, due to the way the switch stores logical port bandwidth, the NMS may have to round down non-integer maximum logical port bandwidth values to the nearest kbs value. For most applications, this does not cause any problems. However, if you need to run 100% line rate traffic through a policed PVC where you have rounded values, policing may cause minor cell loss.

Example

If you send 100% line rate traffic over an ATM DS3 interface that uses HCS direct mapping, the cells arrive at a rate equal to 44209.694 kbs or 104268.15 cps. Because of NMS rounding, the maximum PCR you can provision for this PVC is 104266. If you enable UPC on this PVC, approximately two cells every second are lost. For these cases, you may want to either adjust the traffic rate or disable UPC for this circuit.

Modifying Logical Port Bandwidth

You can modify logical port bandwidth on UNI and NNI logical ports even after you configure PVCs on this port. However, if you reduce the logical port bandwidth such that the new value is not sufficient to support all of the PVCs traversing the port, the available bandwidth enters a negative state. The PVC remains active until it has to be reestablished (i.e., trunk reroute, IOM reboot). If at this time the logical port does not have enough bandwidth to support the PVC, the PVC remains inactive due to insufficient bandwidth.

CBX 500 SP Thread Bandwidth Available for Logical Ports

The NMS and CAC enforce the SP fabric thread bandwidth such that each SP fabric thread is limited to 599.040 Mbs. This enforcement is done to ensure service is guaranteed even when two IOMs are placed on the same SP fabric thread. The 599.040 Mbs. number is derived from the maximum user cell bandwidth which the OC12/STM4 interface supports. (OC12/STM4 physical layer bandwidth is 622.080 Mbs., but the maximum user traffic that any OC12/STM4 port can support is 599.040 Mbs.) This 599.040 thread limitation is also derived from the maximum user cell bandwidth which the four OC3/STM1 interfaces support. (OC3/STM1 physical layer bandwidth is 155.020 Mbs., but the maximum user traffic that any OC3/STM1 port can support is 149.76 Mbs.)

For example, this NMS enforcement is noticeable whenever you attempt to provision two OC3 cards on the same SP fabric thread. As you provision logical ports, the NMS subtracts the assigned bandwidth from the 599.040 Mbs. total. After you provision four OC3 logical ports at the maximum 149.76 Mbs. bandwidth value, there is no bandwidth left for the other OC3 card and its logical ports.

Because of this, when you use two cards on the same fabric thread, Ascend recommends you allocate the bandwidth accordingly across all of the IOM ports. Note that you can oversubscribe the logical ports to avoid any negative implications associated with this restriction. You can use the Set QoS Parameters dialog box (accessible from the Add Logical Port dialog box) to oversubscribe a logical port.

About the Oversubscription Factor

The Oversubscription Factor percentage enables you to optimize the number of permanent virtual circuits (PVCs) and switched virtual circuits (SVCs) you can configure on the network by allowing you to oversubscribe the logical ports. If you configure oversubscription for the VBR classes of service, QoS is no longer guaranteed.

To ensure quality of service, monitor the network closely before you modify oversubscription values to exceed the minimum value of 100%. If you adjust the oversubscription percentage, monitor the cell-loss ratio to be sure the new setting does not impact quality of service.

The port bandwidth is reserved at runtime based on the sum of the effective bandwidth of each VC that uses the port. The CAC (Connection Admission Control) algorithm determines effective bandwidth of a virtual circuit (PVC and SVC). For a VBR circuit, the CAC uses the circuit's PCR, SCR, and MBS values. For CBR circuits, the CAC uses the PCR of the circuit. UBR circuits are assigned 100 cps bandwidth for load and reroute purposes, since it is a "Best Effort" service.

Appendix A describes how to tune the CAC to optimize your network. If you tune the CAC properly, you can optimize network resources without adversely affecting quality of service.

PVC routing is determined by either an OSPF algorithm or the network administrator (if you manually define the circuit path). Each time a PVC attempts to come up after configuration, OSPF reserves the required bandwidth on the port. OSPF deducts the amount of reserved bandwidth from the available virtual bandwidth pool for the applicable class of service.

The available virtual bandwidth can become negative in extreme situations. For the VBR-NRT queue, 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. Any PVC that fails during the time of the reroute is considered to be a new PVC when it attempts to come up after the trunk is rerouted.

Since inter-LAN traffic is bursty in nature, not all network traffic uses the network resources at precisely the same time. Basically, the higher you set the oversubscription factor, the less guarantee there is that user data will get through on the port; the trade-off is that you can provision more circuits on that port. If, however, all network traffic attempts to use the network resources at precisely the same time (for example, during multiple file transfer sessions over the same trunk), some traffic may be delayed or even dropped.

If you leave the Oversubscription factor set for the minimum value of 100%, the port delivers all user data for that service class without unanticipated delays or excessive cell loss. A value of 200% effectively doubles the virtual bandwidth available for that service class. (Ascend reserves a certain percentage of bandwidth for network management, routing updates, and other management traffic.)

VP Shaping on the CBX 500

VP shaping provides the ability to shape OPTimum trunk connections at a specified peak cell rate (PCR) while preserving QoS integrity. This feature is useful when multiple OPTimum trunk logical ports exist on one physical port and each trunk must traverse a VP of a fixed PCR. This new feature ensures that the maximum rate of the OPTimum trunk traffic does not exceed the specified PCR. See either the CBX or NavisCore software release notices for appropriate revision levels.

For a CBX 500, you can enable VP shaping only if the host IOM is equipped certain revisions of the ATM flow control processor (FCP) module. You can configure the FCP to either enable flow control or VP shaping; it does not support both functions at the same time. See either the CBX or NavisCore software release notices for appropriate revision levels.

Shaping is performed by assigning each virtual path (VP) a single queue on the FCP. Shaping is only performed on the first VP; The FCP treats the other VPs (used for VP switching) as VCs. One rate (PCR) is provisioned for each VP and the aggregate traffic of each VP is scheduled with this rate. As a result, the aggregate rate of the VP traffic never exceeds the provisioned PVC rate. Four queues are maintained for each shaped VP. The CBR queue has the highest priority, followed by VBR-RT, VBR-NRT, and ABR/UBR queues. VCCs within a shaped VP are mapped to the four queues according to their QoS.

Administrative Tasks

This section describes how to:

- Use templates to define a new logical port
- Delete ATM logical port components, including:
  - Circuits
  - Trunks

Using Templates

If you defined a logical port configuration and saved it as a template (see *Is Template* field), 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 (see Figure 3-1 on page 3-2).
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 ATM Logical Ports

Before you can delete an ATM logical port, verify the following:

Step 1. The logical port is not defined as part of a circuit. If it is, you must first delete this circuit.

Step 2. There are no trunks defined on this logical port.

Step 3. This logical port is not defined as the feeder (ATM UNI DCE/DTE or ATM NNI) for an existing ATM OPTimum trunk logical port.

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

- Circuits
- Trunks
- 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 (see Figure 6-1 on page 6-4).
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 (see Figure 5-2 on page 5-7). 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 the Logical Port

To delete the logical port:

1. Select the switch from which to delete 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 on which the logical port resides, and press the third (right) mouse button to display a popup menu. Select Logical Port. The Set All Logical Ports in PPort dialog box appears (see Figure 3-1 on page 3-2).
4. Select the logical port to delete. Make sure the Loopback field displays NONE.

Make sure this logical port is not the logical port used as the feeder for an ATM OPTimum trunk. If this is the case, either delete the OPTimum trunk logical port or first define another feeder before you delete this logical port.

5. Choose Delete.
6. Choose Close.

3

Configuring CBX 500 or GX 550 Logical Ports

This chapter provides instructions for configuring ATM logical ports on a CBX 500 or a GX 550 switch. For additional configuration information and a description of Ascend's ATM logical port service, see the following chapters:

- For an overview of Ascend's ATM service, see Chapter 2.
- For information about configuring the CBX 500 6-port Frame DS3/E3 module, or information about configuring ATM logical ports on a B-STDX, see Chapter 4.
- For details on configuring an ATM NNI logical port for PNNI routing, see Chapter 17.

Accessing ATM Logical Port Functions

To access the Logical Port functions in NavisCore:

1. Select the switch to which you want to add a logical port.
2. Log in to NavisCore using either a provisioning or operator password.
3. From the Administer menu, select Ascend Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
4. Select the physical port you want to configure and press the third (right) mouse button to display a pop-up menu. Select Logical Port. The following dialog box appears.

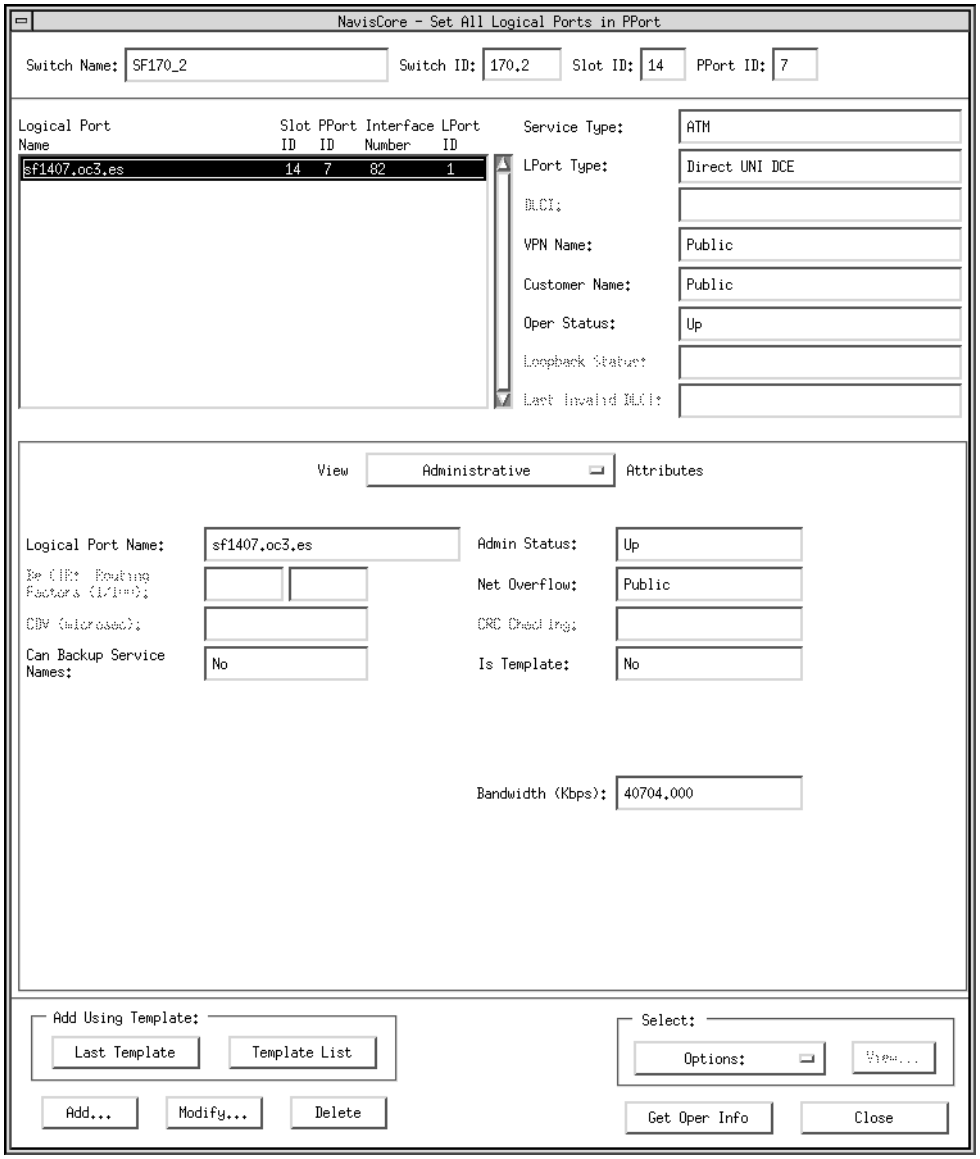

Figure 3-1. Set All Logical Ports in PPort Dialog Box

- To learn more about the Set All Logical Ports in PPort dialog box fields, continue with the following section.
- To begin defining a logical port, proceed to page 3-5.

About the Set All Logical Ports Dialog Box

The Set All Logical Ports In PPort dialog box displays information about existing logical ports and enables you to add a new logical port. It also provides several buttons that you can use to access additional logical port functions, such as add, modify, and delete. Table 3-1 describes dialog box status fields and buttons.

Table 3-1. Set All Logical Ports in PPort Dialog Box Fields and Buttons

Field/Command	Action/Description
Service Type	Displays ATM.
LPort Type	Displays the logical port type: either UNI DCE, UNI DTE, Direct Trunk, or OPTimum Trunk.
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).
View Attributes (option menu)	Displays the appropriate attributes configured for the selected option. See one of the following sections for more information: • "Administrative Attributes" on page 3-9 • "ATM Attributes" on page 3-11 • "ILMI/Signaling/OAM Attributes" on page 3-15 • "Flow Control Processor Attributes" on page 3-21 • "SVC VPI/VCI Range Attributes" on page 3-23 • "Traffic Descriptor Attributes" on page 3-24 • "OPTimum Trunk VPI Range" on page 3-25 • "Defining SVC Attributes" on page 3-31 • "SVC Routing Priorities" on page 3-39

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

Field/Command	Action/Description
Add Using Template	If you have already defined a logical port configuration and saved it as a template, use this option to define a new logical port using similar parameters. See "Using Templates" on page 2-15 for more information.
Add	Adds a new logical port.
Modify	Modifies the selected logical port. The Modify command displays dialog boxes that are similar to those displayed when you Add a logical port; however, you cannot modify the logical port name and the logical port type.
Delete	Deletes the selected logical port. For more information, see page 2-15.
Get Oper Info	Updates the logical port screen with current information from the switch and logical port.
Select: Options menu	Use the Select: Options menu to view logical port options. Once you select an option from this list, choose View to access the information.  The following fields describe the options you can select. *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 9, "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 3-27 for more information. *NTM Parameters* (CBX 500 only) – Displays the network traffic management (NTM) parameters for the selected logical port. *NTM Statistics* (CBX 500 only) – Displays the NTM statistics for the selected logical port. *NDC Statistics* (CBX 500 only) – Displays the network data collection (NDC) statistics for the selected logical port. *IP Parameters* – Accesses the NavisCore IP Navigator logical port functions. *Accounting* (B-STDX and CBX 500 only) – Accesses the NavisXtend Accounting server logical port functions. *Screen Assignments* – Displays the SVC port security screen assignments for the selected logical port. See Chapter 16 for more information.

Defining a Logical Port

The following diagram highlights the process for defining logical ports.

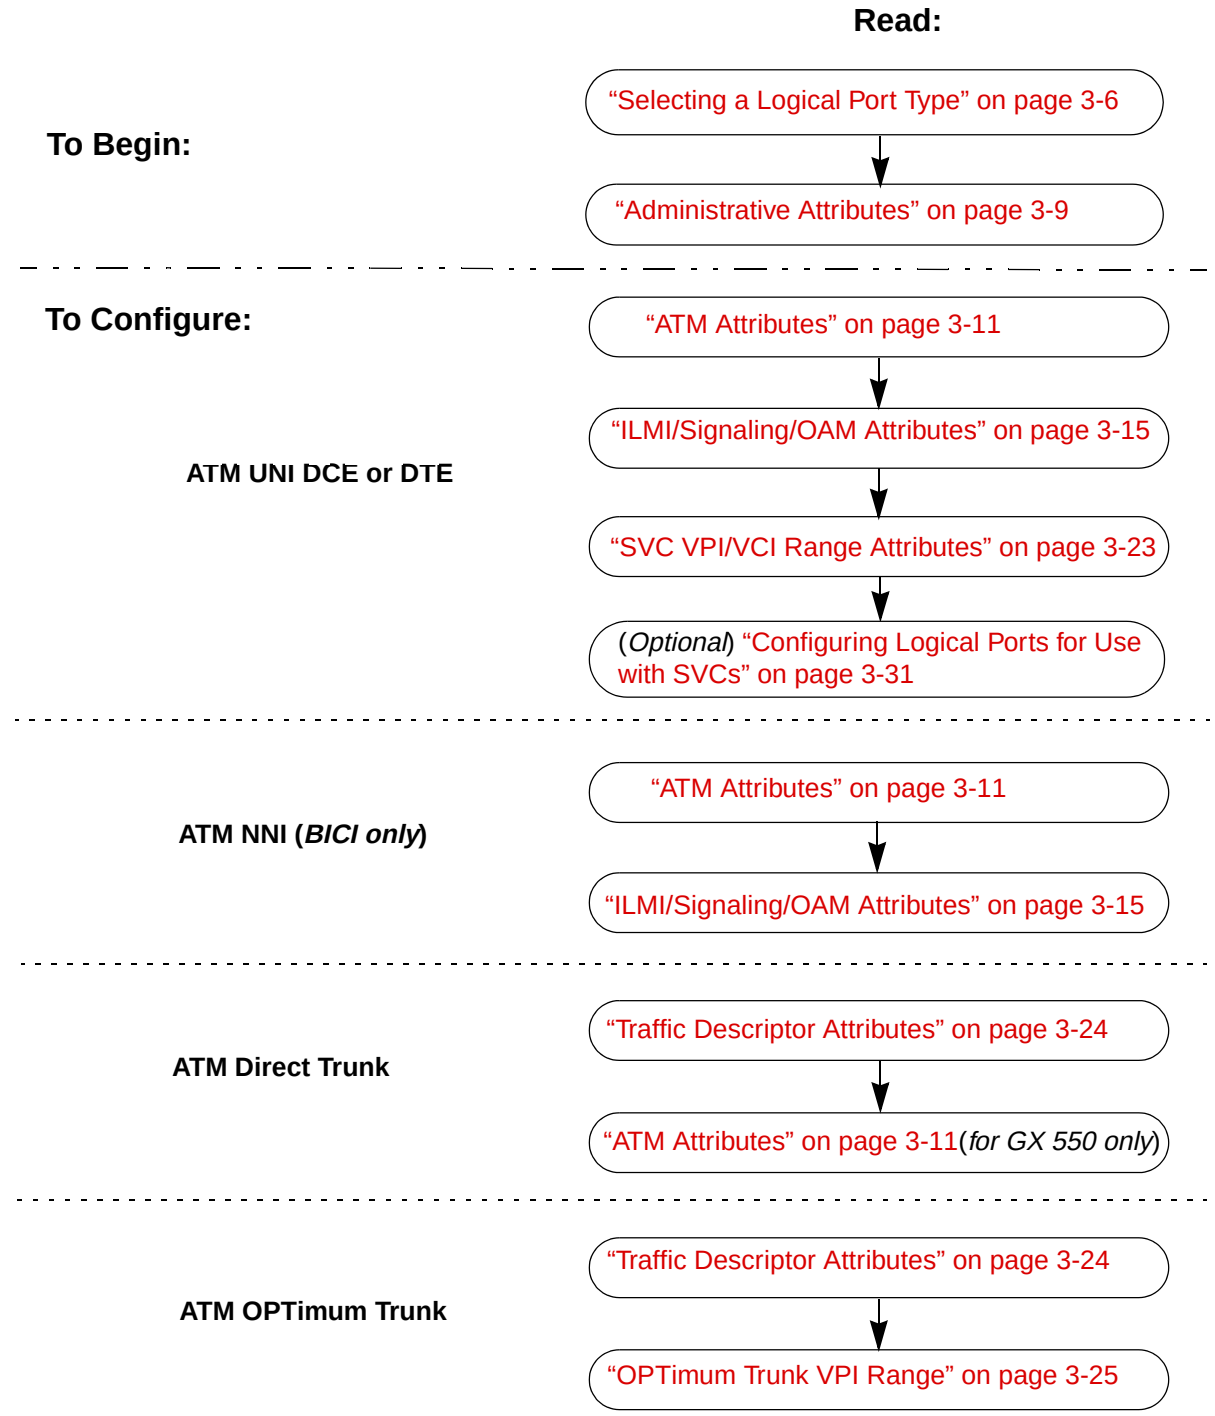

Selecting a Logical Port Type

To select a logical port type:

1. Complete the steps in "Accessing ATM Logical Port Functions" on page 3-2.
2. Choose Add to define a new logical port. The following dialog box appears.

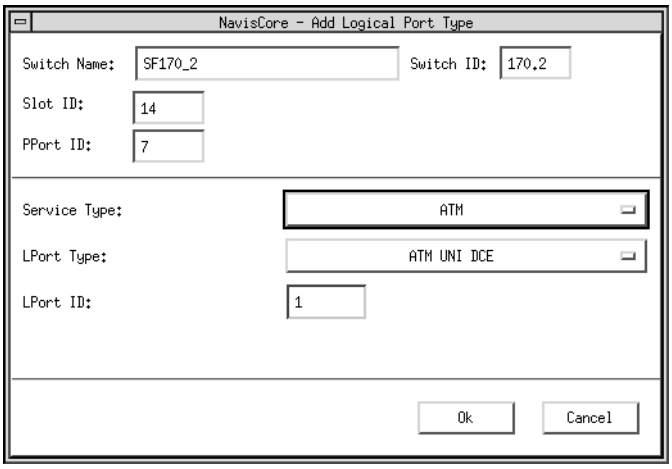

Figure 3-2. Add Logical Port Type Dialog Box

3. Select the required logical port type: ATM UNI DCE, ATM UNI DTE, ATM NNI, ATM OPTimum Trunk, or ATM Direct Trunk.

> For instructions on configuring an ATM NNI logical port for use with the PNNI routing protocol, see Chapter 17.

4. Choose OK. The Add Logical Port dialog box appears. The sample dialog box in Figure 3-3 on page 3-7 shows an ATM UNI DCE logical port.

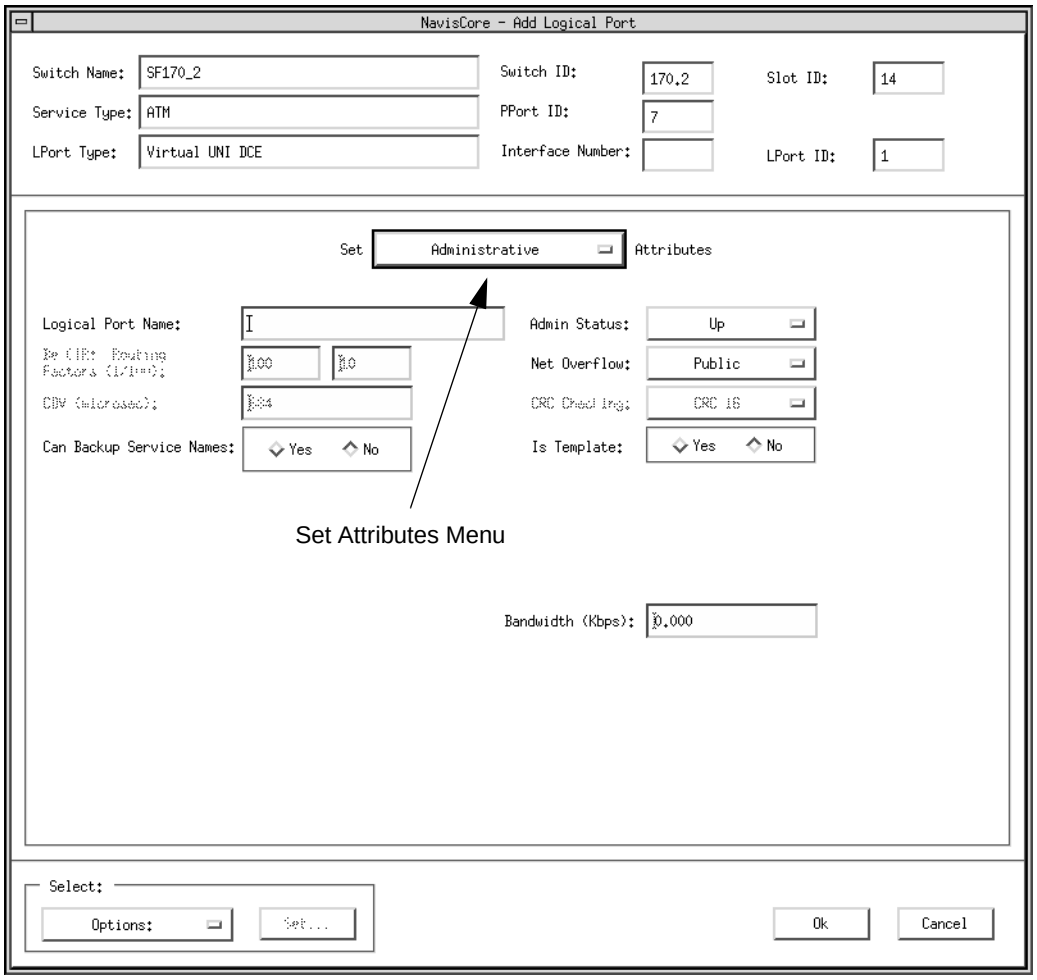

Figure 3-3. Add Logical Port Dialog Box (Virtual UNI Logical Ports)

About the Set Attributes Menu

When you configure logical ports, the Add Logical Port dialog box (Figure 3-3) contains a variety of parameters which you must specify. During this procedure, use the Set Attributes menu on the Add Logical Port dialog box to configure the following:

Administrative — Administrative options, including logical port name, admin status, and bandwidth. See Table 3-2 on page 3-9 to set these attributes.

ATM — ATM-specific options, including the number of valid bits in the VCI and VPI and ATM protocols. You can also enable or disable the Call Admission Control (CAC) or Usage Parameter Control (UPC) functions from this display. See Table 3-3 on page 3-11 to set these attributes.

ILMI/Signaling/OAM — These selections display options that enable you to fine-tune your ATM service. See page 3-15 to set the following attributes:

ILMI – A management information base (MIB) that provides status and communication information to ATM UNI devices and provides for a port keep-alive protocol. This selection also provides an option to configure the traffic characteristics for the ILMI control channel.

Signaling – A signaling protocol that supports the dynamic creation of ATM virtual circuits. To configure SVC signaling parameters, access the Set Logical Port Signaling Tuning Parameters dialog box (Figure 3-6 on page 3-18). This selection also provides an option to configure the traffic characteristics for the signaling control channel.

OAM – A parameter that enables the logical port to generate operations, administration, and maintenance (OAM) alarms.

ATM FCP — Displays options that enable you to configure logical ports for the CBX 500 ATM flow control processor. See Table 3-6 on page 3-21 to set these attributes.

SVC VPI/VCI Range — Configure a separate SVC VPI/VCI address range contained within the PVC VPI/VCI address range. By creating two separate address ranges, you can use one VPI/VCI range for PVCs on a logical port and a different (smaller) VPI/VCI range for SVCs on the same logical port. This addressing scheme enables a CBX 500/GX 550 to interoperate with an SVC-capable CPE that only supports VPI 0 for SVCs. You can set the VPI/VCI range to limit SVCs to VPI 0, while allowing PVCs to utilize the full VPI range. See Table 3-8 on page 3-23 to set these attributes.

Traffic Descriptors — Configure ATM traffic descriptors for the trunk logical ports. See page 3-24 to set these attributes.

OPTimum Trunk VPI Range — Specify the range of VPIs that can be created over an OPTimum trunk. These options work in conjunction with IP Navigator's Mulipoint-to-Point tunneling, a feature that is used to switch IP traffic through an Ascend cloud using ATM VP switching. You also use this option to configure the OPTimum trunk to handle VPCs. See Table 3-8 on page 3-26 to set these attributes.

SVC Parameter Attributes — Define various SVC screening and SVC handling parameters for each logical port on the switch. See page 3-31 to set these attributes.

SVC Routing Priorities — Assign bandwidth and bumping priorities to SVCs based on ingress QoS class. See page 3-39 to set these attributes.

Continue with the following sections to configure these attributes.

Administrative Attributes

From the Set [Administrative] Attributes display (Figure 3-3 on page 3-7), complete the fields described in Table 3-2.

Table 3-2. Set Administrative Attributes Fields

Field	Logical Port	Action/Description
Logical Port Name	All	Enter a unique alphanumeric name for this port. NavisCore uses this name to reference the logical port.
Admin Status	All	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. 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.
Can Backup Service Names	UNI DCE/DTE	Select Yes to configure a logical port for backup service in a fault-tolerant PVC configuration. A fault-tolerant PVC configuration enables a logical port to serve as a backup for any number of active UNI ports. For more information about fault-tolerant PVCs, see Chapter 10, "Configuring Fault-Tolerant PVCs." ***Note:*** *Ascend does not recommend that you configure SVCs on a logical port that is also designated as a backup port in a fault-tolerant PVC configuration.*
Net Overflow	UNI DCE/DTE and NNI	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 (VPN). To assign this logical port to a specific VPN and customer, see page 9-6. 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.
Is Template	All	*(Optional)* Saves these settings as a template to configure another logical port with similar options. To create a template, choose Yes.

Table 3-2. Set Administrative Attributes Fields (Continued)

Field	Logical Port	Action/Description
Bandwidth	All	Enter the amount of bandwidth for this logical port. The default is the amount of bandwidth remaining from the physical clock rate less any logical ports already configured. • If you are defining more than one ATM UNI or NNI logical port type on this port (*Virtual UNI/NNI*), be sure to adjust the bandwidth value to accommodate these virtual ports. • If you are defining an OPTimum cell trunk on this port, configure this UNI logical port with a minimal amount of bandwidth. For specific guidelines on configuring bandwidth with the various physical port types, see page 2-12.
VP Shaping	Virtual UNI/ NNI and OPTimum trunk ports with FCP support	Enables or disables VP shaping. VP shaping provides a method for traffic sent over an Ascend switch through another network to comply with the purchased traffic contract in that other network. The ATM flow control processor functions shape individual cell trunk or virtual UNI/NNI logical port traffic at the configured VP shaping rate. ***Note:*** *See "VP Shaping on the CBX 500" on page 2-14 for more information.*
VP Shaping Rate	Virtual UNI/ NNI and OPTimum trunk ports with FCP support	If you enable VP shaping, enter a value between 100 – *maximum logical port bandwidth* in cells per second. See Table 2-4 on page 2-11 for these values.
CDV	OPTimum trunk	The CDV field enables you to enter a cell delay variation value (in µsecs) that will be added to the Ascend default trunk CDV. For CBR traffic this default is 250 µsecs, and for VBR traffic the value is 500 µsecs. This logical port CDV value is zero by default. If you believe that the path through the network providing the OPTimum trunk connectivity will introduce additional cell delay variation (above the value provided by the Ascend default), enter the appropriate value in this field.

To continue this configuration:

- If this is a UNI or NNI logical port, continue with the following section.
- If this is a trunk logical port type, proceed to "Traffic Descriptor Attributes" on page 3-24.

ATM Attributes

The Set ATM Attributes option is available for UNI and NNI logical port types, as well as for ATM Direct Trunk logical ports residing on the GX 550.

> For the ATM Direct Trunk logical port type, the Number of Valid Bits in VCI field is the only field you need to configure.

Select Set [ATM] Attributes and complete the fields described in Table 3-3.

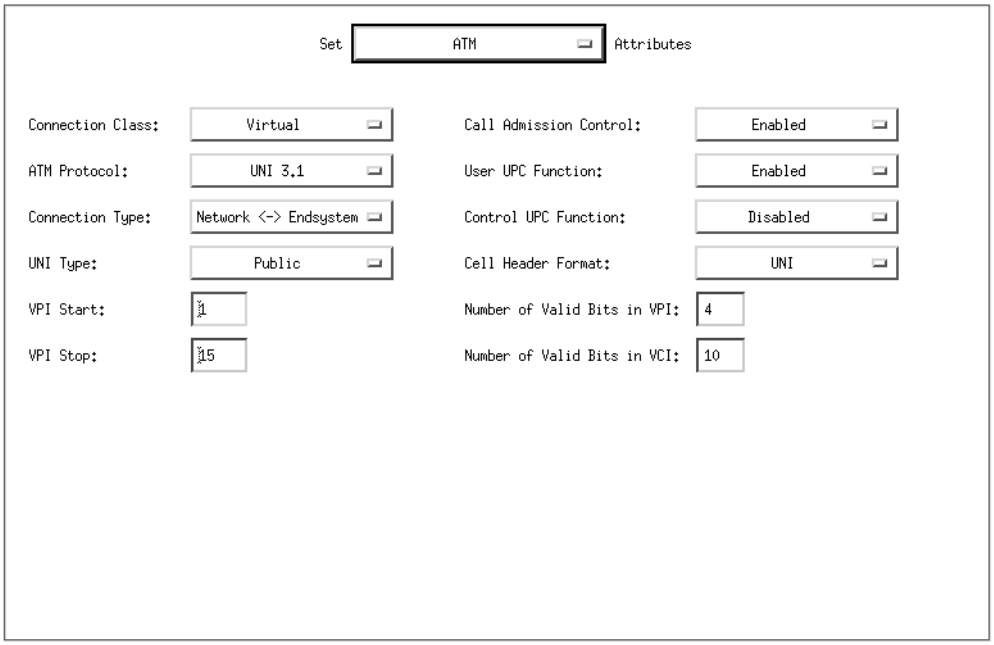

Figure 3-4. Set ATM Attributes (UNI Logical Ports)

Table 3-3. Set ATM Attributes Fields

Field	Logical Port	Action/Description
Connection Class	UNI DCE/DTE, or NNI	Displays the logical port connection type, either direct or virtual. This field is set to Direct when you configure the first UNI/NNI logical port on this physical port. When you configure subsequent UNI/NNI ports on this physical port, this field displays Virtual. For Trunk logical port types, this field defaults to Direct and cannot be changed.

Table 3-3. Set ATM Attributes Fields (Continued)

Field	Logical Port	Action/Description
ATM Protocol	UNI DCE/DTE and NNI	The equipment to which you connect this port must support the protocol you select. The following lists the protocols that Ascend logical ports support: UNI 4.0 UNI 3.1 UNI 3.0 IISP 3.1 IISP 3.0 ITU UNI BICI 1.1 (*NNI only*) PNNI 1.0 (*NNI only*) The default Signaling Tuning parameters are based on the ATM Protocol you select. If you change the Signaling Tuning parameters for this port and later change the UNI version, the default Signaling Tuning parameters for the ATM Protocol you selected will overwrite these changes. For more information on Signaling Tuning parameters, see page 3-18.
Connection Type	UNI DCE	This option lets the switch know if it is attached to another switch or an endsystem. • Select Network <–> Endsystem if this port connects to a router or host. • Select Network <–> Network if this port connects to another ATM switch.
UNI Type	UNI DCE/DTE	Select Public if at least one end of this connection attaches to a public network. Select Private if this connection resides completely within a private network.
VCC VPI Start	Virtual UNI/ NNI	For a virtual UNI/NNI logical port, this field represents the VPI of the control channels (i.e., signaling and ILMI). For more information on VPI Start and Stop values, see page 2-10.
VCC VPI Stop	Virtual UNI/ NNI	To configure this value, use the following formula: `VCC VPI Stop <= (2 numvpibits – 1)` where “numvpibits” equals the value you configure for the Number of Valid Bits in VPI field (page 3-14). For more information on VPI Start and Stop values, see page 2-10.

Table 3-3. Set ATM Attributes Fields (Continued)

Field	Logical Port	Action/Description
Call Admission Control	UNI DCE/DTE and NNI	When enabled (the default), the port rejects a circuit creation request if there is not enough available bandwidth. When disabled, the port attempts to create a circuit even if there is not enough available bandwidth (for VBR Non-Real Time queue only). ***Note:*** *If you disable Call Admission Control for a logical port, you are effectively disabling Ascend's Call Master Connection Admission Control (CAC) function on that logical port. For more information about the CAC function, see Appendix A.*
User UPC Function	UNI DCE/DTE	Enables or disables the Usage Parameter Control (UPC) function for PVCs and SVCs. You can also enable or disable the UPC function for individual PVCs. If you need to enable the UPC function on a per-PVC basis, you must enable the UPC function on the logical port. *Enabled* – (default) Cells that do not conform to the traffic parameters are dropped or tagged as they come into the port. *Disabled* – All traffic, including non-conforming traffic, passes in through the port. If you disable the UPC function on a logical port, quality of service is no longer guaranteed on the network due to the potential for increasing the cell loss ratio on network circuits. For this reason, *Ascend recommends that you leave the UPC function enabled on all logical ports.* For information on UPC traffic parameters, see Chapter 8, "Configuring ATM Traffic Descriptors."
Control UPC Function	UNI DCE/DTE	Enables or disables policing on a user port for control circuits (signaling and ILMI) independent of user traffic. The default is disabled. Enable policing to prevent an attached device from overloading the switch with data on the control circuit. The switch polices the control circuit to pre-defined default traffic characteristics (see Chapter 8). The attached device typically needs to support per-VC shaping on the control channels. ***Note***: *If the attached device is another Ascend switch, do not enable policing since the CBX 500 and GX 550 do not support per-VC shaping on the control channels.*

Table 3-3. Set ATM Attributes Fields (Continued)

Field	Logical Port	Action/Description
NPC Function	NNI	Enables or disables the Network Parameter Control (NPC) function. *Enabled* – (default) Cells that do not conform to the traffic parameters are dropped or tagged as they come into the port. *Disabled* – All traffic, including non-conforming traffic, passes in through the port. If you disable the NPC function on a logical port, quality of service is no longer guaranteed on the network due to the potential for increasing the cell loss ratio on network circuits. For this reason, *Ascend recommends that you leave the NPC function enabled on all logical ports.*
Cell Header Format	All	This field controls the number of VPI bits in the ATM cell header for VPCs on the CBX 500 and VCCs and VPCs on the GX 550. Select UNI to use a range of 0 through 8. Select NNI to use a valid bits in VPI range of 0 through 12. See page 2-8 for more information.
Number of Valid Bits in VPI	UNI DCE/DTE and NNI	Specify a value that is within the valid range for either the NNI or UNI call header format. The default is 4. See page 2-8 for details.
Number of Valid bits in VCI	UNI DCE/DTE and NNI	Specify a value that is compatible with the desired VCI range on the port. The default is 10. See page 2-8 for details.
	GX 550 Direct Trunk only	This field affects the amount of connection entry resource that is reserved for VCs that traverse this trunk endpoint. The default value of 10 translates into a value of 2^10 or a minimum of 1024 connection entries being reserved for VCs on the trunk. The default value of 10 is the most efficient usage of the connection entry resource.

When you finish configuring the ATM Attributes, proceed to the following section.

ILMI/Signaling/OAM Attributes

The ILMI/Signaling/OAM attributes option is only available for ATM UNI and NNI logical port types. ATM UNI logical ports support ILMI, Signaling, OAM, and Proxy; NNI logical ports support OAM only.

For more information about ILMI and Signaling, see page 2-3. The fields on this dialog box also enable you to configure *optional* proxy signaling features for a UNI logical port. See Chapter 13, "Configuring SVC Proxy Signaling," for instructions on using proxy signaling.

Select Set [ILMI/Signaling/OAM] Attributes and complete the fields as described in Table 3-4 on page 3-16.

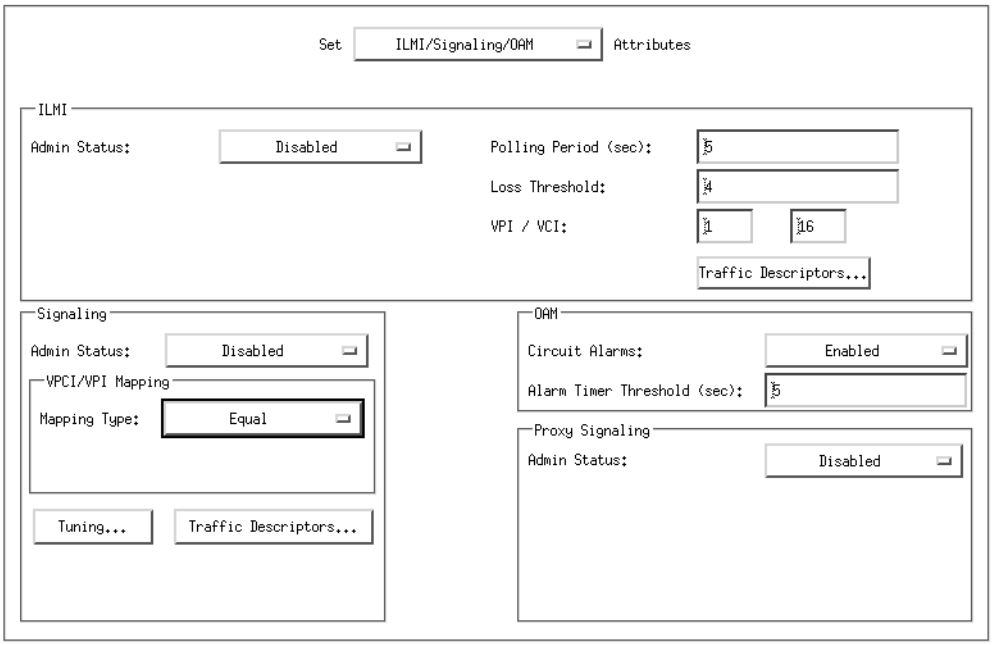

Figure 3-5. Set ILMI/Signaling/OAM Attributes (UNI Logical Ports)

Table 3-4. Set ILMI/Signaling/OAM Attributes Fields

Field	Action/Description
ILMI Attributes *(UNI DCE/DTE logical ports only)*	
Admin Status	Choose Enabled to reserve a percentage of bandwidth in the VBR-NRT QoS class for ILMI. You can override default values for bandwidth and QoS class by choosing the Traffic Descriptors button to assign traffic descriptors for the ILMI channel. When ILMI is Disabled (default), this bandwidth is not reserved. If the attached device cannot run ILMI, leave ILMI disabled. For information about ILMI support, see page 2-3. ***Note:*** *To use line loopback diagnostics, you must disable ILMI support. See the* NavisCore Diagnostic and Troubleshooting Guide *for more information.*
Traffic Descriptors	Choose this command to access the Set ILMI Traffic Descriptors dialog box. This option enables you to modify the traffic characteristics for the control channel. This feature is known as *configurable control channel.* See page 8-9 to complete the fields on this dialog box.
Polling Period (sec)	Specify the polling period (T) for an ILMI poll. The switch generates an ILMI poll every (T) seconds. The default is 5 seconds.
Loss Threshold	Specify the number of times (K) the logical port will issue an ILMI poll before the link is considered down. If no responses are seen in K x T seconds, the link is considered down. The default is 4.
VPI/VCI	Enter the ID of the virtual path (VPI) or virtual channel (VCI) you want to use for ILMI polling. The default is 0 for VPI, 16 for VCI.
DTE Prefix Screen Mode (DTE ports)	When a DTE port receives network prefixes from an external network, you can perform various levels of screening on them against the list of prefixes configured on the node and/or port. Select one of the following options: *Accept All* – No screening occurs; accepts all prefixes. *Node Prefix* – (default) Accepts only network prefixes that partially or fully match a configured node prefix. *Port Prefix* – Accepts only network prefixes that partially or fully match a configured port prefix. *Node or Port Prefix* – Accepts only network prefixes that partially or fully match either a configured node prefix or a configured port prefix. *Reject All* – Rejects all network prefixes received from an external network. For more information about node and port prefixes, see Chapter 11, “About SVCs.”

Table 3-4. Set ILMI/Signaling/OAM Attributes Fields (Continued)

Field	Action/Description
Signaling Attributes *(UNI DCE/DTE logical ports only)*	
Admin Status	Choose enabled to reserve a percentage of bandwidth in the VBR-NRT QoS class to support the UNI signaling protocol. You can override default values for bandwidth and QoS class by choosing the Traffic Descriptors button to assign traffic descriptors for the signaling channel. Use the default setting (Disabled) if you will only use this logical port for PVCs (that is, you will not create SVCs on the port).
Mapping Type (*Virtual UNI/NNI only*)	Use this button to configure the virtual path connection identifier (VPCI). See Chapter 13 for more information on VPCI mapping. *Equal* – (default) The VPI equals the VPCI. *Positive Offset* – The VPI of the corresponding circuit equals the VPCI plus the value you enter. 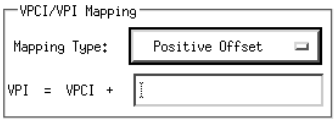 *Negative Offset* – The VPI of the corresponding circuit equals the VPCI minus the value you enter. *Mapping Type Table* – This option has additional functions that are used with Proxy Signaling. For information on configuring them, see Chapter 13.
Traffic Descriptors	Choose this command to access the Set Signaling Traffic Descriptors dialog box. This option enables you to modify the traffic characteristics for the control channel. This feature is known as *configurable control channel*. See "Defining Traffic Descriptor Attributes" on page 8-9 to complete the fields on this dialog box.
Tuning	Choose the Tuning command to display the Set Logical Port Signaling Tuning Parameters dialog box. For information about Tuning parameters, see page 3-18.
OAM Attributes	
Circuit Alarms	Select Enabled (default) to allow this logical port to generate OAM alarms. The switch uses these alarms to signal when the circuits have gone down. Select Disabled to disable OAM alarms on this logical port.
Alarm Timer Threshold (sec)	Before generating an OAM alarm, the switch waits until the circuit has been down for the time period you specify in this field. The default is 5 seconds.

When you finish configuring the ILMI, Signaling, and OAM attributes for this logical port, do one of the following:

- If the logical port you are configuring supports *optional* ATM Flow Control Processor functions, continue with "Flow Control Processor Attributes" on page 3-21.

- If this is a UNI logical port, continue with "SVC VPI/VCI Range Attributes" on page 3-23.
- If this is an NNI logical port, proceed to "Completing the Logical Port Configuration" on page 3-27.

Setting Logical Port Signaling Tuning Parameters

This section describes how to modify the signaling parameters for an ATM UNI logical port. For more information on signaling, see page 2-4.

To modify the signaling tuning parameters:

1. From the Add Logical Port dialog box (Figure 3-5 on page 3-15), choose the Tuning command in the Signaling box. The following dialog box appears.

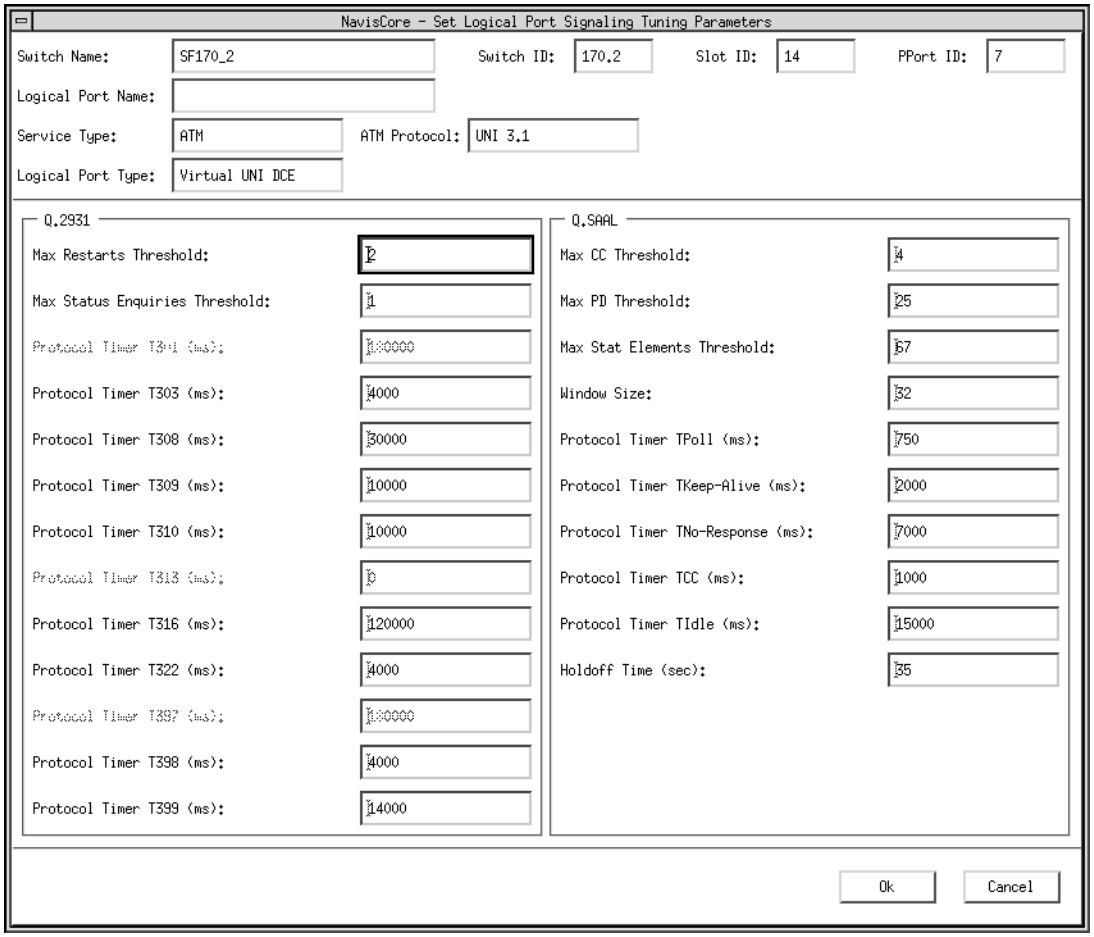

Figure 3-6. Set Logical Port Signaling Tuning Parameters

Use the Set Logical Port Signaling Tuning Parameters dialog box to set the signaling thresholds and timers and the Q.SAAL protocol data unit (PDU) thresholds and timers. In general, you should not change the default values. The displayed defaults are based on the ATM protocol you selected for the logical port (see page 3-12).

2. Complete the fields in Figure 3-6 using the information in Table 3-5. All timer field values are specified in milliseconds (1/1000ths of a second).

Table 3-5. Set Logical Port Signaling Tuning Fields

Field	Description
Signaling	
Max Restarts Threshold	The maximum number of restarts to send without a response. The default is 2.
Max Status Enquiries Threshold	The maximum number of status enquiries that can be unacknowledged before the SVC is dropped. The default is 1.
Protocol Timer T301	How long to wait for a CONNECT after ALERTING has been received. The default is 180,000 milliseconds (180s). (UNI 4.0, Q.2931/Q.2971 protocol only.)
Protocol Timer T303	How long to wait for a response after a SETUP protocol data unit (PDU) has been sent. The default is 4000ms.
Protocol Timer T308	How long to wait for a response after a RELEASE PDU has been sent. The default is 30000ms.
Protocol Timer T309	If Q.SAAL is down, how long to wait before SVCs are dropped. The default is 10000ms for the UNI 3.1 ATM protocol and 90000ms for UNI 3.0.
Protocol Timer T310	How long to wait for the next response after a CALL PROCEEDING PDU has been received. The default is 10000ms.
Protocol Timer T313	How long to wait for a response after a CONNECT PDU has been sent. This function defaults to 4000ms for DTE logical ports; it is disabled for DCE logical ports.
Protocol Timer T316	How long to wait for a response after a RESTART PDU has been sent. The default is 120000ms.
Protocol Timer T322	How long to wait for a response after a STAT ENQUIRY PDU has been sent. The default is 4000ms.
Protocol Timer T397	How long to wait for an ADD PTY ACK after PTY ALERTING has been received. The default is 180,00 milliseconds (180s). (UNI 4.0, Q.2931/Q.2971 protocol only.)
Protocol Timer T398	How long to wait for a response after a DROP PTY PDU has been sent. The default is 4000ms.
Protocol Timer T399	How long to wait for a response after an ADD PTY PDU has been sent. The default is 14000ms.

Table 3-5. Set Logical Port Signaling Tuning Fields (Continued)

Field	Description
Q.SAAL	
Max CC Threshold	The maximum number of transaction retries for control PDUs. The default is 4.
Max PD Threshold	The maximum number of data PDUs without a POLL. The default is 25.
Max Stat Elements Threshold	The maximum number of missing elements in a STATUS PDU. The default is 67.
Window Size	The maximum number of unacknowledged PDUs that can exist at any time. The default is 32. If you decrease this value, peer signaling slows down.
Protocol Timer TPoll (ms)	How often a poll is sent when the Q.SAAL is active. The default is 100ms if this port uses the UNI 3.0 or IISP 3.0 ATM protocol; the default is 750ms for all others.
Protocol Timer TKeep-Alive (ms)	How often a poll is sent when the Q.SAAL is in the transient state. The default is 2000ms.
Protocol Timer TNoResponse (ms)	The maximum amount of time that can pass without a STATUS PDU being received. The default is 7000ms.
Protocol Timer TCC (ms)	The retry time for control PDUs. The default is 1000ms.
Protocol Timer TIdle (ms)	How often a poll is sent when Q.SAAL is idle. This parameter does not apply to UNI 3.0 connections. The default is 15000ms.
Holdoff Time (sec)	The ATM signaling holdoff timer holds off the re-establishment of the ATM signaling connection after you modify a physical or logical port or after a physical port alarm is detected. This mechanism essentially converts SAAL reset conditions in to SAAL failure conditions (also described in Q.2931). The default is 35 sec.

3. When you finish, choose OK to return to the Add Logical Port dialog box (Figure 3-3 on page 3-7). Continue with page 3-17 to configure the remaining OAM attributes.

Flow Control Processor Attributes

Cascade Communications Resource Management (CCRM) cells are a subset of the ATM Forum's *ATM Traffic Management, Version 4.0*, Available Bit Rate (ABR) Resource Management (RM) cells. Backward Congestion Message (BCM) cells provide interoperability with other manufacturers' ATM switches.

If you change the CLP0+1, Discard, or Congestion parameter values, the switch object and the IOM turn yellow, indicating that the switch is "Marginal." You must PRAM Sync the card to resolve this condition. The PRAM Sync command enables you to correct inconsistencies between the NMS database and CBX 500 PRAM. See the *NavisCore NMS Getting Started Guide* for PRAM Sync instructions.

To configure flow-control mechanisms at the logical port level, select Set [ATM FCP] Attributes and complete the fields described in Table 3-6.

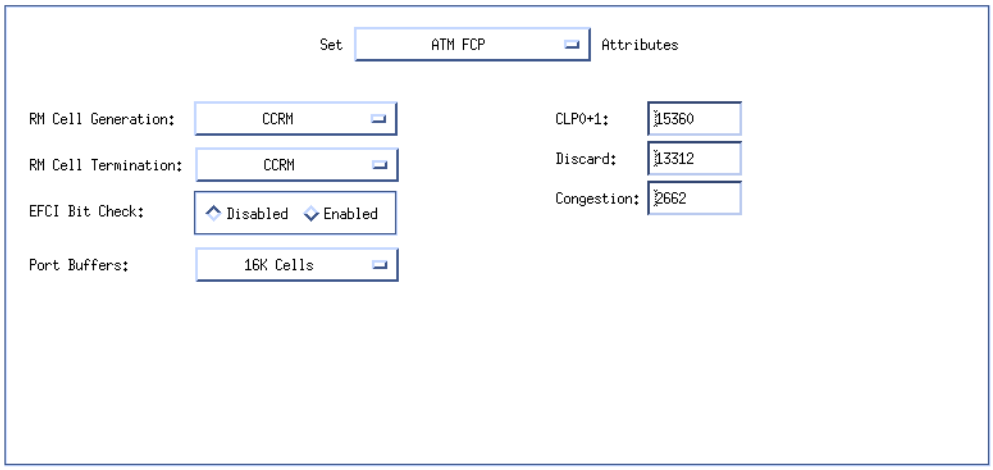

Figure 3-7. Set ATM FCP Attributes

Table 3-6. ATM FCP Attributes

Field	Action/Description
RM Cell Generation	Select the type of RM cell to generate for the VC, either CCRM or BCM. Use the No Loop option (default) to configure the VC to generate no RM cells. (See "CCRM Cell Generation" on page D-7 for information.)
RM Cell Termination	Select the type of RM cell to terminate for the port, either CCRM (default), or CCRM and BCM. (See "BCM Closed-Loop Flow Control" on page D-8 and "Terminating CCRM and BCM Cells" on page D-10 for information.)

Table 3-6. ATM FCP Attributes (Continued)

Field	Action/Description
EFCI Bit Check	Select either *Enabled* or *Disabled* (default). The EFCI Bit Check enables you to support control loops across switches that do not have the ATM Flow-Control Processor installed. These switches mark the EFCI bit in data cells to indicate network congestion. If this option is enabled, the ATM Flow-Control Processor reviews the EFCI bits in the cell stream when it generates a backward RM cell.
Port Buffers	Select the number of desired cell buffers per port. Port buffers enable you to configure the number of cell buffers for each port. The entire 64K-cell buffer can be divided among the ports on an IOM. Options include: 1K, 2K, 4K, 8K, 16K, 32K, and 64K. (Table D-3 on page D-14 lists the defaults.)
CLP0+1	Enter the desired value for the CLP0+1 threshold buffer. The CLP0+1 threshold enables you to reserve buffers before the maximum buffer capacity is reached. (Table D-3 on page D-14 lists the defaults.)
Discard	Enter the desired value for the Global Discard threshold buffer. Global Discard buffers enable you to reserve buffers for cell discard. (Table D-3 on page D-14 lists the defaults.)
Congestion	Enter the desired value for the Congestion threshold. You can configure the Congestion threshold to allow for some margin before the Global Discard buffer threshold is reached. This margin compensates for some of the closed-loop, flow-control delay in the network prior to discarding cells. (Table D-3 on page D-14 lists the defaults.)

- If this is a UNI logical port, continue with "SVC VPI/VCI Range Attributes" on page 3-23.
- If this is an NNI logical port, proceed to "Completing the Logical Port Configuration" on page 3-27.

SVC VPI/VCI Range Attributes

The SVC VPI/VCI Range attributes option is only available for UNI logical port types. The VPI/VCI address range fields allow you to design a VPC VPI/VCI or VCC VPI/VCI address range to match the capability of the equipment attached to this port.

Select Set [SVC VPI/VCI Range] Attributes and complete the fields described in Table 3-7.

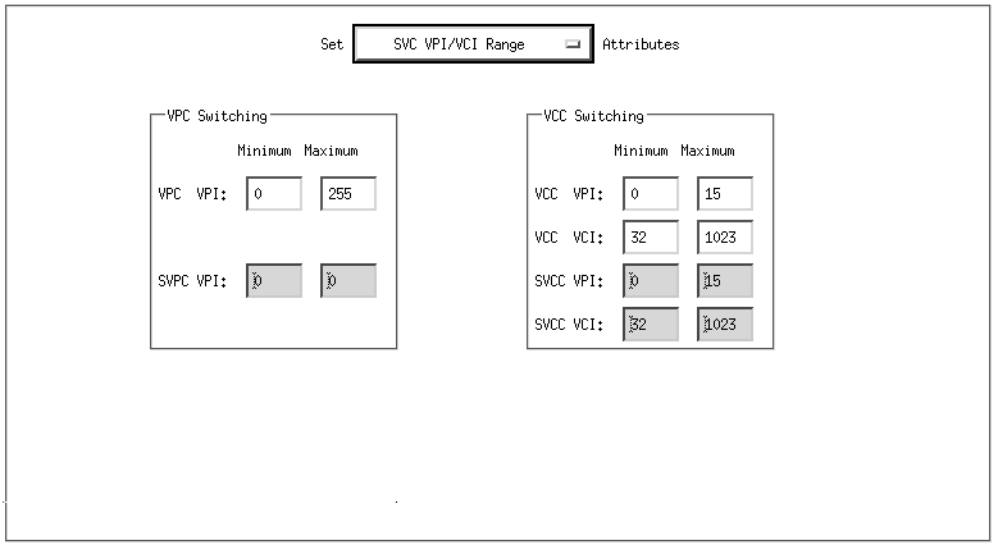

Figure 3-8. Set SVC VPI/VCI Range

Table 3-7. SVC VPI/VCI Range Attributes

Field	Action/Description
VPC VPI	Displays the VPI range for a VPC.
SVPC VPI (*UNI 4.0 only*)	Enter the range of SVPC VPI values. *Direct UNI* – This range corresponds to the cell header format (see page 3-14). For UNI cell header types, the range is from 0 - 255; for NNI, the range is from 0 - 4095. *Virtual UNI/NNI* – This range corresponds to the configured VPI Start/Stop values (see page 3-12).
VCC VPI	Displays the VPI range for a PVC.
VCC VCI	Displays the VCI range for a PVC.

Table 3-7. SVC VPI/VCI Range Attributes (Continued)

Field	Action/Description
SVCC VPI	Enter the range of SVCC VPI values. *Direct UNI* – This range corresponds to the value you entered for Number of Valid Bits in VPI (see page 3-14). *Virtual UNI/NNI* – This range corresponds to the configured VPI Start/Stop values (see page 3-12).
SVCC VCI	Enter the range of SVCC VCI values. *Direct UNI* – This range corresponds to the value you entered for Number of Valid Bits in VCI (see page 3-14). *Virtual UNI/NNI* – This range corresponds to the configured VPI Start/Stop values (see page 3-12).

When you finish configuring these attributes, if you plan to configure SVC addresses for this logical port continue with "Configuring Logical Ports for Use with SVCs" on page 3-31. Otherwise, continue with "Completing the Logical Port Configuration" on page 3-27.

Traffic Descriptor Attributes

The Set Traffic Descriptor Attributes option is only available for ATM Direct and OPTimum Trunk logical port types. This option enables you to modify the traffic characteristics for the configurable control channel.

Select Set Traffic Descriptor Attributes to configure Node and Card traffic descriptors.

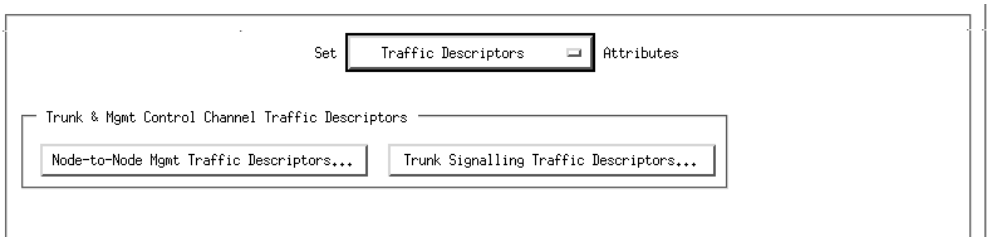

Figure 3-9. Set Traffic Descriptor Attributes

1. Choose Node-to-Node Mgmt Traffic Descriptors. The Set Logical Port Node-to-Node Management Traffic Descriptor dialog box appears.
2. See page 8-9 to complete the fields on this dialog box.

To use the Add Traffic Descriptor command to define a new traffic descriptor, see "Defining Network-wide Traffic Descriptors" on page 8-7.

3. Choose OK to return to the Add Logical Port dialog box.
4. Choose Trunk Signaling Traffic Descriptors.
5. Repeat Step 2 and Step 3 to configure trunk signaling traffic descriptors.

To continue this trunk logical port configuration, do one of the following:

- If this is an OPTimum Trunk logical port, continue with the section, "OPTimum Trunk VPI Range."
- If this is a Direct Trunk logical port, proceed to "Completing the Logical Port Configuration" on page 3-27.

OPTimum Trunk VPI Range

Select Set OPTimum Trunk Attributes and configure the VPI range for an OPTimum trunk as described in Table 3-8.

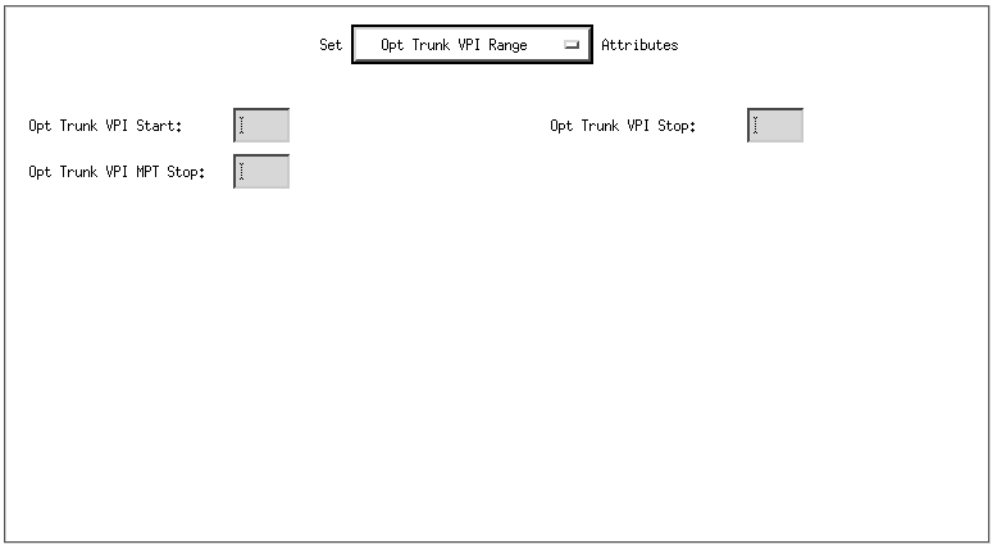

Figure 3-10. Set OPT Trunk VPI Range Attributes

Table 3-8. OPT Trunk VPI Range Attributes

Field	Action/Description
Opt Trunk VPI Start	The VPI start and VPI stop values specify the range of VPIs that can be created over this OPTimum trunk. The maximum allowable range is 0 – 255 for a UNI logical port and 0 – 4095 for an NNI logical ports. Since you can specify more than one OPTimum trunk on the same physical link, make sure the total number of VPIs allowed on each trunk does not exceed these limits. ***Note:*** *Be sure that the range you specify for the OPT Trunk VPI Start/Stop values falls within the VPI/VCI start/stop range defined for the feeder logical port.*
Opt Trunk VPI Stop	
Opt Trunk VPI MPT Stop	To use an ATM OPTimum cell trunk for MPT traffic, enter a VPI MPT stop value that specifies which part of the VPI range is dedicated to MPT traffic. For example: VPI start = 2 VPI MPT stop = 10 VPI stop = 15 VPIs 3 – 10 are dedicated to MPT traffic only; VPIs 2 and 11 – 15 are used for other VPs. The default value, zero (0), prevents MPT traffic from using this OPTimum cell trunk. ***Note:*** *The range of VPIs configured for MPT traffic must be identical on both sides of the trunk. For more information about MPT traffic, see the* NavisCore IP Navigator Configuration Guide.

Completing the Logical Port Configuration

Complete the following steps to select additional options for this new logical port:

1. From the Add Logical Port dialog box, use the Select: Options: menu to review additional options. Choose Set to configure this information. Table 3-9 describes these options.

Table 3-9. Add Logical Port Option Menu Commands

Option	Action/Description
QoS Parameters	To review default QoS parameters and, if necessary, modify these defaults, see "Setting Quality of Service Parameters" on page 3-27.
NTM Parameters (*Optional - CBX only*)	To configure network traffic management (NTM) parameters for this logical port, see the *NavisCore Diagnostics and Troubleshooting Guide*.
Accounting Parameters (*Optional - B-STDX /CBX*)	Enables you to configure NavisXtend Accounting server parameters.

2. Choose OK. The Set All Logical Ports in PPort dialog box reappears (Figure 3-1 on page 3-2).
3. (*Optional*) To configure this logical port for a specific VPN and customer, see "Configuring a Logical Port for VPN" on page 9-6.
4. Choose Close to return to the Set Physical Port attributes dialog box. Then choose Cancel to return to the Switch Back Panel dialog box.

Setting Quality of Service Parameters

This section describes how to set the 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. Ascend recommends you set the logical port QoS fixed and dynamic options before you provision circuits. Under certain conditions, if you change the bandwidth from dynamic to fixed after you provision circuits, one or more QoS classes (including CBR) may display negative bandwidth. For more information about QoS, see "About Quality of Service" on page 8-3.

Table 3-10 lists the default QoS parameters. The switch routes circuits depending on the routing metric you select for the logical port.

Table 3-10. Default Quality of Service Values for ATM UNI Logical Ports

Service Type	Bandwidth Allocation	Routing Metric	Oversubscription Factor
CBR	Dynamic	Admin Cost	100%
VBR-RT	Dynamic	Admin Cost	100%
VBR-NRT	Dynamic	Admin Cost	100%
ABR/UBR	Dynamic	Admin Cost	100%

To modify these settings, from the Select: Options: menu (see page 3-27) select QoS Parameters and choose Set. The following dialog box appears.

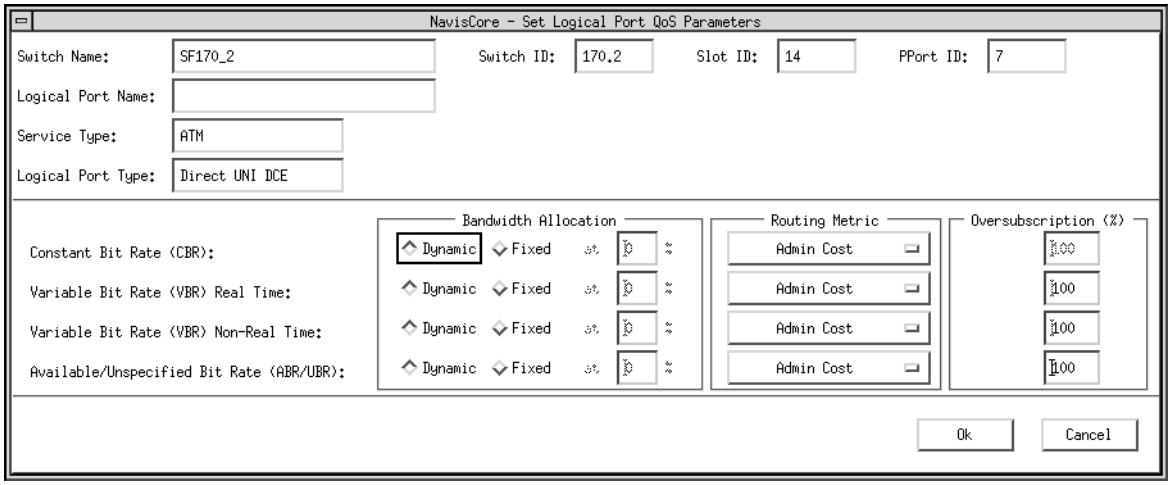

Figure 3-11. Set Logical Port QoS Parameters

To set the QoS parameters:

1. Configure the Bandwidth Allocation for each service class as follows:

   Dynamic — Select Dynamic to enable 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 — Select Fixed to specify 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 CBR or VBR 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 ABR/UBR service class to Fixed, you are guaranteeing that amount of service (at a minimum) for the UBR queue, provided that the VBR queues are not oversubscribed. 100 cells/sec. of bandwidth is allocated for ABR/UBR connections.

If you have service classes set to Dynamic, any remaining bandwidth percentage is allocated to those service classes as needed. For example, if CBR is Fixed at 30%, ABR/UBR is Fixed at 25%, and the two VBR classes are set to Dynamic, the remaining 45% of bandwidth will be dynamically allocated between the two VBR service classes.

2. The switch routes circuits depending on the routing metric you select for the logical port. Routing metrics apply only if the port is configured as UNI DCE, UNI DTE, or NNI logical port. Select one of the following Routing Metrics for each class of service.

Cell Delay Variation (CDV) — This routing metric is only applicable to the CBR and VBR-RT queues. A circuit originating from a queue with the CDV routing metric will find the lowest CDV path to its destination (this is not necessarily the shortest path or the path with the least number of hops). The CDV route is determined from CDV values that are known for the direct and OPTimum trunks.

Admin Cost — A circuit originating from a queue with the Admin Cost routing metric looks for the lowest cost route to its destination (this is not necessarily the shortest path or the path with the least number of hops). The switch determines this route by summing the Admin Costs of each of the direct and OPTimum trunks in the route.

End-to-End Delay — You can configure this routing metric for all service classes. A circuit originating from a queue using the end-to-end delay routing metric finds the path with the lowest end-to-end delay (this is not necessarily the shortest path or the path with the least number of hops). The end-to-end delay is measured between the trunk endpoint interfaces at the time the trunk is initialized.

3. (*Optional*) Specify the Oversubscription Factor percentage for each class of service (except CBR and UBR, which are set to 100% and cannot be modified). This value must be between 100% and 1000%.

If you leave these values set to 100%, Ascend's Call Master Connection Admission Control (CAC) algorithm ensures that the switch packs circuits on a port without experiencing data loss or losing quality of service. (UBR circuits do not use the CAC algorithm.)

After monitoring your network, if 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. For more information on the oversubscription factor, see page 2-13.

4. Choose OK to return to the Add Logical Port dialog box.

Configuring Virtual ATM UNI/NNI Logical Ports

You can create a virtual ATM UNI DCE/DTE or ATM NNI logical port on any physical port on which you have already defined a direct UNI logical port.

If you need to configure an ATM Virtual NNI logical port using PNNI 1.0 routing, see Chapter 17.

To add a virtual ATM UNI/NNI logical port:

1. Complete the steps in "Accessing ATM Logical Port Functions" on page 3-2. Make sure you access a physical port on which you have already defined a UNI logical port.
2. Choose Add to define a new logical port. The Add Logical Port dialog box (Figure 3-2 on page 3-6) appears.
3. Select the LPort Type, either ATM UNI DCE, ATM UNI DCE, or ATM NNI.
4. Choose OK. The Add Logical Port dialog box reappears. The LPort Type field displays a virtual ATM UNI DCE/DTE (or NNI) logical port type.
5. Continue with the instructions beginning with Table 3-2 on page 3-9 to configure attributes for this virtual UNI/NNI logical port.

Configuring Logical Ports for Use with SVCs

If you plan to use SVCs in your network, there are two additional Set Attributes functions you must configure. You only configure these SVC attributes for ATM UNI DCE and DTE logical port types. For more information about SVCs, see Chapter 11, "About SVCs" and Chapter 12, "Configuring SVC Parameters."

- The Set SVC Parameters option enables you to define various SVC screening and handling parameters for each logical port on the switch. These parameters enable you to define SVC addresses for this logical port. Continue with the following section to configure these attributes.
- The Set SVC Routing 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.

Defining SVC Attributes

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [SVC Parameters] Attributes and complete the fields shown in Figure 3-12.

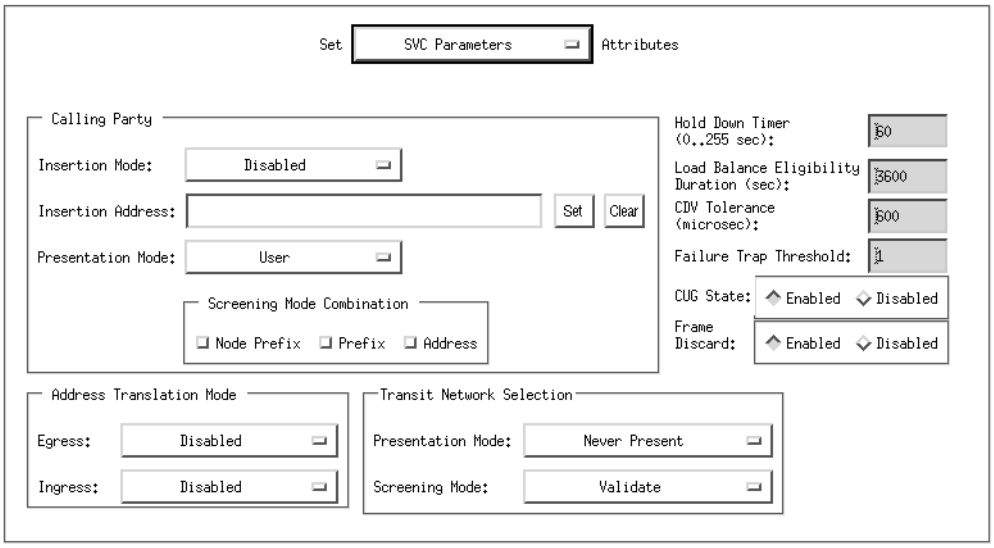

Figure 3-12. Set SVC Attributes

The following sections describe how to define these attributes.

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

For calling party screening to occur, set the Insertion Mode 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.

1. Select one of the following Insertion Mode options:

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 3-34.
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 button to the right of the Insertion Address field. The Set Insertion Address dialog box appears.

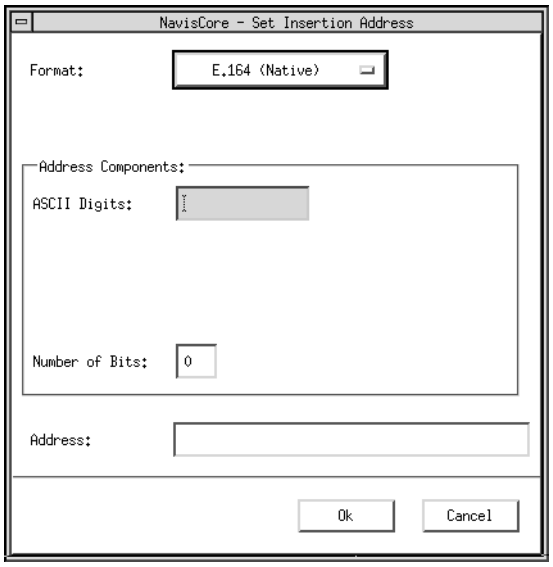

Figure 3-13. Set Insertion Address Dialog Box

The calling party insertion address is not used to route SVCs to this port. To use the calling party insertion address to route SVCs to this port, configure the address (or a prefix corresponding to the address) on this port. For more information, see "Configuring SVC Port Addresses" on page 12-24.

3. Select the appropriate SVC Port Address Format. See the following list of applicable sections for instructions. Then proceed to "Presentation Mode" on page 3-34.
   - For Native E.164 Addresses, see page 12-26.
   - For DCC or ICD AESA addresses, see page 12-27.
   - For E.164 AESA addresses, see page 12-28.
   - For Custom AESA addresses, see page 12-29.

Presentation Mode

In the Set [SVC] Attributes dialog box (Figure 3-12 on page 3-31), 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

In the Set [SVC] Attributes dialog box (Figure 3-12 on page 3-31), 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.

Address Translation Mode Parameters

1. In the Set [SVC] Attributes dialog box (Figure 3-12 on page 3-31), select one of the following Egress Address Translation Mode options:

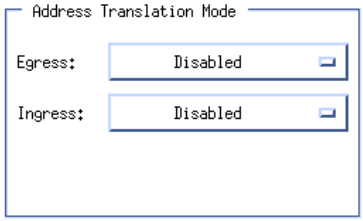

Disabled — No address translation occurs on egress from the logical port.

Tunnel — Select this option if the call is being routed through another network that is using a different address domain (see Figure 3-14). If the calling party address matches a port prefix and the port prefix has a gateway address defined, substitute the local gateway address for the calling party address, and substitute the remote gateway address for the called party address on egress from the logical port. The original addresses are then carried as a sub-address. If you select this option, you should also select Tunnel for the Ingress Address Translation Mode.

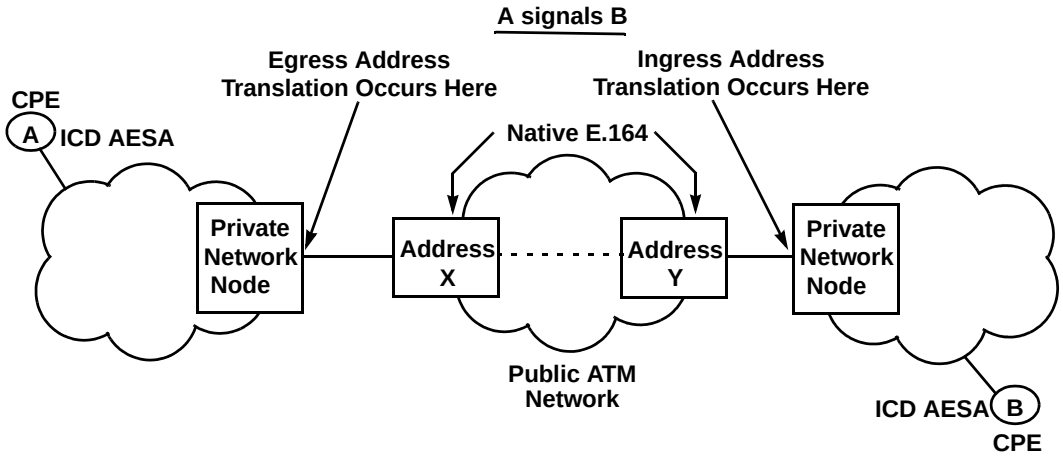

Figure 3-14. Tunnelling Through a Public Network

E.164 Native to AESA — Select this option to convert native E.164 addresses to E.164 AESA format. With this option, the HO-DSP, ESI, and SEL octets of the AESA address are filled with zeros at the network's egress logical port. Also, leading zeros and the trailing Fh are added to the IDP portion. For example, the native E.164 address 5085551234 would be converted to AESA E.164 address 45-000005085551234F-00000000-000000000000-00.

E.164 AESA to Native — Select this option to convert E.164 AESA addresses to native E.164 format. If you select this option, the AFI, HO-DSP, ESI, and SEL octets of the address are removed at the network's egress logical port. Also, all leading zeros and the trailing Fh in the IDP portion of the address are removed. For example, the E.164 AESA address 45-000005085551234F-1A2B3C-0000050F0601-00 would be converted to the native E.164 address 5085551234.

Replace — Select this option if the SVC is being routed into an attached network that is using a different address domain (see Figure 3-15). With this option, the calling party address is replaced with the local gateway address and the called party address is replaced with the remote gateway address at the network's egress logical port.

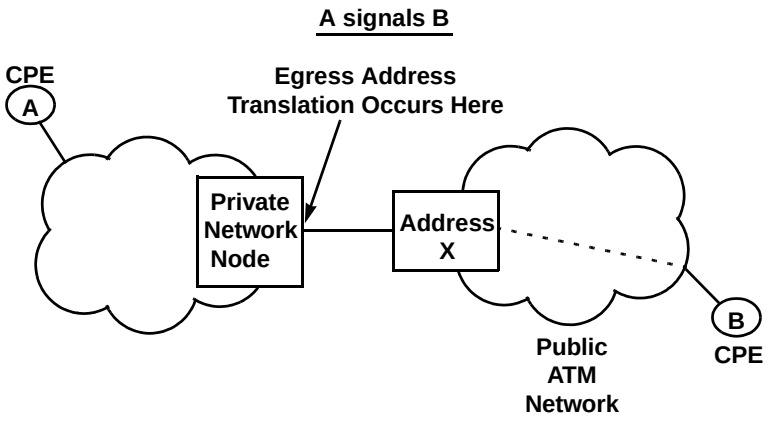

Figure 3-15. Calling Into a Public Network

For more information on egress address translation, see "About Address Translation" on page 11-9.

2. Select one of the following Ingress Address Translation Mode options. These options should match those specified for the Egress Address Translation Mode.

   Tunnel — Select this option if a sub-address is present in the SETUP message, to promote it to the address information element at the ingress port.

   E.164 Native to AESA — Select this option if you selected E.164 AESA to Native as the Egress Address Translation Mode. If you select this option, the AFI, HO-DSP, ESI, and SEL octets of the address are removed at the network's ingress logical port. Also, all leading zeros and the trailing Fh in the IDP portion of the address are removed. For example, the E.164 AESA address 45-000005085551234F-1A2B3C-0000050F0601-00 would be converted to the native E.164 address 5085551234.

E.164 AESA to Native — Select this option if you selected E.164 Native to AESA as the Egress Address Translation Mode. With this option, the HO-DSP, ESI, and SEL octets of the AESA address are filled with zeros at the network's ingress logical port. Also, leading zeros and the trailing Fh are added to the IDP portion. For example, the native E.164 address 5085551234 would be converted to AESA E.164 address 45-000005085551234F-00000000-000000000000-00.

For more information on ingress address translation, see "About Address Translation" on page 11-9.

Transit Network Selection

In the Set [SVC] Attributes dialog box (Figure 3-12 on page 3-31), configure the Transit Network Selection options:

Option	Description
Presentation Mode	Select the egress presentation mode for the selected logical port. Options include: *Never* – (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 12-6 on page 12-37) 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 (presentation mode is 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.

Additional SVC Configuration Options

The Set [SVC] Attributes dialog box (Figure 3-12 on page 3-31) provides the following additional SVC options:

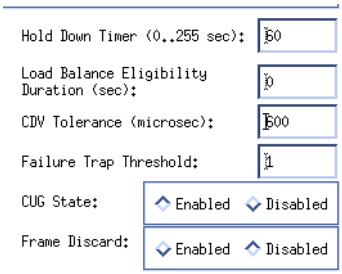

Use Table 3-11 to review the remaining SVC Attributes dialog box options.

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

Table 3-11. Additional SVC Configuration Options

Field	Action/Description
Hold Down Timer (00 . . 255 sec)	Enter the number of seconds to wait before the network initiates SVC 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 (sec)	Enter the number of seconds an SVC must be established before it 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.
CDV Tolerance (microsec)	Configure the cell delay variation tolerance (CDVT). The UPC uses this value to police the requested traffic descriptor. Enter a value between 1 - 65535 µsec which represents cell delay tolerance. The default is 600 µsec.
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 100 SVC failures must occur in a 15-minute window in order to to trigger a trap. If you enter zero (0), the switch never generates a failure trap.
CUG State	Select enable to allow CUG processing for this logical port.

Table 3-11. Additional SVC Configuration Options (Continued)

Field	Action/Description
Frame Discard	When you enable Frame Discard (default), the network performs early packet discard (EPD) and partial packet discard (PPD) on traffic that traverses SVCs that utilize this logical port. This field affects both the CBX 500 EPD/PPD functionality and the CBX 500 and GX 550 output buffer EPD/PPD functionality. Note that if the incoming SVC includes the ATM Adaptation Layer (AAL) parameter information element (IE), there are cases where the information in the AAL IE overrides the logical port setting. This only occurs when requesting a non-UBR AAL 1 and AAL 3/4 connection. For all other cases, including those where AAL 5 is a user-defined AAL or no AAL IE is signaled in, the logical port setting will be in effect. In cases where the incoming SVC does not include the AAL IE or includes a user defined AAL IE, you may wish to disable frame discard as user traffic may be unintentionally discarded if the AAL type is not compatible with EPD/PPD.

SVC Routing Priorities

From the Add Logical Port dialog box (Figure 3-3 on page 3-7), select Set [SVC Routing Priorities] Attributes and complete the fields described in Table 3-12 on page 3-40. When you finish, continue with the instructions in "Completing the Logical Port Configuration" on page 3-27.

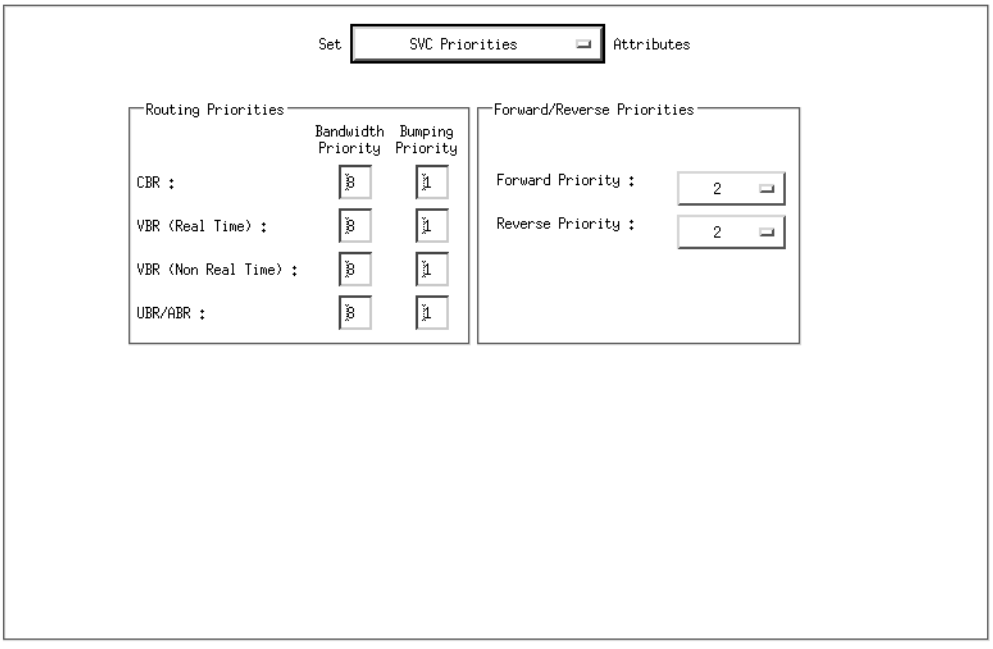

Figure 3-16. Set SVC Routing Priorities

Table 3-12. SVC Routing 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 indicates the highest priority. See Appendix E for more information.
Bumping Priority	For each of the QoS queues, specify a number from 0 through 7 where 1 is the default and indicates the highest priority. See Appendix E for more information.
Forward Priority	Forward priority applies to the caller to callee direction of an SVC. When a particular service category's output queue becomes congested, it must discard cells. This logical port attribute sets the discard priority for the SVC in the forward direction. The lower the number, the higher the priority. Set this attribute from 1 (high priority) to 3 (low priority).
Reverse Priority	Reverse priority applies to the callee to caller direction of an SVC. When a particular service category's output queue becomes congested, it must discard cells. This logical port attribute sets the discard priority for the SVC in the reverse direction. The lower the number, the higher the priority. Set this attribute from 1 (high priority) to 3 (low priority).

4

Configuring ATM Logical Ports on the B-STDX

This chapter describes how to configure logical ports for ATM services on a B-STDX. Most I/O modules in the B-STDX perform a type of “frame-based” ATM switching. The ATM CS DS3/E3 and ATM IWU OC3 modules are capable of performing ATM cell-switching. Keep in mind that the dialog boxes that appear while you configure logical ports display different attributes depending on the type of module, frame- or cell-based.

In addition, the CBX 500 supports a frame-based 6-port DS3 IOM that is also capable of providing frame-based ATM switching. ATM logical ports for this module are configured the same as a B-STDX frame-based module.

For information about the basic elements of ATM service, see one of the following sections in Chapter 2, “About ATM Logical Ports”:

- “Using Interim Local Management Interface (ILMI)” on page 2-3
- “Virtual Paths and Virtual Channels” on page 2-7
- “About the Oversubscription Factor” on page 2-13
- “Administrative Tasks” on page 2-15

About ATM Logical Ports

The following sections describe the types of logical ports you can configure using either a B-STDX module or the frame-based 6-port DS3 module for the CBX 500. For an outline of the logical port types that each module supports, see Table 4-1 on page 4-6.

ATM UNI DCE

The ATM UNI DCE logical port type configures the logical port to communicate with an ATM CPE over ATM PVCs. The Ascend switch acts as an access concentrator feeding multiple Frame Relay and/or ATM PVCs to the CPE via the logical port.

ATM UNI DTE

The ATM UNI DTE logical port type configures the logical port to communicate with an ATM switch over ATM PVCs. The Ascend switch acts as an access concentrator feeding multiple Frame Relay and/or ATM PVCs to the ATM network via the logical port.

ATM Direct Trunk/Direct Cell Trunk

An ATM direct cell trunk (“direct trunk”) logical port type supports the transmission of virtual path connection (VPC) data. Like OPTimum cell trunks, direct trunks have no Ascend header. A unique VPI/VCI identifies the circuit and control traffic using a separate channel. A cell trunk uses a virtual path through the ATM cloud as a channel. When configuring a direct trunk, no DTE feeder is required. Since the direct trunk uses all of the physical port’s bandwidth, you can only configure one direct trunk logical port type on a single physical port; no other logical port types can be configured on this port.

Direct trunks enable you to create a trunk between either two B-STDX switches, or between a CBX 500 or GX 550 and a B-STDX. This logical port type enables a single OSPF routing domain in a mixed network that includes both B-STDX and CBX 500 or GX 550 switches.

The following modules support direct trunk connections between B-STDX switches:

- ATM CS
- ATM IWU
- ATM UNI DS3/E3

The following modules support direct trunk connections between a B-STDX and either a CBX 500 or GX 550 switch:

- ATM CS
- ATM IWU

ATM OPTimum Cell Trunk

An OPTimum cell trunk is a virtual path that supports the transmission of virtual circuit connection (VCC) data. Virtual circuits may be established between any B-STDX and CBX 500, or GX 550 user interface via B-STDX frame/cell trunks and CBX 500 or GX 550 cell trunks.

An OPTimum cell trunk establishes a single OSPF routing domain in a mixed network that includes both B-STDX and CBX 500 or GX 550 switches. Routing decisions allow frame-based traffic to traverse either frame- or cell-based trunks. Cell-based traffic is restricted to routes that traverse Direct Cell trunks. OPTimum Cell trunks have no Ascend trunk header. A unique VPI/VCI identifies the circuit and control traffic using a separate channel. A cell trunk uses a virtual path through the ATM cloud as a trunk.

You can configure an ATM OPTimum trunk to create a switch-to-switch Ascend trunk through a public data network (PDN) into another Ascend network. The Ascend OPTimum trunk allows private enterprises to purchase low-cost, public-carrier services as the trunk between two Ascend switches, rather than use a more expensive leased-line service.

An OPTimum cell trunk enables you to create a trunk between either two B-STDX switches or between a CBX 500 or GX 550 and a B-STDX. The following modules support OPTimum cell trunk connections between B-STDX switches:

- ATM CS DS3/E3
- ATM IWU
- ATM UNI DS3/E3

The following modules support OPTimum cell trunk connections between a B-STDX and either a CBX 500 or GX 550 switch:

- ATM CS DS3/E3
- ATM IWU

OPTimum Frame Trunk

An ATM Open Packet Trunking (OPTimum) logical port enables you to use public ATM networks as trunk lines between two Ascend switches. You can configure an ATM OPTimum trunk logical port to:

- Connect to a peer Ascend switch over an ATM PVC.
- Connect to a peer Ascend switch over an ATM PVC, using an ATM DSU.
- Multiplex Frame Relay PVCs and SMDS “connections” over the ATM PVC.

You can configure an ATM OPTimum trunk to create a switch-to-switch Ascend trunk through a public data network (PDN) into another Ascend network. The Ascend OPTimum trunk allows private enterprises to purchase low-cost, public-carrier services as the trunk between two Ascend switches, rather than use a more expensive leased-line service.

Use this logical port type to do the following:

- Optimize performance and throughput in situations where both ends are connected by Ascend switches.
- Enable the logical port to communicate with an Ascend switch peer over an ATM PVC.
- Multiplex multiple Frame Relay PVCs and SMDS “connections” over the ATM PVC.

Network Interworking for Frame Relay Network-to-Network Interface

The Network Interworking for Frame Relay Network-to-Network Interface (FR NNI) logical port type provides the following access:

- Enables an ATM broadband circuit to interconnect two Frame Relay networks.
- Enables the logical port to communicate with a peer Frame Relay switch over an ATM PVC.
- Multiplexes multiple Frame Relay PVC segments over the ATM PVC.
- Supports many-to-one connection multiplexing.
- Facilitates inter-LATA FR NNI connections.
- Supports Frame Relay/ATM PVC Network Interworking Implementation Agreement FRF.5.

Setting the Number of Valid Bits in VPI/VCI for B-STDX

The Number of Valid Bits setting applies to the VPI and VCI range that you can use for VCCs (both PVCs and SVCs). The default values of VPI = 4 and VCI = 8 mean that you can use VCCs over the range of VPI = 0 – 15 (4 bits of VPI) and a VCI range of VCI = 32 – 255 (bits of VCI). The values have no effect on VPCs, which you can provision anywhere over the VPI = 0 – 255 range.

The valid range for the *number of valid bits in VPI* field is 0 – 6; the valid range for the *number of valid bits in VCI* field is 6 – 12. You may have to adjust these values in the following situations:

- In cases where the required VPI/VCI(s) of the attached devices are outside the VPI = 0 – 15 and VCI = 32 – 255 range the default values provide.
- If you use this logical port as a feeder for OPTimum trunks, the VPI value limits the number of OPTimum trunks you can create on this physical port. The VCI value limits the number of circuits you can route over each OPTimum trunk.

  This OPTimum trunk/circuit trade-off is shown by the following formulas, where *P* represents the value in the Valid Bits in VPI field, and *C* represents the value in the Valid Bits in VCI field:

  Maximum virtual paths = $2^P - 1$
  Maximum virtual channels = $2^C - 32$
  $P + C \leq 12$

  Keep in mind that the default values and range for this setting are different from the CBX 500/GX 550. For an overview of virtual paths and channels, see page 2-7.

When you configure an OPTimum trunk between two endpoints, the OPTimum trunk logical ports must match the VPI of the VPC that provides the connectivity between the two switches. The VPI range for the VPI/VCI valid bits setting for each endpoint must accommodate this VPI.

Using VP Shaping

The VP Shaping feature provides a method of enabling Ascend switch traffic sent to a customer network to comply with the customer's purchased traffic contract. By using VP Shaping, all circuits assigned to the shaper are set to the configured SCR, PCR, and MBS rates. The ATM CS and ATM IWU support VP Shaping.

For ATM Direct and OPTimum Cell Trunk logical ports, you can only configure shaper attributes when VP shaping is selected (see Figure 4-9 on page 4-26). Once you select Shaping Type = VP, the pull-down list bar is enabled. For ATM OPTimum Frame Trunk logical ports, you can select both VP and VC shaping attributes from the pull-down list.

Keep in mind that you must first configure the ATM CS DS3/E3 or IWU physical port shaper attributes before you can specify these attributes for the logical port. For information about configuring VP shaping on the ATM CS or IWU physical ports, see the *NavisCore Physical Interface Configuration Guide.*

I/O Modules for ATM Services

You can configure most ATM logical port types for B-STDX I/O modules. Table 4-1 lists any exceptions.

Table 4-1. I/O Modules for ATM Services

I/O Module Type	ATM Logical Port Support
Frame-based I/O Modules	
8-port Universal I/O	ATM UNI-DCE, ATM UNI-DTE, OPTimum Frame trunk, ATM FR NNI
4-port 24 Channel T1	ATM UNI-DCE, ATM UNI-DTE
4-port 30 Channel E1	ATM UNI-DCE, ATM UNI-DTE
2-port HSSI	ATM UNI-DCE, ATM UNI-DTE, OPTimum Frame trunk, ATM FR NNI
10-port DSX-1	ATM UNI-DCE, ATM UNI-DTE, OPTimum Frame trunk, ATM FR NNI
1-port channelized DS3	ATM UNI-DCE, ATM UNI-DTE, OPTimum Frame trunk, ATM FR NNI
DS3-1-0	ATM UNI-DCE, ATM UNI-DTE, OPTimum Frame trunk, ATM FR NNI
4-port unchannelized T1	ATM UNI-DCE

Table 4-1. I/O Modules for ATM Services (Continued)

I/O Module Type	ATM Logical Port Support
4-port unchannelized E1	ATM UNI-DCE
6-port DS3 for CBX 500	ATM UNI-DCE, ATM UNI-DTE, OPTimum Frame trunk, ATM FR NNI
ATM-based I/O Modules	
ATM DS3/E3 UNI	ATM UNI-DTE, Direct trunk, OPTimum Cell trunk, OPTimum Frame trunk, ATM FR NNI ***Note:*** *Because the ATM DS3/E3 UNI card is not cell-based, you configure the same logical port attributes as the frame-based cards.*
ATM CS	ATM UNI-DCE, ATM UNI-DTE, Direct trunk, OPTimum Cell trunk, OPTimum Frame trunk, ATM FR NNI
ATM IWU	ATM UNI-DCE, ATM UNI-DTE, Direct trunk, OPTimum Cell trunk, OPTimum Frame trunk, ATM FR NNI

Configuring Ports for ATM DXI/FUNI and ATM Services

Low-speed ATM *Data Exchange Interface/Frame User-to-Network Interface* (DXI/FUNI) service enables an Ascend switch to interoperate between Frame Relay and ATM technology on a single platform. ATM DXI/FUNI is a Frame-based protocol that is designed to map easily to Frame Relay. Ascend supports ATM DXI/FUNI, Mode 1A features for the ATM DXI/FUNI standard. These features include:

- Provisioning for up to 938 virtual connections per card
- Support for AAL Type 5 data packaging only
- Frame sizes up to 8192 octets (DTE DSU)
- 16-bit frame checking sequence between the DTE and the DCE

I/O Modules for ATM Interworking Services

You can configure ATM FR NNI logical port services on the following modules:

- 8-port Universal I/O
- 2-port HSSI
- 10-port DSX-1
- 1-port channelized DS3
- 1-port channelized DS3-1-0
- 12-port E1
- 6-port DS3 Frame (*CBX 500 only*)

ATM-based I/O Modules

- ATM UNI DS3/E3
- ATM CS DS3/E3
- ATM IWU OC3

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" on page 4-20 for information about configuring mono- and multi-class services.

Table 4-2 shows the maximum mono-class service threshold values you can configure for each card type.

Table 4-2. 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

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

Card Type	56-Byte Buffers	Bytes
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	9324 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.

About ATM Logical Port Functions

Chapter 3 contains instructions for accessing the Set All Logical Ports in PPort dialog box. This chapter also provides an overview of the fields and functions contained in this dialog box.

- To access the Set All Logical Ports in PPort dialog box, complete the steps in "Accessing ATM Logical Port Functions" on page 3-2.
- To review information about this dialog box, see "About the Set All Logical Ports Dialog Box" on page 3-3.

About the Set Attributes Menu

When you define a new logical port, the Add Logical Port dialog box contains a variety of parameters that you must specify. During the process, use the Set Attributes menu on the Add Logical Port dialog box to configure them.

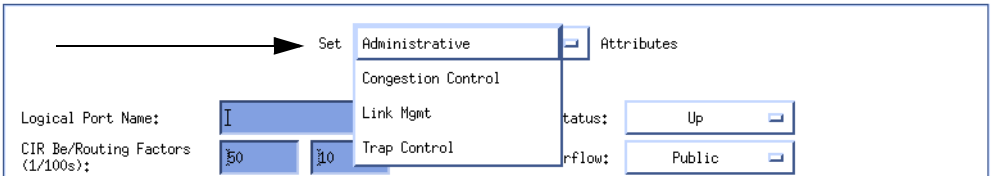

Figure 4-1. Set Attributes Option Menu

Table 4-3 describes the various options you can configure for each ATM logical port type. Keep in mind that some options are only available for certain card types.

Table 4-3. Set Attributes Menu

Option Menu	Description	Logical Port Type	Card Types
Administrative	Determines the number of channels allocated to each port and sets the admin status, net overflow, and bandwidth parameters. The administrative attribute fields may vary depending on the type of service and card type.	All	All
Congestion Control	Sets the threshold parameters (mild, severe, and absolute) that determine how the switch responds to congestion in the network.	• UNI DCE/DTE • Interworking for FR NNI	Frame-based cards (see Table 4-1 on page 4-6)
Trap Control	Sets the congestion threshold percentage in which traps are generated, and the number of frame errors per minute on each logical port.	All	Frame-based cards (see Table 4-1 on page 4-6)
Priority Frame	Specifies the service class type that the logical port can support. The valid values and their corresponding definitions are mono-class and multi-class. When you configure the port for mono-class operation, it only supports VFR-nrt mode; when configured for multi-class operation, it can support CFR, VFR-RT, VFR-NRT, and UFR services.	All	Frame-based cards (see Table 4-1 on page 4-6)

Table 4-3. Set Attributes Menu (Continued)

Option Menu	Description	Logical Port Type	Card Types
ATM	Sets the ATM parameters, which include the number of valid bits in the VPI/VCI, the ATM protocol, and the UNI type.	UNI DCE/DTE	ATM CS and IWU cards
ILMI/Signaling/ OAM	Specifies the ILMI and OAM parameters. *ILMI* – Management information base (MIB) that provides status and communication information to ATM UNI devices and provides a port keep-alive protocol. *OAM* – Sets the alarm functions that generate operations, administration, and maintenance (OAM) alarms. ***Note:*** *The B-STDX does not support signaling functions.*	UNI DCE/DTE	ATM CS and IWU cards
Discard/ Congestion Mapping	Provides support for configurable mapping of DE/CLP and FECN/EFCI bits for both ingress and egress traffic.	• Direct trunk • OPTimum cell trunk • OPTimum frame trunk • Interworking for FR NNI	ATM CS and IWU cards
Link Mgmt	Sets the DCE and DTE polling timer and interval values.	Interworking for FR NNI	All
OPTimum Trunk VPI Range	Specifies the VPI of an OPTimum cell trunk.	OPTimum cell trunk	All

Selecting an ATM Logical Port Type

1. Complete the steps in "Accessing ATM Logical Port Functions" on page 3-2.
2. Choose Add to define a new logical port. The following dialog box appears.

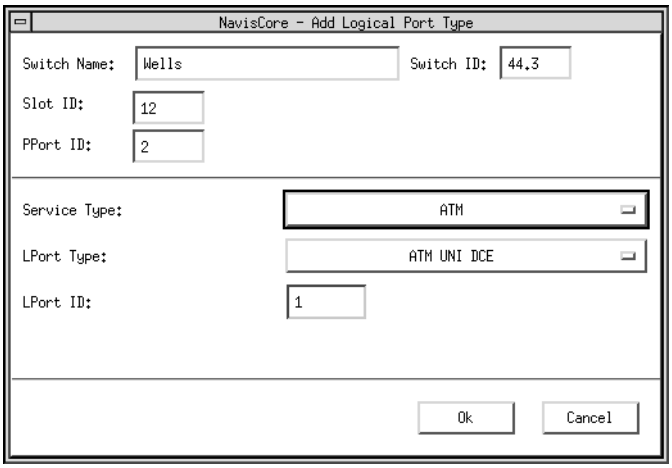

Figure 4-2. Add Logical Port Type Dialog Box

3. Select a logical port type. See Table 4-4 for specific instructions.

Table 4-4. ATM Logical Port Configurations

Logical Port Type	See...
ATM UNI DTE or ATM UNI DCE	"Defining ATM UNI DCE/DTE Logical Ports" on page 4-13.
ATM OPTimum Cell Trunk	"Defining ATM Direct Trunk and OPTimum Cell Trunk Logical Ports" on page 4-24.
ATM Direct Trunk	"Defining ATM Direct Trunk and OPTimum Cell Trunk Logical Ports" on page 4-24.
ATM OPTimum Frame Trunks	"Defining ATM OPTimum Frame Trunk Logical Ports" on page 4-33.
ATM Network Interworking for FR NNI	"Defining ATM Network Interworking for Frame Relay NNI Logical Ports" on page 4-37.

Defining ATM UNI DCE/DTE Logical Ports

To define an ATM UNI-DCE or ATM UNI-DTE logical port:

1. Complete the Add Logical Port Type dialog box (Figure 4-2) fields described in Table 4-5.

Table 4-5. Add Logical Port Type (UNI DCE/DTE) Fields

Field	Action/Description
Service Type	Select ATM.
LPort Type	Select ATM UNI DCE (Network Side) or ATM UNI-DTE.
Logical Port ID	For a T1 module, enter a number between 1 and 24. For an 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.

2. Choose OK. The following dialog box appears.

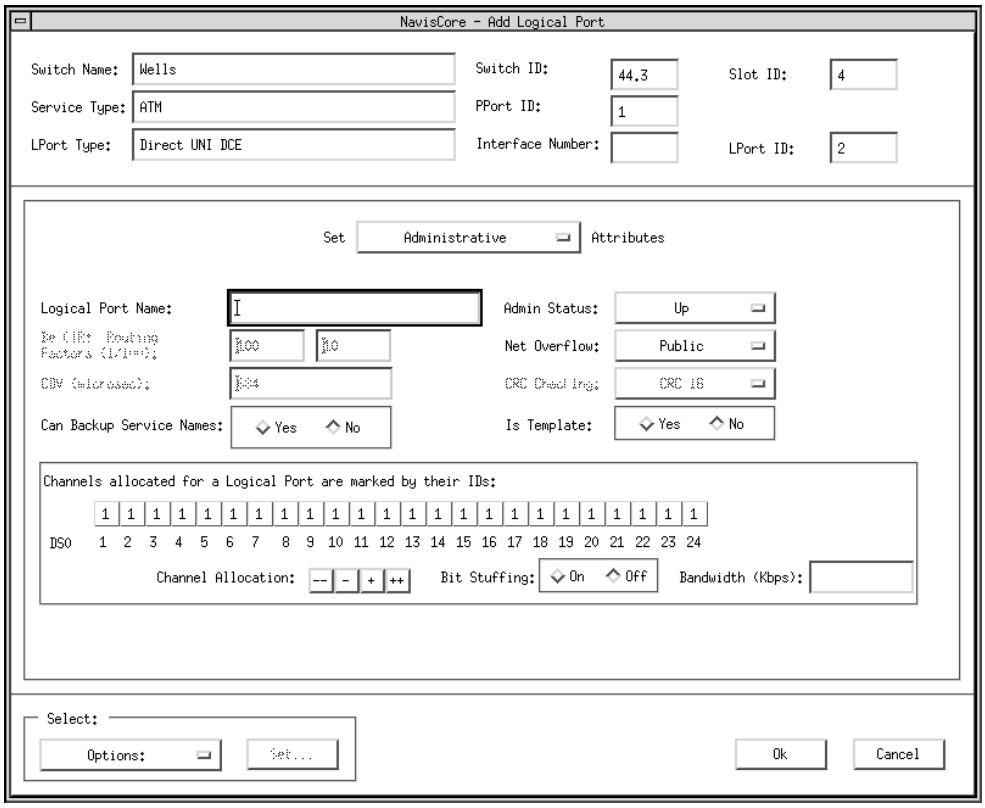

Figure 4-3. Set Administrative Attributes (ATM UNI DCE/DTE)

Administrative Attributes

From the Set Attributes menu in Figure 4-3, complete the fields described in Table 4-6.

Table 4-6. 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.
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 Virtual Private Networks."
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. 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	Set the Net Overflow parameters to one of two modes: *Public* – (Default) Enables the circuit to use public trunks during traffic overflow or trunk failure conditions. *Restrict* – Restricts trunks to their own virtual private network. See "Configuring a Logical Port for VPN" on page 9-6 for more information.
CRC Checking (HSSI 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*) You can save these settings as a template, which you can use again to quickly configure a logical port with the same options. To create a template, choose Yes in the *Is Template* field. See page 4-14 for more information.

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

Field	Action/Description
Channels allocated for a Logical Port are marked by their IDs	Specify the DS0 (for T1) or TS0 (for E1) channel(s) allocation for this 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. Use the following Channel Allocation editing buttons to select/deselect channels: -- 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 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.

See Table 4-7 to continue this configuration:

Table 4-7. Configuring UNI DCE/DTE Attributes

For ATM CS or IWU cards see . . .	For Frame-based cards see . . .
"ATM Attributes" on page 4-21	"Congestion-Control Attributes" on page 4-16
"ILMI and OAM Attributes" on page 4-23	"Trap-Control Attributes" on page 4-18
	"Priority Frame" on page 4-20

When you finish setting the attributes in Table 4-7, continue with the section, "Completing the Logical Port Configuration" on page 4-44.

Congestion-Control Attributes

Select Set [Congestion Control] Attributes and complete the fields described in Table 4-8 on page 4-17.

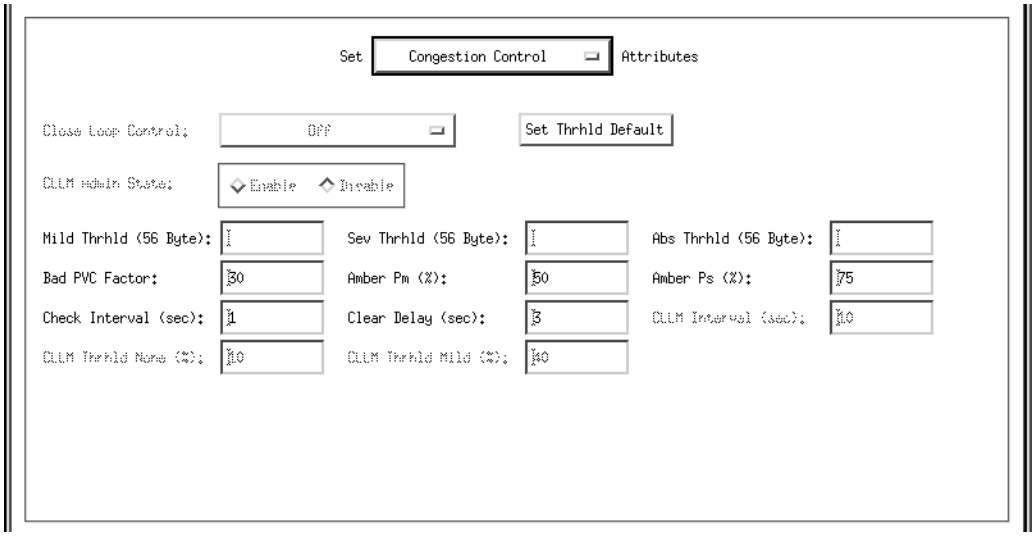

Figure 4-4. Set Congestion Control Attributes

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.

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

Table 4-8. Set Congestion Control Attributes Fields

Field	Action/Description
Set Thrhld Default	Set the Mild, Severe, and Absolute threshold settings to the default settings.
Mild Thrshld (56 Byte) Sev Thrshld (56 Byte) Abs Thrshld (56 Byte)	Accept the defaults or enter values for the mild, severe, and absolute threshold fields. ***Note**: If you set threshold parameters on a T1/E1 card, the default values do not appear until you set the bit stuffing and bandwidth allocation. See Table 4-6 on page 4-14 for more information.*
Bad PVC Factor	Enter a value between 0 – 32 to determine 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 (%)	Enter a Pm% value. This value controls the reduction percentage of Be when mild congestion occurs. 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 an interval that 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 that determines the number of seconds in which the switch monitors the trunk's non-congestion state. The default is 3 seconds.

Trap-Control Attributes

Select Set [Trap Control] Attributes and complete the fields described in Table 4-9.

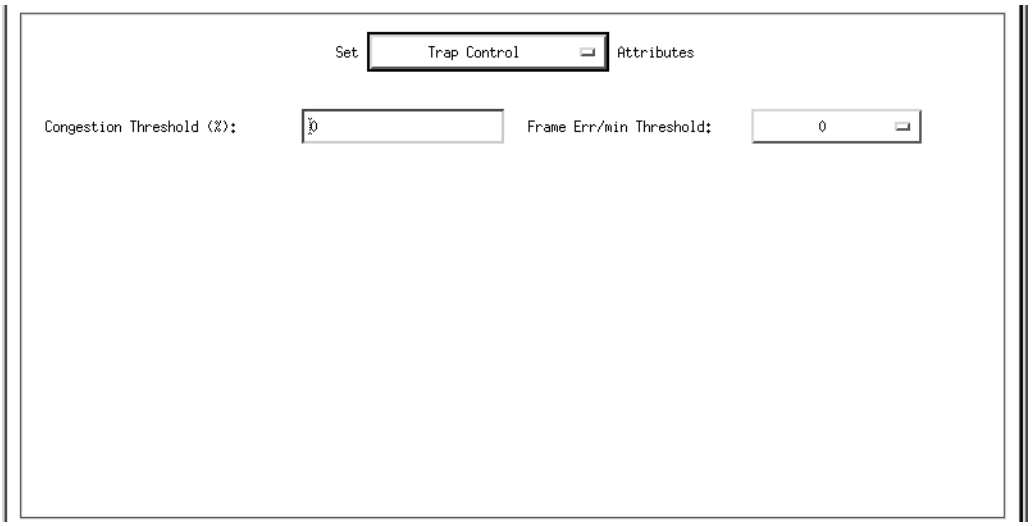

Figure 4-5. Set Trap Control Attributes

Table 4-9. Set Trap Control Attributes Fields

Fields	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: *Zero* – (Default) Disables the congestion threshold. If you enter zero, no traps are generated for this logical port. *Low* – Generates a trap at the first sign of congestion. *High* – Only generates traps for serious network congestion.

Table 4-9. Set Trap Control Attributes Fields (Continued)

Fields	Action/Description
Frame Err/Min Threshold	Enter a value from 0 to 16384 to configure the frame error threshold 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: *Zero* – (Default) Disables this feature, which prevents traps from being generated for this logical port. *Low* – Port is sensitive to frame errors. *High* – Only generates traps when a significant number of frame errors occurs within a one-minute period.
SMDS PDU Violation Threshold (0-255) (*OPTimum Frame 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* – Only issue traps when there is a significant number of SMDS PDU violations.
SMDS PDU Violation Traps (*OPTimum Frame 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 mean 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

Select Set [Priority Frame] Attributes and complete the fields described in Table 4-10. For more information on priority frame features, see "Logical Port Congestion Thresholds" on page 4-8.

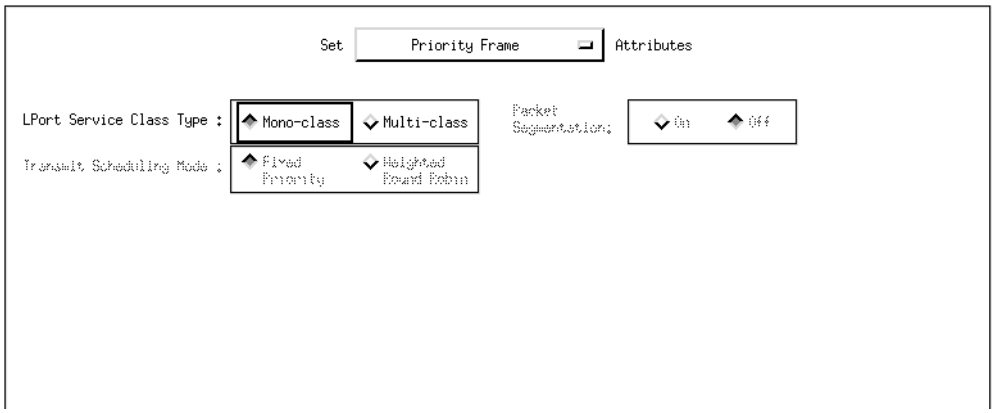

Figure 4-6. Set Priority Frame Attributes

Table 4-10. Set Priority Frame Attributes

Field	Action/Description
LPort Service Class Type	*Mono-class* – (Default) PVC traffic maps to a VFR-NRT service class. *Multi-class* – Allows PVC traffic to utilize all ATM services classes. You must also specify Transmit Scheduling Mode.
Transmit Scheduling Mode	The transmit scheduling mode determines the method used to schedule transmission among the service classes. If you select the Multi-class LPort Service Class, select one of the following queue management options: *Fixed Priority* – Empties the CFR queue first and then the VFR-RT, VFR-NRT, and UFR queues in fixed order. *Weighted Round Robin* – Empties the CFR queue first, VFR-RT and VFR-NRT queues in weighted order, and the UFR queue last.
Allow Vfr-Rt Negative (*Trunk logical port types only)*	This field becomes available if you select the Multi-class Lport Service Class. If you choose enable, the trunk can be oversubscribed. This option is useful in cases where a trunk has failed and PVCs must be rerouted to a new trunk. When this happens, trunk bandwidth can become negative and service may be slow, but PVCs stay up. With this option disabled (*default*), PVCs from the failed trunk will not be rerouted and remains down; however, existing trunk bandwidth and service remains stable.

When you finish setting these attributes, continue with the section, "Completing the Logical Port Configuration" on page 4-44.

ATM Attributes

The Set ATM Attributes option is only available for the ATM IWU and ATM CS cards. Select Set [ATM] Attributes and complete the fields described in Table 4-11.

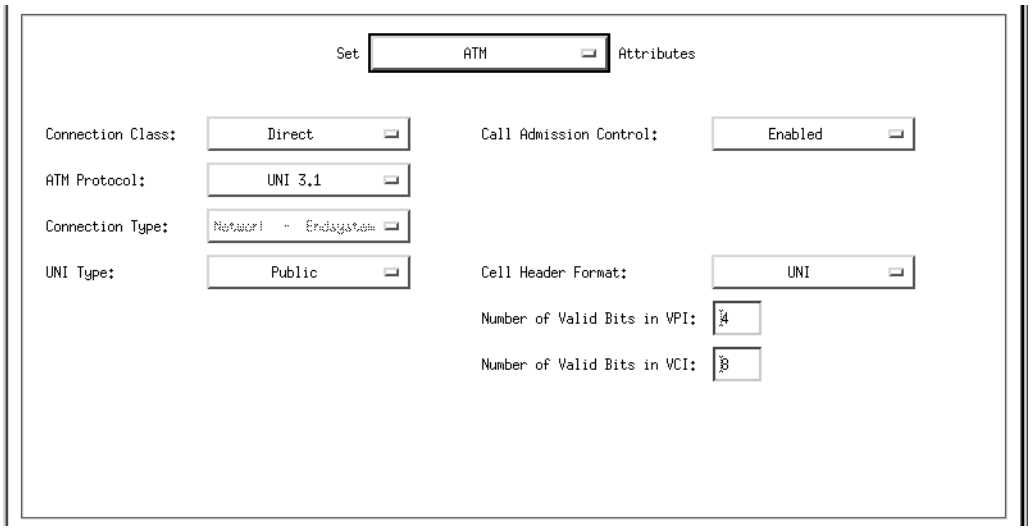

Figure 4-7. Set ATM Attributes

Table 4-11. Set ATM Attributes Fields

Field	Action/Description
Connection Class	Defaults to Direct and cannot be changed.
ATM Protocol	Select the ATM protocol. Options include: • UNI 3.1 (*default*) • UNI 3.0.
UNI Type	Select Public if at least one end of this connection attaches to a public network. Select Private if this connection resides completely within a private network.
Connection Type	Select Network <-> Network if this port connects to another switch or an endsystem. Select Network <-> Endsystem if this port connects to a router or host (UNI-DCE ports only).

Table 4-11. Set ATM Attributes Fields (Continued)

Field	Action/Description
Number of Valid Bits in VPI	Set the number of valid bits in VPI. This field applies to Virtual Channel Connections (VCCs) only. Specify the number of bits used in the ATM cell header for storing the Virtual Path Identifier (VPI). The total of both number of valid bits in VPI/VCI values cannot exceed 12. The default of 4 is recommended; this setting enables you to configure 15 OPTimum trunks, with up to 223 virtual channels on a given virtual path. The valid range for the VPI field is 0-6. For more information see page 4-5.
Number of Valid Bits in VCI	Set the number of valid bits in VCI. This field applies to Virtual Channel Connections (VCCs) only. Specify the number of bits used in the ATM cell header for storing the Virtual Channel Identifier (VCI). The total of both number of valid bits in VPI/VCI values cannot exceed 12. The default of 8 is recommended; this setting enables you to configure 15 OPTimum trunks, with up to 223 virtual channels on a given virtual path. The valid range for the VCI field is 6-12. For more information about setting these values, see page 4-5.
Call Admission Control	Set the call admission control parameter to: *Enabled* – (Default) Port rejects a circuit creation request if there is not enough available bandwidth. *Disabled* – Port attempts to create a circuit even if there is not enough available bandwidth (for VBR Non-Real Time and UBR queues only). ***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.*

ILMI and OAM Attributes

The Set ILMI/Signaling/OAM Attributes option is only available for the ATM IWU and ATM CS cards.

Select Set [ILMI/Signaling/OAM] Attributes and complete the fields described in Table 4-12.

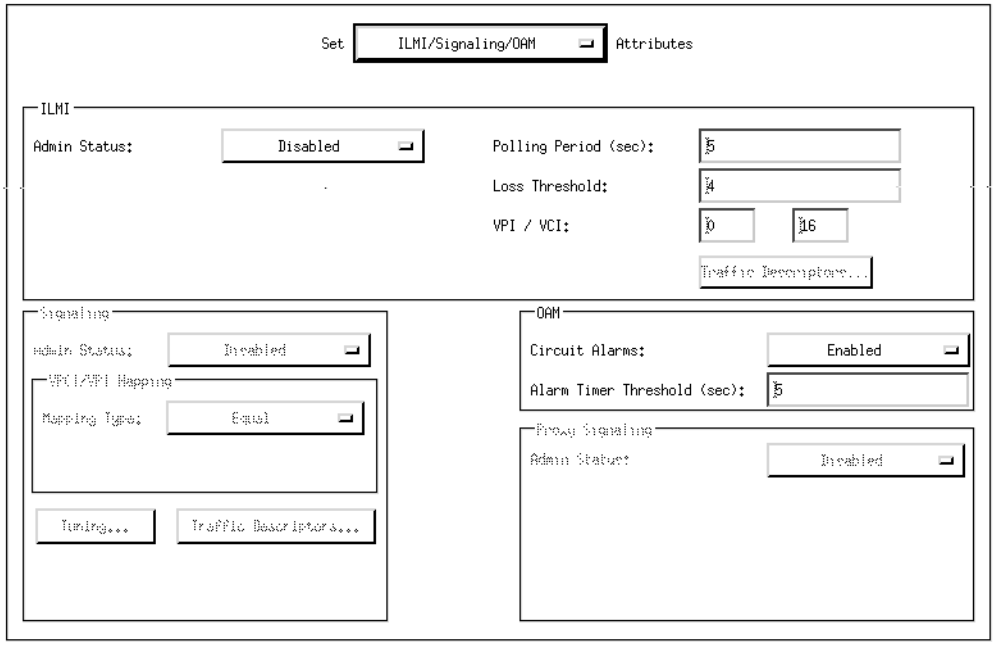

Figure 4-8. Set ILMI/OAM Attributes

Table 4-12. Set ILMI and OAM Attributes Fields

Field	Action/Description
Admin Status	Enable or disable the admin status. *Enabled* – Provides ILMI support. When ILMI is *Disabled* (default), the logical port state is the same as the physical port state. For information about ILMI support, see "Using Interim Local Management Interface (ILMI)" on page 2-3. *Disabled* – Disables ILMI support. If you are using line loopback diagnostics, you must disable ILMI support. See the *NavisCore Diagnostics and Troubleshooting Guide* for more information on loopbacks.
Polling Period (sec)	Specify the polling period (T) for an ILMI poll. The switch generates an ILMI poll every (T) seconds. The default is *5 seconds*.
Loss Threshold	Specify the number of times (K) the logical port will issue an ILMI poll before the link is considered down. If no responses are seen in K x T seconds, the link is considered down. The default is 4.

Table 4-12. Set ILMI and OAM Attributes Fields (Continued)

Field	Action/Description
VPI/VCI	Enter the virtual path ID (VPI) or virtual channel ID (VCI) you want to use for ILMI polling. The default is 0 (VPI) or 16 (VCI).
Circuit Alarms	Set the circuit alarm status. *Enabled* – (Default) Allows this logical port to generate OAM alarms. The switch uses these alarms to signal when the circuits have gone down or come back up. *Disabled* – Disables OAM alarms on this logical port.
Alarm Timer Threshold (sec)	Set the alarm timer threshold (in seconds). The switch waits until the circuit has been down for the time period you specify in this field before generating an OAM alarm. The default is *5 seconds*.

To complete a logical port configuration for an ATM CS or IWU card, proceed to "Completing the Logical Port Configuration" on page 4-44.

Defining ATM Direct Trunk and OPTimum Cell Trunk Logical Ports

This section describes how to configure an ATM Direct trunk or ATM OPTimum Cell Trunk. These logical port types are only available for the ATM CS, IWU, and DS3/E3 UNI card.

Using Direct Cell Trunks

To configure a direct cell trunk, use the following sequence:

Step 1. Configure the physical port you want to use for the direct trunk (see the *NavisCore Physical Interface Guide*). You can configure direct trunks on any of the following ATM modules:

- ATM CS
- ATM IWU
- ATM DS3/E3 UNI

Step 2. Configure an ATM Direct trunk logical port on the physical port.

Step 3. Configure the trunk (see Chapter 5, "Configuring Trunks").

Using ATM OPTimum Cell Trunks

To implement the OPTimum cell trunk, first configure a UNI-DTE feeder logical port on the same physical port. (CS/IWU ports can also use a UNI DCE port for this purpose.)

To configure an OPTimum cell trunk, use the following sequence:

Step 1. Configure the physical port you want to use for the OPTimum trunk (see the *NavisCore Physical Interface Guide*).

Step 2. Configure a UNI-DTE feeder logical port (page 4-13) for the OPTimum trunk on this physical port. Assign this logical port a minimum amount of bandwidth.

Step 3. Configure the OPTimum trunk logical port on the same physical port. Assign the remaining bandwidth to this logical port (if there is only one OPTimum trunk configured on the physical port).

Step 4. Configure the trunk (see Chapter 5, “Configuring Trunks”).

To configure an ATM Direct Trunk or ATM OPTimum Cell Trunk logical port:

1. Complete the Add Logical Port dialog box fields as follows:

Service Type — Select ATM. The available ATM Logical port types appear.

LPort Type — Select ATM OPTimum Cell Trunk or ATM Direct Trunk.

2. Choose OK. The following dialog box appears.

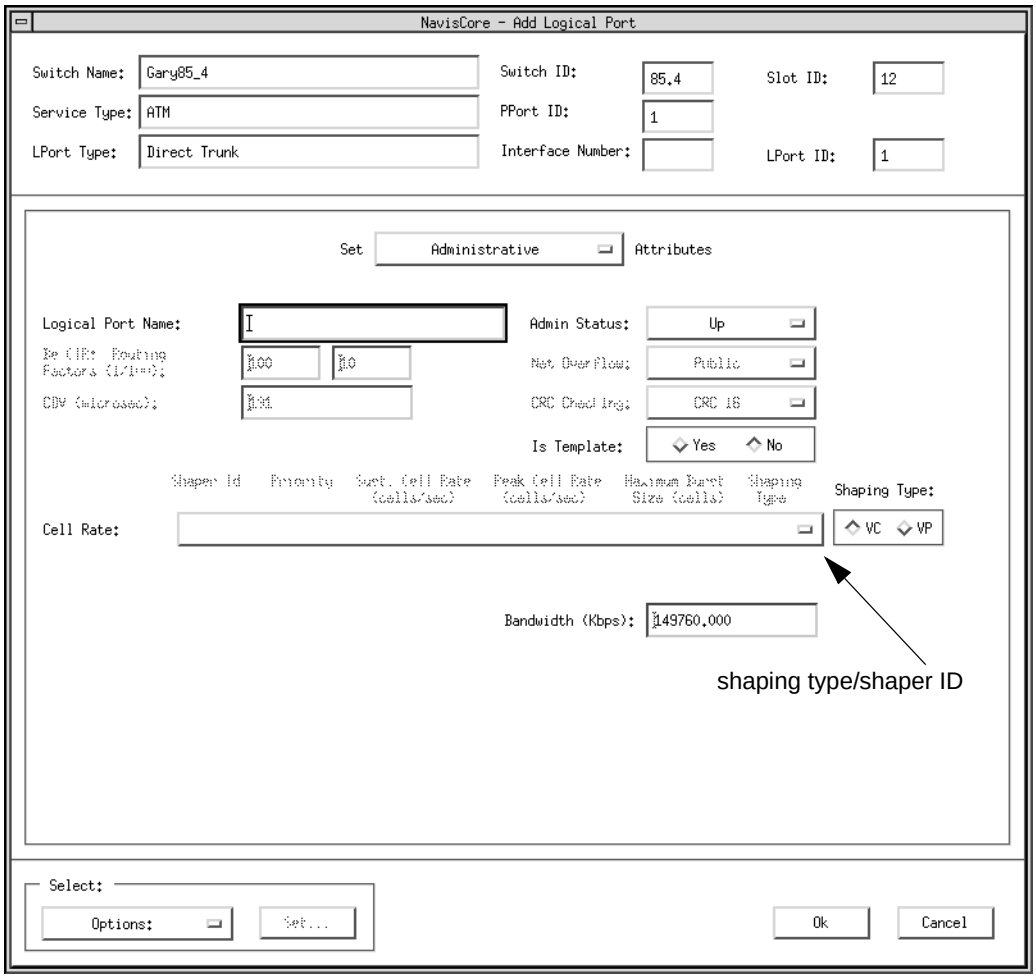

Figure 4-9. Set Administrative Attributes (Direct/OPTimum Cell, ATM CS/IWU Card)

Administrative Attributes

Complete the administrative attributes fields as described in Table 4-13. Note that Table 4-13 lists all possible administrative attributes for a Direct and OPTimum Cell logical port. The attributes vary depending on your type of I/O module.

Table 4-13. Direct and OPTimum Direct Cell Administrative Attributes

Field	Card/Logical Port Type	Action/Description
Logical Port Name	All	Enter an alphanumeric logical port name (up to 32 characters in length) to assign this port.
Admin Status	All	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.
CDV (microsec)	ATM CS/ IWU card OPTimum Direct Cell Trunks	Specify the maximum cell delay variation (in µsecs) for this logical port. This value applies only to CBR traffic, and it specifies the maximum variation in time delays between cells going out of this logical port. The default value is 684 for ATM IWU ports, or 191 for ATM CS ports. To change the default, you need to know the maximum CDV for PVCs on the port, as well as the hardware traffic requirements at the opposite end of the connection.
Net Overflow	All	Set the Net Overflow parameters to either: *Public* – (Default) Enables the circuit to use public trunks during traffic overflow or trunk failure conditions. *Restrict* – Restricts trunks to their own virtual private network. See "Configuring a Logical Port for VPN" on page 9-6 for more information.
Is Template	All	(*Optional*) You can save these settings as a template, which you can 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.

Table 4-13. Direct and OPTimum Direct Cell Administrative Attributes (Continued)

Field	Card/Logical Port Type	Action/Description
Bandwidth (Kbps)	All	Enter the amount of bandwidth you want to configure for this logical port. The default is the amount of bandwidth available on the port. If you are configuring more than one OPTimum Cell trunk on this logical port, enter the appropriate amount of bandwidth to reserve for the OPTimum trunk you are currently configuring. Leave enough bandwidth for any remaining OPTimum trunks you need to configure on this logical port. ***Note:*** *Multipoint VCCs are limited over OPTimum trunks. You cannot configure more than one circuit leaf of a point-to-multipoint circuit on the same logical port, since multiplexing cannot be performed at the port level. If you configure more than one OPTimum trunk on a logical port, set up a circuit leaf for a given circuit root on only one of the OPTimum trunks.*
Shaping Type	ATM CS/IWU card	The Add Logical Port dialog box defaults to a VC Shaping Type (see Figure 4-9 on page 4-26). Using VC Shaping, the shaper pick list is grayed out and the switch uses a method of dynamically selecting a shaper for each circuit routed over the cell trunk. To use the default VC shaping method, at least one VC shaper must exist in the shaper range 1 – 5, at least one in the range 6 – 10, and at least one in the range 11 – 15. To enable VP shaping, select Shaping Type VP, then select a Shaper ID. For more information about VP Shaping, see "Setting the Number of Valid Bits in VPI/VCI for B-STDX" on page 4-5.
Shaper ID	ATM CS/IWU card	If you enable VP shaping, choose a value from 1 through 15 to specify the Shaper ID and associated priority, sustainable cell rate, peak cell rate, and maximum burst size values (see Figure 4-9). Make sure you assign only one trunk logical port per shaper ID. Assigning more than one trunk logical port to a given shaper ID decreases circuit performance.

If you are configuring an ATM DS3/E3 UNI module, Figure 4-9 displays the following fields in place of the Shaping Type/Shaper ID. To configure these fields, use the descriptions in Table 4-14. Otherwise, continue with the instructions in Table 4-15 on page 4-29.

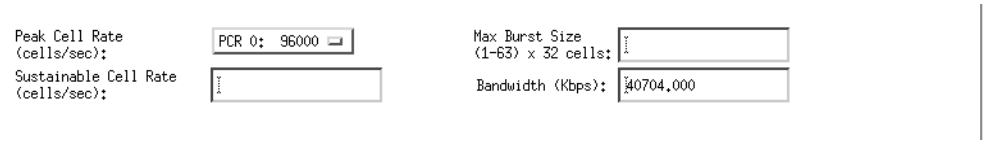

Figure 4-10. ATM DS3 UNI PCR/SCR/MBS Attributes

Table 4-14. ATM DS3 UNI PCR/SCR/MBS Attributes

Field	Card/Logical Port Type	Action/Description
Peak Cell Rate (cells/sec)	ATM UNI DS3/E3 cards	Specify the maximum allowed cell transmission rate (expressed in cells per second). The PCR defines the shortest time period between cells and provides the highest guarantee that network performance objectives (based on cell loss ratio) will be met.
Sustainable Cell Rate (cells/sec)	ATM UNI DS3/E3 cards	Specify the maximum average cell transmission rate that is allowed over a given period of time on a given circuit. SCR allows the network to allocate sufficient resources (but fewer resources than would be allocated based on PCR) for guaranteeing that network performance objectives are met. This parameter applies only to VBR traffic; it does not apply to CBR or UBR traffic.
Maximum Burst Size (cells)	ATM UNI DS3/E3 cards	Specify the maximum number of cells that can be received at the peak cell rate. This allows a burst of cells to arrive at a rate higher than the SCR. If the burst is larger than anticipated, the additional cells are either tagged or dropped. This parameter applies only to VBR traffic; it does not apply to CBR or UBR traffic.

See Table 4-15 to continue this configuration:

Table 4-15. Configuring Direct Trunk/OPTimum Cell Trunk Attributes

For ATM CS or IWU cards see . . .	For ATM DS3/E3 UNI cards see . . .
"Discard/Congestion Mapping" on page 4-30	"Trap-Control Attributes" on page 4-18
"OPTimum Trunk VPI Range" on page 4-32	"Priority Frame" on page 4-20
	"OPTimum Trunk VPI Range" on page 4-32

When you finish setting these attributes, continue with the section, "Completing the Logical Port Configuration" on page 4-44.

Discard/Congestion Mapping

Discard/Congestion mapping attributes enable you to select discard and congestion priority bit mappings on data sent to and from the logical port. Egress mapping takes place just before the ATM interface transmits the data; ingress mapping takes place after the ATM interface receives the data. These options provide support for configurable mapping of the DE/CLP and FECN/EFCI bits.

Select Set [Discard/Congestion Mapping] Attributes and complete the fields described in Table 4-16 on page 4-31. When you finish, continue with "Completing the Logical Port Configuration" on page 4-44.

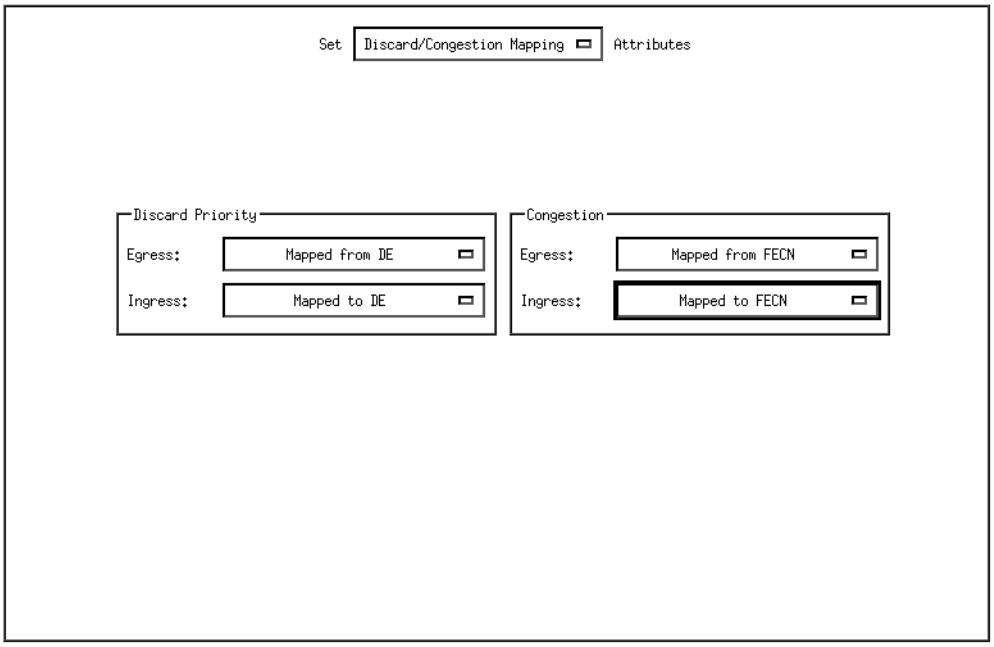

Figure 4-11. Discard/Congestion Mapping Attributes Dialog Box

Table 4-16. Discard/Congestion Mapping Fields

Field	Action/Description
Discard Priority – Egress	Select one of the following options: *Mapped from DE* – (Default) The value of the Discard/Priority bit is used to set the CLP bit when the frame is segmented into cells. The mapping is done just before the frame is presented to the hardware for segmentation. The Discard/ Priority bit is a product of the ingress data stream's Discard/Priority bit setting and whatever modifications are made to this bit due to rate enforcement processing. *Always 0* – The value of the CLP bit is always set to zero (0) for all cells transmitted on this trunk. *Always 1*– The value of the CLP bit is always set to 1 for all cells segmented from all frames transmitted on this trunk.
Discard Priority – Ingress	Select one of the following options: *Mapped to DE* – (Default) The value of the CLP bit received in the cells that make up the ingress frame is transferred directly to the internal Discard/Priority bit. The Discard/Priority bit is transferred with the frame to the egress card for subsequent transmission. If the egress packet format is frame relay, then the Discard/Priority bit is included in the Q.922 header as the DE bit; if the egress packet format is ATM then the CLP bit is set from the Discard/Priority bit. *Not mapped* – The value of the Discard/priority bit is always set to zero (0), ignoring the CLP setting received in the frame. This setting is transferred to the egress card. ***Note:*** *These Discard/Priority settings are used by the rate enforcement and congestion control processing on the egress card. The egress card may change the Discard/Priority bit due to congestion or rate enforcement.*
Congestion – Egress	Select one of the following options: *Mapped from FECN* – (Default) The value of the congestion bit (FECN) is used to set the EFCI bit when the frame is segmented into cells. The mapping takes place just before the frame is presented to the hardware for segmentation. *Always 0* – The value of the EFCI bit is always set to zero (0) for all cells segmented from all frames transmitted on this trunk.

Table 4-16. Discard/Congestion Mapping Fields (Continued)

Field	Action/Description
Congestion – Ingress	Select one of the following options: *Mapped to FECN* – (Default) The value of the EFCI bit received in the cells that comprise the ingress frame is transferred directly to the internal congestion bit (FECN). The congestion bit is transferred with the frame to the egress card for subsequent transmission. If the egress packet format is frame relay, then the congestion bit is included in the Q.922 header as the FECN bit; if the egress packet format is ATM, then the EFCI bit is set from the congestion bit. Note that the egress card can modify the congestion bit due to congestion. *Always 0* – The value of the congestion bit is always set to zero (0), ignoring the setting of the EFCI bit in the cells that comprise the frame. This setting is forwarded to the egress card along with the frame.

OPTimum Trunk VPI Range

Select Set OPTimum Trunk Attributes and configure the VPI range for an OPTimum trunk as described in Table 4-17.

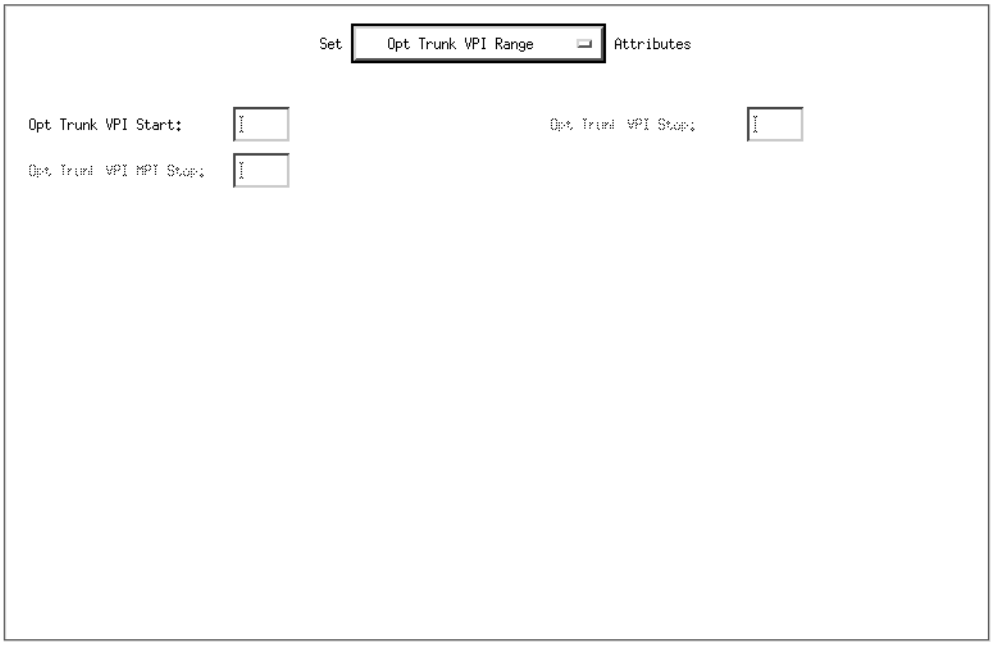

Figure 4-12. Set OPT Trunk VPI Range Attributes

Table 4-17. OPT Trunk VPI Range Attributes

Field	Description
Opt Trunk VPI Start	Enter a number from 0 – *nnnn* to identify the virtual path for the ATM logical port. This is the VPI used for all circuits routed over this OPTimum trunk. Entering a value of 0 enables 4096 circuits to be routed over the trunk. The range of valid VPI values depends upon the number of valid VPI bits you set for the ATM UNI feeder port. For more information, see page 4-5.
Opt Trunk VPI Stop	This value is not used in the B-STDX ATM service.
Opt Trunk VPI MPT Stop	This value is not used in the B-STDX ATM service.

When you finish setting these attributes, continue with the section, "Completing the Logical Port Configuration" on page 4-44.

Defining ATM OPTimum Frame Trunk Logical Ports

To implement the OPTimum frame trunk, first configure a UNI-DTE feeder logical port on the same physical port. To configure an OPTimum frame trunk, use the following sequence:

Step 1. Configure the physical port you want to use for the OPTimum frame trunk (see the *NavisCore Physical Interface Guide*).

Step 2. Configure a UNI-DTE feeder logical port for the OPTimum trunk on this physical port (page 4-13). Assign this logical port a minimum amount of bandwidth.

Step 3. Configure the OPTimum frame trunk logical port on the same physical port. Assign the remaining bandwidth to this logical port.

Step 4. Configure the trunk (see Chapter 5, "Configuring Trunks").

To configure an ATM OPTimum Frame trunk logical port:

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

   Service Type — Select ATM. The available ATM Logical Port Types appear.

   LPort Type — Select ATM OPTimum Frame Trunk.

   VPI — Enter a number from 0 – *nnnn* to identify the virtual path for the ATM logical port. This is the VPI used for all circuits routed over this OPTimum trunk. Entering a value of 0 enables 4096 circuits to be routed over the trunk.

The range of valid VPI values depends upon the number of valid VPI bits you set for the ATM UNI feeder port. For more information, see page 4-5.

VCI — If this logical port resides on the ATM DS3/E3 UNI, enter a value from 32 to 255. Otherwise, enter a value in the range of 32 - *xxx* where *xxx* is determined by the Number of Valid Bits in VCI setting on the feeder port (see page 4-5).

Make sure the number you enter matches the VCI value of the equipment connected to this port. You may have received this value from the ATM network provider.

You must provision a VPC in another ATM network between two Ascend switches. This VPC acts like a physical line. Specify the VPI of this VPC in the Virtual Path ID field.

2. Choose OK. The following dialog box appears.

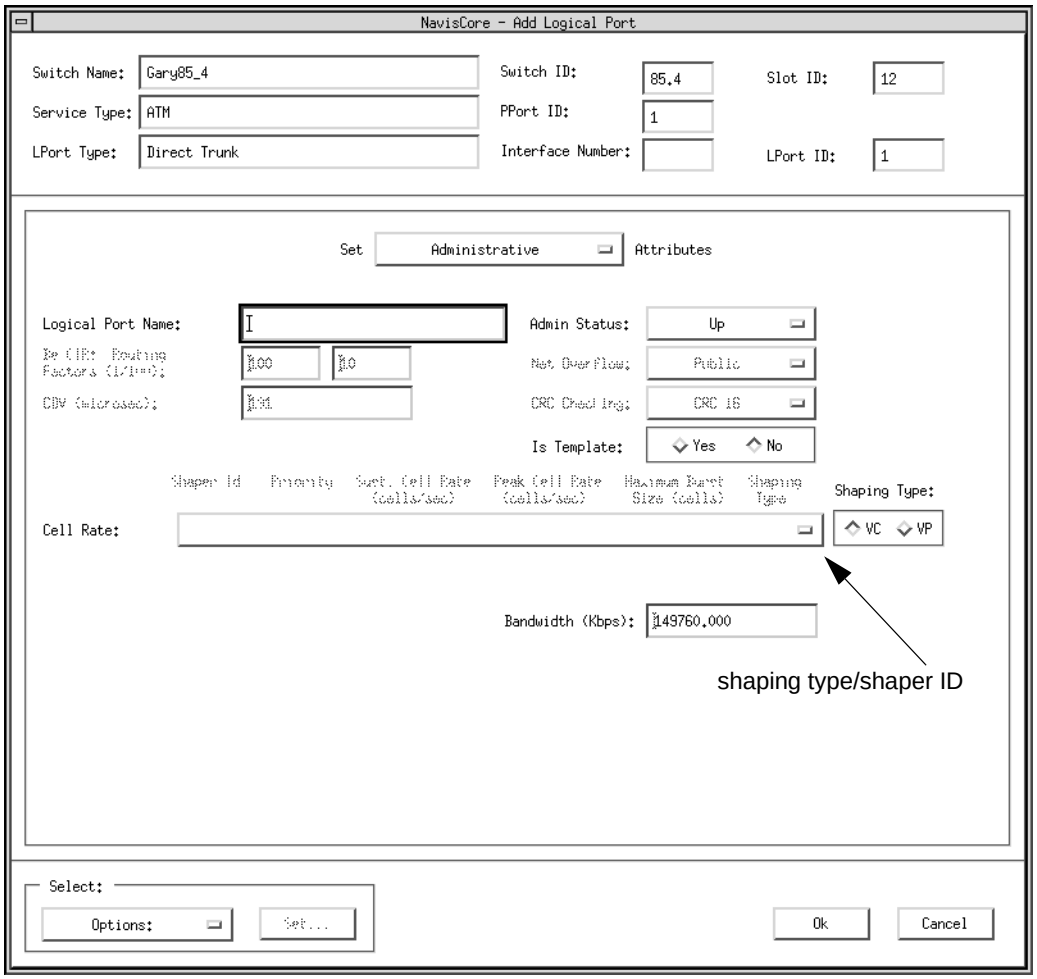

Figure 4-13. Set Administrative Attributes (OPTimum frame, ATM CS/IWU card)

3. Complete the dialog box fields as described in Table 4-18.

Table 4-18. Administrative Attributes Fields

Field	Card Type	Action/Description
Logical Port Name	All	Enter an alphanumeric logical port name (up to 32 characters in length) to assign this port.
Admin Status	All	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.
Net Overflow	All	Set the Net Overflow parameters to one of two modes: *Public* – (Default) Enables the circuit to use public trunks during traffic overflow or trunk failure conditions. *Restrict* – Restricts trunks to their own virtual private network. See "Configuring a Logical Port for VPN" on page 9-6 for more information.
Can Backup Service Names	All	*(Fault-tolerant PVC only)* To configure a logical port for backup service in a fault-tolerant PVC configuration, select Yes. For more information, see Chapter 10.
Is Template	All	(*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.
Shaper ID	ATM CS/IWU	(*See Figure 4-13*) Choose a value from 1 through 16 to specify the Shaper ID and associated priority, sustainable cell rate, peak cell rate, and burst tolerance values. This list provides both VP and VC shaping attributes. Make sure you assign only one trunk logical port per shaper ID. Assigning more than one trunk logical port to a given shaper ID decreases circuit performance. For more information, see "Setting the Number of Valid Bits in VPI/VCI for B-STDX" on page 4-5.
Bandwidth (kbps)	All	Enter the amount of bandwidth you want to configure for this logical port. The default is the amount of bandwidth available on the port. If you are configuring more than one OPTimum frame trunk on this logical port, enter the appropriate amount of bandwidth for the OPTimum trunk you are currently configuring. Remember to leave bandwidth for any remaining OPTimum trunks you need to configure on this logical port.

If you are configuring an ATM DS3/E3 UNI module, Figure 4-13 on page 4-34 displays the following fields in place of the Shaping Type/Shaper ID. To configure these fields, use the descriptions in Table 4-19. Otherwise, continue with the instructions in Table 4-20 on page 4-37.

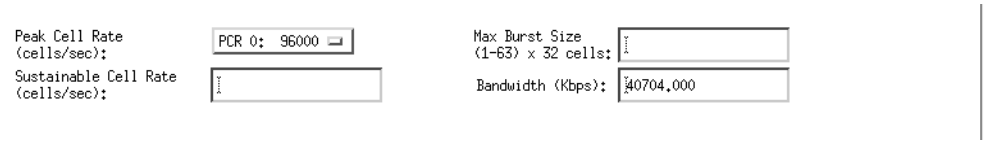

Figure 4-14. ATM DS3 UNI PCR/SCR/MBS Attributes

Table 4-19. ATM DS3 UNI PCR/SCR/MBS Attributes

Field	Card/Logical Port Type	Action/Description
Peak Cell Rate (cells/sec)	ATM UNI DS3/E3	(*See Figure 4-14*) Select a PCR between 0 and 7. This value represents the maximum allowed cell transmission rate (expressed in cells per second). It defines the shortest time period between cells and provides the highest guarantee that network performance objectives (based on cell loss ratio) will be met.
Sustainable Cell Rate (cells/sec)	ATM UNI DS3/E3	(*See Figure 4-14*) Specify the maximum average cell transmission rate that is allowed over a given period of time on a given circuit. The SCR allows the network to allocate sufficient resources (but fewer resources than would be allocated based on PCR) for guaranteeing that network performance objectives are met. This parameter applies only to VBR traffic; it does not apply to CBR or UBR traffic.
Maximum Burst Size (cells)	ATM UNI DS3/E3	(*See Figure 4-14*) Specify the maximum number of cells that can be received at the Peak Cell Rate. MBS allows a burst of cells to arrive at a rate higher than the SCR. If the burst is larger than anticipated, the additional cells are either tagged or dropped. This parameter applies only to VBR traffic; it does not apply to the CBR or UBR traffic.

See Table 4-20 to continue this configuration:

Table 4-20. Configuring OPTimum Frame Trunk Logical Ports

For ATM CS or IWU cards see . . .	For Frame-based cards see . . .
"Discard/Congestion Mapping" on page 4-30	"Trap-Control Attributes" on page 4-18
	"Priority Frame" on page 4-20

When you finish setting these attributes, continue with the section, "Completing the Logical Port Configuration" on page 4-44.

Defining ATM Network Interworking for Frame Relay NNI Logical Ports

Before you can configure an ATM Network Interworking for FR NNI logical port, first configure an ATM UNI DTE feeder logical port with a minimal amount of bandwidth on the same physical port.

To define an ATM Network Interworking for FR NNI logical port:

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

   Service Type — Select ATM.

   LPort Type — Select ATM Network Interworking for FR NNI.

   VPI — Enter the VPI of the ATM VCC that carries the NNI data.

   VCI — Enter the VCI of the ATM VCC used to carry the NNI data. (NNI is a single ATM circuit that can be used to carry a single Frame Relay circuit or many Frame Relay circuits multiplexed over a single ATM circuit.)

2. Choose OK. The following dialog box appears.

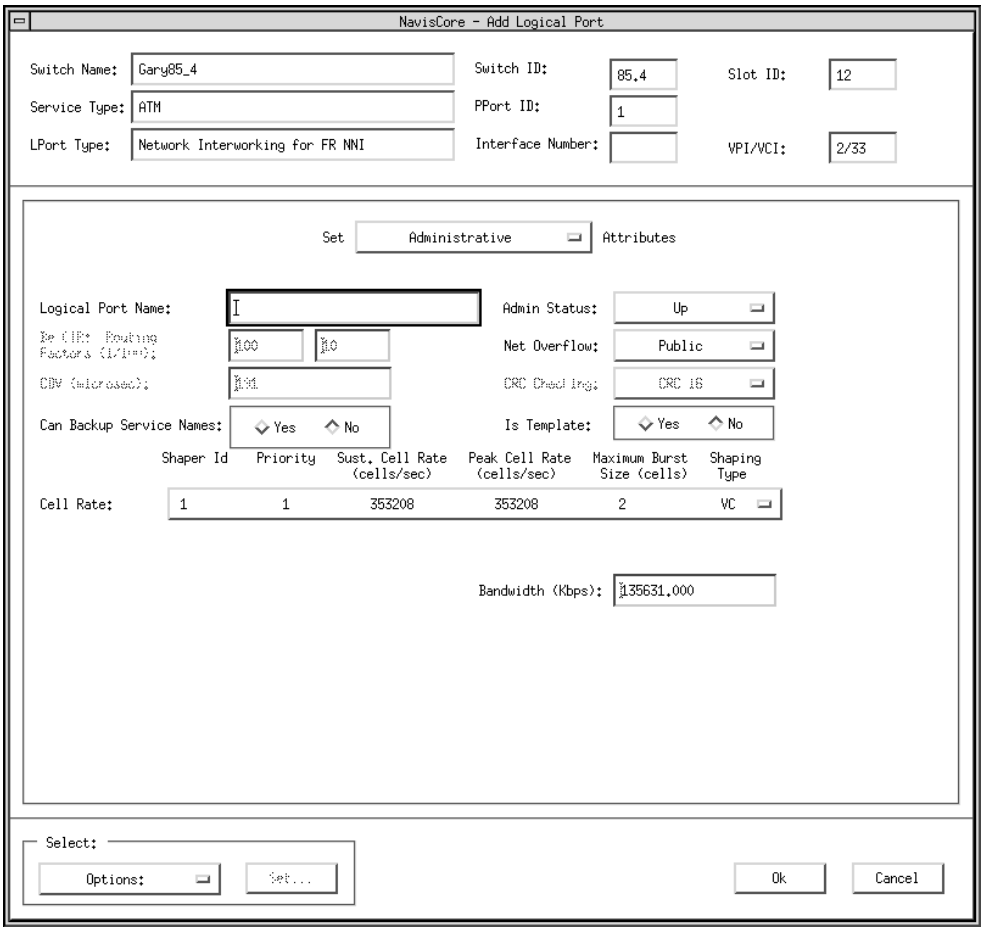

Figure 4-15. Set Administrative Attributes (ATM CS/IWU)

See Table 4-21 to continue this configuration:

Table 4-21. Configuring Frame Relay NNI Attributes

For ATM CS or IWU cards see . . .	For Frame-bsed cards see . . .
"Administrative Attributes" on page 4-27	"Administrative Attributes" on page 4-27
"Link Management Attributes" on page 4-39	"Congestion-Control Attributes" on page 4-16
"Trap-Control Attributes" on page 4-18	"Link Management Attributes" on page 4-39
"Discard/Congestion Mapping" on page 4-41	"Trap-Control Attributes" on page 4-18
	"Priority Frame" on page 4-20

When you finish setting these attributes, continue with the section, "Completing the Logical Port Configuration" on page 4-44.

Link Management Attributes

Select Set [Link Mgmt] Attributes and complete the fields described in Table 4-22.

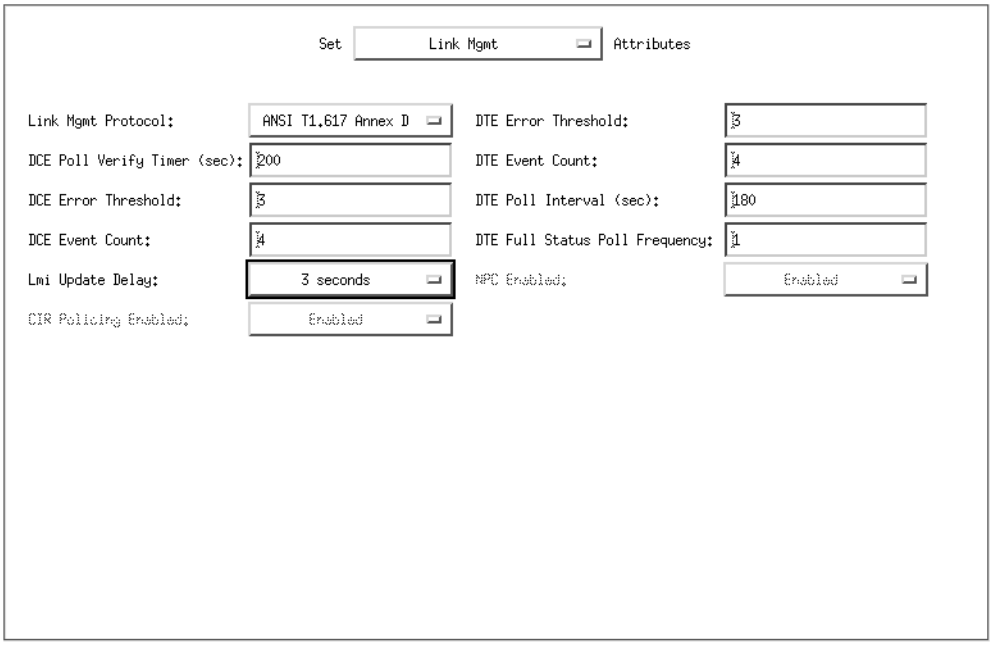

Figure 4-16. Link Management Attributes Dialog Box

Table 4-22. Link Management Attributes Fields

Field	Action/Description
Link Mgmt Protocol	Select the link management protocol used by the Frame Relay equipment connected to this port. 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 customer premise equipment (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 to disable link management for troubleshooting purposes.

Table 4-22. Link Management Attributes Fields (Continued)

Field	Action/Description
DCE Poll Verify Timer	Specify the value of the T392 timer, which sets the length of time the network should wait between status enquiry messages. If the network does not receive a status enquiry message within the number of seconds you specify, the network records an error. The default value is 200 seconds. Increase this value if the DTE device has a poll frequency that is greater than or equal to the DCE Poll Verify Timer. Decrease this value if the DTE's poll frequency is less than or equal to 1/2 of the DCE poll verify timer.
DCE Error Threshold	Specify the DCE error threshold. This parameter is used with the DCE Events Count (N393) parameters. The Local Management protocol monitors the number of events you specify for the DCE Event Count. If the number of events found in error exceeds the DCE Error Threshold you specify, the link is declared inactive. The default value is 3.
DCE Event Count[1]	Specify the DCE event count. This field specifies 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 enquiry message or expiration of the T392 timer. For example, use the default DCE Error Threshold value of three and the default DCE Event Count value of four. If three (N392) of the last four (N393) events are bad, the link is declared inactive. The link remains inactive until the network receives four consecutive error-free events.
LMI Update Delay	Set a timer from 1 to 9 seconds to enable asynchronous LMI updates. The default is three (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.
DTE Error Threshold	Specify an error threshold. This parameter is used with the DTE Events Count (N393) parameter. The Local Management protocol monitors the specified number of events for the DTE Event Count. If the number of events found in error exceeds the specified DTE Error Threshold, the link is declared inactive. The default value is 3.

Table 4-22. Link Management Attributes Fields (Continued)

Field	Action/Description
DTE Event Count[1]	Specify the number of events in a sliding window of events monitored by the network. The default is four. An event is the receipt of a valid or invalid status inquiry message or expiration of the T392 timer. For example, use the default DTE Error Threshold value of 3 and the default DTE Event Count value of 4. If three (N392) of the last four (N393) events are bad, the link is declared inactive. The link remains inactive until the network receives four consecutive error-free events.
DTE Poll Interval (sec)	Specify the number of seconds between the transmission of status enquiry messages. Set the DTE poll interval to a value that is less than the DCE poll verify timer on the attached device. (This value must be greater than 1/2 the value of the DCE poll verify timer.) The default is 180 seconds for one-to-one mapping.
DTE Full Status Poll Frequency	Specify the number of T391 polling cycles between full status enquiry messages. Reduce this value to absorb more bandwidth, since the more frequent full status requests increase overhead. The default value is one for one-to-one mapping.

[1] 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.

Discard/Congestion Mapping

Discard/Congestion mapping attributes enable you to select discard and congestion priority bit mappings on data sent to and from the logical port. Egress mapping takes place just before the ATM interface transmits the data; ingress mapping takes place after the ATM interface receives the data. These options provide support for configurable mapping of the DE/CLP and FECN/EFCI bits.

Select Set [Discard/Congestion Mapping] Attributes and complete the fields described in Table 4-23 on page 4-42.

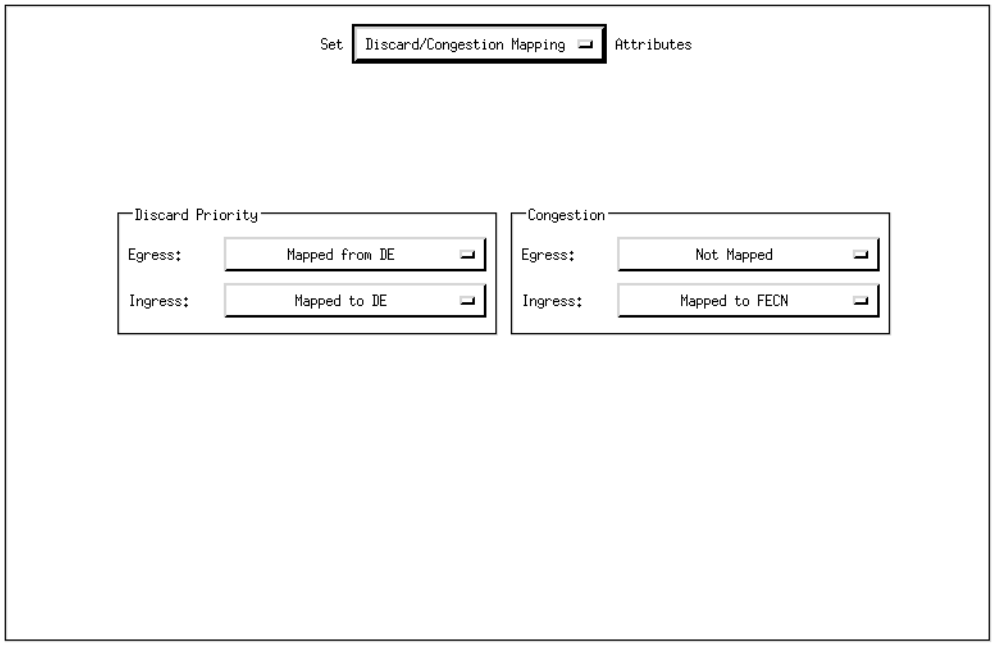

Figure 4-17. Discard/Congestion Mapping Attributes Dialog Box

Table 4-23. Discard/Congestion Mapping Fields

Field	Action/Description
Discard Priority – Egress	Select one of the following options: *Mapped from DE* – (Default) The value of the Discard/Priority bit is used to set the CLP bit when the frame is segmented into cells. The mapping is done just before the frame is presented to the hardware for segmentation. The Discard/Priority bit is a product of the ingress data stream's Discard/Priority bit setting and whatever modifications are made to this bit due to rate enforcement processing. *Always 0* – The value of the CLP bit is always set to 0 for all cells transmitted on this trunk. *Always 1*– The value of the CLP bit is always set to 1 for all cells segmented from all frames transmitted on this trunk.

Table 4-23. Discard/Congestion Mapping Fields (Continued)

Field	Action/Description
Discard Priority – Ingress	Select one of the following options: *Mapped to DE* – (Default) The value of the CLP bit received in the cells that make up the ingress frame is transferred directly to the internal Discard/Priority bit. The Discard/Priority bit is transferred with the frame to the egress card for subsequent transmission. If the egress packet format is frame relay, then the Discard/Priority bit is included in the Q.922 header as the DE bit; if the egress packet format is ATM then the CLP bit is set from the Discard/Priority bit. *Not mapped* – The value of the Discard/priority bit is always set to 0, ignoring the CLP setting received in the frame. This setting is transferred to the egress card. ***Note:*** *These Discard/Priority settings are used by the rate enforcement and congestion control processing on the egress card. The egress card may change the Discard/Priority bit due to congestion or rate enforcement.*
Congestion – Egress	Set to Not mapped. The value of the EFCI bit is always set to 0 for all cells segmented from all frames transmitted on this trunk.
Congestion – Ingress	Set to Mapped to FECN. The value of the EFCI bit received in the cells that comprise the ingress frame is transferred directly to the internal Congestion bit (FECN bit). The Congestion bit is transferred with the frame to the egress card for subsequent transmission. If the egress packet format is frame relay, then the Congestion bit is included in the Q.922 header as the FECN bit; if the egress packet format is ATM, then the EFCI bit is set from the Congestion bit. Note that the egress card can change the Congestion bit due to congestion.

Completing the Logical Port Configuration

Use the following steps to complete this configuration:

1. (*Optional*) If this is an ATM UNI, OPTimum Cell Trunk, or Direct Trunk logical port type on an CS/IWU card, you can configure quality of service (QoS) parameters.

   From the Add Logical Port dialog box, use the Select: Options: menu to select QoS. Choose Set.

   

   See "Setting Quality of Service Parameters" on page 3-27 for more information.

2. Choose OK. The Set All Logical Ports in PPort dialog box reappears (Figure 3-1 on page 3-2).

3. (*Optional*) To configure this logical port for a specific VPN and customer, see "Configuring a Logical Port for VPN" on page 9-6.

4. Choose Close to return to the Set Physical Port attributes dialog box. Then choose Cancel to return to the Switch Back Panel dialog box.

Configure the remaining logical ports for this physical port. See the instructions in this chapter for the specific type of logical port you want to add.

- When you finish configuring a DTE logical port, you can add either an ATM OPTimum cell trunk logical port or an OPTimum frame trunk logical port.
- To configure logical ports on another physical port, select the port, then see the appropriate section for the logical port type you want to configure.
- After you configure both logical port endpoints for a trunk, you can add the trunk line connection between them. See Chapter 5, "Configuring Trunks," for more information.
- After you configure both endpoints of an ATM UNI logical port connection, you can add PVCs between the logical port endpoints. See Chapter 6, "Configuring ATM PVCs," for more information.

5

Configuring Trunks

An Ascend trunk enables two Ascend switches to pass data to each other and exchange internal control messages such as OSPF, SNMP, and others. This chapter describes how to configure an Ascend trunk. In addition, the following sections describe how you can manage trunk traffic:

- "About Administrative Cost" on page 5-2 describes how to configure trunk parameters to route circuits over the trunk which has the lowest administrative cost.
- "About Link Trunk Protocol" on page 5-3 describes how to configure keep alive (KA) control frames.
- "About Virtual Private Networks" on page 9-2 describes how to dedicate trunks to specific customers to guarantee performance and security.
- "About APS Trunk Backup" on page 5-4 describes how to use the CBX 500 and GX 550 optical cards to provide automated trunk backup in cases of equipment failure. These cards include:
  - OC3/STM-1
  - OC12/STM-4
  - OC48/STM-16 (*GX 550 only*)
- "About Trunk Backup for the B-STDX" on page 5-5 describes how to configure *manual* trunk backup for the B-STDX.

About Administrative Cost

You can manage trunk traffic by defining the trunk's administrative cost. Circuits that route data based on administrative cost are created on the path with the lowest administrative cost. You can assign an administrative cost value from 1-65534. The lower the administrative cost of the path, the more likely the path will be chosen when a PVC or SVC routed on administrative cost needs to be created.

The CBX/GX routes circuits based on the routing metric you select for the ATM UNI or NNI logical port endpoint: Admin Cost, CDV (cell delay variation), or End-to-End Delay. See "ATM Direct Trunk" on page 2-6 for more information.

The switch manages circuits as follows:

- When you first define a circuit, the circuit looks for a path that has enough available virtual bandwidth to handle the circuit's effective bandwidth.
- If the circuit finds more than one path with enough available virtual bandwidth, the circuit chooses the path with the lowest administrative cost. This assumes that administrative cost is the designated routing metric. For the UNI or NNI logical port endpoint, if you designate CDV or end-to-end delay as the routing metric, the circuit chooses the trunk(s) with the lowest CDV or end-to-end delay.

The switch automatically reroutes circuits around a failed trunk or switch. If a circuit cannot find a path with sufficient bandwidth, the circuit remains in an inactive state until the bandwidth becomes available.

To establish a PVC, the switch establishes the circuit in the direction of higher-numbered node to lower-numbered node. If the PVC is on the same switch, the switch establishes the circuit in the direction of higher-numbered slot (or port) to lower-numbered slot (or port). In cases where the logical port endpoints use different routing metrics (not recommended), the PVC uses the routing metric that is associated with the higher-numbered element.

About 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, which results in a completed transaction. The request and reply frame formats are identical.

Trunk Delay

Figure 5-1 illustrates the process of KA 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.

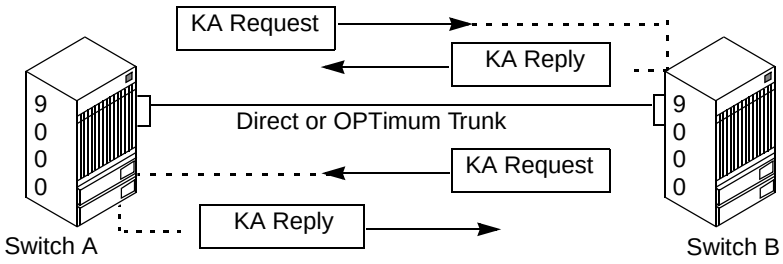

Figure 5-1. Trunk Delay - OSPF Metric and Keep Alive Messaging

Keep Alive Threshold

The Keep Alive Threshold field in the Set All Trunks dialog box 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.

Static and Dynamic Delay

The Static and Dynamic Delay fields in the Set All Trunks dialog box represent the measured one-way delay in units of 100 microseconds. The static delay is measured upon trunk initialization and is updated only when the trunk state changes from down to up. The static delay value is used in conjunction with the end-to-end delay routing metric as a means of allowing users to route circuits over trunks with the lowest end-to-end delay.

The dynamic delay is measured 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.

If you use the Set All Trunks dialog box (Figure 5-2 on page 5-7) to view attributes for a selected trunk, and you notice that the static and dynamic delay values do not match, you can modify the static delay value to match the dynamic delay. To do this:

1. Choose Modify to access the Modify Trunk dialog box (this is similar to Figure 5-4 on page 5-13).
2. Edit the static delay value.
3. Choose OK to accept the change.

If the trunk reinitializes for any reason, the static delay value you inserted when you modified the trunk is automatically replaced by whatever static delay value is measured at the time the trunk reinitialized.

About APS Trunk Backup

The CBX 500 and GX 550 optical cards provide an automatic protection switching (APS) function that enables you to designate a primary (*working*) port and a backup (*protection*) port. These cards include:

- OC3/STM-1
- OC12/STM-4
- OC48/STM-16 (*GX 550 only*)

You can use APS functions to automate trunk backup. If an equipment failure occurs, APS provides a backup physical port. APS Trunk Backup eliminates bandwidth reservation for the backup trunk. You define two trunks between two CBX/GX switches, but only one trunk's worth of bandwidth is required across the Sonet Transport cloud.

Note that on the CBX 500, this feature is only supported on new versions of the OC3/STM-1 IOM. See either the NavisCore or CBX 500 switch software release notice for OC3/STM-1 revision level requirements. To determine whether or not your IOM supports this capability, double-click the switch object to display the Back Panel dialog box; then double-click the OC3/STM-1 IOM to display the View Card Attributes dialog box. The Switch Software Capability and Hardware Capability fields should indicate "APS".

Only one trunk is active at any given moment. If the primary trunk fails (times out), the APS manager detects the link down message and performs an APS switchover to the backup port. APS Trunk Backup provides both facility failure and equipment failure protection.

This feature requires that you configure APS attributes for two working ports that reside on two different switches; these ports act as primary ports. For each of these working ports, select a second port to act as the protection (backup) port. If the working port fails, trunk traffic switches over to the protection port.

If you configure APS Trunk Backup on the CBX, you must select a backup port which resides on a different card in the same switch to act as the protection port; with the GX 550, you can select a protection port on the same card on the same card as the working port. However, this configuration is not recommended.

About Trunk Backup for the B-STDX

The B-STDX switch platform also provides a *manual* trunk backup option. This feature enables you to set up one or more backup trunks to replace a primary trunk. If a 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 B-STDX I/O module.

Using the Add Trunk dialog box (Figure 5-4 on page 5-13), you can configure the trunk type as either primary or backup. 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 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 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.

Configuring B-STDX Trunk Backup

To use trunk backup:

1. Access the Add trunk dialog box (see Figure 5-4 on page 5-13).
2. Define a trunk that has a Trunk Type of *Primary.*
3. Specify all of the primary trunk field values shown in Table 5-3 on page 5-15. 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 enabled Automatic Trunk Backup (see Step 4 on page 5-6).

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 (see Table 5-3 on page 5-15).
2. The system resumes using the primary trunk after it is in service for the period of time specified in the Trunk Restoration Threshold field (see Table 5-3 on page 5-15).

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.

Accessing Trunk Functions

To access the Set All Trunks dialog box, from the Administer menu, select Ascend Parameters ⇒ Set All Trunks.

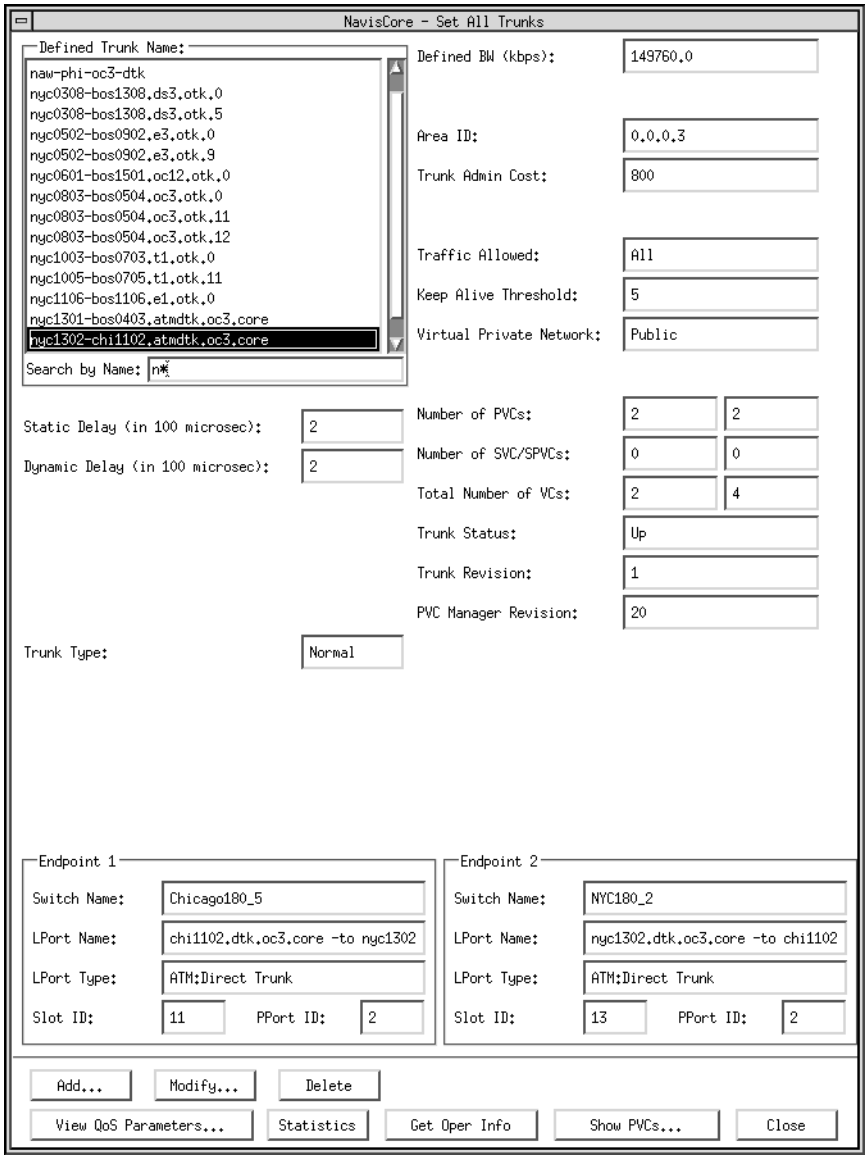

Figure 5-2. Set All Trunks Dialog Box

- To view a list of configured trunks, position the cursor in the Search by Name field and press Return.
- To learn more about the Set All Trunk dialog box, continue with the following section.
- To begin defining a trunk, proceed to page 5-11.

The Set All Trunks Dialog Box

The Set All Trunks dialog box displays information about the configured options for the trunk you select in the Defined Trunk Names list. To use a wildcard search to find a specific circuit name or alias, 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 5-1 describes these dialog box status fields and commands.

Table 5-1. Set All Trunks Dialog Box Status 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)	Displays the bandwidth in Kbps for the selected trunk line.
Area ID	Areas are collections of networks, hosts, and routers used for IP routing. The area ID identifies the area. 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. If a trunk is in Area 1 and the OSPF Backwards Compatibility option (which is set through IP Navigator) is set to Yes, external routes are not advertised across the link. For a detailed description of OSPF areas, and how to use IP Navigator 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. For more information, see "About Administrative Cost" on page 5-2.
Traffic Allowed	Displays one of the following options, which designates the type of traffic allowed on this trunk: *All* – Trunk can carry SVC, PVC, and network management traffic. *Mgmt Only* – Trunk can carry only network management traffic, such as SNMP communication between a switch and the NMS. *Mgmt & Address Restricted* – Trunk can carry PVCs and network management traffic. This trunk option does not support SVC addressing information. If this is the only trunk between two nodes and you configure this option for it, then you effectively prevent SVC traffic from traversing this trunk.

Table 5-1. Set All Trunks Dialog Box Status Fields and Commands (Continued)

Field/Command	Action/Description
Keep Alive Threshold	Displays the number of seconds the trunk protocol will continue to exchange keep alive (KA) control frames without getting a response from the remote node, before bringing the trunk down.
Virtual Private Network	Displays the virtual private network name.
Number of PVCs	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.
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* – One-way communication exists between the two switches. *Two-way* – Bidirectional communication exists between the two switches. *Exchange Start* – The two switches are about to exchange the 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 state changes 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 5-1. Set All Trunks Dialog Box Status Fields and Commands (Continued)

Field/Command	Action/Description
Trunk Type	Displays the type of trunk backup services this trunk provides. Options include: *Normal* – Indicates that this trunk offers no backup service. *Primary* – Indicates that this trunk will act as the main trunk connection in a backup service. *Backup* – Indicates that this is the trunk to which traffic will be diverted in the event of primary trunk failure.
Switch Name	Displays the Ascend switch name on either side of the trunk line.
LPort Name	Displays the logical port name at each endpoint of the trunk.
Lport Type	Displays the configured logical port type.
Slot ID	Displays the slot number where the I/O module containing the selected port is installed.
PPort ID	Displays the physical port ID number on which the logical port is configured.
Add	Defines a new trunk.
Modify	Modifies a trunk definition.
Delete	Deletes the selected trunk.
View QoS Parameters	Displays the Show Logical Port QoS Parameters dialog box for each logical port endpoint of the trunk.
Statistics	Displays the summary statistics for the selected trunk.
Get Oper Info	Displays the status of the selected trunk. A message appears in the Trunk Status field. See page 5-9 for a description of these messages.
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* *GX 550 switch software 01.00.xx.xx.*
Close	Exits the Set All Trunks dialog box.

Defining a Trunk

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

The trunk definition is a three-step sequence:

Step 1. Configure a trunk logical port type.

For a CBX 500/GX 550, see "Defining a Logical Port" on page 3-5.

For a B-STDX, see one of the following sections:

- "Defining ATM Direct Trunk and OPTimum Cell Trunk Logical Ports" on page 4-24
- "Defining ATM OPTimum Frame Trunk Logical Ports" on page 4-33

Step 2. Define a trunk configuration between the two switches. Begin with Step 1 on page 5-12.

Step 3. Create a trunk line connection on the network map to represent the trunk. See page 5-17.

To define a trunk between two Ascend switches:

1. See "Accessing Trunk Functions" on page 5-7 for instructions on accessing the Set All Trunks dialog box (Figure 5-2 on page 5-7).
2. Choose Add. The following dialog box appears.

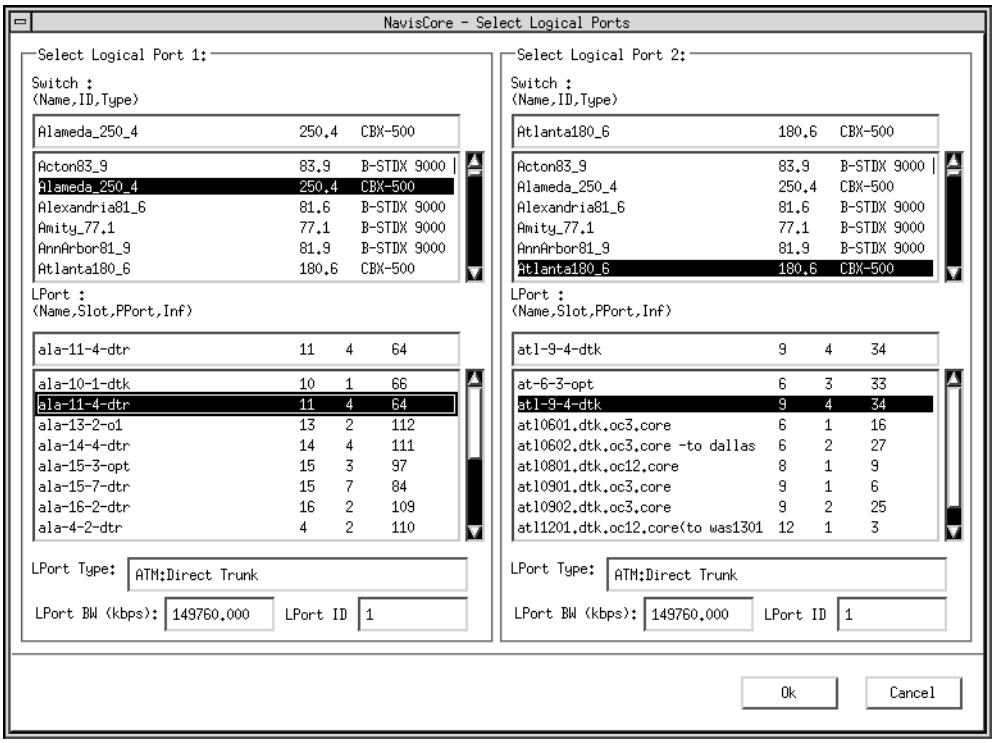

Figure 5-3. Select Logical Ports Dialog Box

3. Provide the following information for both Logical Port 1 and Logical Port 2:
   a. Select the name of the switch where logical port 1 resides, then select the name of the switch where logical port 2 resides.
   b. Select the name of logical port 1, then select the name of logical port 2.
   c. Review the LPort Type field. Both endpoints must use the same logical port type (either ATM:Direct Trunk and ATM:OPTimum Cell Trunk configurations).

When you configure an OPTimum trunk or virtual UNI between two endpoints, the logical ports must match the VPI of the VPC that provides the connectivity between the two switches. The VPI range for the VPI/VCI valid bits setting for each endpoint must accommodate this VPI.

   d. Review the LPort BW field. The bandwidth for each logical port endpoint must be the same.

4. Choose OK. The Add Trunk dialog box appears, displaying the parameters for both logical ports in the trunk configuration.

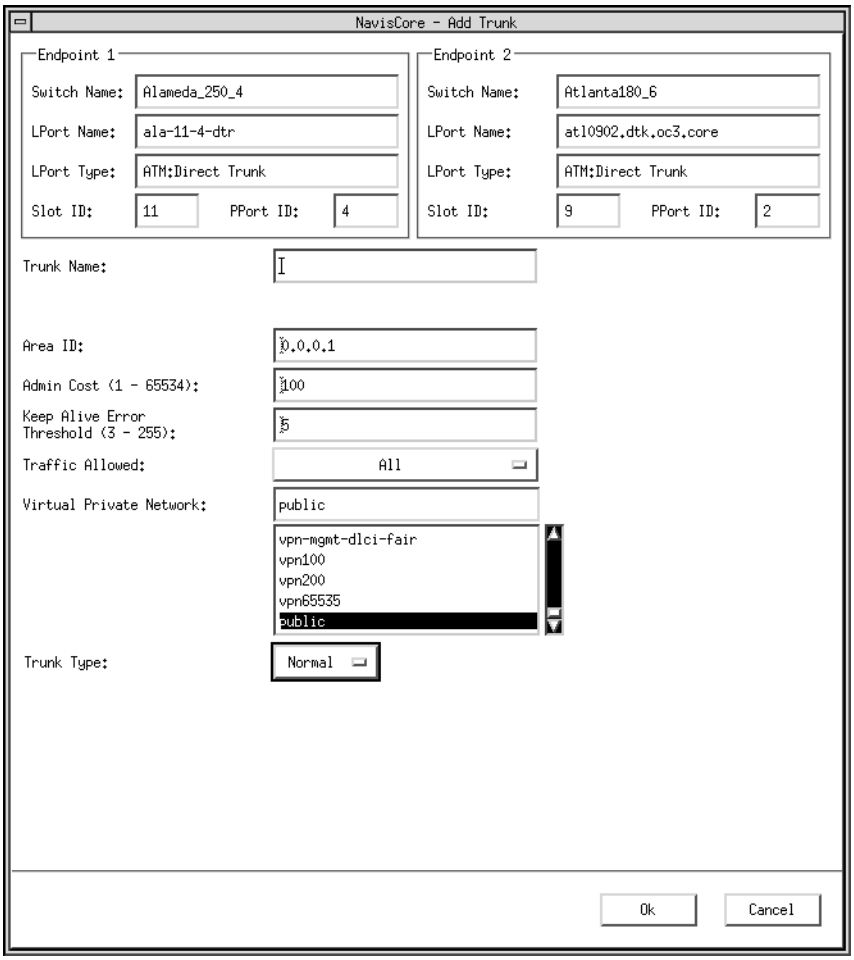

Figure 5-4. Add Trunk Dialog Box

5. Use Table 5-2 on page 5-14 to complete the Add Trunk dialog box.

Table 5-2. Add Trunk Dialog Box Fields

Field	Action/Description
Trunk Name	Enter a unique alphanumeric name to identify the trunk.
Area ID	Areas are collections of networks, hosts, and routers used for IP routing. The area ID identifies the area. Enter the area ID (x.x.x.x) for the destination area for this endpoint. 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. 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. For a detailed description of OSPF areas, and how to use IP Navigator to configure multiple OSPF areas, see the *NavisCore IP Navigator Configuration Guide.*
Admin Cost	Enter a value (from 1 - 65534) that defines the cost of using this trunk for a virtual circuit when a virtual circuit is being dynamically created on the switch. For guidelines, see "About Administrative Cost" on page 5-2.
Keep Alive Error Threshold	Enter a value between 3 and 255 seconds to define the KA threshold. The default is 5 seconds. Service is disrupted if you modify this value once the trunk is online. For more information about this parameter, see page 5-3.
Traffic Allowed	Specify one of the following options to designate the type of traffic allowed on this trunk: *All* – Trunk can carry SVC, PVC, and network management traffic. *Mgmt Only* – Trunk can carry only network management traffic, such as SNMP communication between a switch and the NMS. *Mgmt & Address Restricted* – Trunk can carry PVCs and network management traffic. This trunk option does not support SVC addressing information. If this is the only trunk between two nodes and you configure this option for it, then you effectively prevent SVC traffic from traversing this trunk.
Virtual Private Network	Select a VPN name. The default is Public. For more information about VPNs, see Chapter 9, "Configuring Virtual Private Networks."
Trunk Type	If you are configuring APS trunk backup for a CBX/GX switch, follow the instructions in "Configuring APS Trunk Backup" on page 5-20. To configure Trunk Backup for the B-STDX, select one of the following trunk types: *Normal* – Indicates that this trunk offers no backup service. *Primary* – Indicates that this trunk will act as the main trunk connection. To configure trunk backup features, you must first configure the Primary trunk. Continue with the instructions on page 5-15. *Backup* – Indicates that this is the trunk to which traffic will be diverted in the event of primary trunk failure. Once you configure a Primary trunk, continue with the instructions on page 5-16 to designate a Backup trunk.

6. (*Optional*) If you plan to use the B-STDX trunk backup feature, continue with the instructions in "Using B-STDX Trunk Backup."

7. When you finish defining the trunk attributes, choose OK to complete the trunk configuration. Choose Close to return to the network map.
8. Continue with "Creating a Trunk Line Connection" on page 5-17.

Using B-STDX Trunk Backup

Complete the steps in one of the following sections depending on the trunk type. Keep in mind that you must first configure a Primary trunk before you designate a backup trunk.

- For specific details on implementing the CBX/GX switch platform's APS trunk backup options, see "Configuring APS Trunk Backup" on page 5-20.
- For an overview of B-STDX trunk backup, see "About Trunk Backup for the B-STDX" on page 5-5.

Configuring a Primary Trunk

1. If you selected *Primary* as the Trunk Type the system displays the fields shown in Figure 5-5. Complete the fields described in Table 5-3.

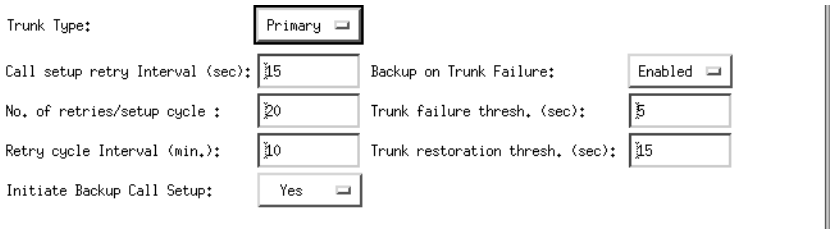

Figure 5-5. Add Trunk Dialog Box (Primary Trunk)

Table 5-3. 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 an 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	Choose Enabled to use trunk backup.

Table 5-3. Add Primary Trunk Fields (Continued)

Field	Action/Description
Trunk Failure thresh. (sec)	Specify the number of seconds. The default is 5 seconds. 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. This field prevents a switch-over to a primary trunk that has only been temporarily restored.

2. Choose OK to complete the configuration. Choose Close to return to the network map.

Configuring the Backup Trunk

1. Once you select a Backup trunk type of primary, the Add Trunk dialog box displays the *Primary Trunk of the backup* field. Select the name of the primary trunk. For a B-STDX, you configure up to eight different backup trunks for each primary trunk.
2. Choose OK to complete the configuration. Choose Close to return to the network map.

Creating a Trunk Line Connection

After you define the trunk configuration between two switches, you can 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 following dialog box appears.

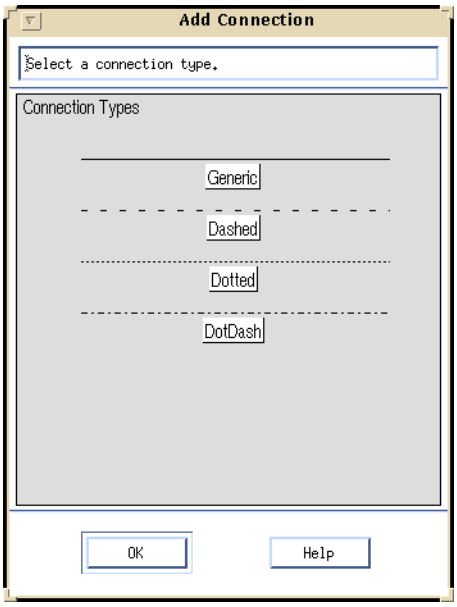

Figure 5-6. Add Connection Dialog Box

2. Select a connection symbol from the palette.
3. With the Add Connection dialog box open, create a trunk line between the two Ascend switches on the network map by clicking on the first switch object (source symbol) and then the second switch object (destination symbol).

4. The following dialog box appears.

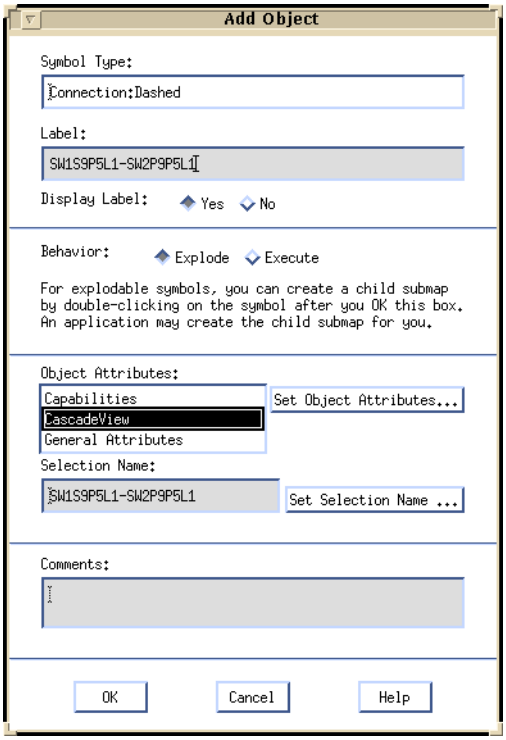

Figure 5-7. Add Object Dialog Box

5. Complete the Add Object dialog box fields as follows:

Symbol Type — Displays the type of connection you are adding to the network map.

Label — Enter the trunk name you specified on the Add Trunk dialog box.

Display Label — Select Yes to have the label (name) appear beneath the switch object on the network map. Select No if you do not want the label to appear on the map.

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 following dialog box appears.

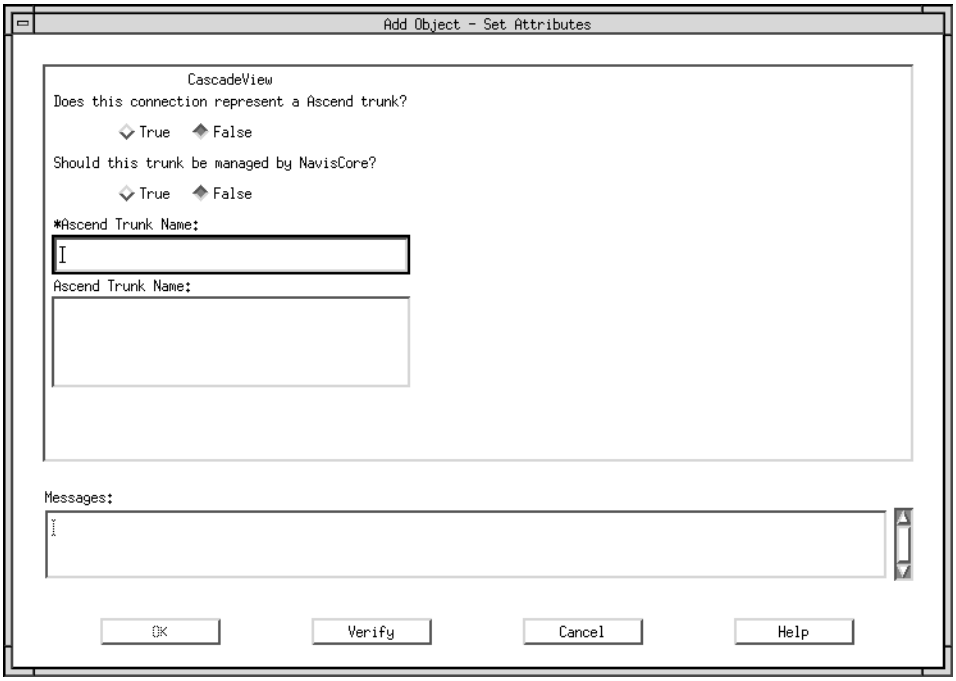

Figure 5-8. Add Object - Set Attributes Dialog Box

6. Complete the Add Object – Set Attributes dialog box fields as follows:

   Does this connection represent an Ascend Trunk? — Select True.

   Should this trunk be managed by NavisCore? — Select True.

   Ascend Trunk Name — Select the trunk name from the Ascend Trunk Name list. The selected trunk name appears in the *Ascend Trunk Name field.

7. Choose Verify to confirm your selections and then choose OK to return to the Add Object dialog box.

8. Choose OK to return to Add Connection dialog box. Then choose OK to return to the network map. A trunk line appears between the two switches on the map.

Configuring APS Trunk Backup

The steps in this section require you to configure two physical ports (either OC3/STM-1, OC12/STM-4, or OC48/STM-16) located on two different switches. For each of these *working* ports, select a port that resides on a different card in the same switch to act as the protection port. You then configure each of these physical ports with a direct trunk logical port. If a working port fails, trunk traffic is diverted to the protection port. See "About APS Trunk Backup" on page 5-4 for more information.

Complete the following sections to configure APS trunk backup.

Defining Physical Port Attributes

1. Select the switch for which you want to configure the *working* physical port.
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 dialog box appears.
4. Select the physical port you want to configure and choose Attrs. The following dialog box appears.

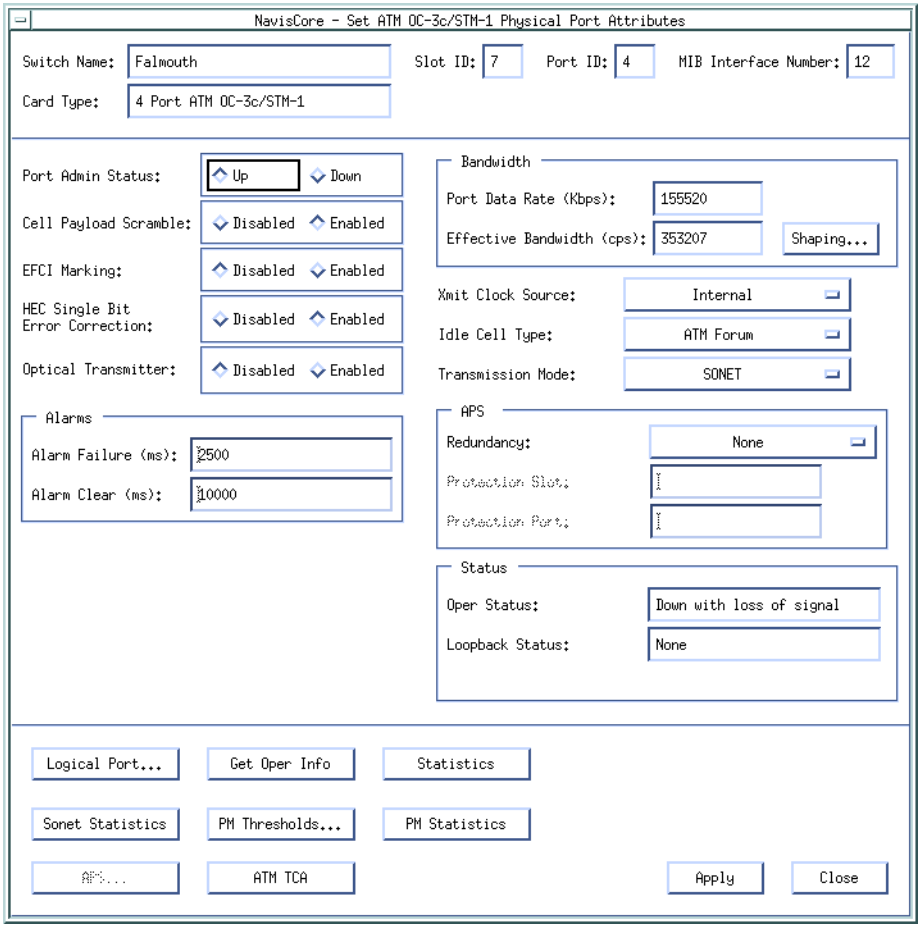

Figure 5-9. Set ATM OC3/STM-1 Ports Physical Port Attributes Dialog Box

5. Complete the dialog box fields as described in Table 5-4.

Table 5-4. Set ATM OC3/STM-1 Physical Port Attributes Fields

Field	Action/Description
MIB Interface Number	Displays the MIB interface number for the physical port. The software assigns a unique number to each physical port on the switch.
Port Admin Status	Set this option to Up to enable immediate port access. Set the Admin Status to Down to save the configuration in the database without activating the port or to take the port offline to run diagnostics. Each time you modify the Port Admin Status, choose Apply to send the change to the switch. ***Note:*** *Changing the Port Admin Status to down sets the physical port operational state to down, but this action does not result in an APS switchover; if you admin down a physical port, user data will be disrupted if the port is active.*

Table 5-4. Set ATM OC3/STM-1 Physical Port Attributes Fields (Continued)

Field	Action/Description
Cell Payload Scramble	Enables (default) or disables the Cell Payload Scramble function. The Cell Payload Scramble function prevents user data from being misinterpreted (that is, it prevents ATM cell header alienation).
EFCI Marking	The explicit forward congestion indicator (EFCI) determines if congestion (or impending congestion) exists in a node. The default is *disabled*. If Enabled, the congested node modifies the EFCI bit in the ATM cell header to indicate congestion. If the equipment connected to the CBX 500 can use the EFCI bit to adjust its transmission rate, it may lower the connection cell rate to relieve the congestion. EFCI is only set in the UBR queue and affects all connections in this physical port's queue. Disable this option if you do not need EFCI marking on this physical port.
HEC Single Bit Error Correction	Enables or disables single bit header error control (HEC) on a per-port basis. When the framer is operating in the default mode of single bit error correction enabled, the framer corrects the single bit errors, but does not count them. Disable the single bit error correction function on the framer to determine how many errors are occurring on the physical port.
Optical Transmitter	This field is a safety feature intended to prevent personal injury when you repair/replace the module or connect cables to the module. By default, this option is disabled. This disables the transmit laser or LED for this port so it cannot transmit incoming traffic. Enable this option to transmit incoming traffic out of this port. ***Note:*** *When you disable the transmit laser, the CPE or switch at the other end of the connection reports a red port alarm to indicate signal loss. Disabling the transmit laser may cause an APS switchover; if you disable the optical transmitter, user data will be disabled if the port was active.* ***WARNING****: Before you remove the optical cable, set this field to disabled. If the optical connectors are exposed, the transmit laser beam can cause personal injury.*
Alarm Failure (ms)	Enter a value between 0 to 65535 ms to determine how long the switch waits before declaring a physical layer problem (i.e., loss of signal) a real failure. The default value of 2500 ms (2.5 seconds) means the switch "soaks" the physical layer alarm for 2.5 seconds before declaring the physical port down. A value of 0 ms means the physical port goes down immediately following any physical layer failure. If you set the value lower than the default of 2.5 seconds, the switch takes the physical port down due to any transient failure in the transmission path; for a port that provides trunk connectivity, this may cause unnecessary rerouting of circuits.
Alarm Clear (ms)	Enter a value between 0 to 65535 ms to determine how long the switch waits once a failure is cleared before declaring a physical layer problem (i.e., loss of signal) resolved. The default value of 10000 ms (10 seconds) means the switch waits 10 seconds after the alarm clears before declaring the physical port up. A value of 0 ms means the physical port comes back up as soon physical layer failure alarm clears. If you set the value lower than the default of 10 seconds, the switch may declare the physical port up before the transmission path is stabilized.

Table 5-4. Set ATM OC3/STM-1 Physical Port Attributes Fields (Continued)

Field	Action/Description
Port Data Rate (Kbps)	Represents the raw physical data rate of the port. Due to the bandwidth lost as a result of the ATM layer to physical layer mapping, this number is always greater than the actual cell rate that can be transmitted out the port. The actual rate of cell transmission is dependent on the method of ATM layer mapping used.
Effective Bandwidth (cps)	Represents the actual cell transmission rate the physical port uses. By default, the physical port transmits cell traffic at the maximum rate supported on the physical interface. However, you can use the Shaping command to select a transmission rate that is lower than the maximum rate.
Xmit Clock Source	Specify the transmit clock source. *Internal* – (Default) The IOM/BIO internal timing source provides the clock source to this port. The IOM/BIO Clock Source setting in the Set IOM/BIO Card Attributes dialog box determines the internal clock source. *Loop-Timed* – The clock source is derived from the signal coming into this port.
Idle Cell Type	Allows you to specify the type of cell that is used to fill the gaps between user data cells that are transmitted out of the physical port. The physical port receive function is not affected by this option (both ITU and ATMF are recognized and processed by the physical port receiver). Select one of the following options: *ATM Forum* (default) – The fill cell will have a header of 00 00 00 00 55 and a payload of 6A (for all 48 bytes). CLP=0 in the cell header. *ITU* – The fill cell will have a header of 00 00 00 01 52 and a payload of 6A (for all 48 bytes). CLP=1 in the cell header.
Transmission Mode	Enables you to designate individual ports for either SONET (OC3/OC12/OC48) or SDH (STM-1/STM-4/STM-16). For OC3 modules, you can configure OC3 framing on one port and STM-1 framing on another. • Select *SONET* (default) to configure the port for North America • Select *SDH* to configure the port for International
Redundancy	Select *APS with Trunk Backup*.
Protection Slot	Enter the Slot ID of the protection port. Be sure to select a port that resides on a different card in the same switch to act as the protection port.
Protection Port	Enter the Port ID of the protection port.
Oper Status	Indicates the operational status of the physical port (Up or Down). If this field is blank, the IOM/BIO did not respond to a status request.
Loopback Status	Displays the port's loopback status if you enabled diagnostic loopback tests. The default is None. See the *NavisCore Diagnostic and Troubleshooting Guide* for details.

Defining APS Attributes

1. From the Set Physical Port Attributes dialog box (Figure 5-9 on page 5-21), choose *APS*. The following dialog box appears.

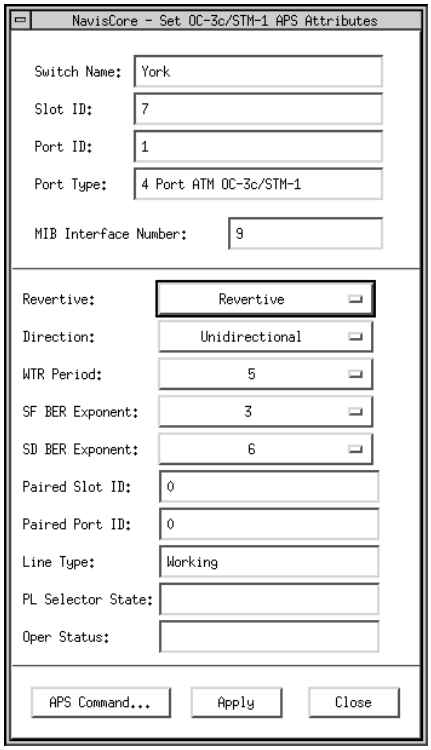

Figure 5-10. Set APS Attributes Dialog Box

2. Complete the APS attributes as described in Table 5-5.

Table 5-5. Set APS Attributes Fields

Field	Action/Description
Revertive	Select *Nonrevertive.* This option enables traffic to continue to pass over the protection port until you use the APS Command to transfer back to the working port. (This is the only option available with APS trunk backup.)
Direction	Select *Bidirectional.* With this option, when one end detects a line defect, it signals the other end to switch. Thus, both ends switch in tandem. (This is the only option available with APS trunk backup.)
WTR Period	Set to 5 minutes. This is the period of time the port waits once the automatically initiated switch condition (i.e., SD, SF, LOF, AIS, or LOS) clears before it transfers traffic back to the working port.

Table 5-5. Set APS Attributes Fields (Continued)

Field	Action/Description
SF BER Exponent	Set the signal fail (SF) bit error rate (BER) exponent. Values can range from 3 to 5; the default is 3. The port uses this value to compute the signal fail threshold. When the BER exceeds $10^{-threshold}$, the port detects the line failure and transfers traffic to the protection port. The port transfers traffic based on how you configured the Direction option (either unidirectional or bidirectional). ***Note**: The clear condition for SF BER is a bit error rate of less than 10^{-7}.*
SD BER Exponent	Set the signal degrade (SD) BER exponent. Values can range from 6 to 9; the default is 6. The port uses this value to compute the signal degrade threshold. When the BER exceeds $10^{-threshold}$, the port detects the line failure and transfers traffic to the protection port. The port transfers traffic based on how you configured the Direction option (either unidirectional or bidirectional). ***Note:** The clear condition for SD BER is a bit error rate of less than 1/10th the current SD threshold.*
Paired Slot ID	Displays the slot location ID of the card on which the protection port resides. If you view this information for a protection port, it displays the slot ID of the working port.
Paired Port ID	Displays the port ID of the protection port. If you view this information for a protection port, it displays the port ID of the working port.
Line Type	Indicates whether the port is a working line or a protection line.
PL Selector State	Indicates the protection line selector state. “Selected” means the protection line is active; “Released” indicates the working line is active. If either port in an APS-pair is capable of receiving data, the Oper Status field on the Set/Show Physical Port Attributes dialog box should display Up.
Oper Status	Indicates the operational status of the port, either Up or Down.

3. Choose Apply to save these settings and Close to return to the Set Physical Port Attributes dialog box (Figure 5-9 on page 5-21).
4. To exit, choose Apply and then OK to save the physical port attributes and create an SNMP SET command to send to the switch. Choose Cancel to exit.
5. To define an APS configuration on a corresponding module located in a *different* switch:
   - Repeat Step 1 through Step 5 in the section, “Defining Physical Port Attributes” beginning on page 5-20.
   - Then, repeat Step 1 through Step 4 in the section, “Defining APS Attributes” beginning on page 5-24.
6. Continue with the following section to define ATM Direct Trunk logical ports.

Defining ATM Direct Trunk Logical Ports for APS

Use the following steps to create an ATM Direct Trunk logical port for each working and protection port.

1. Select the switch on which the first working/protection port pair resides. (You will define an ATM Direct Trunk logical port for each of these physical ports.)
2. From the Administer menu, select Ascend Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
3. Select the working port (of the APS pair) and press the third (right) mouse button to display a pop-up menu. Select Logical Port. The Set All Logical Ports in PPort dialog box appears (Figure 3-1 on page 3-2).
4. Choose Add to display the Add Logical Port dialog box (Figure 3-2 on page 3-6).
5. Select *ATM Direct Trunk* as the logical port type.
6. Choose OK. The Add Logical Port dialog box reappears (Figure 3-3 on page 3-7).
7. Use the instructions in Table 3-2 on page 3-9 to set the Administrative Attributes.
8. (*Optional*) Select Set Traffic Descriptor Attributes to configure Node and Card traffic descriptors. The fields in Figure 5-11 appear.

Figure 5-11. Set Traffic Descriptor Attributes

   a. Choose Node-to-Node Mgmt Traffic Descriptors. The Set Logical Port Node-to-Node Management Traffic Descriptor dialog box appears.
   b. See "Defining Traffic Descriptor Attributes" on page 8-9 to complete the fields on this dialog box.

> To use the Add Traffic Descriptor command to define a new traffic descriptor, see "Defining Network-wide Traffic Descriptors" on page 8-7.

   c. Choose OK to return to the Add Logical Port dialog box.
   d. Choose Trunk Signaling Traffic Descriptors.
   e. Repeat Step b and Step c to configure trunk signaling traffic descriptors.
9. When you finish, choose OK to return to the Set All Logical Ports in PPort dialog box.

10. Choose Close to return to the Set Physical Port Attributes dialog box. Choose Cancel to return to the Switch Back Panel dialog box.
11. To configure an ATM Direct Trunk logical port for the protection port, repeat Step 3 through Step 10.
12. Choose Close to return to the network map.
13. Select the switch on which the second working/protection port pair resides and complete Step 2 though Step 11 beginning on page 5-26 to define ATM Direct Trunk logical ports for the second APS pair.
14. Continue with the following section to define the Primary and Backup trunks.

Defining APS Primary and Backup Trunks

Complete the steps in the following sections to configure primary and backup trunks.

Configure the Primary Trunk

1. From the Administer menu, select Ascend Parameters ⇒ Set All Trunks. The Set All Trunks dialog box (Figure 5-2 on page 5-7) appears.
2. Choose Add. The Select Logical Ports dialog box appears (Figure 5-3 on page 5-12).
3. Select the name of the switch where the first working port resides, then select the name of the switch where the second working port resides.
4. For each switch endpoint, select the ATM Direct Trunk logical port that resides on the working port.
5. Choose OK.
6. Complete the fields as described in Table 5-2 on page 5-14. Be sure to select the Trunk Type *Primary.*
7. Use the instructions in Table 5-3 on page 5-15 to complete the additional fields that appear for a Primary Trunk.
8. Choose OK to complete this configuration and Close to return to the network map.

Configure the Backup Trunk

1. From the Administer menu, select Ascend Parameters ⇒ Set All Trunks. The Set All Trunks dialog box (Figure 5-2 on page 5-7) appears.
2. Choose Add. The Select Logical Ports dialog box appears (Figure 5-3 on page 5-12).
3. Select the name of the switch where the first protection port resides, then select the name of the switch where the second protection port resides.
4. For each switch endpoint, select the ATM Direct Trunk logical port that resides on the protection port.
5. Choose OK.
6. Complete the fields as described in Table 5-2 on page 5-14. Be sure to select the Trunk Type *Backup*.
7. When you select a Trunk Type of Backup, the dialog box displays the *Primary Trunk of the backup* field. Select the name of the primary trunk you configured using the corresponding APS working ports.
8. Choose OK to complete this configuration. Choose Close to return to the network map.

The APS Trunk Backup configuration is now complete.

6

Configuring ATM PVCs

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

- Point-to-Point
- Point-to-Multipoint

In addition, this chapter explains how to manually define PVCs and use the Move Circuit function.

Reliable Scalable Circuit

The NavisCore Reliable Scalable Circuit feature (set to *On* by default) improves PVC configuration reliability. This feature is set to *On* by default. 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 on this feature's reported error types, see Appendix F, "Reliable Scalable Circuit."

Disabling the Reliable Scalable Circuit Feature

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

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

Setting the VPI/VCI Values for PVCs

For each PVC you configure, you must specify a value from 0-*nnnn* to represent the Virtual Path Identifier (VPI) for the PVC (see page 6-11). The maximum value that you can specify is based on the Valid Bits in VPI that is configured for the logical port, as follows:

Maximum value = $2^P - 1$

where *P* is the value in the Valid Bits in VPI field. For example, if you entered 5 in the Valid Bits in VPI field, the maximum value is 31 ($2^5 - 1 = 31$) which would give you up to 32 virtual paths (numbered 0-31). See page 2-8 for details on setting the Valid Bits in VPI.

If you are defining a Virtual Channel Connection (VCC), you must also specify a value to represent the Virtual Channel Identifier (VCI) for an ATM circuit (see page 6-11). The maximum value that you can specify is based on the Valid Bits in VCI value that is configured for the logical port, as follows:

Maximum value = $2^C - 1$

where C is the value in the Valid Bits in VCI field. For example, if you entered 6 in the Valid Bits in VCI field, the maximum VCI value you can enter is 63 (which would give you 32 virtual channels, numbered 32 to 63).

These VPI/VCI range restrictions only apply to VCCs. You can provision a Virtual Path Connection (VPC) to any value in the VPI=0-255 range. In addition, if the logical port uses the NNI cell header format, you can provision VPCs over the 0 - 4095 range. For more information on the Valid Bits in VPI/VCI fields, see page 2-8.

The VPI/VCI combination must be unique at each circuit endpoint (including multipoint circuits). As a result, since a VPC has access to all valid VCIs, a VCC or multipoint circuit that uses a VPI that is already assigned to a VPC cannot be established, nor can a VPC be established if the selected VPI is already assigned to a VCC or multipoint circuit.

Accessing the Set All PVCs On Map Dialog Box

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.

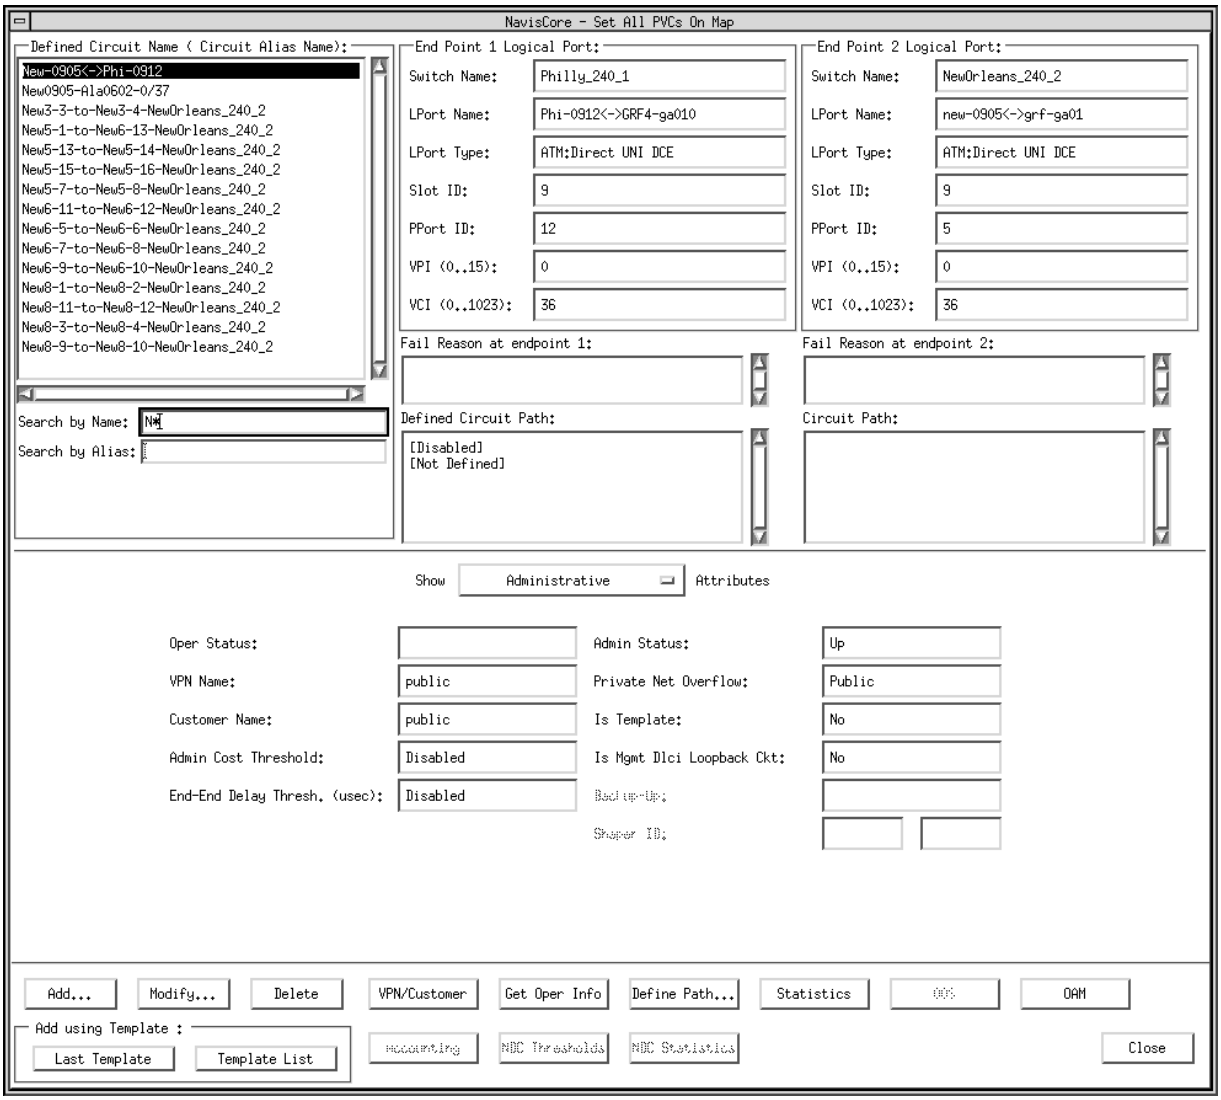

Figure 6-1. Set All PVCs On Map Dialog Box

- To view a list of configured circuits, position the cursor in the Search by Name field and press Return.
- To learn more about the Set All PVCs on Map dialog box, continue with the following section.
- To begin defining a PVC, proceed to page 6-8.

About the Set All PVCs On Map Dialog Box

The Set All PVCs On Map dialog box displays configured option information for the circuit name/circuit alias you select. To use a wildcard search to find a specific circuit name or alias, 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 these dialog box status fields and commands.

Table 6-1. Set All PVCs On Map Dialog Box Status Fields and Commands

Field/Command	Action/Description
Defined Circuit Name	Displays a listing 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 circuit endpoint. *LPort Name* – Displays the name of the logical port at each circuit endpoint. *LPort Type* – Displays the configured logical port type at each circuit endpoint. *Slot ID* – Indicates the physical slot number where the I/O module containing the selected port is installed. *PPort ID* – Displays the ID number of the physical port on which the selected logical port is configured. *VPI (0..nnnn)* – Displays the VPI for the selected circuit endpoint. *VCI (32..nnnn)* – Displays the VCI for the selected circuit endpoint. For more information on the Valid Bits in VPI/VCI fields, see page 2-8.
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.
Actual Circuit Path	Displays the actual path that OSPF selected for this circuit to use to get to its destination.

Table 6-1. Set All PVCs On Map Dialog Box Status Fields and Commands (Continued)

Field/Command	Action/Description
Show [Administrative] Attributes menu	Displays the appropriate attributes configured for the selected option. See one of the following sections for more information: *For ATM PVCs:* "Administrative Attributes" on page 6-12 "User Preference Attributes" on page 6-15 "Traffic Type Attributes" on page 6-13 "Frame Discard Attributes" on page 6-18 "Extended QoS Parameter Attributes" on page 6-19 *For Interworking PVCs* "Administrative Attributes" on page 6-12 "User Preference Attributes" on page 6-30 "Traffic Type Attributes" on page 6-26
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 explanations. *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 logical port.) *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.
Add	Adds a new PVC.
Modify	Modifies the selected circuit. The Modify command displays dialog boxes that are similar to those displayed for Add circuit; however, you cannot modify the circuit name, logical port endpoints, circuit type, or VPI/VCI values from this dialog box.
Delete	Deletes the selected circuit.
VPN/Customer	Assigns the selected circuit to a specific VPN and customer name.

Table 6-1. Set All PVCs On Map Dialog Box Status Fields and Commands (Continued)

Field/Command	Action/Description
Get Oper Info	Displays the selected circuit's current operational status in the Oper Status field. See page 6-6 for an explanation of these status messages.
Defined Path	Manually defines a circuit path (see page 6-34).
Statistics	Displays the summary statistics for the selected circuit.
OAM	Runs the Operations, Administration, and Management loopback diagnostics for the selected circuit.
Add using Template	If you have already defined a circuit configuration and saved it as a template, use this command to define a new circuit using similar parameters. • Choose Last Template to apply the last template you used to establish a circuit in this NMS session. • Choose Template List to display a list of previously defined templates for this map.
Accounting	Accesses the NavisXtend Accounting server functions for a PVC.
NDC Thresholds (*CBX 500 only*)	Displays the configured network data collection (NDC) thresholds for the selected circuit.
NDC Statistics (*CBX 500 only*)	Displays the NDC statistics for the selected circuit.
Close	Exit this dialog box and return to the network map.

Defining a PVC Connection

To set up a PVC connection between two UNI or NNI logical ports:

1. See page 6-4 for instructions for accessing the Set All PVCs On Map dialog box (Figure 6-1 on page 6-4).
2. From the Set All PVCs On Map dialog box, choose Add. The following dialog box appears.

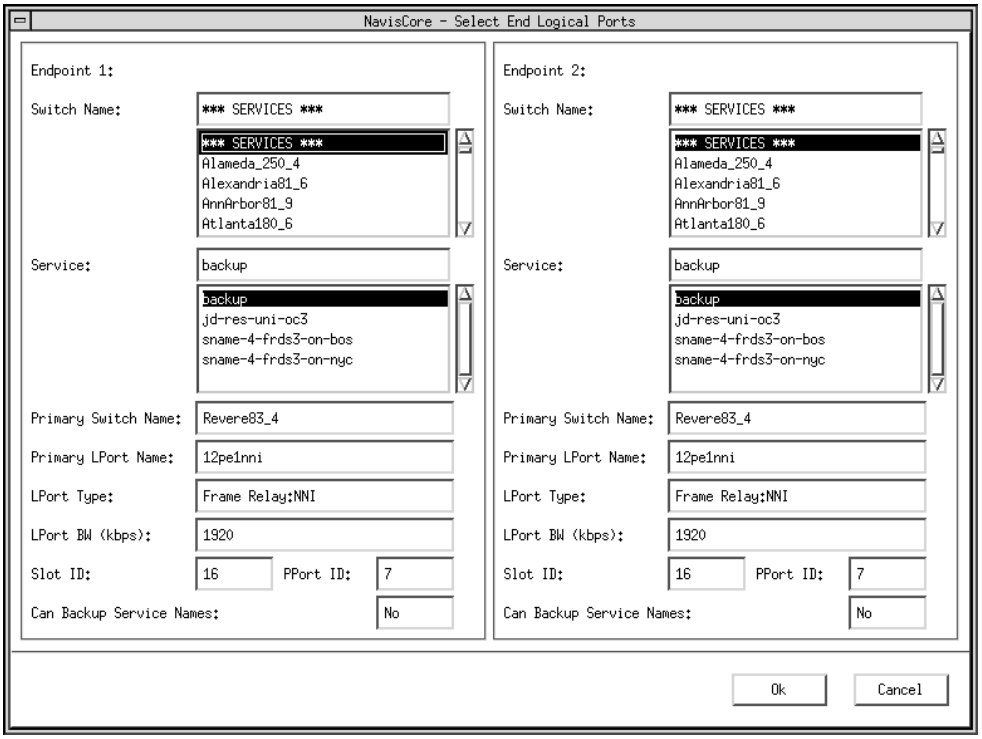

Figure 6-2. Select End Logical Ports Dialog Box

3. Configure Endpoint 1 and Endpoint 2 as follows:

For a fault-tolerant PVC configuration

For more information about fault-tolerant PVCs, see Chapter 10.

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

b. Select a service name from the list. You can only configure a fault-tolerant PVC for ATM UNI DCE and UNI DTE logical port types.

c. Continue with Step 4 on page 6-9.

For a standard circuit configuration

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

b. Select the name of the logical port for Endpoint 1, then select the name of the logical port for Endpoint 2. Note that if you enable the VPN/Customer view function (see page 9-7), only logical ports that belong to the VPN or customer you select appear in this list.

c. Continue with Step 4.

4. The Select End Logical Ports dialog box displays the following information for both Endpoint 1 and Endpoint 2:

Primary Switch Name — (*fault-tolerant PVC only*) Displays the name of the switch on which the primary (active) logical port resides.

Primary LPort Name — (*fault-tolerant PVC only*) Displays the name of the primary (active) logical port endpoint.

LPort Type — Displays the logical port type for the selected logical ports.

LPort Bandwidth — Displays the logical port bandwidth for the selected logical ports. At each endpoint, logical ports may have different bandwidth.

Slot ID — Displays the I/O slot (number) where the cards for the selected logical ports reside.

PPort ID — Displays the port ID numbers for the selected logical ports.

5. Choose OK. The Add PVC dialog box appears.

Figure 6-3. Add PVC Dialog Box

Continue with one of the following sections according to the ATM service you are configuring:

- If both endpoints provide ATM services, continue with the following section, "Configuring an ATM Service PVC."
- If one endpoint provides Frame Relay services, continue with "Configuring Frame Relay-to-ATM Service Interworking Circuits" on page 6-20.

Configuring an ATM Service PVC

To configure attributes for this type of PVC, define the following parameters for each of the circuit's two endpoints.

VPI (0 .. nnn) — Enter a value from 0-*nnnn* to represent the Virtual Path Identifier for the PVC. The maximum value you can enter is based on the valid bits in VPI that are configured for the logical port. Note that zero is not a valid value for a management PVC. See page 6-3 for information about setting this value.

VCI (32..nnnn) — (*for VCCs only*) Depending on the circuit configuration, enter a value to represent the Virtual Channel Identifier for an ATM PVC. See page 6-3 for information about setting this value.

About the Set Attributes Menu

When you configure a PVC, the dialog box provides detailed parameters that you need to specify for each endpoint. During this procedure, you 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.

NDC — Defines CBX 500 Network Data Collection (NDC) functions, which can detect any violation of PVC service subscription parameters, and establish trends in network traffic patterns and loads. See the *NavisCore Diagnostic and Troubleshooting Guide* for information about NDC functions.

Frame Discard — Defines UNI 4.0 signaling frame discard features for the forward and/or backward direction. The method of achieving frame discard depends on the implementation of early packet discard/partial packet discard (EPD/PPD) in your network. Your equipment must support frame discard.

Extended QoS Parameters — Defines the ATM Forum TM 4.0 Extended QoS Parameters. This selection enables you to define cell delay variation (CDV) and cell loss ratio (CLR) in the forward and reverse direction.

Continue with the following sections to configure these parameters.

Administrative Attributes

Select Set [Administrative] Attributes and complete the fields described in Table 6-2.

Table 6-2. Administrative Attributes

Field	Action/Description
Circuit Name	Enter any unique, alphanumeric name to identify the circuit. Do not use parentheses and asterisks. This name must be unique to the entire map.
Circuit Alias	(*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 NavisXtend Report Generator. Enter any unique, alphanumeric name to identify the circuit. Do not use parentheses and asterisks. This name must be unique to the entire map.
Admin Status	Select Up (default) to activate the circuit at switch startup, or Down if you do not want to activate the circuit at switch startup.
Circuit Type	Specify whether the circuit is a Virtual Path Connection (VPC) or Virtual Channel Connection (VCC, the default). If you select VPC, the VCI field is set to 0 and cannot be changed. A VPC enables a network that interfaces with an OPTimum trunk to accept circuits with this VPI and any of its valid VCIs.
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 9. 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). *Restrict* – 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)*	You can save these settings as a template to configure another PVC with similar options. To create a template, choose Yes in the "Template" field. The default value is No. See "Using Templates" on page 6-48 for more information.
Mgmt Loop Back Ckt	If you choose Yes, this PVC configuration will be included in the NMS initialization script file. This file contains all the SNMP set requests necessary to replicate the entire switch configuration. Once you download the configuration file to the switch, the PVC can be used to establish NMS-to-switch connectivity. The Mgmt Loop Back Ckt option is especially useful in some Management DLCI configurations. The default value is No.

Table 6-2. Administrative Attributes (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 in a value of 1000, the PVC will not be routed over a path whose total admin cost exceeds 1000. The NMS calculates the total admin cost for a path by using the sum of the admin cost for each trunk in the path. The valid range for this field is 1 – 65534.
End-End Delay Threshold (CTD)	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 in a value of 500 µsec., the PVC will not be routed over a path whose total end-to-end delay exceeds 500 µsec. 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. The valid range for this field is 0 – 167777214 µsec. ***Note:*** *The value you enter should reflect your network topology. If a PVC typically traverses high speed trunks, set the delay rate lower. You need to increase the delay if the PVC uses low-speed trunks.*

Traffic Type Attributes

Select Set [Traffic Type] Attributes (see Figure 6-4) to specify Traffic Descriptor settings for forward and reverse traffic. For more information about using ATM traffic descriptors, see Chapter 8, "Configuring ATM Traffic Descriptors."

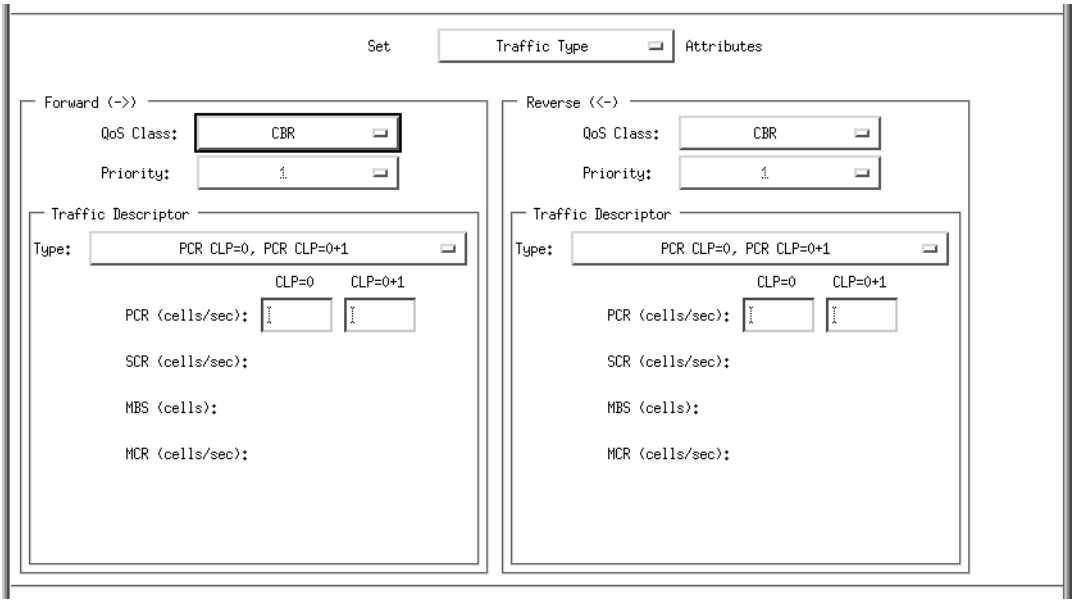

Figure 6-4. Set Traffic Type Attributes

Forward traffic is traffic from Endpoint 1 to Endpoint 2, and reverse traffic is from Endpoint 2 to Endpoint 1. Complete the fields described in Table 6-3 to set traffic type attributes in each direction.

Table 6-3. Traffic Type Attributes

Field	Action/Description
QoS Class (Fwd/Rev)	Select the Quality of Service class for forward and reverse traffic. The forward and reverse QoS classes do not have to match. The QoS class determines which traffic descriptors you can select. For more information on QoS, see Table 8-1 on page 8-3. ***Note:*** *For a CBX 500 that uses the flow control processor, RM cells are sent in the backward direction. As a result, they assume the QoS class of the other direction.*
Priority (Fwd/Rev) (*VBR-NRT and VBR-RT QoS classes on CBX/GX only*)	Select both the forward and reverse circuit priority, where 1 is high priority, 2 is medium priority, 3 is low priority, and 4 is lowest priority. (Note that for a B-STDX endpoint the priority range is from 1 – 3 only.) The forward and reverse circuit priority values do not have to match. CBR QoS class priority is set to 1.
Traffic Descriptor Type	Select from the following traffic descriptor options: *PCR CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes PCR CLP=0. If so, specify the PCR in cells per second for high-priority traffic (i.e., the CLP=0 cell stream). *PCR CLP=0+1 (cells/sec)* – Specify the PCR in cells per second for the combined high- and low-priority traffic (i.e., the CLP=0+1 aggregate cell stream). *SCR CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes SCR CLP=0. If so, specify the SCR in cells per second for the combined high-priority traffic (i.e., the CLP=0 cell stream). *SCR CLP=0+1 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes SCR CLP=0+1. If so, specify the SCR in cells per second for the combined high- and low-priority traffic (i.e., the CLP=0+1 aggregate cell stream). *MBS CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes MBS CLP=0. If so, specify the MBS (in cells per second) for the combined high-priority traffic (i.e., the CLP=0 cell stream). *MBS CLP=0+1 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes MBS CLP=0+1. If so, specify the MBS (in cells per second) for the combined high- and low-priority traffic (i.e., the CLP=0+1 cell stream). *MCR CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes MCR CLP=0. If so, specify the MCR (in cells per second) for the combined high-priority traffic (i.e., the CLP=0 cell stream). ***Note:*** *While the MCR traffic descriptor is only applicable to a CBX 500 with an FCP, this attribute is offered as a selection on non-CBX endpoints. This is because even though one or both endpoints may not be on a CBX with FCP, the PVC might traverse a CBX 500 FCP trunk. In this case, the provisioned attribute is used.*

Table 6-3. Traffic Type Attributes (Continued)

Field	Action/Description
Shaper (*B-STDX CS/IWU endpoint only*)	If this circuit carries ATM cell traffic, use the default of NO SHAPER. If this circuit carries frame relay traffic, select one of the configured shapers in the following pull-down menu. 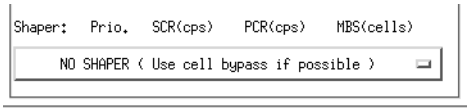 These shapers correspond to the traffic shapers configured for the physical port on which this logical port resides. For information about physical port traffic shaping, see the *NavisCore Physical Interface Configuration Guide*.

User Preference Attributes

Select Set [User Preference] Attributes (see Figure 6-5) and complete the fields described in Table 6-4 on page 6-16.

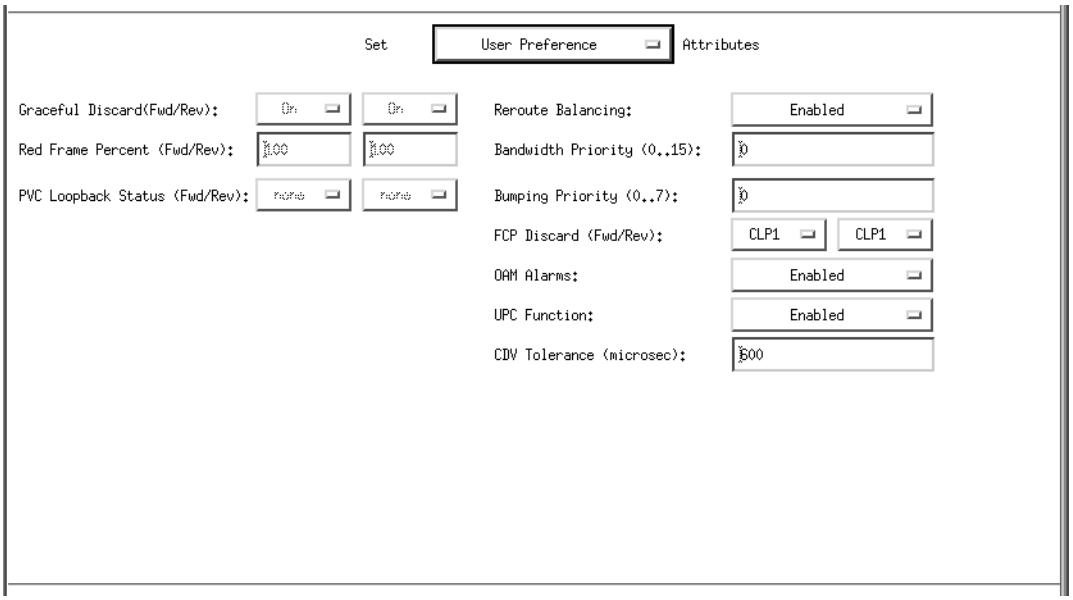

Figure 6-5. Set User Preference Attributes

Table 6-4. User Preference Attributes

Field	Action/Description
Reroute Balancing	When enabled (default), the PVC conforms to the configured reroute tuning parameters. This means that when the PVC reroutes during trunk failure, it will migrate back to its original trunk at a rate and time determined by the configured reroute tuning parameters. When disabled, the PVC ignores the switch tuning parameters. For more information, see the *NavisCore NMS Getting Started Guide*.
Bandwidth Priority	Set a value from 0 through 15 where 0 is the default and indicates the highest priority. See Appendix E, "Priority Routing," for more information.
Bumping Priority	Set a number from 0 through 7 where 0 is the default and indicates the highest priority. See Appendix E, "Priority Routing," for more information.
FCP Discard	Displays only if you selected a QoS class that supports FCP Discard. Select one of the following options: *CLP1* – You can provision selective CLP1 discard for UBR, ABR, and VBR-NRT PVCs. If the current cell causes the queue for a PVC to exceed the discard thresholds, and the cell has CLP set to 1, the cell is discarded. Note that EPD is not performed in this case. *EPD* – Early Packet Discard. The ATM Flow-control processor can perform EPD for UBR, ABR, and VBR-NRT PVCs. If you select this option, when a cell causes the queue for a PVC to exceed the discard thresholds, the VC enters the EPD state. The cells in the current packet of the VC are admitted to the queue. However, when the end of the current packet is detected, all of the cells in the next packet are discarded for that PVC. See "ATM Flow-Control Discard Mechanisms" on page D-16 for more information. ***Note:*** *While the frame discard attribute is only applicable to a CBX 500 with an FCP, this attribute is offered as a selection on non-CBX endpoints. This is because even though one or both endpoints may not be on a CBX with FCP, the PVC might traverse a CBX 500 FCP trunk. In this case, the provisioned attribute is used.*
OAM Alarms (*CBX/GX and ATM CS/IWU modules only*)	Set to Enabled (default) to use OAM alarms on this circuit. Set to Disabled to disable OAM alarms on this circuit. When enabled, the switch sends OAM F5 or F4 AIS (alarm indicator signal) cells out of each UNI logical port endpoint to indicate that the circuit is down. ***Note:*** *For a management PVC, this field is set to disabled and cannot be changed.*

Table 6-4. User Preference Attributes (Continued)

Field	Action/Description
UPC Function (*PVCs with CBX/GX endpoints only*)	Enables (default) or disables the Usage Parameter Control (UPC) function. When you enable UPC, the circuit tags or drops cells as they come into the port that do not conform to the configured traffic descriptors. When you disable UPC, the circuit allows all traffic, including non-conforming traffic, into the port. As a result, when you disable UPC, quality of service is no longer guaranteed for circuits in the network due to the potential for increasing the cell loss ratio because of port congestion. For this reason, *Ascend recommends that you enable the UPC function on all circuits.* For information about UPC traffic parameters, see Chapter 8, "Configuring ATM Traffic Descriptors." ***Note:*** *To use the UPC function for individual circuits, verify that the UPC function is enabled for both logical port endpoints on which you will define the circuit. Enabling UPC at the circuit level has no effect if you did not enable UPC at the logical port level. UPC is enabled by default for both logical ports and circuits.*
CDV Tolerance (*PVCs with CBX/GX endpoints only*)	Configure the Cell Delay Variation Tolerance (CDVT). The UPC uses this value to police the requested traffic descriptor. Valid values are between 1 - 65535 µs. The default is 600 µs.
Graceful Discard (Fwd/Rev) (*ATM UNI endpoint on frame-based card*)	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. ***Note:*** *For the ATM UNI DS3/E3, if you set this value for shaping purposes, the switch code ignores the PCR, SCR, and MBS values calculated from Set Traffic Descriptor Attributes (Figure 6-9 on page 6-26); the switch instead picks the highest PCR queue available and sets the SCR to that PCR.*
Red Frame Percent (Fwd/Rev) (*ATM UNI endpoint on frame-based card*)	Set this value only if Graceful Discard is set to *On*. See "Graceful Discard" on page 6-22 for more information. The Red Frame Percent limits the number of red frames the network is responsible to deliver.
PVC Loopback Status (Fwd/Rev) (*ATM UNI endpoint on frame-based card*)	Displays the current loopback state. If **None** is not displayed in the PVC Loopback Status field, do not attempt to modify or delete the selected circuit. See the *NavisCore Diagnostic and Troubleshooting Guide* for more information about loopback testing.

If both ATM endpoints reside on a CBX 500 or GX 550 switch, proceed to the following section, "Frame Discard Attributes." Otherwise, continue with "Completing the PVC Configuration" on page 6-20.

Frame Discard Attributes

This option menu (see Figure 6-6) only appears if both endpoints reside on either a CBX 500 or GX 550 switch. Select Set [Frame Discard] Attributes and complete the fields described in Table 6-5.

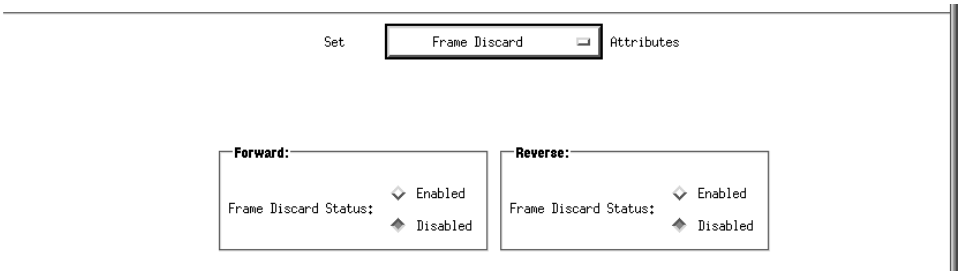

Figure 6-6. Set Frame Discard Attributes

You should only enable frame discard if the traffic traversing the VC is encapsulated using AAL5. If frame discard is enabled on VCs that are not using AAL5 encapsulation, all traffic traversing the VC may be discarded.

Table 6-5. Frame Discard Attributes

Field	Description
Forward Frame Discard Status	Select Enable to turn on the physical port output buffer early packet discard/partial packet discard (EPD/PPD) function for this particular PVC. When enabled, AAL5 traffic that is traversing the PVC will be subject to EPD/PPD when physical port congestion is experienced.
Reverse Frame Discard Status	

Extended QoS Parameter Attributes

This option menu (see Figure 6-7) only appears if both endpoints reside on either a CBX 500 or GX 550 switch. Select Set [Extended QoS Parameters] Attributes and complete the fields described in Table 6-6.

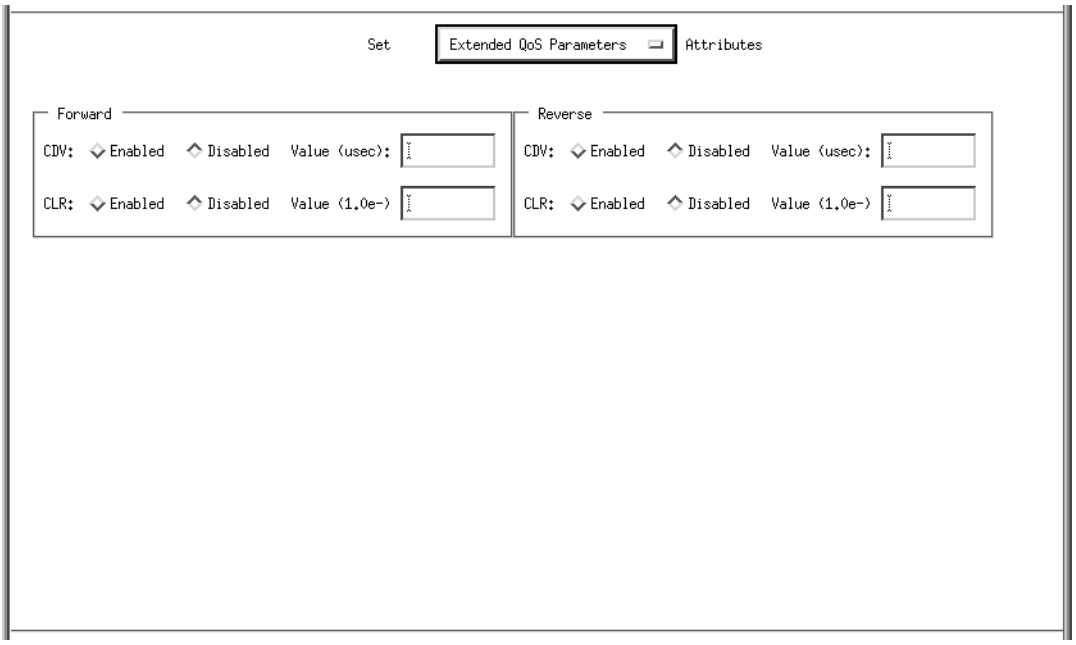

Figure 6-7. Set Extended QoS Parameter Attributes

Table 6-6. Extended QoS Parameter Attributes

Field	Action/Description
Forward/ Reverse CDV	If you enable this option, the PVC will not be routed over a path whose total cell delay variation (CDV) exceeds the entered value. If you enable this field and enter in a value of 1000 µsec, the PVC will not be routed over a path whose total CDV exceeds 1000 µsec. The total CDV for a path is calculated by summing the CDV for each trunk in the route. The valid range for this field is 1 – 167777214 µsec.
Forward/ Reverse CLR	If you enable this option, the PVC will not be routed over a path if the cell loss ratio (CLR) of one of the trunks exceeds the entered value. If you enable this field and enter in a value of 10, the PVC will not be routed over a path that has one or more trunks with a CLR worse than 1.0 e $^{-10}$. The CLR for a trunk is based on the CAC objective for the host switches. The valid range for this field is 1.0e^{-1} to 1.0e^{-12}. Enter a value between 1 and 12.

Completing the PVC Configuration

Use the following steps to complete the circuit configuration.

1. *(Optional)* To configure CBX 500 Network Data Collection parameters for this circuit, select Set [NDC] Attributes. For more information, see the *NavisCore Diagnostic and Troubleshooting Guide*.
2. *(Optional)* To configure NavisXtend Accounting Server parameters for this circuit, choose the Accounting button. For more information, see the *NavisXtend Accounting System Administrator's Guide.*
3. Choose OK to define the circuit parameters. The Set All PVCs on Map dialog box reappears (Figure 6-1 on page 6-4).

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 F for more information.

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

Configuring Frame Relay-to-ATM Service Interworking Circuits

This section describes how to configure the following circuits for B-STDX ATM services.

Frame Relay-to-ATM Service Interworking — This service uses a circuit with a Frame Relay logical port at one endpoint and an ATM logical port at the other endpoint. The circuit uses a 10-bit address called a *data link connection identifier* (DLCI). DLCIs identify the logical endpoints of a virtual circuit and have local significance only.

ATM Data Exchange Interface/Frame User-to-Network Interface (DXI/FUNI) — This service uses a circuit with an ATM logical port defined on a frame-based I/O module, such as the 8-port Universal I/O module. The circuit is identified by a 4-bit *virtual path identifier* (VPI) and a 6-bit *virtual channel identifier* (VCI). Circuits on the ATM DS3/E3 module use an 8-bit VCI.

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

A *virtual channel* (VC) is a connection between two communicating ATM devices. A VC may consist of a group of several ATM links, customer premise equipment (CPE) to central-office switch, switch-to-switch, and switch-to-user equipment.

Frame Relay to ATM Service Interworking enables a Frame Relay device to connect to an ATM user device over a common WAN backbone. Frame Relay to ATM Service Interworking provides a seamless communication between ATM and Frame Relay networks or end-user devices.

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 configured level (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 under various congested states or congestion conditions. Table 6-7 describes the designators (traffic colors) and discard policy.

Table 6-7. 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 for delivering. 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 on the rate-enforcement discard process, see the *Networking Services Technology Overview Guide*.

Rate Enforcement Schemes

Rate enforcement schemes provide more flexibility, increased rate enforcement accuracy, and improved switch performance. You configure the rate enforcement scheme in the Add PVC dialog box under the Traffic Type attributes (see Table 6-10 on page 6-27).

Table 6-8 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 6-8. Rate Enforcement Schemes

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

Defining Interworking PVCs

To define a circuit for Frame Relay-to-ATM service interworking:

1. Follow Step 1 through Step 5 beginning on page 6-20 to select the PVC endpoints.
2. The Add PVC dialog box displays the fields shown in Figure 6-8.

Figure 6-8. Set Administrative Attributes Dialog Box (FR-ATM Service IW)

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

Table 6-9. Set Administrative Attributes Fields

Field	Action/Description
DLCI *(Frame Relay endpoint)*	Enter a unique DLCI for this logical port.
VPI (0..*nnnn*) *(ATM endpoint)*	Enter a value from 0 to *nnnn* to represent the virtual path identifier for an ATM circuit. The maximum value you can enter is based on the valid bits in VPI that are configured for the logical port. Note that zero is not a valid value for a management PVC. See page 6-3 for information about setting this value. • For an ATM CS or ATM IWU module, the VPI range depends on the number of VPI bits selected on the physical port. See the *NavisCore Physical Interface Configuration Guide* for more information. • For an ATM UNI DS3/E3 module, the number of VPI bits is set to four; the VPI range is 0 – 15.
VCI (32..*nnn*) *(ATM endpoint only, VCCs only)*	Enter a value to represent the Virtual Channel Identifier for an ATM circuit. See page 6-3 for information about setting this value. When you configure the ATM circuit: • On a frame-based I/O module, enter a value from 32 to 63. • On an ATM-based I/O module (such as the ATM DS3 module), enter a value from 32 to 255. • On an ATM CS or ATM IWU module, the total number of bits available for the VPI and VCI is 12 bits. For example, if the VPI is set to 1, there are 11 bits available for the VCI. If the VPI is set to 2, there are 10 bits available for the VCI. ***Note**: If you are configuring the VCI on an ATM CS or ATM IWU module, the VCI range depends on the number of VPI bits selected on the physical port. See the* NavisCore Physical Interface Configuration Guide *for more information.*
Administrative Attributes Fields	
Circuit Name	Enter any unique, continuous, alphanumeric name to identify the circuit. Do not use parentheses and asterisks. You can use hyphens.
Circuit Alias	(*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 NavisXtend Report Generator. Enter any unique, alphanumeric name to identify the circuit. Do not use parentheses and asterisks. This name must be unique to the entire map.
Admin Status	Select Up (default) to activate the circuit at switch startup, or Down if you do not want to activate the circuit at switch startup.

Table 6-9. Set Administrative Attributes Fields (Continued)

Field	Action/Description
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 9. 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). *Restrict* – 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*)	You can save these settings as a template to configure another PVC with similar options. To create a template, choose Yes in the "Is Template" field. The default value is No. See "Using Templates" on page 6-48 for more information.
Mgmt Loopback Ckt	If you choose Yes, this PVC configuration will be included 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.
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 in a value of 1000, the PVC will not be routed over a path whose total admin cost exceeds 1000. The total admin cost for a path is calculated by summing the admin cost for each trunk in the path. The valid range for this field is 1 – 65534.
End-End Delay Threshold (CTD)	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 in a value of 500 µsec., the PVC will not be routed over a path whose total end-to-end delay exceeds 500µsec. 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.*

Using the Set Attributes button, complete the following sections:

- "Traffic Type Attributes" on page 6-26
- "User Preference Attributes" on page 6-30

Traffic Type Attributes

Select Set [Traffic Type] Attributes to specify Traffic Descriptor settings for forward and reverse traffic.

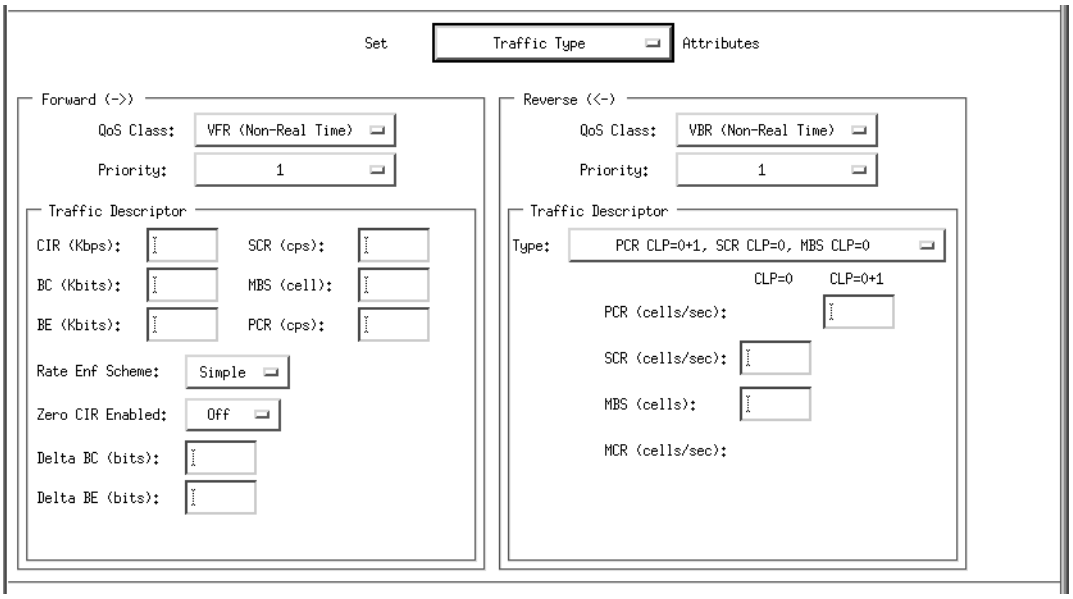

Figure 6-9. Add PVC-Set Traffic Type Dialog Box (FR-ATM Service IW)

In the example shown in Figure 6-9, configure the fields beneath Forward (–>); then configure the fields beneath Reverse (<–). The attributes that appear depend upon the endpoint type, either Frame Relay or ATM. Use Table 6-10 to configure these attributes.

Table 6-10. Set Traffic Type Attributes Fields (Traffic Parameters)

Field	Action/Description
Frame Relay Endpoint Traffic Parameters	
QoS Class (Fwd/Rev)	Select one of the following Frame Relay Class of Service values: *VFR (Real-Time)* – Variable Frame Rate (VFR). Used for packaging special delay-sensitive applications, such as packet video, which require low cell-delay variation between endpoints. *VFR (Non-Real Time)* – Handles packaging for transfer of long, bursty data streams over a pre-established ATM connection. This service is also used for short, bursty data, such as LAN traffic. CPE protocols adjust for any delay or loss incurred through the use of VFR-NRT. *UFR* – Unspecified Frame Rate. Primarily used for LAN traffic. The CPE should compensate for any delay or lost cell traffic. *ABR* – Available Bit Rate. Primarily used for LAN traffic. The CPE compensates for any delay or lost cell traffic. Select this option if the PVC will traverse a CBX 500 cloud that uses an FCP.
Priority	Select both the forward and reverse circuit priority, where 1 is high priority, 2 is medium priority, and 3 is low priority. The forward and reverse circuit priority values do not have to match.
CIR (Kbps) (Committed Information Rate)	Enter the CIR rate in Kbps at which the network transfers data under normal conditions. Normal conditions refer to a properly designed network with ample bandwidth and switch capacity. The rate is averaged over a minimum increment of the Committed Rate Measurement interval (Tc).
SCR (cps)	Displays the sustainable cell rate that is calculated from the CIR value you enter.
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.
MBS (cell)	Displays the maximum burst size (MBS) that is calculated from the BC value you enter.
BE (Kbits) (Excess Burst Size)	Enter the maximum amount of uncommitted data, in Kbits, the network will attempt to deliver during a specified time interval, Tc. Tc is calculated BC/CIR. The network treats this data as Discard Eligible (DE) data. ***Note**: For ATM UNI DS3/E3 modules, if the sum of BC + BE is greater than the value of MBS, you will get an error. If you set BC = CIR and BE to zero, traffic shaping is disabled on the ATM side of the circuit and MBS is forced to equal 32.*
PCR (cps)	Displays the peak cell rate that is calculated from the BE value you enter.

Table 6-10. Set Traffic Type Attributes Fields (Traffic Parameters) (Continued)

Field	Action/Description
Rate Enf Scheme	Select *Simple* (default) or *Jump*. The configurable rate enforcement scheme provides more flexibility, increased rate enforcement accuracy, and improved switch performance. See "Rate Enforcement Schemes" on page 6-22 for more information.
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 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)	The maximum number of bits the network agrees to transfer over the circuit (as committed bits) during the measurement interval provided there are positive committed bit (BC) credits before receiving the frame, but negative BC credits after accepting the frame. Set the number of Delta BC bits for this circuit between 0 - 53535 (*default* 53535).
Delta BE (bits)	The maximum number of bits the network agrees to transfer over the circuit (as excess bits) during the measurement interval provided there are positive excess bit (BE) credits before receiving the frame, but negative BE credits after accepting the frame. Set the number of Delta BE bits for this circuit between 0 - 53535. (*default* 53535).
ATM Endpoint Traffic Parameters	
QoS Class (Fwd/Rev)	Select the Quality of Service class for forward and reverse traffic. The forward and reverse QoS classes do not have to match. The QoS class determines which traffic descriptors you can select. For more information on QoS, see Table 8-1 on page 8-3. ***Note:** For a CBX 500 that uses the flow control processor, RM cells are sent in the backward direction. As a result, they assume the QoS class of the other direction.*
Priority (Fwd/Rev) (*VBR-NRT and VBR-RT QoS classes on CBX/GX only*)	Select both the forward and reverse circuit priority, where 1 is high priority and 4 is lowest priority. The forward and reverse circuit priority values do not have to match.

Table 6-10. Set Traffic Type Attributes Fields (Traffic Parameters) (Continued)

Field	Action/Description
Traffic Descriptor Type	Select from the following traffic descriptor options: *PCR CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes PCR CLP=0. If so, specify the PCR in cells per second for high-priority traffic (i.e., the CLP=0 cell stream). *PCR CLP=0+1 (cells/sec)* – Specify the PCR in cells per second for the combined high- and low-priority traffic (i.e., the CLP=0+1 aggregate cell stream). *SCR CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes SCR CLP=0. If so, specify the SCR in cells per second for the combined high-priority traffic (i.e., the CLP=0 cell stream). *SCR CLP=0+1 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes SCR CLP=0+1. If so, specify the SCR in cells per second for the combined high- and low-priority traffic (i.e., the CLP=0+1 aggregate cell stream). *MBS CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes MBS CLP=0. If so, specify the MBS (in cells per second) for the combined high-priority traffic (i.e., the CLP=0 cell stream). *MBS CLP=0+1 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes MBS CLP=0+1. If so, specify the MBS (in cells per second) for the combined high- and low-priority traffic (i.e., the CLP=0+1 cell stream). *MCR CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes MCR CLP=0. If so, specify the MCR (in cells per second) for the combined high-priority traffic (i.e., the CLP=0 cell stream). ***Note:*** *While the MCR traffic descriptor is only applicable to a CBX 500 with an FCP, this attribute is offered as a selection on non-CBX endpoints. This is because even though one or both endpoints may not be on a CBX with FCP, the PVC might traverse a CBX 500 FCP trunk. In this case, the provisioned attribute is used.*
Shaper (*B-STDX CS/IWU endpoint only*)	If this circuit carries ATM cell traffic, use the default of NO SHAPER. If this circuit carries frame relay traffic, select one of the configured shapers in this pull-down menu. 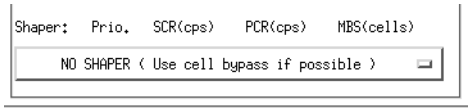 These shapers correspond to the traffic shapers configured for the physical port on which this logical port resides. For information about physical port traffic shaping, see the *NavisCore Physical Interface Configuration Guide*.

User Preference Attributes

Select Set [User Preference] Attributes and complete the fields described in Table 6-11.

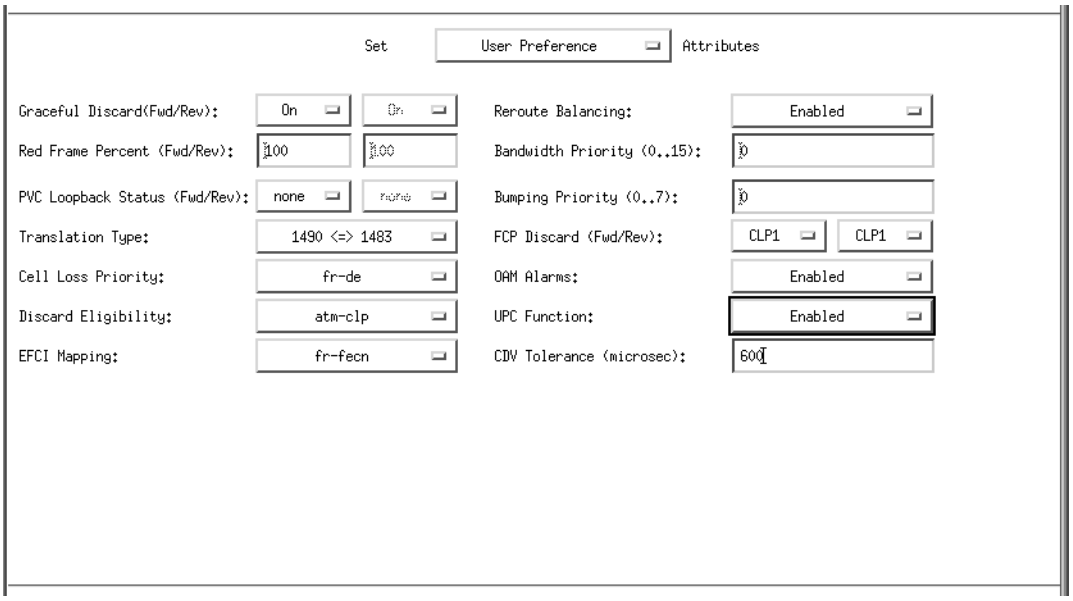

Figure 6-10. Set User Preference Attributes Dialog Box (FR-ATM Service IW)

Table 6-11. Set User Preference Attributes Fields

Field	Action/Description
Graceful Discard (Fwd/Rev) (*PVCs with frame relay UNI endpoints only*)	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 cannot discard this packet unless Graceful Discard is set to On. *On* – Forwards some red packets if there is no congestion. *Off* – Immediately discards red packets. ***Note:*** *For the ATM UNI DS3/E3, if you set this value for shaping purposes, the switch code ignores the PCR, SCR, and MBS values calculated from Set Traffic Descriptor Attributes (Figure 6-9 on page 6-26); the switch instead picks the highest PCR queue available and sets the SCR to that PCR.*
Red Frame Percent (Fwd/Rev) (*PVCs with frame relay UNI endpoints only*)	Set this value only if Graceful Discard is set to *On*. See "Graceful Discard" on page 6-22 for more information. The Red Frame Percent limits the number of red frames the network is responsible to deliver.

Table 6-11. Set User Preference Attributes Fields (Continued)

Field	Action/Description
PVC Loopback Status (Fwd/Rev)	Displays the current loopback state. If **None** is not displayed in the PVC Loopback Status field, do not attempt to modify or delete the selected circuit. See the *NavisCore Diagnostic and Troubleshooting Guide* for more information about loopback testing.
Translation Type	*(ATM endpoint only)* Select the ATM Translation Type protocol. Options include: *None* – Each end of the circuit uses the 1490 protocol. *RFC 1490 ⇔ 1483* – If you have a Frame Relay logical port on Endpoint 1 and an ATM logical port on endpoint 2. *RFC 1483 ⇔ 1490* – If you have an ATM logical port on Endpoint 1 and a Frame Relay logical port on Endpoint 2.
Cell Loss Priority	Specify the CLP setting. The CLP bit (cell loss priority) is in each cell's cell header. Options include: *fr-de* (ATM CS and IWU modules only) – Sets the CLP bit to the same value as the Frame Relay frame discard eligible (DE) bit on all ATM cells. This maps the DE bit to CLP. *0* – Sets the CLP bit to 0. *1* – Sets the CLP bit to 1.
Discard Eligibility	Select one of the following settings: *atm-clp* (ATM CS and IWU modules only) – Sets the CLP bit received in last cell of the frame to Frame Relay frame DE bit. *0* – Sets the DE to 0. *1* – Sets the DE to 1.
EFCI Mapping	Choose fr-fecn (*default*) to map the EFCI bit on the ATM endpoint to the frame relay FECN bit; choose 0 to ignore EFCI to FECN bit mapping.
Reroute Balancing	When enabled (default), the PVC conforms to the configured reroute tuning parameters. This means that when the PVC reroutes during trunk failure, it will migrate back to its original trunk at a rate and time determined by the configured reroute tuning parameters. When disabled, the PVC ignores the switch tuning parameters. For more information, see the *NavisCore Getting Started Guide*.

Table 6-11. Set User Preference Attributes Fields (Continued)

Field	Action/Description
Bandwidth Priority	Set a value from 0 through 3 where 0 is the default and indicates the highest priority. See Appendix E, "Priority Routing," for more information.
Bumping Priority	Set a number from 0 through 7 where 0 is the default and indicates the highest priority. See Appendix E, "Priority Routing," for more information.
FCP Discard	Displays only if you selected a QoS class that supports FCP Discard. Select one of the following options: *CLP1* – You can provision selective CLP1 discard for UBR, ABR, and VBR-NRT PVCs. If the current cell causes the queue for a PVC to exceed the discard thresholds, and the cell has CLP set to 1, the cell is discarded. Note that EPD is not performed in this case. *EPD* – Early Packet Discard. The ATM Flow-control processor can perform EPD for UBR, ABR, and VBR-NRT PVCs. If you select this option, when a cell causes the queue for a PVC to exceed the discard thresholds, the VC enters the EPD state. The cells in the current packet of the VC are admitted to the queue. However, when the end of the current packet is detected, all of the cells in the next packet are discarded for that PVC. See "ATM Flow-Control Discard Mechanisms" on page D-16 for details. ***Note:*** *While the frame discard attribute is only applicable to a CBX 500 with an FCP, this attribute is offered as a selection on non-CBX endpoints. This is because even though one or both endpoints may not be on a CBX with FCP, the PVC might traverse a CBX 500 FCP trunk. In this case, the provisioned attribute is used.*
OAM Alarms (*CBX/GX and ATM CS and IWU module endpoints only*)	Choose *Enable* to allow this circuit to generate OAM alarms to indicate whether the circuit is up or down. These alarms send a signal to the logical port whenever the circuit goes down or comes back up.

Table 6-11. Set User Preference Attributes Fields (Continued)

Field	Action/Description
UPC Function *(PVCs with CBX/GX endpoints only)*	Enables (default) or disables the Usage Parameter Control (UPC) function. When you enable UPC, the circuit tags or drops cells as they come into the port that do not conform to the configured traffic descriptors. When you disable UPC, the circuit allows all traffic, including non-conforming traffic, into the port. As a result, when you disable UPC, quality of service is no longer guaranteed for circuits in the network due to the potential for increasing the cell loss ratio because of port congestion. For this reason, *Ascend recommends that you enable the UPC function on all circuits.* For information about UPC traffic parameters, see Chapter 8, "Configuring ATM Traffic Descriptors." ***Note:*** *To use the UPC function for individual circuits, verify that the UPC function is enabled for both logical port endpoints on which you will define the circuit. Enabling UPC at the circuit level has no effect if you did not enable UPC at the logical port level. UPC is enabled by default for both logical ports and circuits.*
CDV Tolerance *(PVCs with CBX/GX and ATM CS and IWU module endpoints only)*	Enter a value between 1 - 65535 µs to define the Cell Delay Variation Tolerance (CDVT). The UPC uses this value to police the requested traffic descriptor. The default is 600 µs.

4. Choose OK to accept the circuit parameters and send the configuration information to the switch (provided the switch is communicating with the NMS).
5. The Set All PVCs dialog box reappears.
6. (*Optional*) To configure this PVC for a specific VPN and customer, see page 9-8.
7. 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 the OSPF algorithm's 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 Circuits on Map dialog box (Figure 6-1 on page 6-4) appears.
2. Select the circuit for which you want to manually define the circuit path. Choose Define Path. The following dialog box appears.

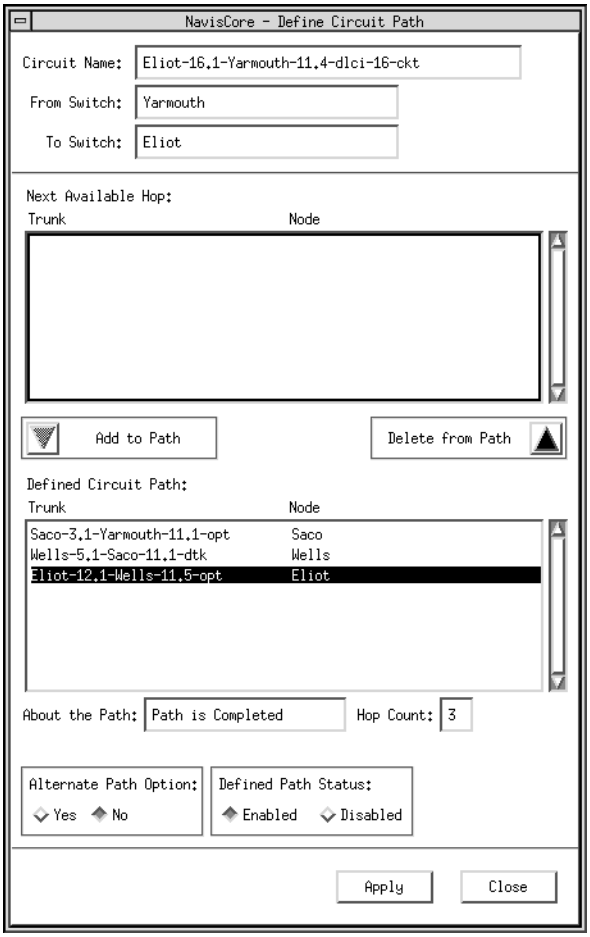

Figure 6-11. Define Circuit Path Dialog Box

3. Complete the dialog box fields as described in Table 6-12 on page 6-35.

Table 6-12. Define Circuit Path Fields

Field	Action/Description
Next Available Hop	Displays a listing of the available hops (e.g., trunk-node pairs). • Select the trunk-node pair through which you want to route the circuit. When there are multiple trunks between two nodes, select Any Trunks to cause OSPF to decide which is the best path to use at any given time. • Choose Add to Path. The trunk-node selection is added to the Trunk/Node field, which displays all selected hops.
Alternate Path Option	Select either Yes or No to define whether OSPF should route the circuit if the manual route fails. • Select Yes to enable OSPF to route the circuit based on the best available path. • Select No to disable the circuit from being rerouted; the circuit remains down until the defined path becomes available.
Defined Path Status	Select Enabled to route the circuit based on the manual route defined. Select Disabled to route the circuit based on the network's OSPF algorithm.

4. Choose Apply and then choose OK.

Moving Circuits

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

This function has the following restrictions:

- You should not move a circuit that is currently in use because it may lose traffic.
- You cannot move a circuit for which you have manually defined a circuit path.
- The VPI/VCI must be unique to the destination logical port.
- The Move Circuit function fails if the number of circuits moved exceeds the maximum allowed for the I/O card.

To move a circuit:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Circuits ⇒ Move Circuit Endpoint. The following dialog box appears.

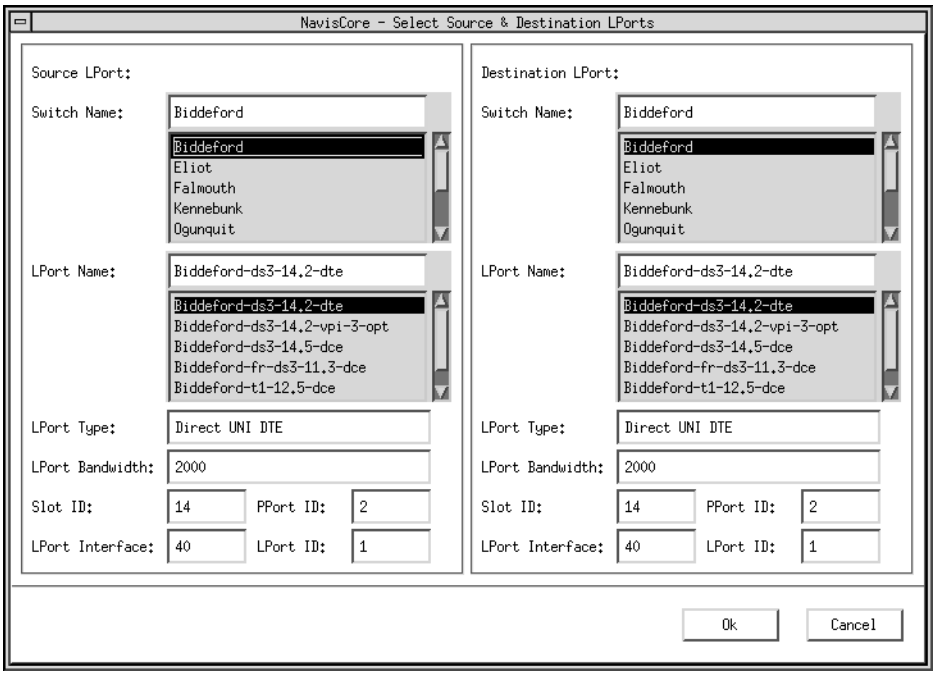

Figure 6-12. Select Source & Destination LPorts Dialog Box

2. To select the Source LPort that contains the circuit you want to move:
   a. Select the Switch Name.
   b. Select the LPort Name.
3. To select the Destination LPort to which you want to move the circuit:
   a. Select the Switch Name.
   b. Select the LPort Name.

4. Choose OK. The following dialog box appears. This dialog box displays the circuits that have the source logical port as an endpoint.

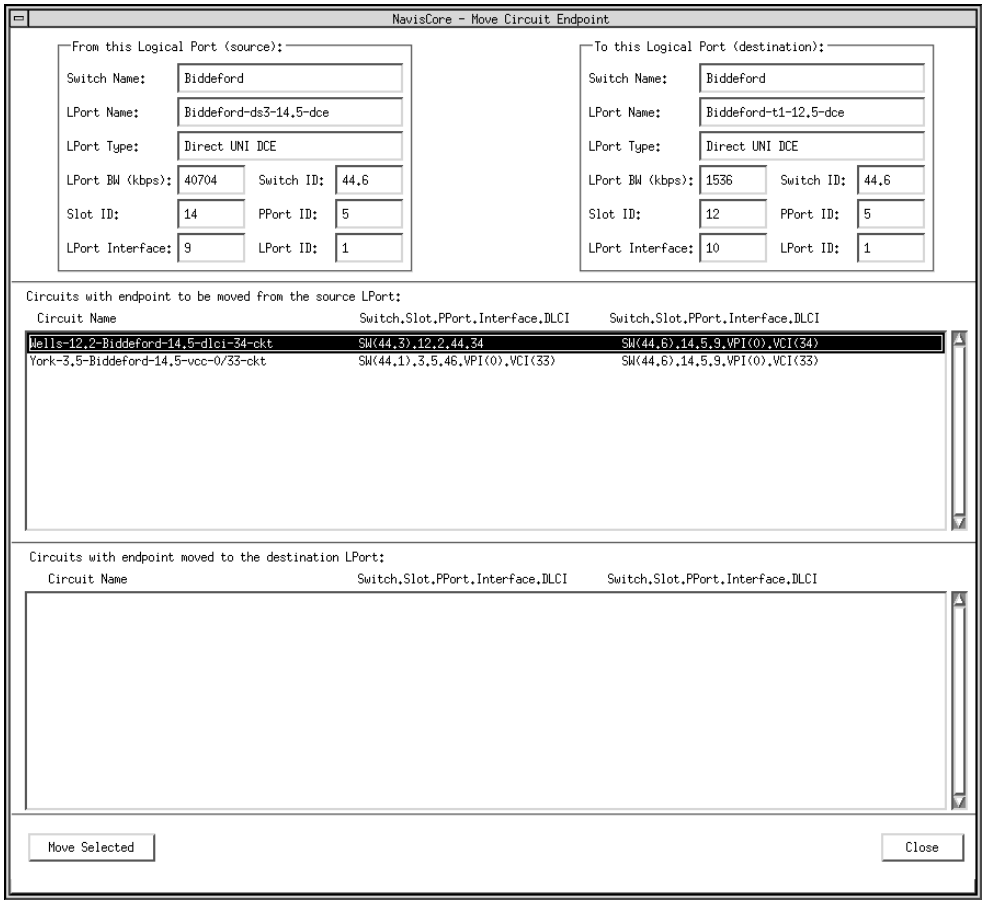

Figure 6-13. Move Circuit Dialog Box

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

6. Choose Move Selected (or Move All to move more than one circuit). The selected circuit appears in the *Circuits with endpoint moved to the destination LPort* list.

7. When you finish, choose Close.

Configuring Point-to-Multipoint Circuits

A point-to-multipoint (PMP) circuit consists of the originating point (circuit root), and endpoints (circuit leafs). The endpoints of a given PMP circuit can be on any switch in the network map, and on any number of switches (that is, the endpoints do not have to terminate on the same switch).

To access the Set All Point-to-Multiple-Point Circuit Roots dialog box, from the Administer menu, select Ascend Parameters ⇒ Set All Circuits ⇒ Point-to-Multipoint Circuits. The following dialog box appears.

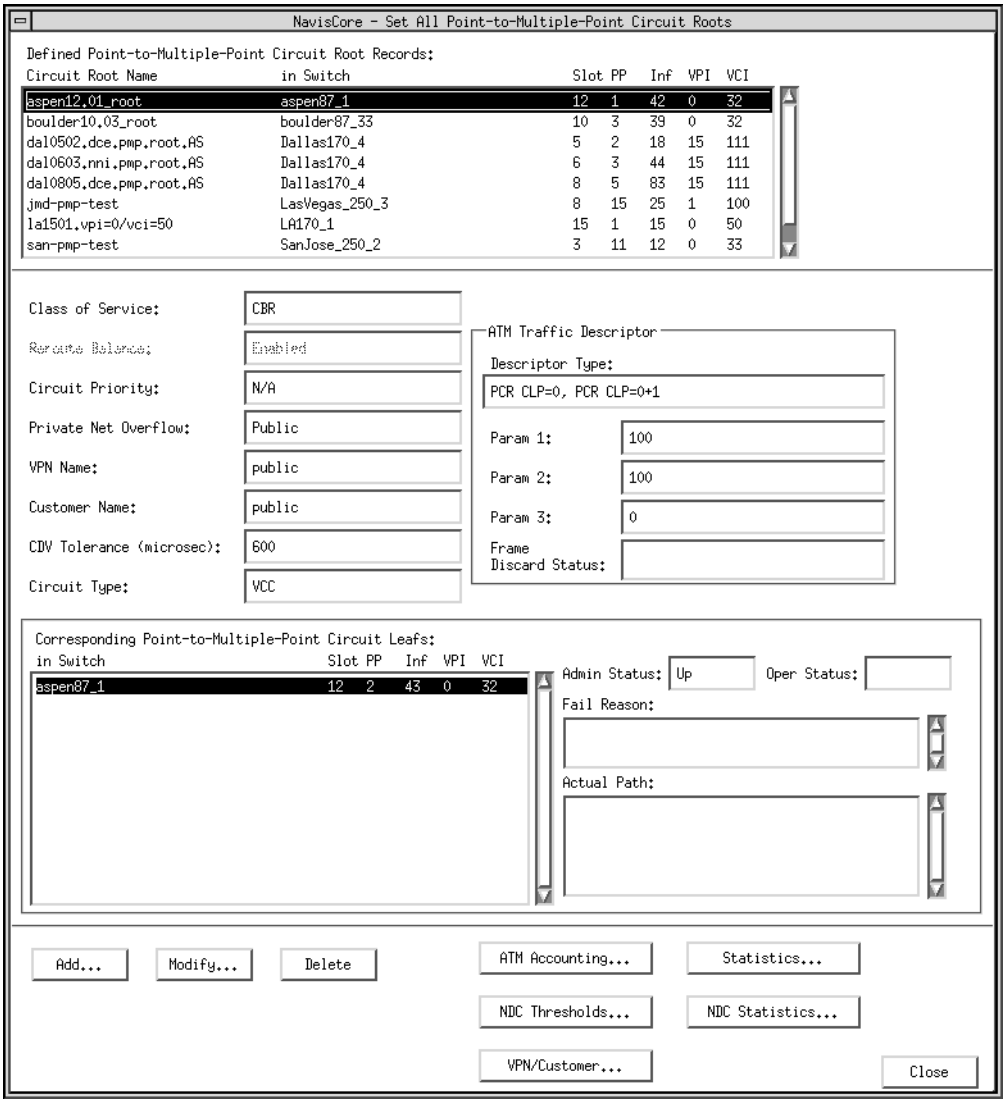

Figure 6-14. Set All Point-to-Multiple-Point Circuit Roots Dialog Box

The *Defined Point-to-Multiple-Point Circuit Root Records* box at the top of the screen lists any existing PMP circuit roots. The *Corresponding Point-to-Multiple-Point Circuit Leafs* box at the bottom of the screen lists any existing circuit leafs (endpoints) for the selected circuit root.

Table 6-13 displays the following information for each root and leaf:

Table 6-13. Set All Point-to-Multipoint Dialog Box Fields and Buttons

Field/Command	Action/Description
in Switch	The switch ID and switch name on which the root or leaf resides.
Slot	The physical slot for the IOM on which the root or leaf was created.
PP	The number of the physical port on which the root or leaf was created.
Inf	The MIB interface number for the logical port on which the root or leaf was created.
VPI	The virtual path ID of the logical port assigned to the root or leaf.
VCI	The virtual channel ID of the logical port assigned to the root or leaf.
Class of Service	Displays the QoS class (CBR, VBR-RT, VBR-NRT, or UBR) for the PMP circuit.
Reroute Balance	Shows whether reroute balancing is Enabled (default) or Disabled for this PMP circuit. See page 6-16 for information about reroute balancing.
Circuit Priority	Displays the priority of the circuit: 1 (High), 2 (Medium), 3 (Low), or 4 (Lowest). See page 8-3 for more information on circuit priority.
Circuit Type	Displays whether the circuit is a VCC (virtual channel connection) or VPC (virtual path connection).

Table 6-13. Set All Point-to-Multipoint Dialog Box Fields and Buttons (Continued)

Field/Command	Action/Description
ATM Traffic Descriptor	Displays the circuit's traffic descriptor(s) settings. The number of values displayed depends on the traffic descriptor combination that was selected for the circuit. For example, if you selected the combination PCR CLP=0+1, SCR CLP=0, and MBS CLP=0, three values are displayed: • The first value (Param 1) is the PCR for CLP=0+1 • The second value (Param 2) is the SCR for CLP=0 • The third value (Param 3) is the MBS for CLP=0 If, however, you selected the PCR CLP=0+1 combination, only one value is displayed: the PCR for CLP=0+1.
Add	Adds a new point-to-multipoint circuit.
Modify	Modifies the selected point-to-multipoint circuit root. The Modify command displays dialog boxes that are similar to those displayed for Add point-to-multipoint; however, you cannot modify the circuit name, logical port endpoints, circuit type, or VPI/VCI values from this dialog box.
Delete	Deletes the selected point-to-multipoint circuit root.
Accounting	Accesses the NavisXtend Accounting Server functions for a PVC.
Statistics	Displays the summary statistics for the selected point-to-multipoint circuit.
NDC Thresholds (*CBX 500 only*)	Displays the configured network data collection (NDC) thresholds for the selected circuit.
NDC Statistics (*CBX 500 only*)	Displays the NDC statistics for the selected circuit.
VPN/Customer	Assigns the selected point-to-multipoint circuit root to a specific VPN and customer name. You must do this before you create PMP circuit leafs.

Defining a Point-to-Multipoint Circuit Root

To configure the originating point (circuit root) for a point-to-multipoint circuit:

1. Choose Add. The following dialog box appears.

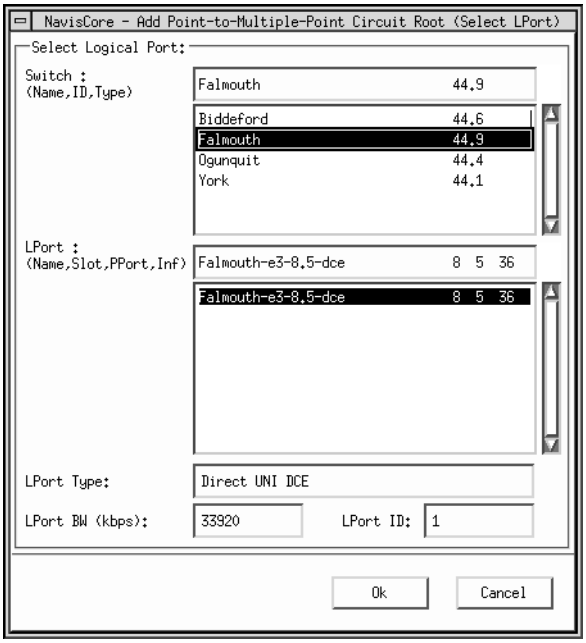

Figure 6-15. Add Point-to-Multiple-Point Circuit Root (Select LPort) Dialog Box

2. In the Switch list box, select the switch on which the originating point of the circuit will reside. The list contains the switch name and switch ID for all switches the NMS can currently access. The selected switch name appears in the text box above the list.
3. In the LPort list box, select the logical port on which the originating point of the circuit will reside. The selected logical port appears in the text box above the list. The list box displays the logical port name, slot ID, physical port number, and MIB interface number (Inf) for all logical ports on the selected switch.

   The LPort list box displays following information for the selected port:

   LPort Type — Displays the type of logical port.

   LPort BW (Kbps) — Displays the total logical port bandwidth for the selected logical port, in kilobits per second (Kbps).

   LPort ID — Displays the logical port ID for the selected logical port.

4. Choose OK to display the Add Point-to-Multiple-Point Circuit Root dialog box.

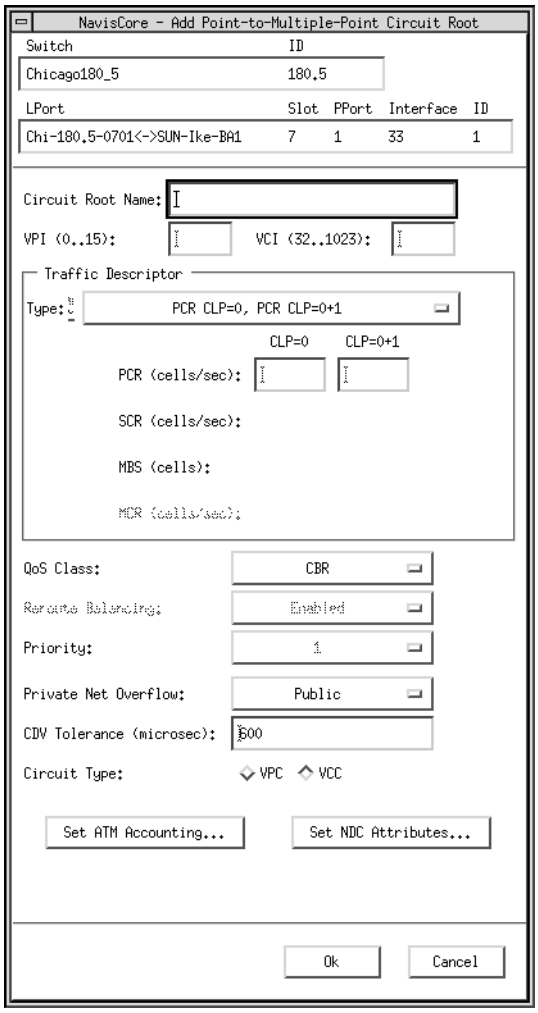

Figure 6-16. Add Point-to-Multiple-Point Circuit Root Dialog Box

5. Configure the fields as described in Table 6-14.

Table 6-14. Add Point-to-Multiple-Point Circuit Root Fields

Field	Action/Description
Circuit Root Name	Enter an alphanumeric name for the circuit root.
VPI (0..15)	Enter a value from 0-*nnnn* to represent the VPI for the PVC. The maximum value you can enter is based on the valid bits in VPI that are configured for the logical port. Note that zero is not a valid value for a management PVC. See page 6-3 for information about setting this value.
VCI (32..1023) (*for VCCs only*)	Depending on the circuit configuration, enter a value to represent the VCI for an ATM PVC. See page 6-3 for information about setting this value.

Table 6-14. Add Point-to-Multiple-Point Circuit Root Fields (Continued)

Field	Action/Description
Traffic Descriptor Type	Select from the following traffic descriptor options: *PCR CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes PCR CLP=0. If so, specify the PCR in cells per second for high-priority traffic (i.e., the CLP=0 cell stream). *PCR CLP=0+1 (cells/sec)* – Specify the PCR in cells per second for the combined high- and low-priority traffic (i.e., the CLP=0+1 aggregate cell stream). *SCR CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes SCR CLP=0. If so, specify the SCR in cells per second for the combined high-priority traffic (i.e., the CLP=0 cell stream). *SCR CLP=0+1 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes SCR CLP=0+1. If so, specify the SCR in cells per second for the combined high- and low-priority traffic (i.e., the CLP=0+1 aggregate cell stream). *MBS CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes MBS CLP=0. If so, specify the MBS (in cells per second) for the combined high-priority traffic (i.e., the CLP=0 cell stream). *MBS CLP=0+1 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes MBS CLP=0+1. If so, specify the MBS (in cells per second) for the combined high- and low-priority traffic (i.e., the CLP=0+1 cell stream). *MCR CLP=0 (cells/sec)* – Displays only if you selected a traffic descriptor combination that includes MCR CLP=0. If so, specify the MCR (in cells per second) for the combined high-priority traffic (i.e., the CLP=0 cell stream). ***Note:*** *While the MCR traffic descriptor is only applicable to a CBX 500 with an FCP, this attribute is offered as a selection on non-CBX endpoints. This is because even though one or both endpoints may not be on a CBX with FCP, the PVC might traverse a CBX 500 FCP trunk. In this case, the provisioned attribute is used.*
FCP Discard (*CBX 500 with FCP only*)	Displays only if you selected a QoS class that supports FCP Discard. Select either the CLP1 or EPD option. (See "ATM Flow-Control Discard Mechanisms" on page D-16 for more information.)
Reroute Balancing	When enabled, circuits use the tuning parameters you defined for the switch. When disabled, switch tuning parameters are ignored for the circuit. For more information, see the *NavisCore NMS Getting Started Guide*.
Priority (*VBR-NRT and VBR-RT QoS classes only*)	Select circuit priority, where 1 is high priority, 2 is medium priority, 3 is low priority, and 4 is lowest priority.
CDV Tolerance (microsec)	Enter a value between 1 - 65535 µsec. to define the Cell Delay Variation Tolerance (CDVT). The UPC uses this value to police the requested traffic descriptor. The default is 600 µsec.

Table 6-14. Add Point-to-Multiple-Point Circuit Root Fields (Continued)

Field	Action/Description
Private Net Overflow	Determines whether this PVC is restricted to trunks of its own VPN or can use public (shared) trunks during overflow conditions. To configure this circuit for a specific VPN and customer, see page 9-8. For more information about VPNs, see page 9-2. 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). *Restrict* – 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.
Circuit Type	Specify whether the circuit is a virtual path connection (VPC) or virtual channel connection (VCC, the default). If you select VPC, the VCI field is set to 0 and cannot be changed.

6. (*Optional - CBX 500*) Choose the Set [NDC] Attributes option and continue the instructions on page 6-45. When you finish, continue with Step 7.
7. Choose OK to return to the Set All PMP Circuit Roots dialog box (Figure 6-14 on page 6-38).

For information about the Set Accounting command, see the *NavisXtend Accounting System Administrator's Guide*.

8. To assign this PMP Circuit to a VPN or customer, choose the VPN/Customer command and see page 9-8.
9. Continue with "Configuring Point-to-Multipoint Circuit Leafs" on page 6-45.

Configuring Point-to-Multipoint Circuit Leafs

To add endpoints to the root circuit:

1. To configure the multiple endpoints of the PMP circuit, from the Set All Point-to-Multiple-Point Circuit Roots dialog box (Figure 6-14 on page 6-38) choose Modify. The following dialog box appears.

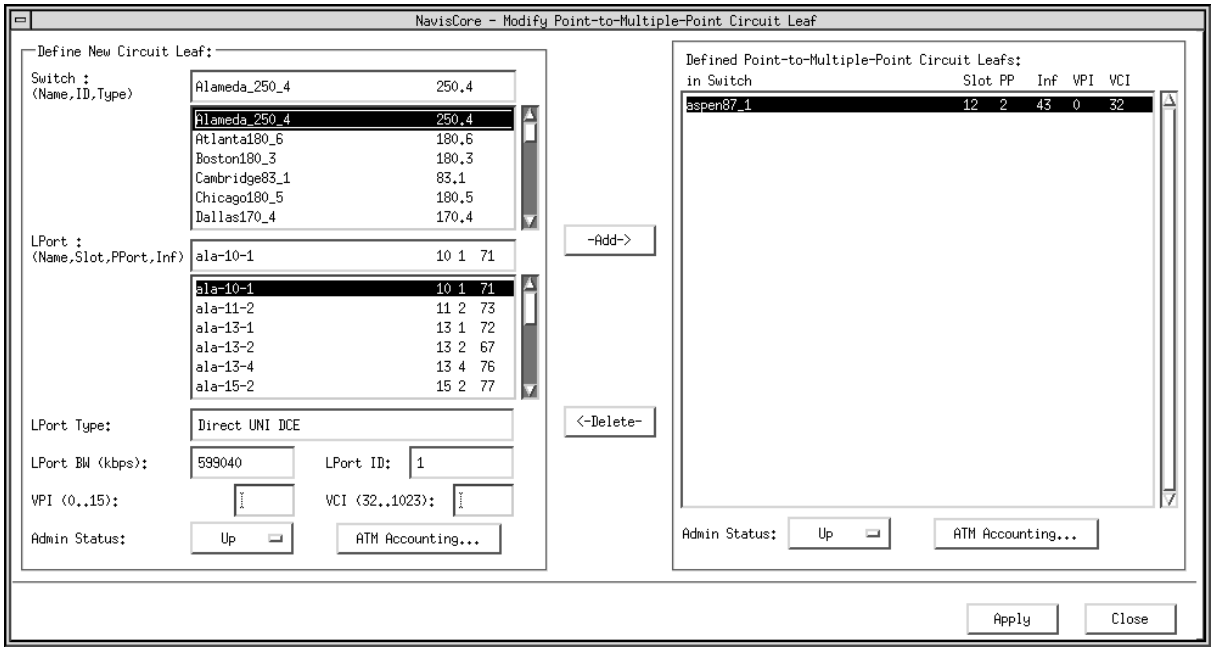

Figure 6-17. Modify PMP Circuit Leaf Dialog Box

The left side of this box enables you to define the endpoints of the PMP circuit. The list on the right shows the endpoints that have already been defined for the selected originating point.

2. To add a new endpoint:
   - **a.** From the Switch list box, select the switch on which to configure the new endpoint. The LPort list box changes to show the logical ports that are configured on the selected switch.
   - **b.** In the LPort list box, select the logical port for the new endpoint.
   - **c.** In the VPI and VCI fields, enter the virtual path ID and virtual channel ID for the PMP circuit as appropriate (i.e., VPCs do not require a VCI).
   - **d.** Set the Admin Status to Up if you want to activate this circuit when the switch comes online. Set the Admin Status to Down if you do not want to activate this circuit when the switch comes online.
   - **e.** Choose -Add-> to add the circuit to the PMP Circuit Leaf list.

f. Repeat Step 2 for each endpoint you want to create for this PMP circuit. When done, go to Step 3.

Do not configure more than one circuit leaf for a given root on the same physical port. If you configure more than one OPTimum trunk on a physical port, only one OPTimum trunk can be used for routing one of the leafs for a given root. For information about the Accounting command, see the *NavisXtend Accounting System Administrator's Guide*.

Figure 6-18 illustrates invalid and valid configuration examples that show how multiplexing cannot occur at the port level.

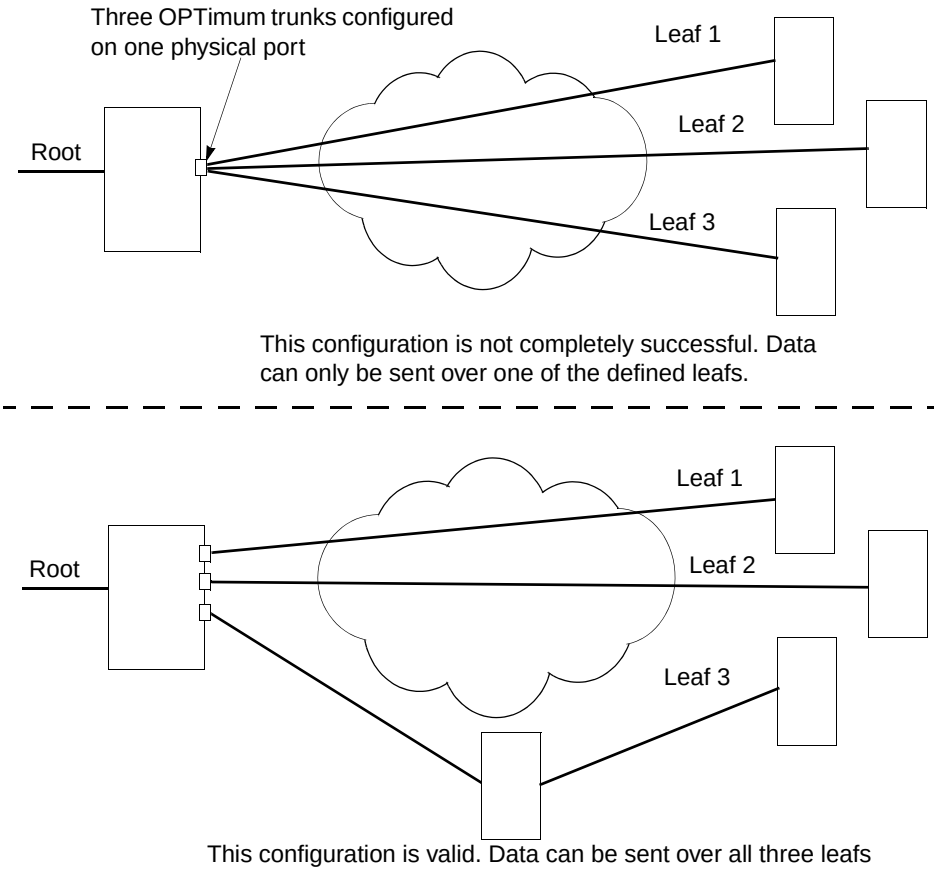

Figure 6-18. Point-to-Multipoint Circuit Example

3. Choose Apply, then choose Close to return to the Set All Point-to-Multiple Point Circuits dialog box.

To define additional PMP circuits and endpoints, repeat Step 1 through Step 3. When you are done adding PMP circuits and endpoints, choose Close to return to the network map.

Deleting a PMP Circuit Root and Leafs

Before you delete the root of a circuit, you must delete all of the circuit's leafs.

To delete a PMP circuit root and/or one or more of the circuit's leafs, as well as the entire circuit:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Point-to-Multipoint Circuits. The Set All Point-to-Multiple Point Circuit Roots dialog box appears (Figure 6-14 on page 6-38).
2. From the Circuit Root Name list, select the PMP circuit root you want to delete.
3. Choose Modify. The Modify PMP Circuit Leaf screen appears.
4. From the Defined PMP Circuit Leafs list, select the leaf you want to delete. Then choose Delete. A confirmation box appears. Choose OK to continue.
5. Repeat Step 4 for each circuit leaf you want to remove from the PMP circuit. If you are deleting the circuit, delete all leafs.
6. When done, choose Apply and then Close.
7. In the Circuit Root Name list, verify that the circuit you want to delete is selected. Also, in the Corresponding PMP Circuit Leafs list, verify that no leafs are listed (that is, they have all been deleted).
8. Choose Delete. A confirmation box appears. Choose OK to delete the circuit root. The circuit is now deleted from the network.
9. Choose Close to return to the network map.

Using Templates

If you defined a circuit configuration and saved it as a template (see *Template* field on page 6-12), 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 (see Figure 6-1 on page 6-4).
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 to delete.
4. Choose Delete.

7

Configuring Management Paths

This chapter explains 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 this same procedure to establish communications between the switch and any IP host (i.e., 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 an ATM 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 logical port (to which the NMS or IP host connects), and the remote switch SP module, the switch that connects the NMS or IP host is not burdened by the traffic traversing the management PVC.

Management VPI/VCI — This is the preferred method if you only use the attached NMS or IP host to transfer information between the host and the local switch. Even though you can use a management VPI/VCI connection to transfer information between the host and remote switch(es), using this method to transfer large amounts of information can have a negative impact on the local switch.

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, CBX, and GX switch platforms.

MPVCs originate at the switch I/O interface: IOP (B-STDX), IOM (CBX), and BIO (GX 550). They terminate at an internal logical port located on the switch processor module (either CP, SP, or NP, respectively). 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[***switch name***]. It uses an interface number (ifnum) of 4093. To form the circuit, connect the MgmtLPort. SW[***switch name***] endpoint to any UNI logical port type. You can configure MPVCs across different switch platforms; for example, B-STDX Frame Relay 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's processor module can send (and receive) IP packets over the MPVC to (and from) the ASE. 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 a UNI logical port.

Defining Physical Port Attributes

1. Select the switch for which you want to configure the ATM 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 dialog box appears.
4. Select the physical port you want to configure and choose Attrs. The Set Physical Port Attributes dialog box appears (see Figure 5-9 on page 5-21 for a sample dialog box).
5. Complete the dialog box fields. See the *NavisCore Physical Interface Configuration Guide* if you need information about changing default values.
6. Choose Apply.

Defining an ATM 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 3-2).
2. Choose Add to display the Add Logical Port dialog box (Figure 3-2 on page 3-6).
3. Select either ATM UNI DCE or DTE as the logical port type.
4. Choose OK. The Add Logical Port dialog box reappears (Figure 3-3 on page 3-7).
5. Use the instructions in Table 3-2 on page 3-9 to set the Administrative Attributes.
6. Use the instructions in Table 3-3 on page 3-11 to set the ATM Attributes.
7. Use the instructions in Table 3-4 on page 3-16 to set the ILMI/Signaling/OAM Attributes.
8. Use the instructions in Table 3-7 on page 3-23 to set the SVC VPI/VCI Range Attributes.
9. Choose OK to return to the Set All Logical Ports in PPort dialog box.
10. Choose Close to return to the Set Physical Port Attributes dialog box. Choose Cancel to return to the Switch Back Panel dialog box.
11. 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 appears (Figure 6-1 on page 6-4).
2. Choose Add. The Select End Logical Ports dialog box appears (Figure 6-2 on page 6-8).
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 appears (Figure 6-3 on page 6-10).
8. Enter the VPI/VCI or DLCI values as follows:
   - For an ATM UNI endpoint, enter a VPI and VCI value. Use the instructions on page 6-11 for CBX/GX ATM endpoints; see page 6-24 for B-STDX endpoints.
   - For a Frame Relay UNI endpoint, use the instructions on page 6-24 to enter a DLCI value.

9. Enter a Circuit Name for the management PVC. You will select this name when you configure the NMS path.
10. Set the remaining PVC attributes as follows:

For an ATM Service PVC see . . .	For an Interworking PVC see . . .
Table 6-2 on page 6-12 to set the Administrative Attributes	Table 6-9 on page 6-24 to set the Administrative Attributes
Table 6-3 on page 6-14 to set the Traffic Type Attributes	Table 6-10 on page 6-27 to set the Traffic Type Attributes
Table 6-4 on page 6-16 to set the User Preference Attributes	Table 6-11 on page 6-30 to set the User Preference Attributes
Table 6-5 on page 6-18 to set the Frame Discard Attributes	

11. *(Optional)* To configure CBX 500 Network Data Collection parameters for this circuit, select Set [NDC] Attributes. For more information, see the *NavisCore Diagnostic and Troubleshooting Guide*.
12. *(Optional)* To configure NavisXtend Accounting Server parameters for this circuit, choose the Accounting button. For more information, see the *NavisXtend Accounting System 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 following dialog box appears.

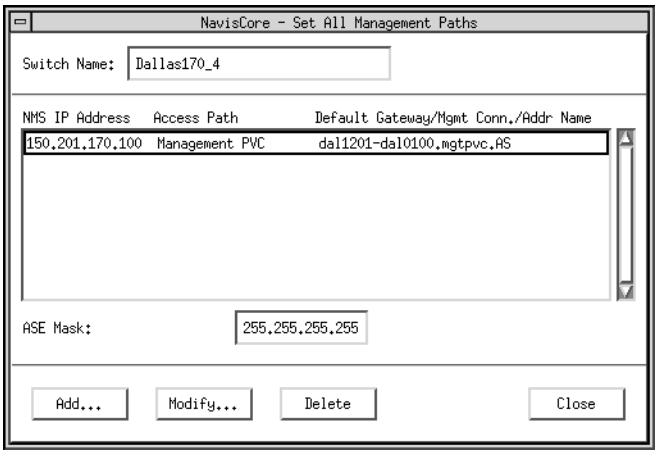

Figure 7-1. Set All Management Paths

3. Choose Add.

4. Select Management PVC as the Access Path. The following dialog box appears.

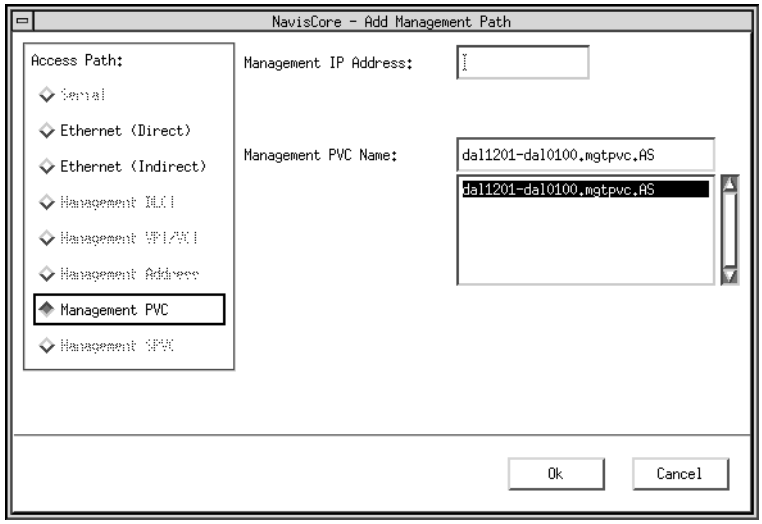

Figure 7-2. Add Management Path

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

6. Select the Management PVC Name you entered in Step 9 on page 7-4.

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

Using Management VPI/VCI

You use a Management VPI/VCI when the NMS connects to the gateway switch via an ATM router or ATM network interface card (NIC). The NMS accesses the gateway switch through this connection. This method of access enables you to monitor the network without the use of an Ethernet module in the switch.

Configuring a Management VPI/VCI

The following sections describe how to configure an NMS Path using a management VPI/VCI. To begin this process:

1. Configure an unused physical port using the instructions in "Defining Physical Port Attributes" on page 7-2.
2. Define a UNI logical port using the instructions in "Defining an ATM UNI Logical Port" on page 7-3.
3. Continue with the instructions in the following section, "Defining the Management VPI/VCI Connection."

Defining the Management VPI/VCI Connection

1. From the Administer menu, select Ascend Parameters ⇒ Set All Management VPI/VCIs. The following dialog box appears.

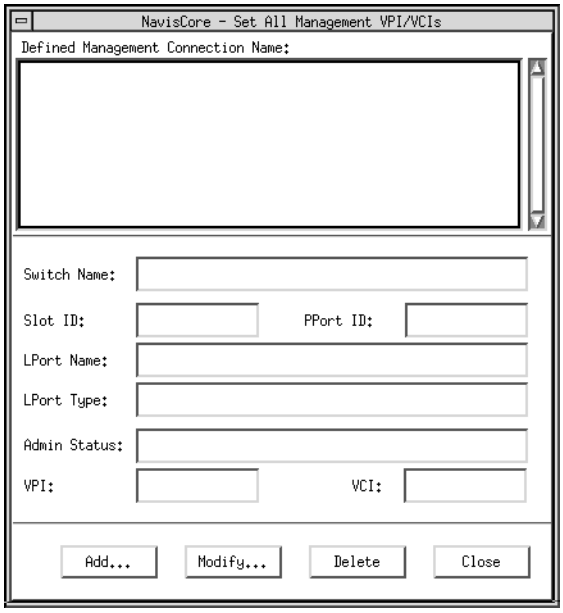

Figure 7-3. Set All Management VPI/VCIs Dialog Box

If you have already configured a Management VPI/VCI, the Set All Management VPI/VCIs dialog box displays the management connection names and configured parameters. From the Set All Management VPI/VCI dialog box, use the Modify or Delete commands to modify or delete Management VPI/VCI configurations.

2. Choose Add. The following dialog box appears.

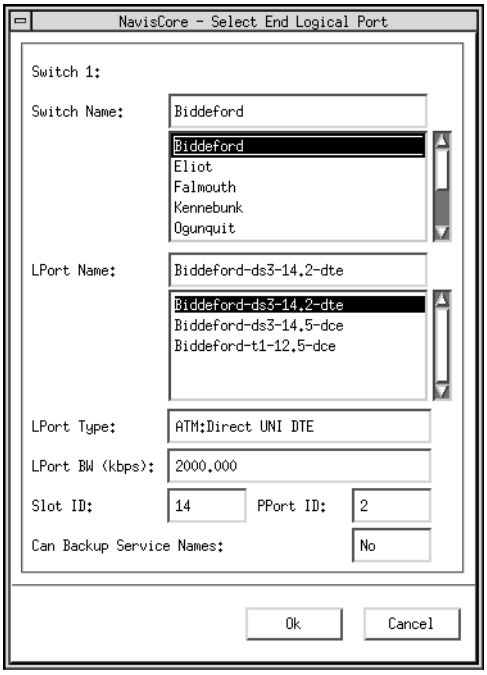

Figure 7-4. Select End Logical Port Dialog Box

3. Complete the dialog box fields as 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 or NIC that serves as the interface for the Network Management VPI/VCI.
LPort Name	Select the name of the logical port configured to access the router or NIC.
LPort Type	Displays the logical port type.
LPort BW (kbps)	Displays the logical port bandwidth.

Table 7-1. Select End Logical Port Fields (Continued)

Field	Action/Description
Slot ID	Displays the I/O slot number in which the I/O card resides.
PPort ID	Displays the port number for the port you are configuring.

4. Choose OK. The following dialog box appears.

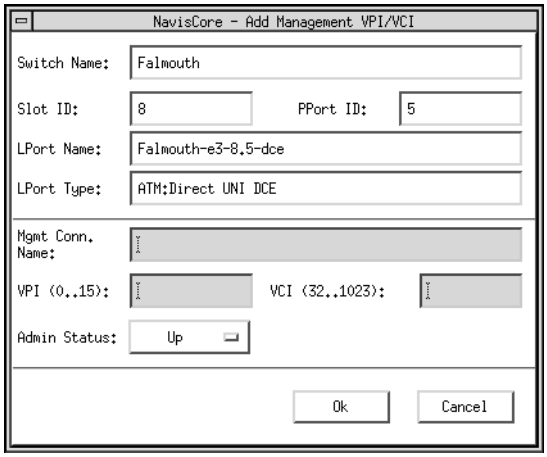

Figure 7-5. Add Management VPI/VCI Dialog Box

5. Complete the Add Management VPI/VCI dialog box fields as described in Table 7-2.

Table 7-2. Add Management VPI/VCI Fields

Field	Action/Description
Switch Name	Displays the name of the switch that connects to the router that serves as the interface for the Network Management VPI/VCI.
Slot ID	Displays the I/O slot (number) in which the I/O card resides.
PPort ID	Displays the port number for the physical port.
LPort Name	Displays the name of the logical port configured for the router.
LPort Type	Displays the logical port type.
Mgmt Conn. Name	Enter a unique, continuous, alphanumeric name to identify the connection. Do not use hyphens, dashes, parentheses, or asterisks.
VPI	Enter the VPI that is used for the connection.

Table 7-2. Add Management VPI/VCI Fields (Continued)

Field	Action/Description
VCI	Enter the VCI that is used for the connection.
Admin Status	Select either Up or Down to define whether the Management VPI/VCI connection 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 appears (Figure 7-1 on page 7-5).
3. Choose Add. The Add Management Path dialog box (Figure 7-2 on page 7-5) appears.
4. Select Management VPI/VCI 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 VPI/VCI Name (Management Conn. Name) you defined in Step 5 on page 7-8.
7. Choose OK.
8. Choose Close to return to the network map.

Defining the Static Route

To complete the Management VPI/VCI configuration, you must enter a static route in the router or NMS workstation to access the internal IP network.

8

Configuring ATM Traffic Descriptors

This chapter describes basic information you need for configuring traffic descriptors. Both the CBX 500 and the GX 550 can use traffic descriptors to define a service contract which guarantees that a specified amount of data is delivered. While the network can still deliver data that exceeds the limits of this traffic contract, this data may be delayed or lost if network resources are unavailable.

The B-STDX switch and the CBX 500 6-port Frame DS3/E3 and 4-port Ethernet modules do not support ATM traffic descriptors.

When you configure a PVC, you select the desired ATM traffic descriptor and enter the appropriate parameter value based on those items provided in the menu selection list. When you configure an SPVC, you first configure the specific traffic descriptor and then assign this traffic descriptor to the SPVC. Alternatively, for SPVCs, you may also choose one of the preconfigured traffic descriptors.

Overview

Configuring a logical port associates ATM traffic descriptors with the logical port control channels. Depending on the type of logical port, these control channels include ILMI, UNI signaling, PNNI routing, trunk protocol, and management traffic control channels. To simplify the provisioning process, you do not have to explicitly select the ATM traffic descriptor needed for the applicable control channel. A default value is always provided. Table 8-5 through Table 8-7 (beginning on page 8-11) describe these default values in more detail.

In most cases, you do not need to change the control channel default traffic descriptors. However, if you wish to have a particular control channel use a different QoS class or a different PCR/SCR/MBS, you have the ability to do so. For example, the default trunk signaling and management control channels that are used on trunks between Ascend switches are assigned to use the CBR QoS class and 5% of the configured logical port bandwidth (2.5% for each of the two channels). If necessary, you can change the QoS class of the trunk signaling channel; you can also change the amount of bandwidth associated with it.

If you plan to change the default values for logical port control channels, keep the following guidelines in mind:

- For control channels between two Ascend switches (which encompasses the trunk signaling control channel and the node-to-node management traffic control channel), the traffic descriptor (TD) values (PCR/SCR/MBS) are only used to calculate the amount of bandwidth reserved by the CAC for this type of traffic. The TD values do not affect traffic shaping on these channels nor do they affect the channel policing (these channels are never policed). You can change the default amount of bandwidth reserved for these control channels if you find the amount unacceptable.

- Starting with CBX 500 release 3.x switch software, the default amount of control channel bandwidth reserved on OC3/STM1 and OC12/STM4 trunks has changed. Previously, trunk control channel traffic was reserved at 5% of the logical port bandwidth (regardless of the media type). In some network scenarios for OC3/STM1 and OC12/STM4 trunks, this value was excessive; see Table 8-7 on page 8-13 for new default reserved bandwidth values.

- For control channels between an Ascend switch and another vendor device (including the ILMI, UNI signaling, and PNNI routing control channels), the TD values calculate both the amount of bandwidth reserved by CAC and the rate at which the control channels are policed.

  Control channels are not policed by default. You enable the UPC/NPC for the particular logical port, and the control channel will be policed at the traffic descriptor rate. Similar to the trunk control channels, the TD values associated with the ILMI, UNI signaling, and PNNI routing control channels do not affect the traffic shaping rate.

About Traffic Descriptors

To define a traffic descriptor, you must select the Quality of Service Class (QoS Class) and traffic descriptor combination to meet your network needs. The following sections describe each of the QoS classes, as well as the various traffic descriptor parameters. For each QoS class, you can select combinations of traffic parameters, which together form a traffic descriptor.

About Quality of Service

ATM supports four service classes to handle the various data types in a network. By selecting the appropriate service class, you can ensure optimal network usage.

Table 8-1 describes each service type class. The numerical value for the QoS Class reflects the ATM Forum definitions.

Table 8-1. Quality of Service Classes

Type	Description	QoS Class
Constant Bit Rate (CBR)	Handles digital information, such as video and digitized voice that is represented by a continuous bit stream. CBR traffic requires guaranteed throughput rates and service levels.	1
Variable Bit Rate (VBR) Real Time	For packaging special delay-sensitive applications, such as packet video, that require low cell delay variation between endpoints.	2
Variable Bit Rate (VBR) Non-Real Time	Handles packaging for transfer of long, bursty data streams over a pre-established ATM connection. This service is also used for short bursty data, such as LAN traffic. CPE protocols adjust for any delay or loss incurred through the use of a VBR non-real time service class.	3
Available Bit Rate/Unspecified Bit Rate (ABR/UBR)	Primarily used for LAN traffic. The CPE should compensate for any delay or lost cell traffic. You can only use this service class in conjunction with the ATM Flow Control Processor.	4

If the network equipment connected to the logical port does not support QoS, select the corresponding Unspecified class of service type. This provides a QoS class of 0.

About Logical Port Quality of Service Parameters

When you configure a logical port, you specify QoS parameters for each class of service (CBR through ABR/UBR). For more information about configuring these parameters, see "Setting Quality of Service Parameters" on page 3-27. The following list summarizes each of these parameters:

Bandwidth Allocation — Configures the amount of bandwidth to allocate on a logical port. You can configure a fixed percentage of bandwidth, or enable the bandwidth to change dynamically according to bandwidth demands.

Routing Metric — Optimizes network resources by routing traffic over the path that best matches the QoS needs of the associated VC. By selecting one of these metrics, you can ensure that a PVC, SVC, or SPVC originating from this logical port follows an efficient routing path to its destination.

Oversubscription Factor — Enables you to provision more PVCs, SVCs, or SPVCs on a given logical port than the amount of supported physical bandwidth. This ability to "oversubscribe" a logical port's bandwidth assumes that not all network resources are in use at the same time. For more information about the oversubscription factor, see page 2-13.

About Traffic Parameters

This section describes network traffic parameters and their associated ATM traffic descriptor combinations. When you create a logical port, PVC, or an SPVC, you can select a traffic descriptor that specifies how the network controls traffic going in the forward and reverse direction on that entity. This traffic descriptor is made up of individual traffic parameters which work together to provide traffic shaping.

Table 8-2 describes the individual traffic parameters.

Table 8-2. Traffic Parameters

Traffic Parameter	Description
CLP=0	Specifies the high-priority cell stream (cells whose Cell Loss Priority bit is set to 0).
CLP=1	Specifies the low-priority cell stream (cells whose Cell Loss Priority bit is set to 1).
CLP=0+1	Specifies the aggregate cell stream (all cells in this circuit whose Cell Loss Priority bit is either 0 or 1).
PCR (Peak Cell Rate)	Peak Cell Rate is the maximum allowed cell transmission rate (expressed in cells per second). It defines the shortest time period between cells and provides the highest guarantee that network performance objectives (based on cell loss ratio) will be met.

Table 8-2. Traffic Parameters (Continued)

Traffic Parameter	Description
SCR (Sustained Cell Rate)	Sustained Cell Rate is the maximum average cell transmission rate that is allowed over a given period of time on a given circuit. It allows the network to allocate sufficient resources (but fewer resources than would be allocated based on PCR) for guaranteeing that network performance objectives are met. This parameter applies only to VBR traffic; it does not apply to CBR or UBR/ABR traffic.
MBS (Maximum Burst Size)	Maximum Burst Size is the maximum number of cells that can be received at the Peak Cell Rate. This allows a burst of cells to arrive at a rate higher than the SCR. If the burst is larger than anticipated, the additional cells are either tagged or dropped. This parameter applies only to VBR traffic; it does not apply to the CBR or UBR traffic.
MCR (Minimum Cell Rate) (CBX 500 with FCP support only)	Minimum cell rate is the rate at which the source switch is always allowed to send data. This parameter only applies to ABR traffic. For more information about Flow Control Processor features, see Appendix D, "Implementing CBX 500 ATM Flow Control."
Tagging	Tagging refers to the method of changing a high-priority cell (CLP=0) to a low-priority cell (CLP=1). This method provides an alternative to simply dropping the cells from the cell stream, when the CLP=0 cell stream is non-conforming.
Best Effort	This option means that the network attempts to deliver traffic that exceeds the limits of the traffic contract. However, there are no guarantees that traffic will be delivered.

The traffic descriptor combination you select determines the number and type of cells that are admitted into a congested queue, and whether or not high-priority cells are tagged as low-priority cells when traffic exceeds the traffic parameter thresholds.

You can configure up to 512 traffic descriptors per switch. The following table lists the traffic descriptors that are available for each QoS class.

Table 8-3. QoS Class Traffic Descriptors

QoS Class	Traffic Descriptor	Description
Constant Bit Rate (CBR) (specified/ unspecified)	PCR CLP=0, PCR CLP=0+1, tagging	Traffic conformance is based on the Peak Cell Rate (PCR) of both the CLP=0 and CLP=0+1 cell streams with Tagging enabled.
	PCR CLP=0, PCR CLP=0+1, no tagging	Traffic conformance is based on the PCR of both the CLP=0 and CLP=0+1 cell streams with no Tagging.
	PCR CLP=0+1	Traffic conformance is based only on the PCR of the CLP=0+1 aggregate cell stream with no Best Effort.

Table 8-3. QoS Class Traffic Descriptors (Continued)

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

When you choose the Forward (or Reverse) Traffic Descriptor combination, select the combination that best describes the traffic characteristics. The Usage Parameter Control (UPC) function uses the traffic parameters to determine the conforming cells of an ATM connection, based on the threshold values for PCR, SCR, and MBS as specified in the service contract. If a traffic descriptor combination is not valid for the service class specified in the Forward (or Reverse) QoS class field, you cannot select it.

For more information on how each traffic descriptor combination affects the cell streams under different traffic conditions, see Appendix B, “ATM Traffic Descriptors.”

Configuring ATM Traffic Descriptors

NavisCore provides the ability to preconfigure a set of network-wide traffic descriptors. When you need to specify traffic information for a logical port or SPVC, you can select a predefined traffic descriptor definition.

The *Configurable Control Channel* feature enables you to define traffic descriptors for control circuits. To do this, you configure traffic descriptor information for the logical port's ILMI, UNI, and PNNI signaling or trunk control channels.

To configure ATM traffic descriptors, do the following:

- To define network-wide traffic descriptors, continue with the following section.
- To specify traffic descriptors for an existing logical port or SPVC, see "Defining Traffic Descriptor Attributes" on page 8-9.

Defining Network-wide Traffic Descriptors

To configure a set of traffic descriptors for use in your network:

1. From the Administer menu, choose Ascend Parameters ⇒ Set All ATM Traffic Descriptors. The following dialog box appears.

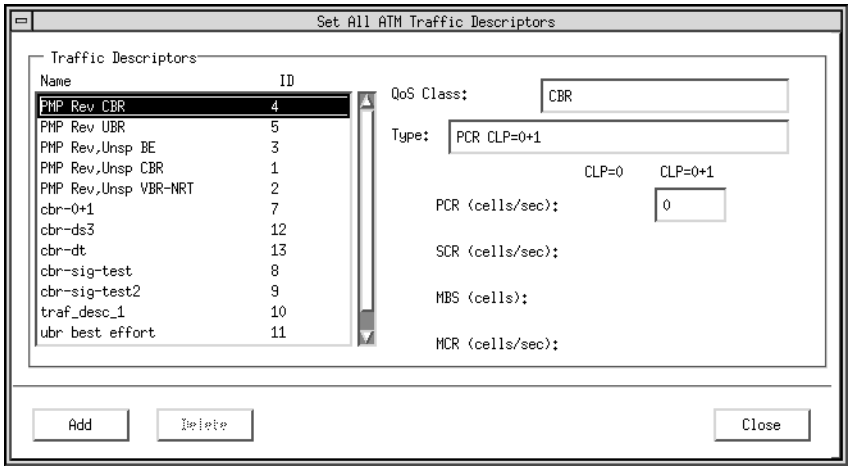

Figure 8-1. Set All ATM Traffic Descriptors Dialog Box

This dialog box displays information about previously configured traffic descriptors. If you need to delete a traffic descriptor, select the name from the list and choose Delete.

2. Choose Add to add a traffic descriptor. The following dialog box appears.

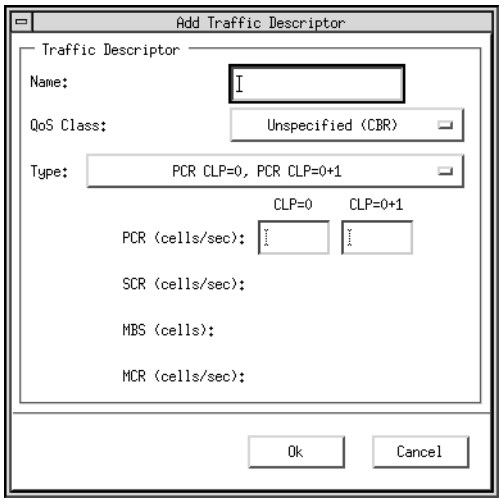

Figure 8-2. Add Traffic Descriptor Dialog Box

3. Enter a name (up to 20 characters) for this character descriptor type.
4. See Table 8-1 on page 8-3 to select the QoS class. Note that your choice of QoS class affects which traffic descriptors are available. If the attached equipment does not support QoS classes other than 0, select only the unspecified service classes.
5. See Table 8-3 on page 8-5 to select the traffic descriptor type.
6. Use the following table to specify the required values in cells per second.

Table 8-4. Traffic Descriptor Types

Traffic Descriptor Type	Description
PCR CLP=0 (cells/sec)	Displays only if you selected a traffic descriptor combination that includes PCR CLP=0. If so, specify the PCR in cells per second for high-priority traffic (i.e., the CLP=0 cell stream).
PCR CLP=0+1 (cells/sec)	Specify the PCR in cells per second for combined high- and low-priority traffic (i.e., the CLP=0+1 aggregate cell stream).
SCR CLP=0 (cells/sec)	Displays only if you selected a traffic descriptor combination that includes SCR CLP=0. If so, specify the SCR in cells per second for combined high-priority traffic (i.e., the CLP=0 cell stream).
SCR CLP=0+1 (cells/sec)	Displays only if you selected a traffic descriptor combination that includes SCR CLP=0+1. If so, specify the SCR in cells per second for combined high- and low-priority traffic (i.e., the CLP=0+1 aggregate cell stream).
MBS CLP=0 (cells/sec)	Displays only if you selected a traffic descriptor combination that includes MBS CLP=0. If so, specify the MBS (in cells per second) for combined high-priority traffic (i.e., the CLP=0 cell stream).

Table 8-4. Traffic Descriptor Types (Continued)

Traffic Descriptor Type	Description
MBS CLP=0+1 (cells/sec)	Displays only if you selected a traffic descriptor combination that includes MBS CLP=0+1. If so, specify the MBS (in cells per second) for combined high- and low-priority traffic (i.e., the CLP=0+1 cell stream).
MCR CLP=0 (cells/sec)	Displays only if you selected a traffic descriptor combination that includes MCR CLP=0. If so, specify the MCR (in cells per second) for combined high-priority traffic (i.e., the CLP=0 cell stream).

7. Choose OK to set the ATM traffic descriptor. The Set All ATM Traffic Descriptors dialog box appears (Figure 8-1 on page 8-7).
8. Choose Add to repeat these steps to create additional ATM traffic descriptors, or choose Close to return to the network map.

Defining Traffic Descriptor Attributes

To assign a traffic descriptor to a logical port or SPVC:

1. See one of the following sections to access the traffic descriptor attributes (see Figure 8-3):
   - For UNI logical ports, see "ILMI/Signaling/OAM Attributes" on page 3-15.
   - For Direct/OPTimum trunk logical ports, see "Traffic Descriptor Attributes" on page 3-24.
   - For SPVCs, see "Adding an SPVC" on page 14-7.
   - For PNNI RCC, see "Configuring an ATM NNI Logical Port" on page 17-11.

2. The dialog box displays following fields:

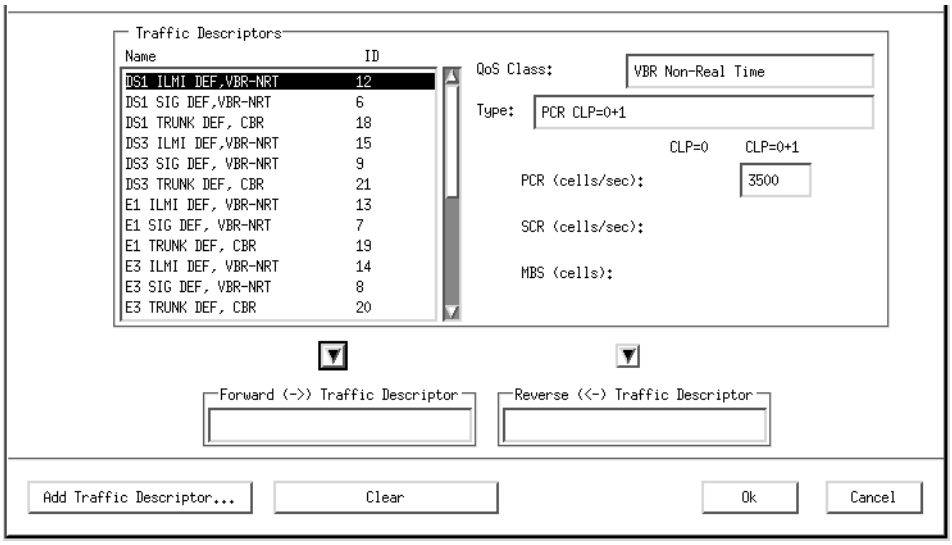

Figure 8-3. Traffic Descriptor Dialog Box Fields

3. Select a traffic descriptor name from the list and click on the Forward Traffic Descriptor arrow to assign this as the forward traffic descriptor.
4. Select a traffic descriptor name from the list and click on the Reverse Traffic Descriptor arrow to assign this as the reverse traffic descriptor.

To create a new traffic descriptor definition, choose the Add Traffic Descriptor command and use the instructions beginning with Step 3 on page 8-8.

5. Choose OK to complete this configuration.

Deleting Traffic Descriptor Definitions

To delete a traffic descriptor definition for either a logical port or SPVC:

1. Access the traffic descriptor attributes. A dialog box similar to Figure 8-3 appears.
2. For each direction, select the traffic descriptor you want to remove.
3. Choose Clear.
4. Choose OK.

Control Channel Traffic Descriptor Defaults

Each type of control channel is initially configured with a set of default traffic descriptors. These defaults specify the QoS class information and cell rate values for each traffic descriptor type. See the following tables to review these defaults:

- Table 8-5, "UNI Signaling Control Channel Traffic Descriptor Defaults"
- Table 8-6, "ILMI Control Channel Traffic Descriptor Defaults," on page 8-12
- Table 8-7, "Trunk Control Channel Traffic Descriptor Defaults," on page 8-13
- Table 8-8, "PNNI Routing Control Channel Traffic Descriptors," on page 8-14

Table 8-5. UNI Signaling Control Channel Traffic Descriptor Defaults

	DS1	E1	E3 (with PLCP)	E3 (with HEC)	DS3 (with PLCP)	DS3 (with HEC)	OC3/ STM1	OC12/ STM4
Type	NoClpScr[1]	NoClpScr	NoClpScr	NoClpScr	NoClpScr	NoClpScr	NoClpScr	NoClpScr
Class	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT
PCR [2] (cell/sec)	3500	4700	72000	80000	96000	106000	365000	1466000
SCR (cell/sec)	42	42	500	500	500	500	2000	8000
MBS (cells)	16	16	16	16	16	16	16	16
Approximate EBW[3] (cell/sec)	52	52	617	617	617	617	2468	9873

[1] The default Type, NoClpScr, represents the following: PCR CLP=0+1, SCR CLP=0+1, MBS CLP=0+1.

[2] If the configured logical port bandwidth is less than the physical port bandwidth, then PCR is 100% of logical port bandwidth.

[3] The approximate EBW (equivalent bandwidth) values are based on the use of the default values with the Ascend CAC in absence of user circuits on the logical port. It is provided here only as an aid to determine how much total bandwidth is reserved for the control channel.The amount reserved changes if you modify the traffic descriptor class or value.

Table 8-6. ILMI Control Channel Traffic Descriptor Defaults

	DS1	E1	E3 (with PLCP)	E3 (with HEC)	DS3 (with PLCP)	DS3 (with HEC)	OC3/ STM1	OC12/ STM4
Type	NoClpScr[1]	NoClpScr	NoClpScr	NoClpScr	NoClpScr	NoClpScr	NoClpScr	NoClpScr
Class	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT	VBR-NRT
PCR [2] (Cells/Sec)	3500	4700	72000	80000	96000	106000	365000	1466000
SCR (Cell/Sec)	21	21	21	250	250	250	1000	4000
MBS (cells)	16	16	16	16	16	16	16	16
Approximate EBW[3] (cell/sec)	26	26	26	309	309	309	1236	4944

[1] The default Type, NoClpScr, represents the following: PCR CLP-0+1, SCR CLP=0+1, MBS CLP=0+1.

[2] If the configured logical port bandwidth is less than the physical port bandwidth, then PCR is 100% of logical port bandwidth.

[3] The approximate EBW (equivalent bandwidth) values are based on the use of the default values with the Ascend CAC in absence of user circuits on the logical port. It is provided here only as an aid to determine how much total bandwidth is reserved for the control channel. The amount reserved changes if you modify the traffic descriptor class or value.

Table 8-7. Trunk Control Channel Traffic Descriptor Defaults

	DS1	E1	E3 (with PLCP)	DS3 (with PLCP)	DS3 (with HEC)	OC3/ STM1 [4]	OC12/ STM4 [4]
Type	NoClpNoScr [1]	NoClpNoScr	NoClpNoScr	NoClpNoScr	NoClpNoScr	NoClpNoScr	NoClpNoScr
Class	CBR	CBR	CBR	CBR	CBR	CBR	CBR
PCR [2,3] (cells/sec)	90	115	2000	2400	2600	CBX 6750 GX 9100	CBX 6750 GX 28000

[1] The default Type, NoClpNoScr, represents the following: PCR CLP=0+1.

[2] If the configured logical port bandwidth is less than the physical port bandwidth, then PCR is 2.5% of logical port bandwidth.

[3] Approximate EBW values are not provided in this case for CBR circuits, EBW=PCR.

[4] For OC3/STM1 and OC12/STM4, the default values are associated with the maximum control channel transmission rate the card type supports. For CBX IOMs, this is 13500cps; for GX BIOs this is 56000.

Both a trunk signaling and a node-to-node management control channel are used on a trunk. This means that when you examine the bandwidth reserved on a trunk that uses these default values, the values that are reserved are equal to the values in this table times two.

Table 8-8. PNNI Routing Control Channel Traffic Descriptors

Value	Description
Type	NoClpScr
Class	VBR-NRT
PCR	906 cps
SCR	453 cps
MBS	171 cells
EBW [1]	645

[1] The approximate EBW (equivalent bandwidth) values are based on the use of the default values with the Ascend CAC in absence of user circuits on the logical port. It is provided here only as an aid to determine how much total bandwidth is reserved for the control channel. The amount reserved changes if you modify the traffic descriptor class or value.

PNNI routing control traffic descriptors are the same across all port types.

9

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 out of a single public network. After creating a VPN name and ID, you then create and associate one or more customer names and IDs with the desired VPN. Once the VPNs and customers have been created in the database, you then assign any desired UNI/NNI logical ports to the particular VPN/customer pairing. In addition, it is also necessary to associate any of the desired public network trunks with a particular VPN.

Any PVCs you create on the UNI/NNI logical ports must be manually associated with the desired VPN/customer pairing. SVCs on the other hand, automatically inherit the VPN/customer pairing 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 9-1 provides a restrictive mode example.

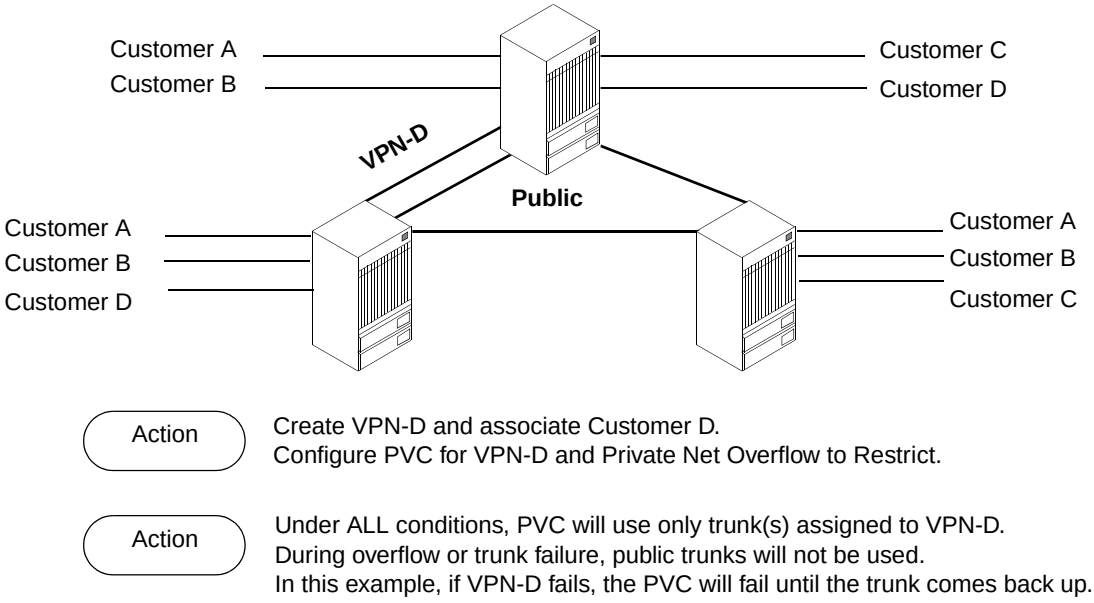

Figure 9-1. VPN Restrictive Mode Example

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 9-2). 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.

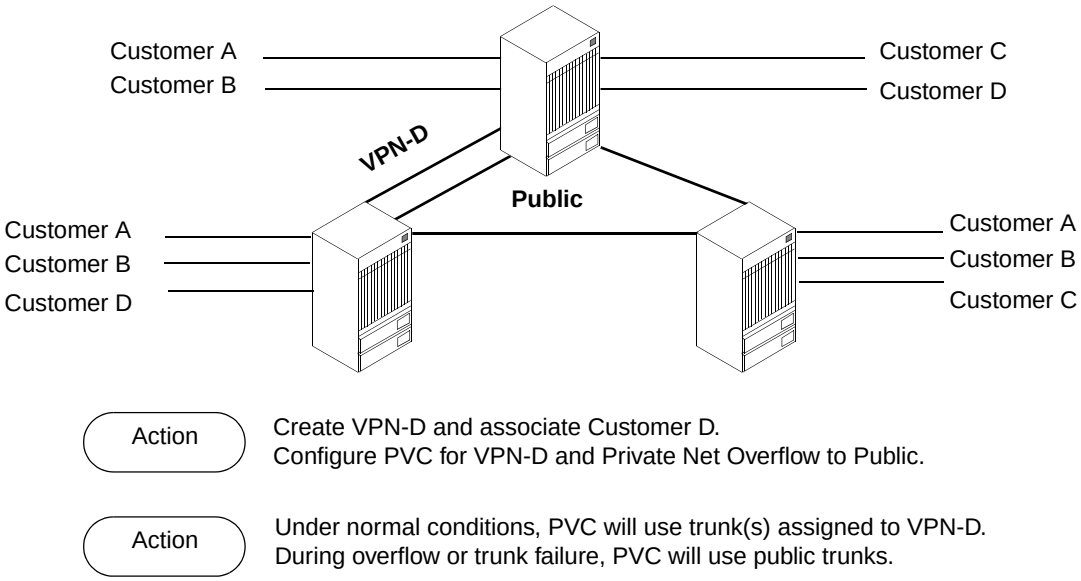

Figure 9-2. VPN Inclusive Mode Example

Configuring a Virtual Private Network

Use the following sequence to set up a VPN:

Step 1. Create the VPN (see page 9-3).

Step 2. Add customers to a specific VPN (see page 9-5).

Step 3. Dedicate a trunk to a specific VPN (see page 5-14).

Step 4. *For SVC traffic*, when you configure the UNI or NNI logical port, specify the net overflow attribute (see page 3-9). Then, dedicate this logical port to a specific VPN and customer (page 9-8).

Step 5. *For PVC traffic*, specify the net overflow attribute for the circuit (page 6-12). Then, dedicate the circuit to a specific VPN and customer (page 9-8).

Creating a VPN

To create a VPN and add customers to this network:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Virtual Private Networks. The following dialog box appears.

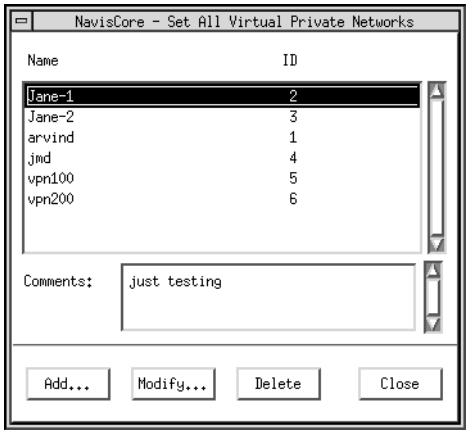

Figure 9-3. Set All Virtual Private Networks Dialog Box

2. Choose Add. The following dialog box appears.

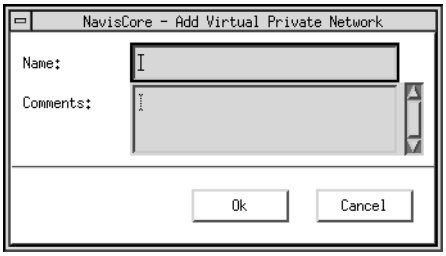

Figure 9-4. Add Virtual Private Network Dialog Box

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 following dialog box appears.

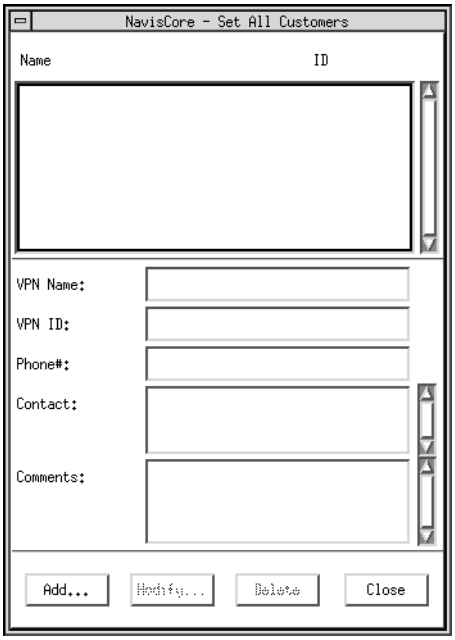

Figure 9-5. Set All Customers Dialog Box

2. Choose Add. The following dialog box appears.

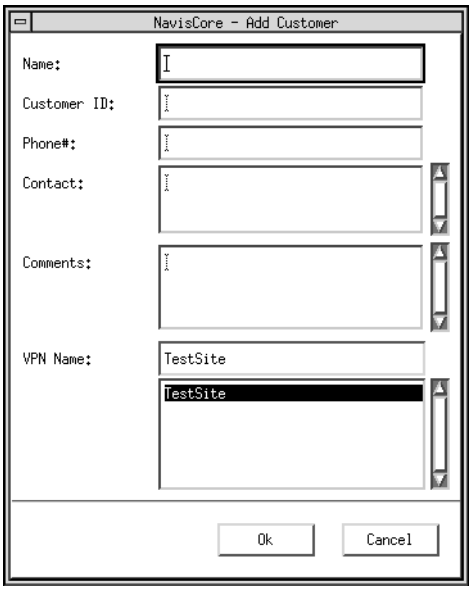

Figure 9-6. Add Customer Dialog Box

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 (see Table 3-2 on page 3-9). 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 3-2 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 button, select VPN/Customer Info and choose Set. The following dialog box appears.

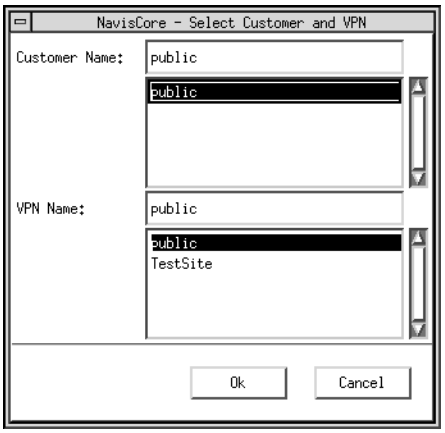

Figure 9-7. Select Customer and VPN Dialog Box

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 6-8) 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 9-6 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 following dialog box appears.

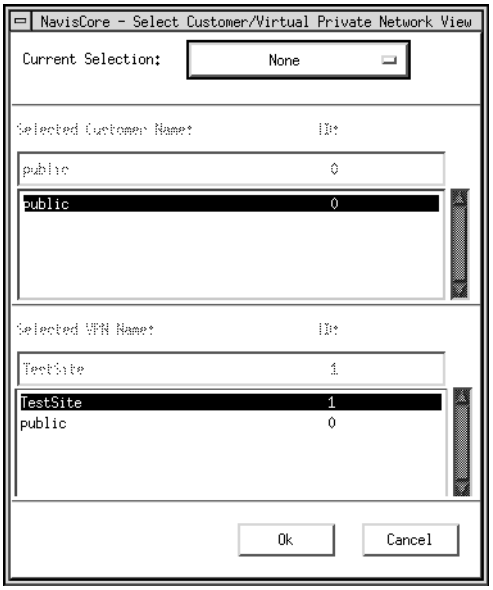

Figure 9-8. Select Customer/Virtual Private Network Dialog Box

2. Use the Current Selection button 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 (see Table 6-2 on page 6-12). This parameter determines whether the 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. See page 6-4 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 9-7 on page 9-6).
4. Select the customer and VPN name.
5. Choose OK.
6. Choose Close to exit.

10

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. The backup port is activated if a primary port fails or if you need to take a primary port off-line. This function is sometimes referred to as *resilient UNI*.

To automate resilient UNI/fault-tolerant PVC functions, you can configure the CBX 500 or GX 550 physical port on which this UNI logical port resides for automatic protection switching (APS). The APS with resilient UNI configuration protects against facility defects and equipment failure as well as I/O module failure. Keep in mind that this feature requires a circuit reroute. Although you can configure fault-tolerant PVCs for the B-STDX, the B-STDX switch platform does not support the APS with resilient UNI feature. See “Configuring APS with Resilient UNI” on page 10-6 for more information.

Configuring Fault-Tolerant PVCs

Use the following sequence to configure fault-tolerant PVCs:

Step 1. Follow the sequence on page 3-5 to define a UNI-DCE or UNI-DTE logical port as a backup port. Choose Yes for the option "Can Backup Service Names" (see page 3-9).

Step 2. Specify a service name for the primary port (see page 10-3).

Step 3. Configure circuits to use a service name as the endpoint (see page 6-8).

Step 4. Activate the backup port (see page 10-5).

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" (see Table 3-2 on page 3-9). 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 assigned service name as the endpoint instead of selecting 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

You define the *service name binding* 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 following dialog box appears, displaying any service names you have already configured.

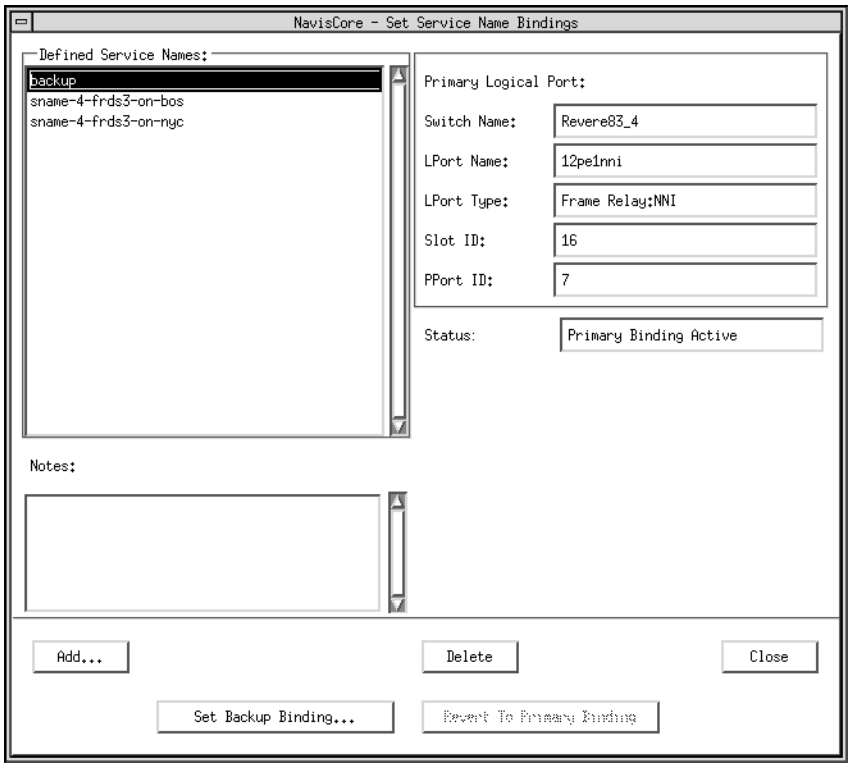

Figure 10-1. Set Service Name Bindings Dialog Box

You can access the following functions from the Set 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.

2. Choose Add. The following dialog box appears.

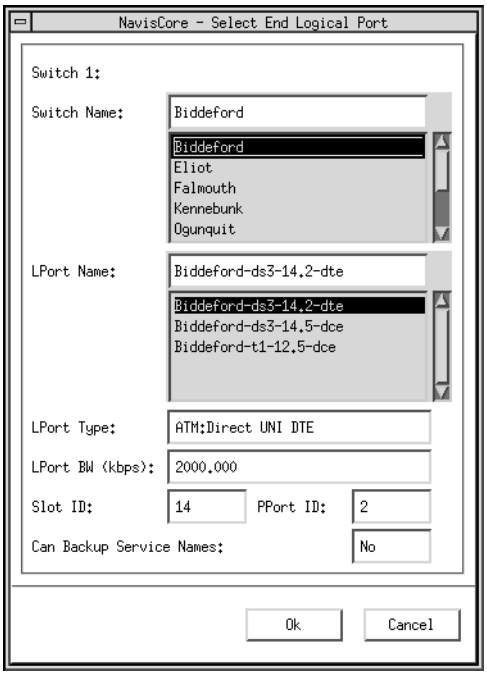

Figure 10-2. Select End Logical Port Dialog Box

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 following dialog box appears.

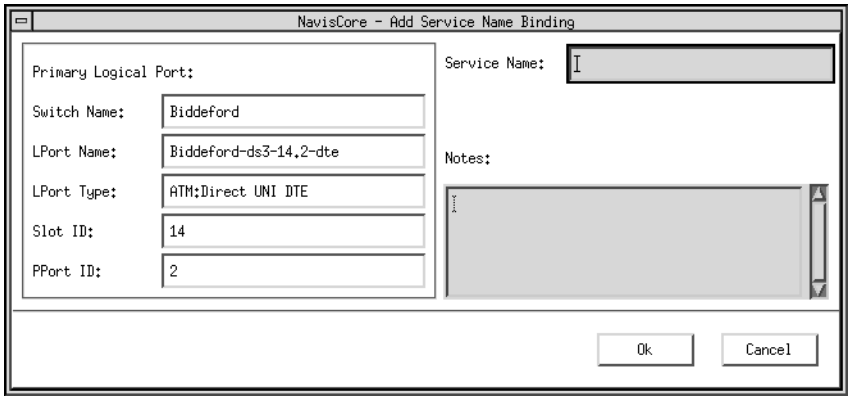

Figure 10-3. Add Service Name Binding Dialog Box

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 "Defining a PVC Connection" on page 6-8 to configure the circuits for fault-tolerant PVCs.

To reroute a PVC if it fails, see the next section, "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 ⇒ Service Name Binding. The Set Service Name Bindings dialog box (Figure 10-1 on page 10-3) appears.
2. Choose Set Backup Binding. The following dialog box appears.

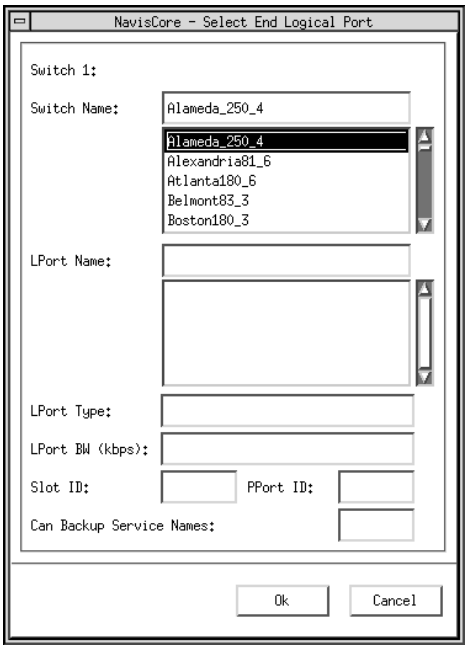

Figure 10-4. Select End Logical Port Dialog Box

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 following dialog box appears, displaying the Service Name that corresponds to the switch and logical port names you selected.

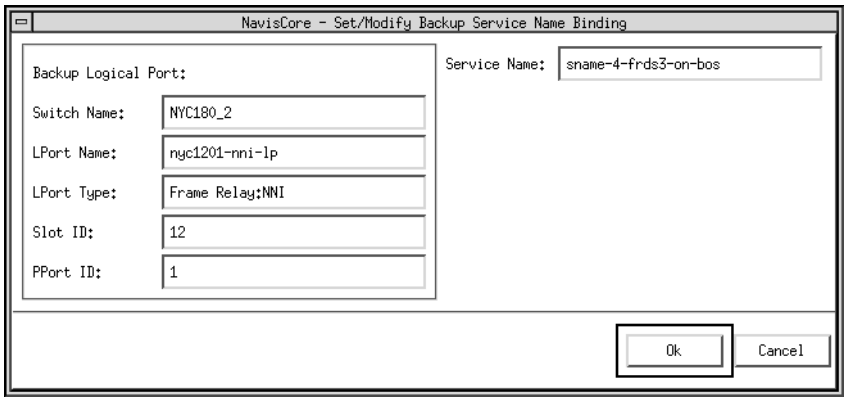

Figure 10-5. Set/Modify Backup Service Name Binding Dialog Box

6. Choose OK. The Set Service Name Bindings dialog box reappears (Figure 10-1 on page 10-3). The Status field should now display the message, Backup Binding Active.

Configuring APS with Resilient UNI

The CBX 500 and GX 550 optical cards provide an automatic protection system (APS) that enables you to designate a primary (*working*) port and a backup (*protection*) port. These cards include:

- OC3/STM-1
- OC12/STM-4
- OC48/STM-16 (*GX 550 only*)

You can use APS functions to automate the basic fault-tolerant PVC/resilient UNI feature. If an equipment failure occurs, the APS provides a backup physical port.

The fault-tolerant PVC/resilient UNI feature requires that you configure APS attributes for two working ports which reside on two different switches; these ports act as working ports. For each of these working ports, select a port that resides on a different card in the same switch to act as the protection port. If the working port fails, the fault-tolerant PVC/resilient UNI software automatically moves circuits to the protection port.

Note that on the CBX 500, this feature is only supported on new versions of the OC3/STM1 IOM. See either the NavisCore or CBX 500 switch software release notice for OC3/STM-1 revision level requirements. To determine whether or not your IOM supports this capability, double-click the switch object to display the Back Panel dialog box; then double-click the OC3/STM1 IOM to display the View Card Attributes dialog box. The Switch Software Capability and Hardware Capability fields should indicate “APS”.

Complete the following sections to configure APS with resilient UNI.

Defining Physical Port Attributes

1. Select the switch for which you want to configure the *working* physical port.
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 dialog box appears.
4. Select the physical port you want to configure and choose Attrs. The Set Physical Port Attributes dialog box appears (Figure 5-9 on page 5-21).
5. Complete the dialog box fields as described in Table 5-4 on page 5-21. Be sure to select *APS Resilient UNI* as the APS Redundancy type.

Defining APS Attributes

1. From the Set Physical Port Attributes dialog box (Figure 5-9 on page 5-21), choose *APS*. The Set APS Attributes dialog box appears (Figure 5-10 on page 5-24).
2. Complete the APS attributes as described in Table 5-5 on page 5-24.
3. Choose Apply to save these settings and Close return to the Set Physical Port Attributes dialog box (Figure 5-9 on page 5-21).
4. To exit, choose Apply and then OK to save the physical port attributes and create an SNMP SET command to send to the switch. Choose Cancel to exit.

   To define an APS configuration on a corresponding optical card located in a *different* switch:

   - Repeat Step 1 through Step 5 in the previous section, “Defining Physical Port Attributes.”
   - Then, repeat Step 1 through Step 4 in the previous section, “Defining APS Attributes.”

   Continue with the following section to define ATM UNI logical ports.

Defining ATM UNI Logical Ports for APS Resilient UNI

Use the following steps to create an ATM UNI logical port for each working and protection port.

1. Select the switch on which the first working/protection port pair resides. (You will define an ATM UNI logical port for each of these physical ports.)
2. From the Administer menu, select Ascend Parameters ⇒ Set Parameters. The Switch Back Panel dialog box appears.
3. Select the working port (of the APS pair) and press the third (right) mouse button to display a pop-up menu.
4. Select Logical Port. The Set All Logical Ports in PPort dialog box appears (Figure 3-1 on page 3-2).
5. Choose Add to display the Add Logical Port dialog box (Figure 3-2 on page 3-6).
6. Select *ATM UNI* as the logical port type.
7. Choose OK. The Add Logical Port dialog box reappears (Figure 3-3 on page 3-7).
8. Use the instructions in Table 3-2 on page 3-9 to set the Administrative Attributes.
9. Use the instructions in Table 3-3 on page 3-11 to set the ATM Attributes.
10. Use the instructions in Table 3-4 on page 3-16 to set the ILMI/Signaling/OAM Attributes.
11. Use the instructions in Table 3-7 on page 3-23 to set the SVC VPI/VCI Range Attributes.
12. Choose OK to return to the Set All Logical Ports in PPort dialog box.
13. Choose Close to return to the Set Physical Port Attributes dialog box. Choose Cancel to return to the Switch Back Panel dialog box.
14. To configure an ATM UNI logical port for the protection port, repeat Step 3 through Step 13 above. When you set the Administrative Attributes for this protections port, be sure to select Yes for the option, *Can Backup Service Names* (see page 3-9).
15. Choose Close to return to the network map.
16. Select the switch on which the second working/protection port pair resides and complete Step 2 through Step 14 above to define ATM UNI logical ports for the second APS pair.

Defining the APS Resilient UNI/Fault-Tolerant PVC

1. Complete the steps described in "Creating Service Names" on page 10-3 to assign a service name to both working/ATM UNI logical port endpoints.
2. Continue with the instructions in "Defining a PVC Connection" on page 6-8 to configure the fault-tolerant PVC.

11

About SVCs

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

The B-STDX does not support ATM SVCs. You also cannot configure ATM SVCs on the 6-port Frame-based DS3 module.

Address Formats

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

ATM End System Address (AESA) formats — AESA formats give service providers using a private ATM network the flexibility to develop an addressing scheme that best suits their network needs; for example, you may find that most CPEs in your network only support a specific AESA address format.

AESA Anycast Formats – AESA Anycast formats give service providers "group address" functionality for each of the AESA address formats. Using the Anycast format, a call is placed to the group address and the network selects one of the members to which the call will be routed. This group address could, for example, represent a group of Internet servers which contain the same information and perform identical functions. It does not matter which of these servers handles the call.

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

The following sections describe these address formats.

ATM End System Address (AESA) Formats

The GX 550 and CBX 500 support four AESA formats:

Data Country Code (DCC) — For DCC AESA addresses, the initial domain identifier (IDI) is a two-byte data country code field that identifies the country in which this address is registered. These country codes are standardized and defined in ISO reference 3166. *DCC Anycast AESA* provides a group address function for this address type.

International Country Designator (ICD) — For ICD AESA addresses, the IDI field contains the international country designator that uniquely identifies an international organization. The British Standards Organization administers these values. *ICD Anycast AESA* provides a group address function for this address type.

E.164 — For E.164 AESA addresses, the IDI field contains an eight-byte E.164 address. This E.164 address uses the international format and consists of up to fifteen decimal digits. *E.164 Anycast AESA* provides a group address function for this address type.

Custom — Custom AESA addresses enable you to use a customized octet structure and a customized authority and format identifier (AFI).

All AESA address formats consist of 20 octets. Each of these address formats contain the following components:

Initial Domain Part (IDP) — Defines the type of address and the regulatory authority responsible for allocating and assigning the Domain Specific Part. There are two subfields: the AFI and IDI fields.

Authority and Format Identifier (AFI) – The AFI part of the AESA address identifies the authority that allocates the DCC, ICD, or E.164 part of the AESA address, as well as the syntax of the rest of the address. The following are default AFIs:

Table 11-1. AFI Default Values

Address Type	AFI
DCC	0x39
DCC Anycast	0xBD
ICD	0x47
ICD Anycast	0xC5
E.164	0x45
E.164 Anycast	0xC3
Custom	A user-specific code for custom prefixes/addresses. (You must know the appropriate code to enter when defining custom prefixes/addresses.)

Initial Domain Identifier (IDI) – A hex code that identifies the sub-authority that has allocated the address. The format depends on the following address types:

Table 11-2. IDI Default Values

Address Type	IDI Description
DCC (including Anycast)	Consists of 2 octets (4 hex digits) that identify the country in which this address is registered. The DCC is generally considered a three digit quantity with a trailing hex “f” semi-octet. For example, the ANSI IDI of 840 is encoded as 0x840f.
ICD (Anycast)	Consists of 2 octets (4 hex digits) that identify an international organization to which this address is registered. The ICD is generally considered a four digit quantity. For example, the US GOSIP IDI of “5” is encoded as 0x0005.
E.164 (Anycast)	Consists of 8 octets in BCD format. (1-15 hex digits, plus a trailing Fh; if less than 15 digits are entered, type leading zeros to fill the 8 octets.) Represents an international E.164 address. For example, the E.164 address of 978-555-1212 is encoded as 0x000009785551212f.

Domain Specific Part — Consists of the HO-DSP, EDI, and SEL fields.

High-Order Domain-Specific Part (HO-DSP) – The authority specified in the AFI/IDI octets determines the format of this field. It identifies a segment of address space that is assigned to a particular user or subnetwork. It should be constructed to facilitate routing through interconnected ATM subnetworks. The general format for each address type is as follows:

Table 11-3. HO-DSP Default Values

Address Type	HO-DSP Description
DCC, ICD (including Anycast)	Consists of 10 octets (20 hex digits)
E.164 (Anycast)	Consists of 4 octets (8 hex digits)
Custom	Consists of 12 octets (24 hex digits)

End System Identifier (ESI) – A 6-octet (12 hex digit) field that uniquely identifies the end system within the specified subnetwork. This is typically an IEEE MAC address.

Selector (SEL) – A 1-octet (2 hex digit) field that is not used for ATM routing, but may be used by the end system.

Figure 11-1 shows how the octets are assigned for each AESA address format. Each octet is equivalent to two hex digits.

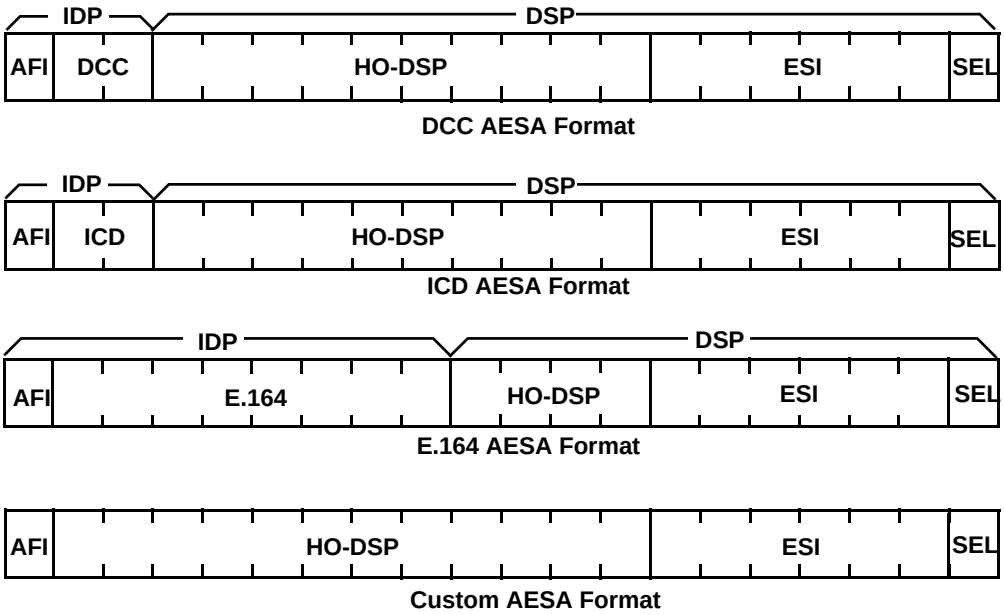

Figure 11-1. AESA Address Formats

Native E.164 Address Format

Native E.164 addresses are the standard Integrated Services Digital Network (ISDN) numbers, including telephone numbers. Native E.164 addresses consist of 1-15 ASCII digits. For example, standard 10-digit United States telephone numbers, such as 508-555-1234, are native E.164 addresses.

Unlike AESA address formats, native E.164 addresses are not broken down into an AFI, HO-DSP, ESI, and SEL portion. When a native E.164 address is translated to E.164 AESA format, the native E.164 address is stored in octets 2-9 of the 20-octet AESA address, while the HO-DSP, ESI, and SEL portions are filled with zeros. Conversely, when an E.164 AESA address is translated to native E.164 address format, the AFI, HO-DSP, ESI, and SEL portions, as well as any leading zeros in the 8-octet AESA E.164 address, are stripped off to produce the native E.164 address.

Designing an Address Format Plan

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

You use address formats to develop a network numbering plan. Using an AESA address, you can design the IDP portion of an address to target a specific network; then use the HO-DSP portion of the address to identify subnetworks within that network, and use the ESI portion to identify a specific end system.

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

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

About Address Registration

Address information in a switch is used to determine call routing; it is also used for calling party screening. When used for route determination, the switch advertises an appropriate subset of its configured node prefixes, port prefixes, and port addresses to all other switches in the network. When used for calling party screening, the switch uses the configured node prefixes, port prefixes, and/or port addresses to determine whether or not the network should accept an SVC request.

To perform these two functions at a UNI, both the user and the network need to know the ATM addresses that are valid at the UNI. Address registration provides a mechanism for address information to be dynamically exchanged between the user and the network, enabling both to determine the valid ATM addresses that are in effect at a UNI. Address registration applies only to UNI ports on which ILMI is enabled (see page 3-15 for instructions on how to enable ILMI on a UNI logical port). Any ILMI-eligible node or port prefix will be transferred from all ILMI-enabled private UNI-DCE ports and all ILMI-enabled public end-system UNI-DCE ports to their peer DTE devices.

Node prefixes are not exchanged from “network-to-network” UNI-DCE ports. Only port prefixes are exchanged from these ports. For address registration to work, attached UNI devices must support ILMI.

ILMI-eligible prefixes include:

- All native E.164 node prefixes
- All 13-octet (104-bit) AESA node prefixes
- All native E.164 port prefixes
- All 13-octet (104-bit) AESA port prefixes

The network side of the UNI provides the network prefix, which consists of the IDP and HO-DSP portions. The user side of the UNI provides the remaining portion of the address, which consists of the IEEE MAC address (the ESI portion) and the SEL portion of an ATM address; this forms the user part of the address. Figure 11-2 shows this addressing scheme.

Native E.164 prefixes sent by the network are concatenated with a NULL user part by the user, and returned to the network as native E.164 addresses. (The prefix and address are identical.)

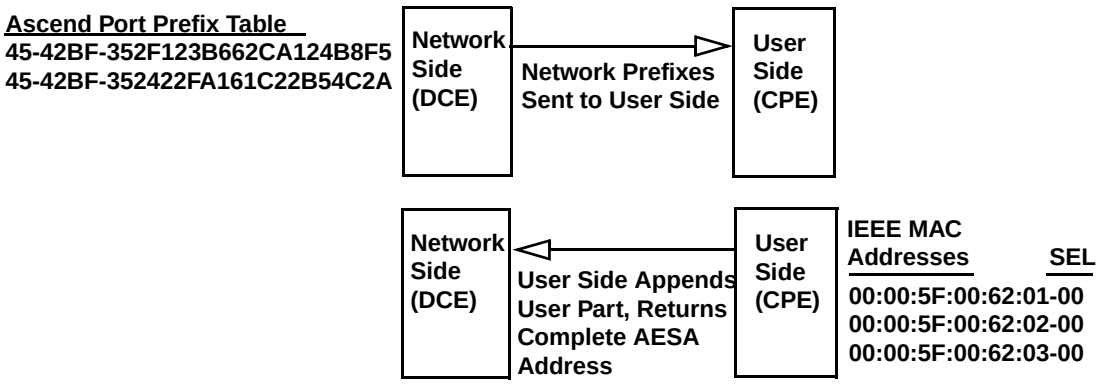

Resulting ILMI Address Table at DCE

45-42BF-352F123B662CA124B8F5-00:00:5F:00:62:01-00
45-42BF-352422FA161C22B54C2A-00:00:5F:00:62:01-00
45-42BF-352F123B662CA124B8F5-00:00:5F:00:62:02-00
45-42BF-352422FA161C22B54C2A-00:00:5F:00:62:02-00
45-42BF-352F123B662CA124B8F5-00:00:5F:00:62:03-00
45-42BF-352422FA161C22B54C2A-00:00:5F:00:62:03-00

Figure 11-2. Address Registration

About Route Determination

The node prefixes, port prefixes, and port addresses configured on network nodes are used to determine the route for a given SVC. A “best match” hierarchy determine the route, 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 the node. For example, if all addresses on the node contain the digits 15085551, you would define this as the node prefix. To allow for address routing, node prefixes should be unique to a switch; if not, the switch has to perform subsequent matching to find a route to the destination.

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

	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 example shows the node to which the SVC request is routed for certain called-party addresses, and describes why the request is routed to that node:

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 is an exact match for 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.

Called Party Address	Node	Reason
5085531111	1 or 2	Since the longest match occurs on both Nodes 1 and 2, the Admin Cost value assigned to port prefix 508553 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 508553.
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 (see "Defining Default Routes for Network-to-Network Connections" on page 12-22).

About Address Translation

This section describes how address translation occurs in various situations and at various points along a network connection. This information applies only if you enable address translation. Also, egress address translation requires matching a called party address to a configured prefix on the egress port.

Calling party and called party addresses are stored as information elements in the SETUP message, which is sent to initiate call setup. In some situations, calling party and called party sub-addresses are also stored as information elements in the SETUP message.

Calling Party Address	**Called Party Address**
Calling Party Subaddress	**Called Party Subaddress**

Egress address translation, when enabled on a network-to-network port, functions as described in Table 11-4 and Table 11-5. The following factors determine how address translation occurs:

- Whether or not local and/or remote gateway addresses are defined on the egress port
- The type of translation (tunnel or replace) selected as the egress address translation mode
- The numbering plan of the signaled calling and called addresses

Calling party and called party processing are independent. Note that in the SETUP message, the called party address is mandatory, while the calling party address is optional. In the case of a native E.164 called party or calling party address, the related sub-address field is always set to null, since the sub-address field cannot carry native E.164 addresses (note that in the tables, if the signaled calling party address is native E.164 format, the calling party sub-address field is always set to null).

Using ingress address translation, the calling party sub-address (if it is not null) overwrites the calling party address at the ingress port, and the called party sub-address (if it is not null) overwrites the called party address.

Table 11-4 shows how calling party addresses are translated at the egress port.

Table 11-4. Calling Party Address Translation at Egress Port

Signaled Address	SETUP Information Element	No Local Gateway Address	Local Gateway Address with Tunnel Option	Local Gateway Address with Replace Option
No calling party	Calling Party Address	Null	Local Gateway Address	Local Gateway Address
	Calling Party Sub-address	Null	Null	Null
AESA calling party	Calling Party Address	Signaled AESA Calling Party Address	Local Gateway Address	Local Gateway Address
	Calling Party Sub-address	Null	Signaled AESA Calling Party Address	Null
Native E.164 calling party	Calling Party Address	Signaled Native E.164 Calling Party Address	Local Gateway Address	Local Gateway Address
	Calling Party Sub-address	Null	Null	Null

Table 11-5 shows how called party addresses are translated at the egress port.

Table 11-5. Called Party Address Translation at Egress Port

Signaled Address	SETUP Information Element	No Remote Gateway Address	Remote Gateway Address with Tunnel Option	Remote Gateway Address with Replace Option
AESA called party	Called Party Address	Signaled AESA Called Party Address	Remote Gateway Address	Remote Gateway Address
	Called Party Sub-address	Null	Signaled AESA Called Party Address	Null
Native E.164 called party	Called Party Address	Signaled Native E.164 Called Party Address	Remote Gateway Address	Remote Gateway Address
	Called Party Sub-address	Null	Null	Null

Examples

The following example diagrams show the state of the SETUP message calling party/called party address and sub-address elements at various points along the connection.

The example diagrams represent the calling party and called party address and sub-address elements as follows:

Calling Party Address	Called Party Address
Calling Party Subaddress	Called Party Subaddress

Example 1
- Egress tunneling enabled on Network 1's egress port
- Ingress tunneling enabled on Network 2's ingress port
- Local Gateway address X configured to a prefix on Network 1's egress port, and the prefix corresponds to B
- Remote Gateway address Y configured to a prefix on Network 1's egress port, and the prefix corresponds to B

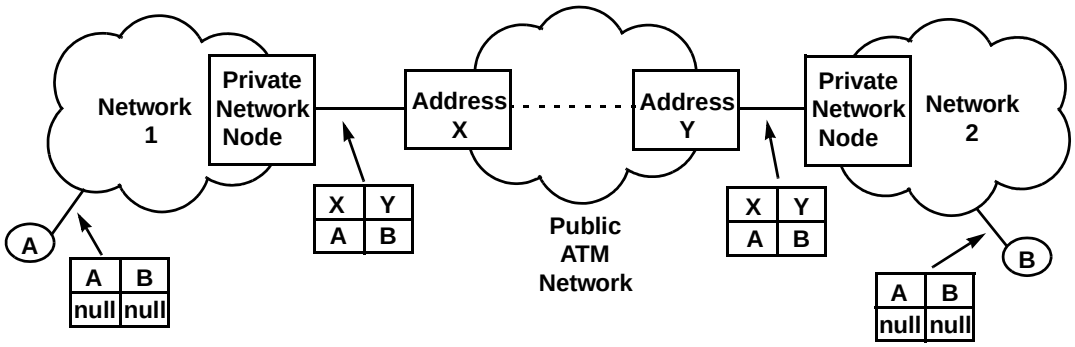

Example 2
- Egress tunneling enabled on Network 1's egress port
- Ingress tunneling enabled on Network 2's ingress port
- No Local Gateway address defined on egress port
- Remote Gateway address Y configured to a prefix on Network 1's egress port, and the prefix corresponds to B

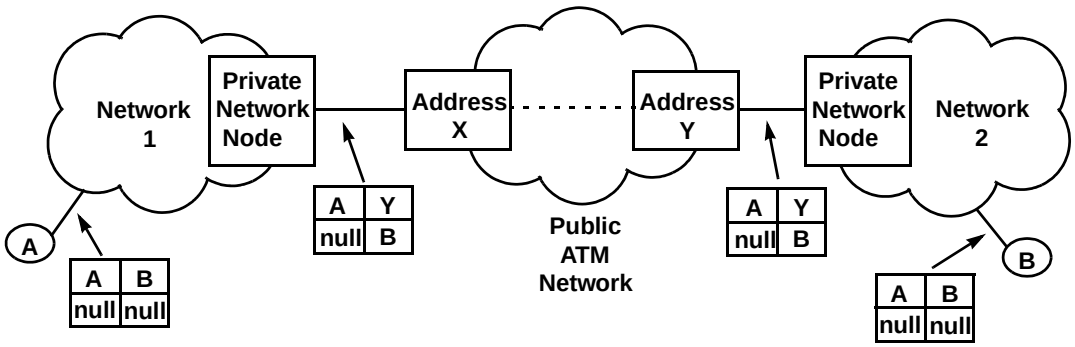

Example 3

- Replace option selected on egress port of Network 1
- Local Gateway address X configured to a prefix on Network 1's egress port, and the prefix corresponds to B

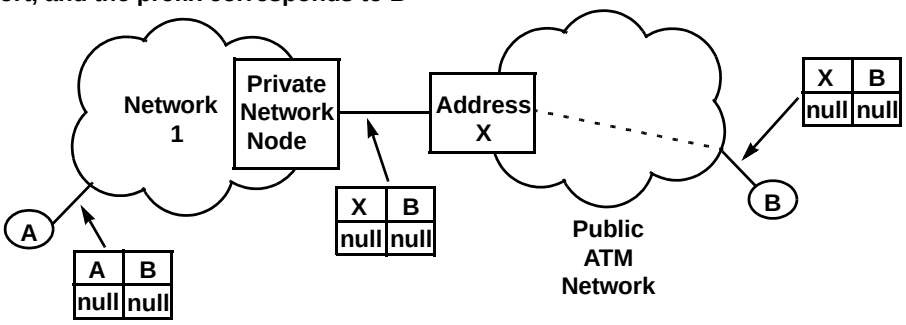

Example 4

- Replace option selected on egress port of Network 1
- Local Gateway address X configured to a prefix on Network 1's egress port, and the prefix corresponds to B
- Remote Gateway address Y configured to a prefix on Network 1's egress port, and the prefix corresponds to B

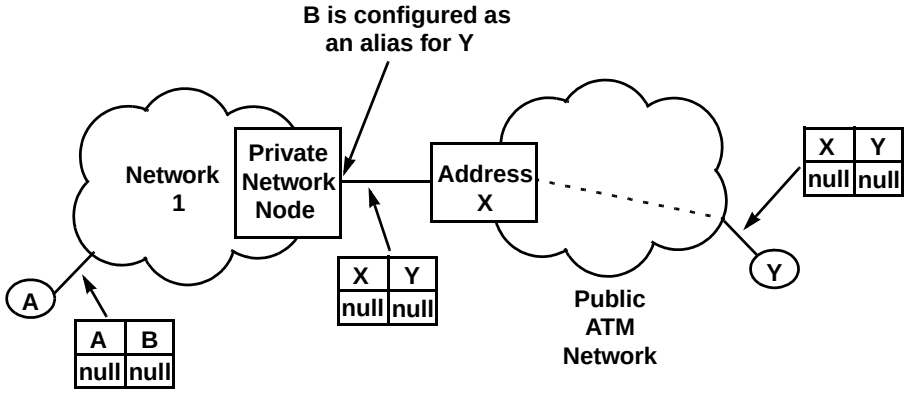

Network ID Addressing

A network ID can be used to identify an inter-exchange carrier (IXC). You can configure network ID addressing on ATM and Frame Relay 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). A network ID enables you to associate a network-to-network connection with a particular IXC (using a route determination ID) and 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 the 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
- Provisional Default TNS

Since routing is performed based on either the signaled/provisioned TNS value or the signaled called party address (not both), you can configure either route determination network IDs or route determination port prefixes/addresses (not both) 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 ATM UNI 3.x, 4.0, IISP, or FRF.4 ports.

You can configure a maximum of 1024 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 1024 minus the sum of port addresses, prefixes, and user parts.

12

Configuring SVC Parameters

This chapter contains procedures to perform the following tasks:

- Configure node and port prefixes to route SVC requests to a specific node or logical port. With node and port prefixes, you may take advantage of address registration.
- Configure the port user part of an address (DTE ports only). Address registration combines the port user part with a node or port prefix to route the SVC request.
- Configure SVC port addresses to route SVC requests to a specific logical port when the attached network device does not support address registration.
- Configure a network ID to uniquely identify an inter-exchange carrier (IXC).

The B-STDX does not support ATM SVCs. You also cannot configure ATM SVCs on the 6-port Frame-based DS3 module.

Configuring Node Prefixes

Node prefixes apply to all ports on the switch and are used for routing aggregation, source address validation, and address registration. 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 the two AFI digits of the AESA address, or 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 AESA or 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 granularity you need to define 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 you do so, you may need to define port prefixes or port addresses to provide more granularity for routing determination. For example, you may define port prefixes 508551, 508552, and 508553 on the first node, and port prefixes 508554, 508555, and 508556 on the second node.

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 following dialog box appears.

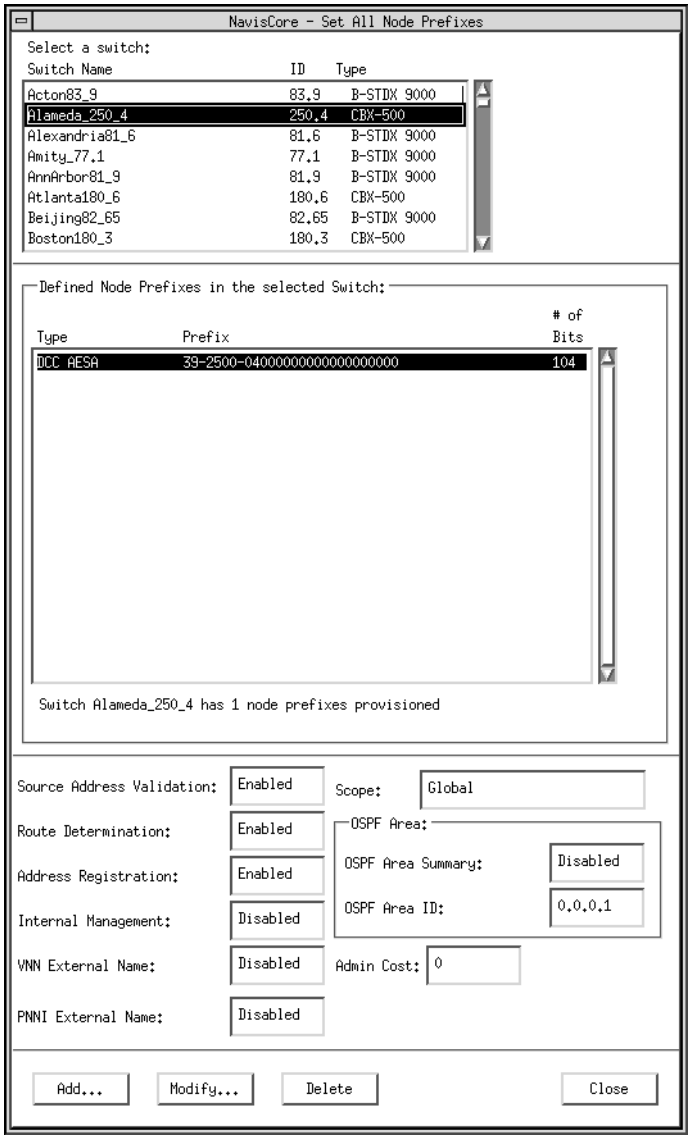

Figure 12-1. Set All Node Prefixes Dialog Box

The top list box (*Select a switch:*) displays all switches that are accessible from this NMS. The bottom list box (*Defined Node Prefixes in the selected Switch:*) diplays all the node prefixes configured on the selected switch. You can display the address and routing options for each node prefix.

2. From the top list box, select the switch on which to configure node prefixes.

3. Choose Add. The Add Node Prefix dialog box appears (see Figure 12-2 on page 12-5).
4. To define a node prefix for a specific format:
   a. See Table 12-1 to select the address format.
   b. Select the scope. Organizational scope defines how far into a hierarchical PNNI domain the switch should advertise this prefix or address. Although this release of NavisCore does not support hierarchies, you can configure this parameter to support future migration. For more information about PNNI, see the *ATM Forum PNNI Specification*.
   c. Continue with the section that corresponds to the address format you select.

Table 12-1. Address Format Descriptions

Format	Description	See . . .
E.164 Native	Standard 1-15 digit Integrated Services Digital Network (ISDN) number, which includes telephone numbers.	"Native E.164 Node Prefix Format" on page 12-5
DCC AESA	Data country code ATM End System Address, which identifies the country in which the address is registered.	"DCC and ICD AESA Node Prefix Format" on page 12-6
DCC Anycast AESA	Provides a group address function using DCC AESA address formats. Use the DCC AESA configuration instructions.	"DCC and ICD AESA Node Prefix Format" on page 12-6
ICD AESA	International country designator ATM End System Address, which identifies the international organization to which the address applies.	"DCC and ICD AESA Node Prefix Format" on page 12-6
ICD Anycast AESA	Provides a group address function using ICD AESA address formats. Use the ICD AESA configuration instructions.	"DCC and ICD AESA Node Prefix Format" on page 12-6
E.164 AESA	E.164 ATM End System Address, which encapsulates a standard 1-15 digit ISDN number, including telephone numbers.	"E.164 AESA Node Prefix Format" on page 12-7
E.164 Anycast AESA	Provides a group address function using E.164 AESA address formats. Use the E.164 AESA configuration instructions.	"E.164 AESA Node Prefix Format" on page 12-7
Custom AESA	ATM End System Address with customized octet structure and customized authority and format identifier (AFI).	"Custom AESA Node Prefix Format" on page 12-9

Native E.164 Node Prefix Format

Complete the following information for the E.164 (Native) format:

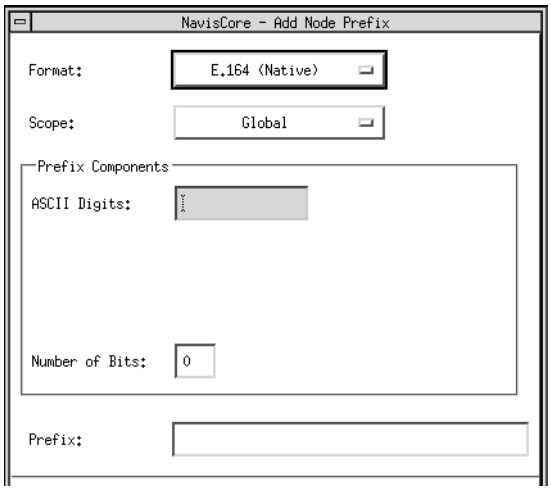

Figure 12-2. Add Node Prefix Dialog Box (E.164 Native Format)

1. In the ASCII Digits field, enter all or part of the 1-15 ASCII digits that represent the E.164 address.

   For example, enter 5085552600 (a standard 10-digit U.S. phone number), or enter a partial number (such as 508). The value you enter is converted to the ASCII hex values that represent each digit in the number (this value is displayed in the Address in Hex field). If you entered 5085552600, it converts to 35303835353532363030. This value is also displayed in the Address in Hex column on the Set All Node Prefixes dialog box (Figure 12-1 on page 12-3).

2. Configure the address routing options using the steps on page 12-10, "Defining Address and Routing Options."

3. Choose OK to return to the Set All Node Prefixes dialog box.

DCC and ICD AESA Node Prefix Format

Complete the following information for the DCC or ICD AESA format:

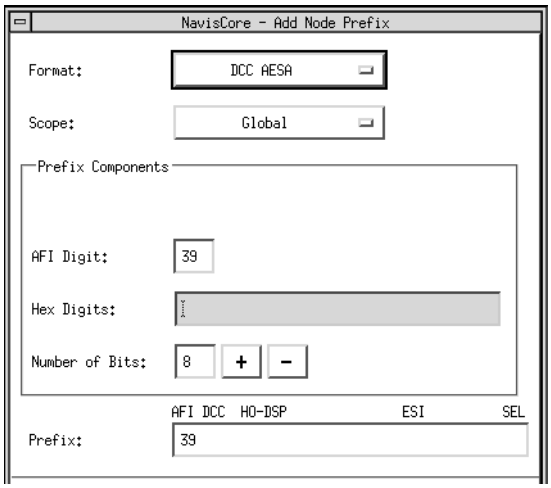

Figure 12-3. Add Node Prefix Dialog Box (DCC or ICD AESA Format)

1. In the Hex Digits field, enter the Data Country Code (DCC) of the country in which the address is registered, or the International Country Designator (ICD) that identifies the international organization to which this address applies. DCCs and ICDs consist of 4 hex digits, and occupy two octets.
2. (*Optional*) Enter the HO-DSP, ESI and SEL portions of the address.

   For information on the appropriate format to use for DCC and ICD addresses, see "ATM End System Address (AESA) Formats" on page 11-2.

To register the AESA address in the attached DTE devices ILMI prefix table, enter exactly the first 13 octets (26 digits) of the AESA address. Address registration occurs only on ILMI-enabled UNI ports with prefixes that have the address registration option set to enabled (see page 12-11).

3. 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. (The value increases by eight with each pair of address digits you type). Click on the — icon to decrease the number of address bits that are checked, thereby enabling the node to perform call screening and call routing down to the bit level. You can decrease the value by 1-7 bits.

For example, if you enter the partial DCC AESA address 39-43BF12AC (which uses 40 bits) as the port prefix, but only need to check the first 38 bits of the port prefix for call screening and call routing (because the last two binary digits are 00), click the — icon until the value in the Number of Bits field is 38.

Address you entered:	**39-43BF12AC**
Address in binary (40 bits):	**00111001-0100001111001111000010010101111100**
Address in binary (38 bits):	**00111001-0100001111001111000010010101111**

Address you entered:	**39-43BF12A8**
Address in binary (40 bits):	**00111001-0100001111001111000010010101111000**
Address in binary (37 bits):	**00111001-0100001111001111000010010101111**

4. Configure the address routing options using the steps on page 12-10, "Defining Address and Routing Options."

5. Choose OK to return to the Set All Node Prefixes dialog box (Figure 12-1 on page 12-3). The new entry appears, along with the preceding AFI (39 or 47), in the bottom list box (*Defined Node Prefixes in the selected Switch*).

E.164 AESA Node Prefix Format

Complete the following information for the E.164 AESA format:

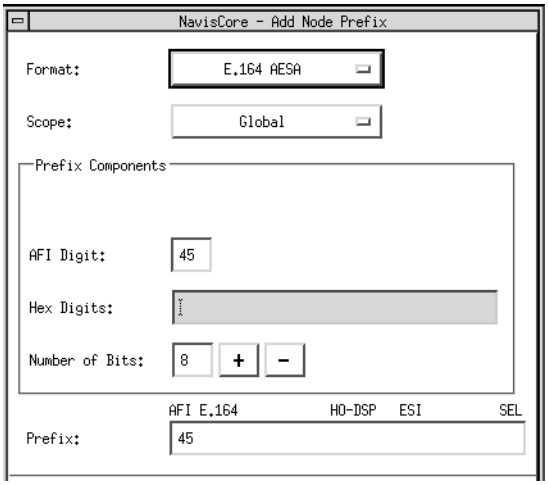

Figure 12-4. Add Node Prefix Dialog Box (E.164 AESA Format)

1. In the Hex Digits field, enter the full or partial E.164 AESA address. Since the IDI portion of the address is 8 octets (16 hex digits), but the E.164 address format is a maximum of 15 digits, you must terminate the IDI portion with Fh.

   If you do not enter all 15 digits of the IDI portion, you must enter leading zeros to fill in the octets. For example, enter 508 as 000000000000508F; enter 508555 as 000000000508555F.

2. If you enter the IDI portion of the address, you can optionally enter the HO-DSP, ESI, and SEL portions. For example, if you enter 000005085551234F, you can then enter all or some of the remaining parts. For information about the appropriate format to use for E.164 AESA addresses, see "ATM End System Address (AESA) Formats" on page 11-2.

To register the AESA address in the attached DTE devices ILMI prefix table, enter exactly the first 13 octets (26 digits) of the AESA address. Address registration occurs only on ILMI-enabled UNI ports with prefixes that have the address registration option set to enabled (see page 12-11).

3. 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. (The value increases by eight with each pair of address digits you enter). Click on the — icon to decrease the number of address bits that are checked, thereby enabling the node to perform call screening and call routing down to the bit level. You can decrease the value by 1-7 bits.

   For example, if you enter the partial E.164 AESA address 45-00000504 (which uses 40 bits) as the port prefix, but only need to check the first 38 bits of the port prefix for call screening and call routing (because the last two binary digits are 00), click the — icon until the value in the Number of Bits field is 38.

Address you entered:	**45-00000504**
Address in binary (40 bits):	**01000101-00000000000000000000010100000100**
Address in binary (38 bits):	**01000101-00000000000000000000101000001**

4. Configure the address routing options using the steps on page 12-10, "Defining Address and Routing Options."
5. Choose OK to return to the Set All Node Prefixes dialog box (Figure 12-1 on page 12-3). The new entry appears, along with the preceding AFI (45), in the bottom list box (*Defined Node Prefixes in the selected Switch*).

Custom AESA Node Prefix Format

Complete the following information for the Custom AESA format:

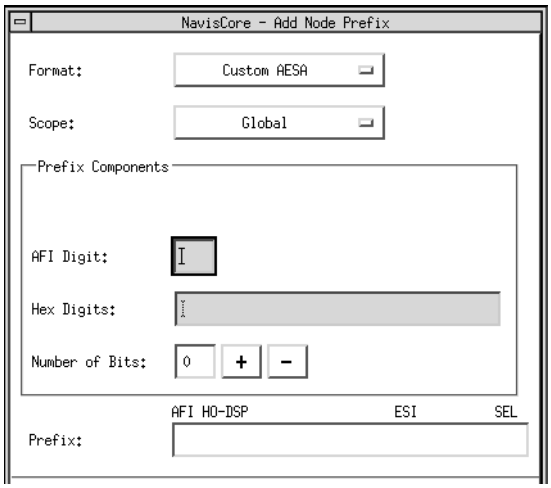

Figure 12-5. Add Node Prefix Dialog Box (Custom AESA Format)

1. In the AFI Digits field, enter the custom AFI you want to use.
2. In the Hex Digits field, enter in the customized address format, starting with the HO-DSP, followed by the ESI and SEL values (in that order).

   This address can be up to 19 octets (38 hex digits) long, with 12 octets used for the HO-DSP, 6 octets used for the ESI, and 1 octet used for the SEL. You do not have to enter the entire address; the HO-DSP, ESI, and SEL entries are optional. However, you must enter the AFI digits. For information about these items, see "ATM End System Address (AESA) Formats" on page 11-2.

To register the AESA address in the attached DTE devices' ILMI prefix table, enter exactly the first 13 octets (26 digits) of the AESA address. Address registration occurs only on ILMI-enabled UNI ports with prefixes that have the address registration option set to enabled (see page 12-11).

3. 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. (The value increases by eight with each pair of address digits you type). Click the − icon to decrease the number of address bits checked, thereby enabling the node to perform call screening and call routing down to the bit level. You can decrease the value by 1-7 bits.

For example, if you enter the partial address 51-43BF12AC (which uses 40 bits) as the port prefix, but only need to check the first 38 bits of the port prefix for call screening and call routing, click the — icon until the value in the Number of Bits field is 38.

Address you entered:	**51-43BF12AC**
Address in binary (40 bits):	**01010001-0100001111001111000100101011110 0**
Address in binary (38 bits):	**01010001-0100001111001111000100101011 11**

Address you entered:	**51-43BF12A8**
Address in binary (40 bits):	**01010001-0100001111001111000100101011100 0**
Address in binary (37 bits):	**01010001-0100001111001111000100101011 1**

4. See the following section, "Defining Address and Routing Options," to configure the address and routing options.
5. Choose OK to return to the Set All Node Prefixes dialog box (Figure 12-1 on page 12-3).

Defining Address and Routing Options

The Add Node Prefix dialog boxes contain fields that allow you to enable or disable the address and routing options.

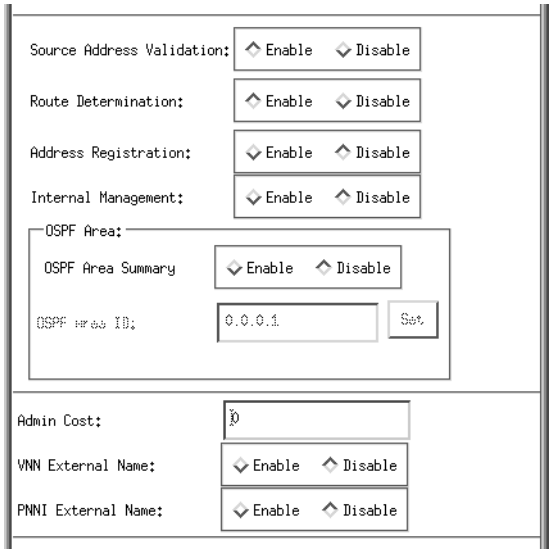

Figure 12-6. Add Node Prefix Address and Routing Fields

Use Table 12-2 to configure these options.

Table 12-2. Add Node Prefix Address and Routing Fields

Field	Action/Description
Source Address Validation	Select enable to validate the calling party address against the node prefix associated with the UNI/NNI 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. Enable this option to use PVC/PVP termination (see page 14-2).
Address Registration	If enabled, this node prefix is used for ILMI address registration for all UNI-DCE “network-to-endsystem” logical ports that support ILMI. You cannot enable this option for AESA node prefixes that are not 13 octets long.
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 represents an area border router. Then enter an OSPF Area ID. For more information, see the *NavisCore IP Navigator Services Guide.*
OSPF Area ID	If you enable OSPF Area Summary, choose Set to select the Area ID. This assigns an OSPF area to a node prefix in cases where the node acts as an area border router. OSPF Area IDs enable the VC manager to determine which way to route the PVC.
Admin Cost	Enter the administrative cost associated with the 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.
VNN External Name	Select enable to advertise this name within the VNN routing domain as an external name. An external name is a name that is reachable within another VNN routing domain. The default, disable, means this name is only reachable within the VNN routing domain.
PNNI External Name	Select enable to advertise this name within the PNNI routing domain as an external name. An external name is a name that is reachable within another PNNI routing domain. The default, disable, means this name is only reachable within the PNNI routing domain.

Configuring SVC Port Prefixes

The Set All 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 following dialog box appears.

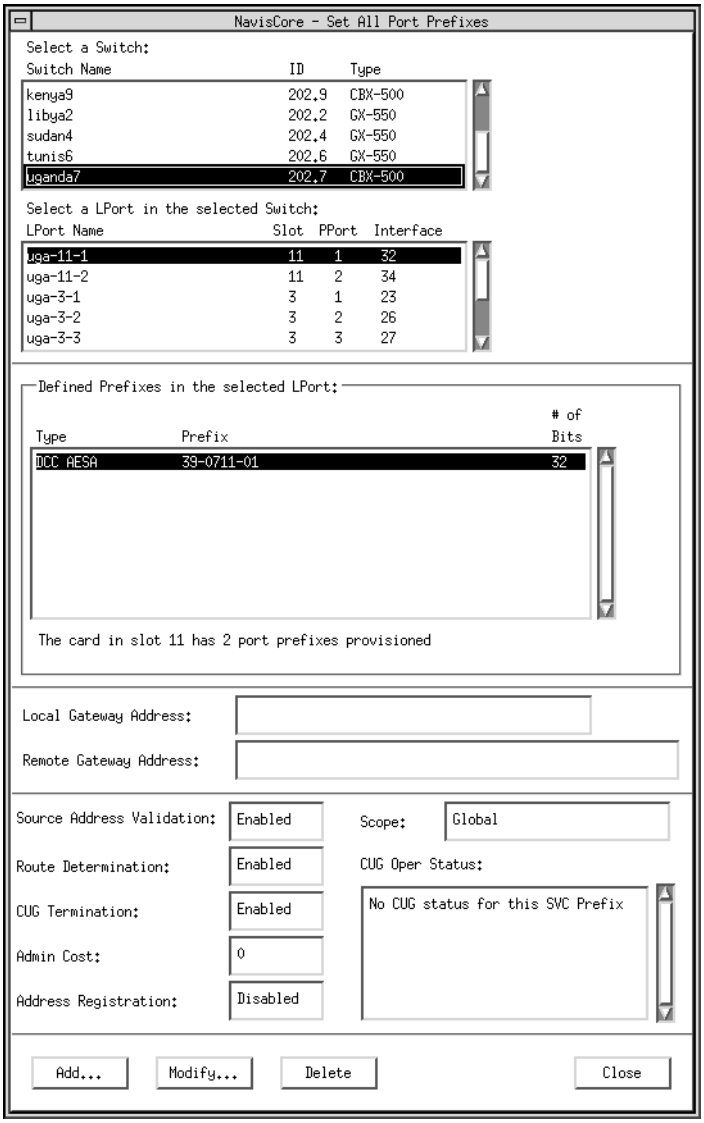

Figure 12-7. Set All Port Prefixes Dialog Box

- The *Select a Switch:* list box displays all switches that this NMS can access.
- The *Select a LPort in the selected Switch:* list box displays the logical ports that are configured for the selected switch, along with the slot, physical port, and MIB interface number for the logical port.
- The *Defined Prefixes in the selected LPort:* list box displays all port prefixes that have already been defined on the selected logical port. You can display configured options for each of these prefixes.

2. Select the switch on which to configure a port prefix.
3. Select the logical port on which to configure a port prefix.
4. Choose Add. The Add Prefix dialog box appears (Figure 12-8 on page 12-14).
5. Select an address format and scope. See page 12-4 for a description of these attributes.
6. Continue with the section that corresponds to the address format you select.

Format	See . . .
E.164 Native	page 12-14
DCC (Anycast) AESA	page 12-15
ICD (Anycast) AESA	page 12-15
E.164 (Anycast) AESA	page 12-16
Custom AESA	page 12-18
Default Route	page 12-22

E.164 Native Prefixes Port Prefix Format

Complete the following information for the E.164 native prefix format:

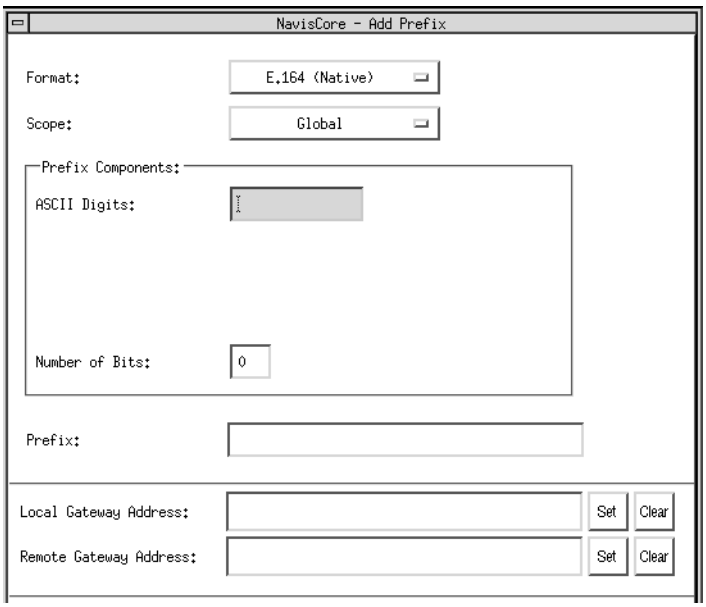

Figure 12-8. Add Prefix Dialog Box (E.164 Native Format)

1. In the ASCII Digits field, enter all or part of the 1-15 ASCII digits that represent the E.164 address.

   For example, enter 5085552600 (a standard 10-digit U.S. phone number), or enter a partial number (such as 508). The value you enter is converted to the ASCII hex values that represent each digit in the number (this value is displayed in the Address in Hex field). If you entered 508555260, it converts to 35303835353532363030. This value is also displayed in the Address in Hex column on the Set All Port Prefixes screen.

2. If the port provides a network-to-network connection, see "Setting the Local and Remote Gateway Address for Port Prefixes" on page 12-20 for instructions. When done, proceed to Step 3.

3. See Table 12-3 on page 12-23 to configure additional port prefix options.

4. Choose OK to return to the Set All Port Prefixes dialog box (Figure 12-7 on page 12-12).

DCC and ICD AESA Port Prefix Format

Complete the following information for the DCC or ICD AESA prefix format:

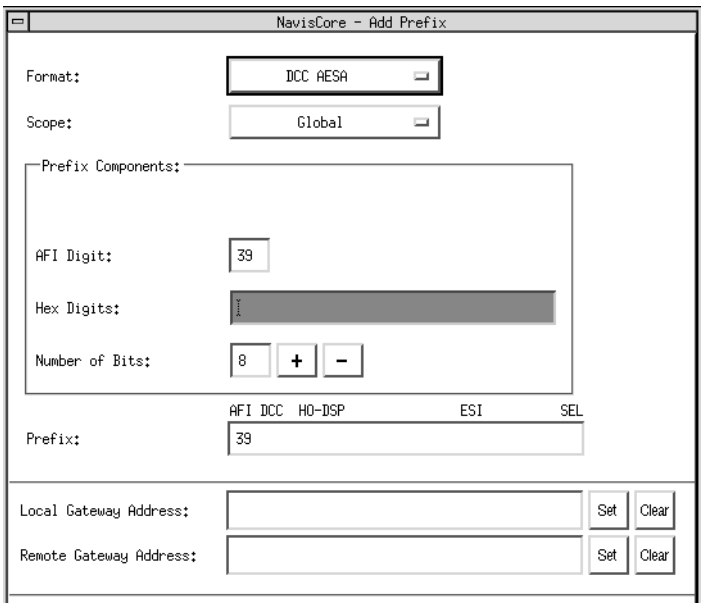

Figure 12-9. Add Prefix Dialog Box (DCC and ICD AESA Format)

1. In the Hex Digits field, enter the Data Country Code (DCC) of the country in which the address is registered, or the International Country Designator (ICD) that identifies the international organization to which this address applies. DCCs and ICDs consist of 4 hex digits, and occupy two octets.
2. (*Optional*) Enter the HO-DSP, ESI and SEL portions of the address.

   For information on the appropriate format to use for DCC and ICD addresses, see "ATM End System Address (AESA) Formats" on page 11-2.

To register the AESA address in the attached DTE devices' ILMI prefix table, enter exactly the first 13 octets (26 digits) of the AESA address. Address registration occurs only on ILMI-enabled UNI ports.

3. As you type the address, the value in the Number of Bits field changes to indicate the number of address bits checked during call screening and call routing. (The value increases by eight with each address digit you type). Click the — icon to decrease the number of address bits checked, thereby enabling the node to perform call screening and call routing down to the bit level. You can decrease the value by 1-7 bits.

For example, if you enter the partial DCC AESA address 39-43BF12AC (which uses 40 bits) as the port prefix, but only need to check the first 38 bits of the port prefix for call screening and call routing (because the last two digits are binary 00), click the — icon until the value in the Number of Bits field is 38.

Address you entered:	**39-43BF12AC**
Address in binary (40 bits):	**00111001-0100001111001111000100101011110 0**
Address in binary (38 bits):	**00111001-0100001111001111000100101011 11**

Address you entered:	**39-43BF12A8**
Address in binary (40 bits):	**00111001-0100001111001111000100101011 1000**
Address in binary (37 bits):	**00111001-0100001111001111000100101011 1**

4. If the port provides a network-to-network connection, see "Setting the Local and Remote Gateway Address for Port Prefixes" on page 12-20 for instructions. When done, proceed to Step 5.
5. See Table 12-3 on page 12-23 to configure additional port prefix options.
6. Choose OK to return to the Set All Port Prefixes dialog box (Figure 12-7 on page 12-12). The new entry appears, along with the preceding AFI (39 or 47), in the *Defined Prefixes in the selected LPort* list box on the bottom half of the screen.

E.164 AESA Port Prefix Format

Complete the following information for the E.164 AESA prefix format:

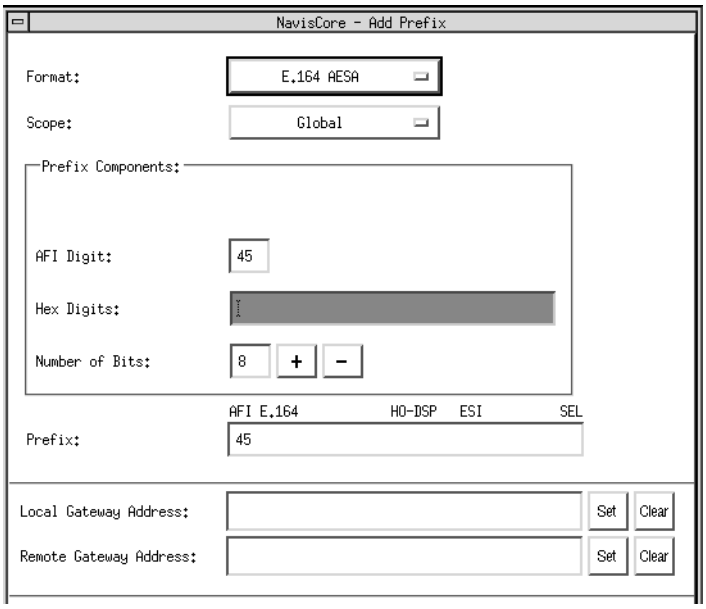

Figure 12-10. Add Prefix Dialog Box (E.164 AESA Format)

1. In the Hex Digits field, enter the full or partial E.164 AESA address. Since the IDI portion of the address is 8 octets (16 hex digits), but the E.164 address format is a maximum of 15 digits, you must terminate the IDI portion with Fh.

   If you do not enter all 15 digits of the IDI portion, you must enter leading zeros to fill in the octets. For example, enter 508 as 000000000000508F; enter 508555 as 000000000508555F.

2. If you enter the IDI portion of the address, you can optionally enter the HO-DSP, ESI, and SEL portions. For example, if you enter the IDI portion as 000005085551234F, you can then enter all or some of the remaining parts. For information on the appropriate format to use for E.164 addresses, see "ATM End System Address (AESA) Formats" on page 11-2.

To register the AESA address in the attached DTE devices' ILMI prefix table, enter exactly the first 13 octets (26 digits) of the AESA address. Address registration occurs only on ILMI-enabled UNI ports.

3. As you type the address, the value in the Number of Bits field changes to indicate the number of address bits checked during call screening and call routing. (The value increases by eight with each pair of address digits you type). Click the — icon to decrease the number of address bits checked, thereby enabling the node to perform call screening and call routing down to the bit level. You can decrease the value by 1-7 bits.

   For example, if you enter the partial E.164 AESA address 45-00000504 (which uses 40 bits) as the port prefix, but only need to check the first 38 bits of the port prefix for call screening and call routing (because the last two binary digits are 00), click the — icon until the value in the Number of Bits field is 38.

Address you entered:	**45-00000504**
Address in binary (40 bits):	**01000101-00000000000000000000010100000100**
Address in binary (38 bits):	**01000101-00000000000000000000101000001**

4. If the port provides a network-to-network connection, see "Setting the Local and Remote Gateway Address for Port Prefixes" on page 12-20 for instructions. When done, proceed to Step 5.
5. See Table 12-3 on page 12-23 to configure additional port prefix options.
6. Choose OK to return to the Set All Port Prefixes dialog box (Figure 12-7 on page 12-12). The new entry appears, along with the preceding AFI (45), in the *Defined Prefixes in the selected LPort* list box.

Custom AESA Port Prefix Format

Complete the following information for the Custom AESA prefix format:

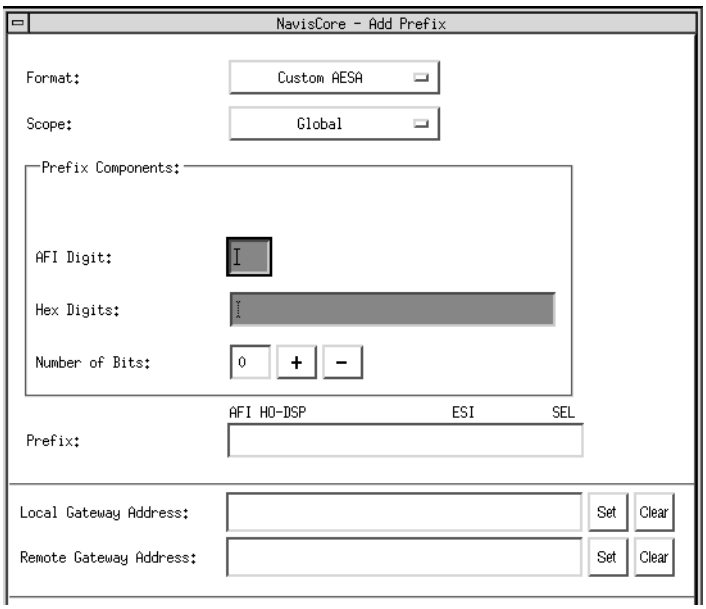

Figure 12-11. Add Prefix Dialog Box (Custom AESA Format)

1. In the AFI Digits field, enter the custom AFI you want to use.
2. In the Hex Digits field, type in the customized address format, starting with the HO-DSP, followed by the ESI and SEL values (in that order). This address can be up to 19 octets (38 hex digits) long, with 12 octets used for the HO-DSP, 6 octets used for the ESI, and 1 octet used for the SEL. You do not have to enter the entire address; the HO-DSP, ESI, and SEL entries are optional. However, the AFI digits are required. For information on these items, see "ATM End System Address (AESA) Formats" on page 11-2.

To register the AESA address in the attached DTE devices' ILMI prefix table, enter exactly the first 13 octets (26 digits) of the AESA address. Address registration occurs only on ILMI-enabled UNI ports.

3. As you type the address, the value in the Number of Bits field changes to indicate the number of address bits checked during call screening/call routing. (The value increases by eight with each address digit you type). Click the — icon to decrease the number of address bits, thereby enabling the node to perform call screening and call routing down to the bit level. You can decrease the value by 1-7 bits.

   For example, if you enter the partial address 51-43BF12AC (which uses 40 bits) as the port prefix, but only need to check the first 38 bits of the port prefix for call screening and call routing, click the — icon until the value in the Number of Bits field is 38.

Address you entered:	**51-43BF12AC**
Address in binary (40 bits):	**01010001-01000011110011110001001010111100**
Address in binary (38 bits):	**01010001-0100001111001111000100101011111**

Address you entered:	**51-43BF12A8**
Address in binary (40 bits):	**01010001-01000011110011110001001010111000**
Address in binary (37 bits):	**01010001-0100001111001111000100101011**1

4. If the port provides a network-to-network connection, see "Setting the Local and Remote Gateway Address for Port Prefixes" on page 12-20 for instructions. When done, proceed to Step 5.
5. See Table 12-3 on page 12-23 to configure additional port prefix options.
6. Choose OK to return to the Set All Port Prefixes dialog box (Figure 12-7 on page 12-12).

Setting the Local and Remote Gateway Address for Port Prefixes

This section describes how to set the optional local and remote gateway addresses for ports that are providing a network-to-network connection. Local and remote gateway addresses are used in conjunction with the egress address translation feature (see page 11-9).

Figure 12-12 shows which addresses to enter as the local and remote gateway addresses for each end of the network-to-network connection.

Figure 12-12. Setting Local and Remote Gateway Addresses

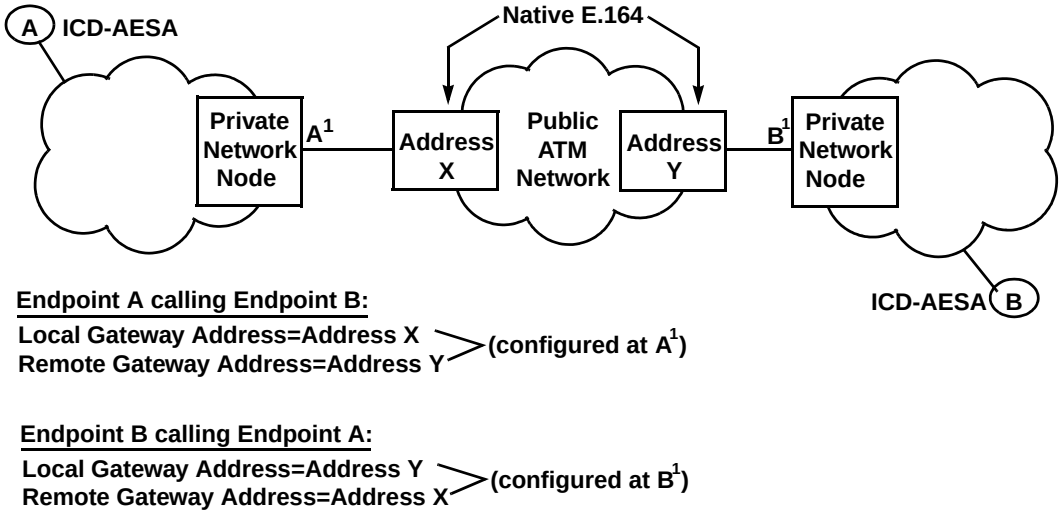

You can configure prefixes on a network-to-network port with the following:

- Null local and remote gateway addresses
- Only a local gateway address
- Only a remote gateway address
- Both a local and a remote gateway address

You need to define gateway addresses for address translation only. For more information on egress address translation, see page 11-9.

To set the local (or remote) gateway address:

1. From the Add Prefix dialog box, choose the Set command. The Set Local (or Remote) Gateway Address dialog box appears.

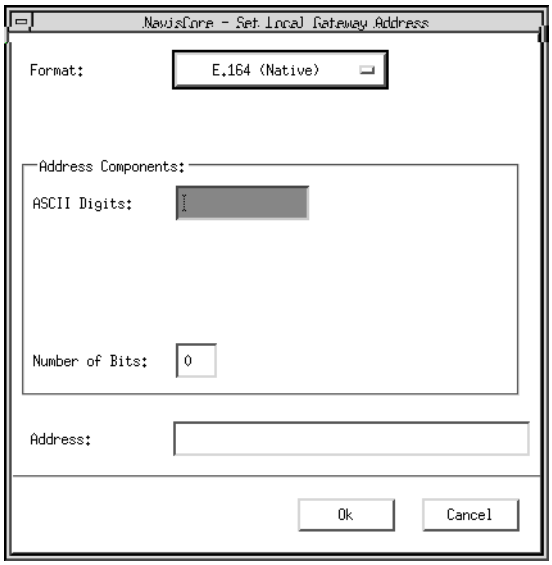

Figure 12-13. Set Local Gateway Address Dialog Box

2. Select the local (or remote) gateway address format.
3. If you specify the local gateway address, enter the address of the public network gateway used to enter the public network. If you specify the remote gateway address, enter the address of the public network gateway used to exit from the public network back to the private network.
4. When done, choose OK to return to the Add Prefix dialog box.
5. To continue, do one of the following:
   - If you are defining an E.164 Native prefix format, proceed to Step 3 on page 12-14.
   - If you are defining a DCC or ICD AESA prefix format, proceed to Step 5 on page 12-16.
   - If you are defining an E.164 AESA prefix format, proceed to Step 5 on page 12-17.
   - If you are defining a Custom AESA prefix format, proceed to Step 5 on page 12-19.

Defining Default Routes for Network-to-Network Connections

For ports being used for 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 calling party address does not match any port prefixes or addresses, it routes the call to the port on which the default route is defined. If more than one port has a default route defined, then the administrative cost value is used to determine the port to which the call is routed.

It is important that you define a port address for calling/called party addresses. Admin cost is not always the criteria for routing a call, because the call can be placed out the same interface on which it was received.

You can define multiple default routes within a node or network. The default route typically applies to network-to-network logical ports (IISP or public UNI DTE).

Complete the following information for a default route:

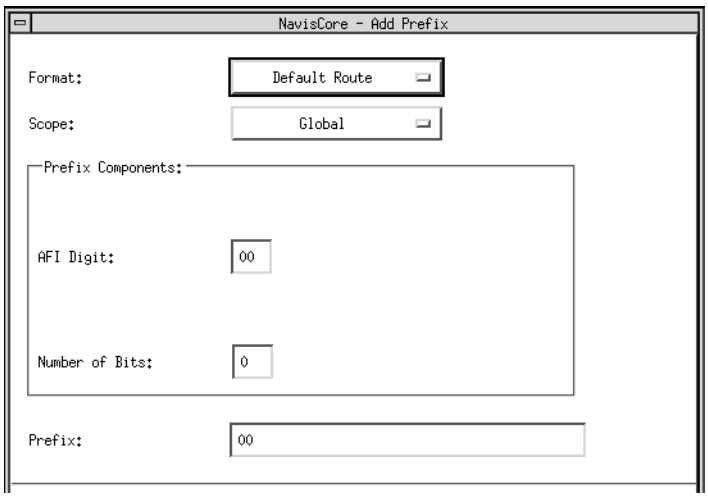

Figure 12-14. Add Prefix Dialog Box (Default Route)

1. In the Format field, select Default Route.
2. See Table 12-3 on page 12-23 to configure additional port prefix options.
3. Choose OK to return to the Set All Port Prefixes dialog box.

Defining Port Prefix Options

When you add a port prefix, the Add Prefix dialog box contains fields which allow you to enable or disable the following options, as shown in Figure 12-15.

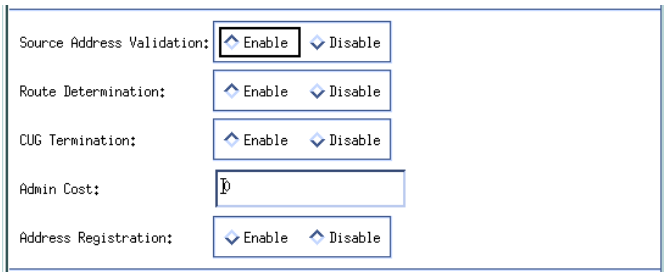

Figure 12-15. Add Port Prefix Option Fields

Use Table 12-3 to configure these options.

Table 12-3. Add Port Prefix Option Fields

Field	Action/Description
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. Enable this option to use PVC/PVP termination (see page 14-2).
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 on CUGs, see Chapter 15, “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.
Address Registration	If enabled, port prefixes are used for ILMI address registration if ILMI is enabled on this logical port. This option cannot be enabled for AESA port prefixes that are not 13 octets long.

Configuring SVC Port Addresses

If the device attached to a given physical port does not support ILMI address registration, or to fully specify an address to use for calling party screening, you can define SVC addresses for all the logical ports on a given physical port. The AESA formats must have full-length address definitions and include all 20 octets (40 hex digits). That is, you must enter the AFI, IDI, HO-DSP, ESI, and SEL portions of the address. (Since ATM routing does not use the SEL portion, you can enter any value for that part of the address.) For native E.164 addresses, you enter the 1-15 digit E.164 address.

Defining an SVC Port Address

To define SVC addresses:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Port Addresses. The following dialog box appears (see Figure 12-16 on page 12-25).

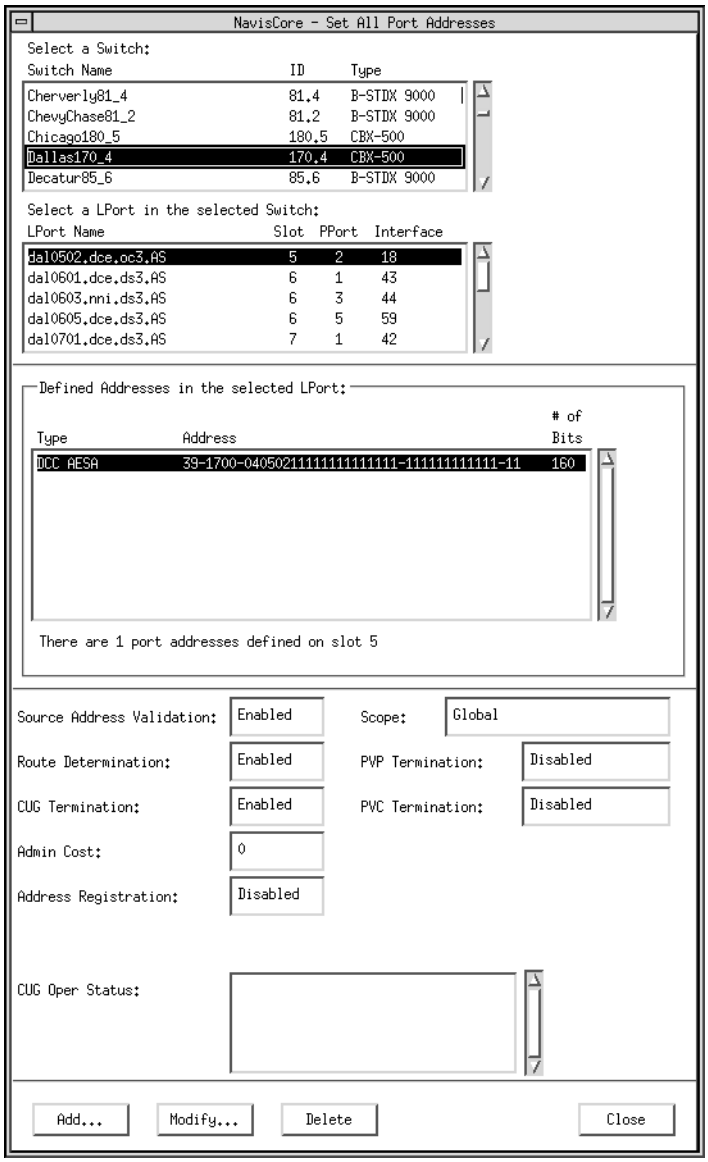

Figure 12-16. Set All Port Addresses Dialog Box

- The top list box (*Select a Switch*) displays all switches that the NMS can access.
- The center list box (*Select a LPort*) displays the logical ports that are configured on the selected switch.
- The bottom list box (*Defined Addresses*) displays any SVC addresses that you already defined on the selected logical port.

2. Select the switch on which to configure SVC addresses.
3. Select the logical port on which to configure SVC addresses.
4. Choose Add. The Add Address dialog box appears.

5. Select an address format and scope. See page 12-4 for a description of these attributes.
6. Continue with the section that corresponds to the address format you select.

Format	See . . .
E.164 Native	page 12-26
DCC (Anycast) AESA	page 12-27
ICD (Anycast) AESA	page 12-27
E.164 (Anycast) AESA	page 12-28
Custom AESA	page 12-29

Native E.164 SVC Addresses

Complete the following information for E.164 (Native) format:

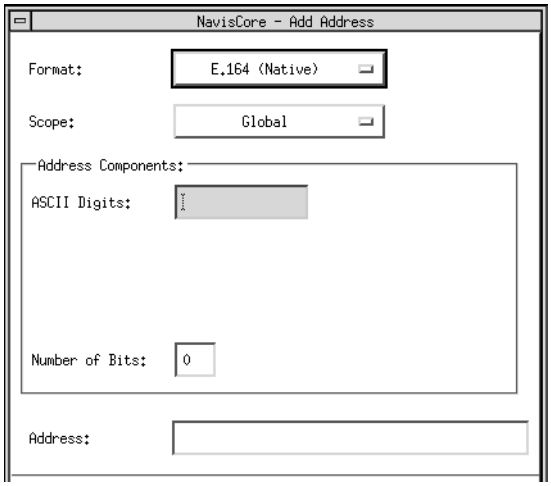

Figure 12-17. Add Address Dialog Box (Native E.164 SVC Address Format)

1. In the ASCII Digits field, enter all of the 1-15 ASCII digits that represent the E.164 address. For example, enter 5085552600 (a standard 10-digit U.S. phone number). The value you enter is converted to the ASCII hex values that represent each digit in the number (this value is displayed in the Address in Hex field).

   For example, 5085552600 converts to 35303835353532363030. The Address in Hex column on the Set All Port Addresses dialog box also displays this value.
2. See Table 12-4 on page 12-30 to configure additional fields.
3. Choose OK to return to the Set All Port Addresses dialog box (Figure 12-16 on page 12-25).

DCC and ICD AESA SVC Addresses

Complete the following information for DCC or ICD AESA format:

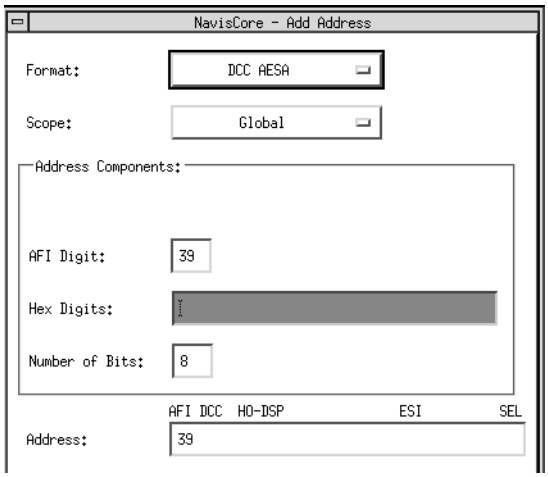

Figure 12-18. Add Address Dialog Box (DCC or ICD AESA Format)

1. In the Hex Digits field, enter the data country code (DCC) of the country in which the address is registered, or the international country designator (ICD) that identifies the international organization to which this address applies. DCCs and ICDs consist of 4 hex digits, and occupy two octets.
2. Enter the appropriate HO-DSP, ESI and SEL values. For information on these items, and the appropriate format to use for DCC and ICD AESA addresses, see "ATM End System Address (AESA) Formats" on page 11-2.
3. See Table 12-4 on page 12-30 to configure additional fields.
4. Choose OK to return to the Set All Port Addresses dialog box (Figure 12-16 on page 12-25).

E.164 AESA SVC Addresses

Complete the following information for the E.164 AESA format:

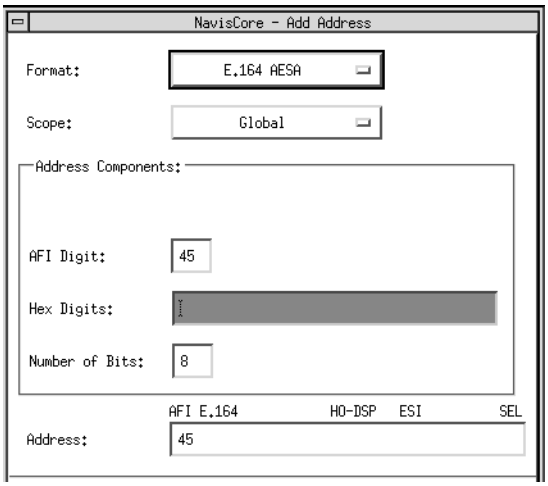

Figure 12-19. Add Address (E.164 AESA Format)

1. In the Hex Digits field, enter the full or partial E.164 AESA address. Since the IDI portion of the address is 8 octets (16 hex digits), but the E.164 address format is a maximum of 15 digits, you must terminate the IDI portion with Fh. For example, enter 5085551234 as 000005085551234F.
2. After you type the IDI portion of the address, enter the appropriate HO-DSP, ESI, and SEL portions to complete the address. For information on the appropriate format to use for E.164 AESA addresses, see "ATM End System Address (AESA) Formats" on page 11-2.
3. See Table 12-4 on page 12-30 to configure additional fields.
4. Choose OK to return to the Set All Port Addresses dialog box (Figure 12-16 on page 12-25).

Custom AESA SVC Addresses

Complete the following information for the Custom AESA format:

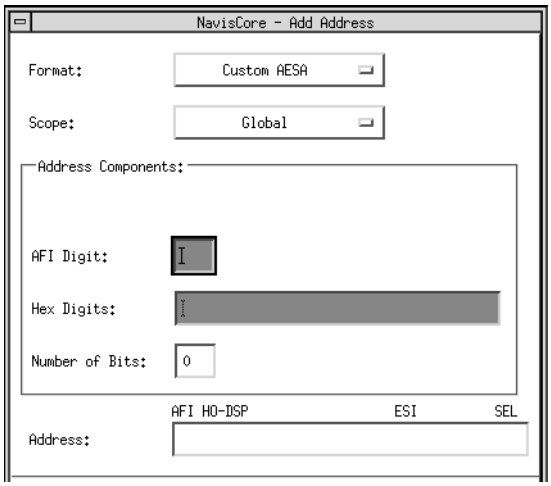

Figure 12-20. Add Address Dialog Box (Custom AESA Format)

1. In the AFI Digits field, enter the custom AFI.
2. In the Hex Digits field, enter the customized address format, starting with the HO-DSP, followed by the ESI and SEL values (in that order).

   This address must be the full 19 octets (38 hex digits) long, with 12 octets used for the HO-DSP, 6 octets used for the ESI, and 1 octet used for the SEL. For information on these items, see "ATM End System Address (AESA) Formats" on page 11-2.
3. See Table 12-4 on page 12-30 to configure additional fields.
4. Choose OK to return to the Set All Port Addresses dialog box (Figure 12-16 on page 12-25).

Defining SVC Port Address Options

The Add Address dialog boxes contain fields that allow you to enable or disable the following options.

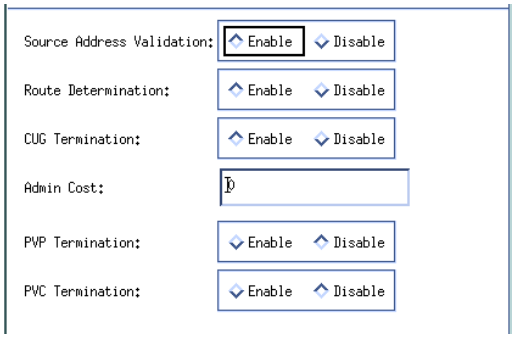

Figure 12-21. Add SVC Port Address Option Fields

Use Table 12-4 to configure the options shown in Figure 12-21.

Table 12-4. Add SVC Port Address Option Fields

Field	Action/Description
Source Address Validation	Select Enable to validate the calling party address against the UNI/NNI port address 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. Enable this option to use PVC/PVP termination (see page 14-2).
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 15, "Closed User Groups."
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.

If you are using soft PVCs in your network, continue with the following section, "Configuring PVP Termination." Otherwise, choose OK to return to the Set All Port Address dialog box (Figure 12-16 on page 12-25).

Configuring PVP Termination

Enable this option to terminate an SPVC to this address on the logical port. If you select Enabled, the Connection ID field appears.

Complete the following fields:

Connection ID — Select Any if you want the network to allocate a VPI for the SPVPC. Select Specific to supply a VPI. Note that if you enable PVC termination, this field is set to Any and cannot be changed.

VPI — Enter the VPI of the logical port on which you want the switch to terminate this SPVPC. The logical port cell header type limits the range of values you can enter: UNI = 255, NNI = 4095.

For more information about SPVCs, see "About Point-to-Point SPVCs" on page 14-4.

Configuring PVC Termination

Enable this option to terminate an SVC (spoofing) or SPVCC to this address on this logical port. If you select Enabled, the Connection ID field appears.

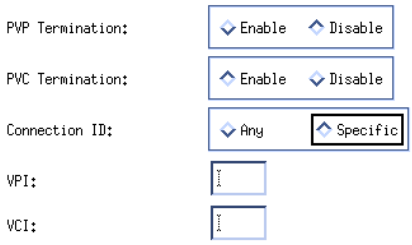

Complete the following fields:

Connection ID — Select Any if you want the network to allocate a VPI/VCI for the spoofed SVC or terminated SPVCC. Select Specific to supply a VPI/VCI value. Note that you cannot select specific if you also enable PVP termination.

VPI/VCI — Enter the VPI/VCI of the logical port on which you want the switch to terminate this SPVCC.

Configuring the Port User Part of the Address

The port user part of an AESA address consists of the ESI and SEL portions of the address. It is used for the DTE (user) ports on a Ascend switch and provides information for the address table on the DCE device attached to the UNI on the public network side (see "About Address Registration" on page 11-6). When the attached DCE device receives prefixes, the user part(s) are concatenated to form full addresses. The full addresses are then written back to the DCE device's ILMI address table.

When you configure the port user part to complete an address connection with the attached DCE device, you can supply any 7-octet value as the user part (it does not have to be a real IEEE MAC address and SEL combination). Also, you should enter any user addresses in your network that you want to make known to the attached public network. To do this, collect Media Access Control (MAC) addresses from attached devices and enter them as user parts at the public UNI port.

You may have to define user parts only on UNI-DTE ports where the device attached to that port expects address registration completion. That is, the attached device is broadcasting its network prefixes to the Ascend port, and expects the Ascend switch to respond with the user part of the address.

Defining a Port User Part

To define the port user part of the address:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Port User Parts. The following dialog box appears.

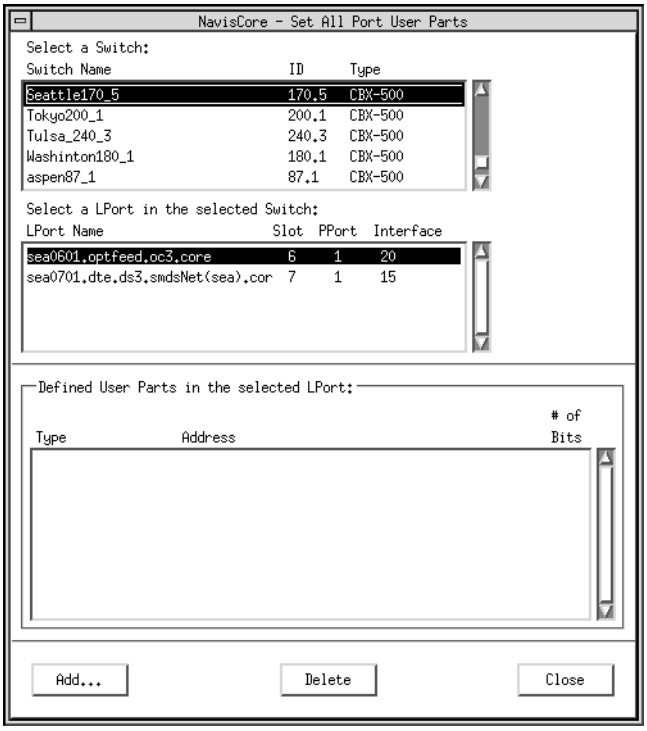

Figure 12-22. Set All Port User Parts Dialog Box

   - The *Select a Switch* list box shows all switches that this NMS can access.
   - The *Select a Logical Port* list box shows the UNI-DTE logical ports that are configured for this switch.
   - The *Defined User Parts* list box displays any user parts you already defined for this logical port.

2. In the Switch Name list box, select the switch on which you want to configure port user parts. The logical ports configured on the selected switch appear in the Lport Name list box.
3. In the Lport Name list box, select the logical port on which you want to configure user parts. If any user parts are already configured for this logical port, they appear in the list box at the bottom of the dialog box.

4. Choose Add. The Add User Part dialog box appears.

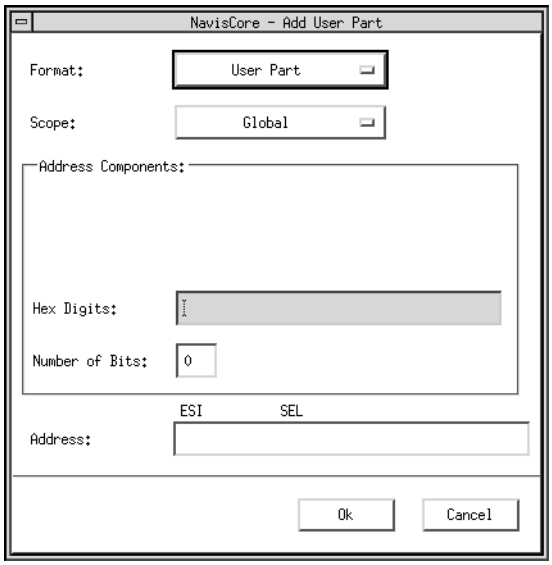

Figure 12-23. Add User Part Dialog Box

5. Enter the 7-octet (14-digit) user part in the Hex Digits field.
6. When done, choose OK to return to the Set All Port User Parts dialog box (Figure 12-22 on page 12-33).

Defining Network ID Parameters

You can add, delete, and modify network IDs. For an overview of network ID features, see page 11-14.

Adding Network IDs

To add a network ID:

1. From the Administer menu, select Ascend Parameters ⇒ Set All SVC Parameters ⇒ Set All Port Network IDs. The following dialog box appears.

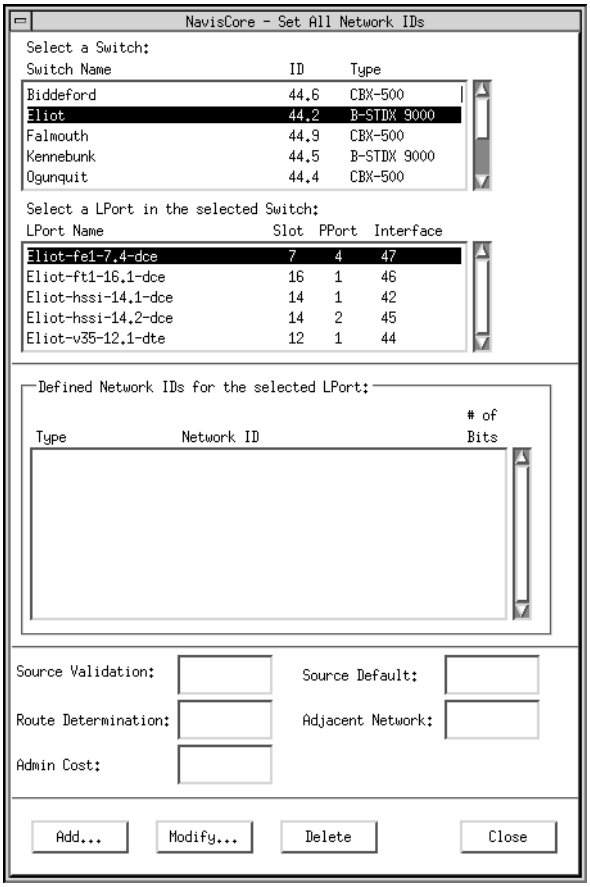

Figure 12-24. Set All Network IDs Dialog Box

Table 12-5 on page 12-36 describes the dialog box fields.

Table 12-5. Set All Network IDs Fields

Field	Action/Description
Select a Switch	Displays the Switch Name, ID, and Type for all existing switches in the network.
Select a Lport in the selected Switch	Displays the Lport Name, Slot, PPort, and Interface for all ATM UNI and NNI, and Frame Relay UNI logical ports 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 logical port.
Source Validation Route Determination Source Default Adjacent Network	Displays Enabled or Disabled for the selected Network ID.
Admin Cost	Displays the admin cost for the selected Network ID.

2. To add a network ID, select a switch name from the *Select a Switch* list box and then select a logical port from the *Select an LPort in the selected Switch* list box. Choose Add. The following dialog box appears.

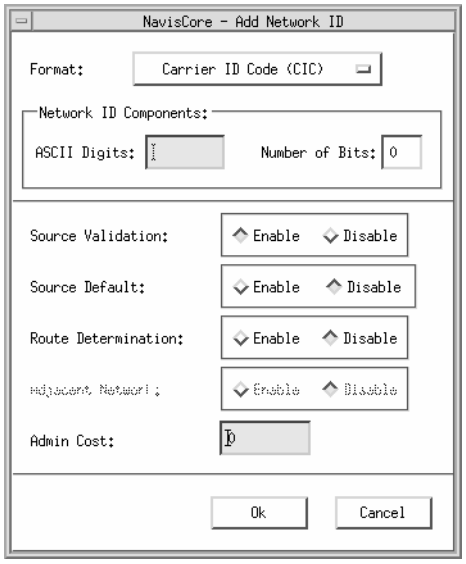

Figure 12-25. Add Network ID Dialog Box

3. Complete the Add network ID fields described in Table 12-6.

Table 12-6. 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 source default 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.
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.
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.
Admin Cost	Enter an admin 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 Network IDs dialog box (Figure 12-24 on page 12-35), select the desired network from the Defined Network IDs for the selected logical port list box.
2. Choose Modify.
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 12-24 on page 12-35), select the desired network from the Defined Network IDs for the selected logical port list box.
2. Choose Delete. The selected network is deleted from the list.
3. Choose Cancel to exit the dialog box.

13

Configuring SVC Proxy Signaling

Switched Virtual Circuit (SVC) proxy signaling is an optional CBX 500/GX 550 feature that enables a single signaling entity to signal on behalf of multiple endpoints. You can use proxy signaling to allow endsystems that do not understand ATM signaling to set up SVCs via a proxy signaling agent (PSA). The PSA performs all signaling functions on behalf of the endsystem, known as the proxy signaling client (PSC).

Before you can configure the PSA and PSC, use the instructions in Chapter 3, "Configuring CBX 500 or GX 550 Logical Ports," to configure the ATM UNI DCE logical ports.

About Proxy Signaling

The following terms are used to define proxy signaling functions:

Proxy Signaling Agent (PSA) — The network port attached to the signaling entity that performs signaling for non-signaling entities. In the Ascend implementation, a PSA is an ATM UNI DCE logical port for which signaling is enabled.

Proxy Signaling Client (PSC) — The network port attached to the endsystem for which a PSA performs signaling duties. In the Ascend implementation, a PSC is an ATM UNI DCE logical port for which signaling is disabled.

You can use proxy signaling to enable high-end ATM equipment to support multiple physical interfaces that share the same ATM address. This application provides high-end equipment with the ability to support connections that have an aggregate bandwidth which exceeds the physical interface line rate. The individual connection(s) must be at a rate that is less than or equal to the line rate.

Proxy signaling enables a “smart device” to signal on behalf of a “dumb” device (see Figure 13-1). It allows high-end devices with multiple network interface cards (NICs) to use a single signaling channel. In this instance, you use proxy signaling to allow a single signaling entity (PSA) to signal on behalf of multiple, non-signaling endsystems (PSC). This application extends the ATM signaling protocol to endsystems that do not necessarily understand ATM signaling.

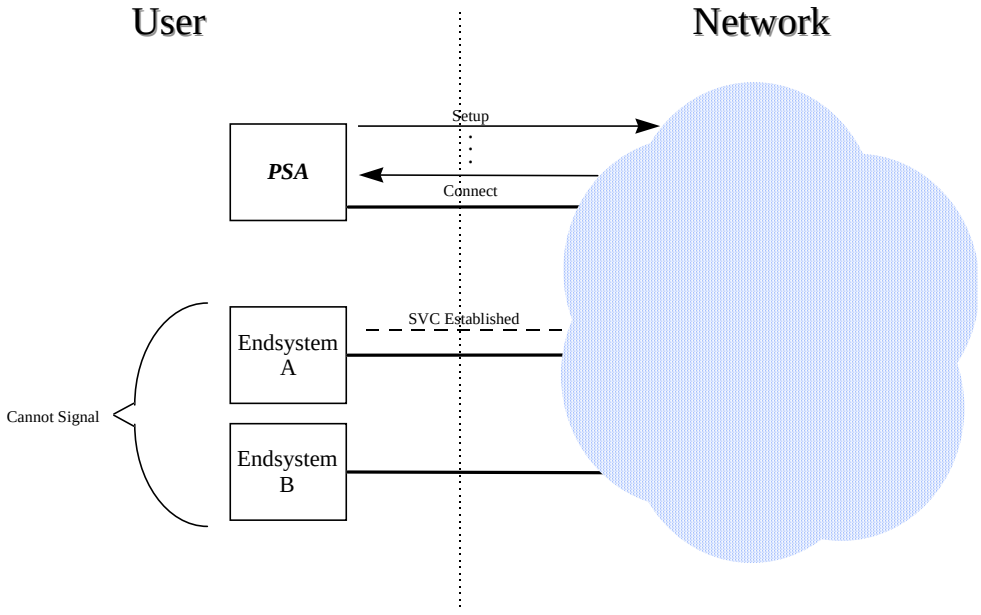

Figure 13-1. Establishing SVC for Endsystem via Proxy Signaling Agent

Proxy Signaling Agent

To define SVC proxy signaling functions, you must first configure the ILMI/Signaling/OAM attributes for an ATM UNI DCE logical port, and set the Proxy Admin Status to Agent. Signaling must be enabled (see "ILMI/Signaling/OAM Attributes" on page 3-15).

Acting as the PSA, the UNI DCE port uses VPI/VPCI mapping to determine if a particular SVC request is destined for a PSC. With UNI 4.0 signaling, VPCI mapping provides an alias that represents the PSC's logical port and VPI address. Each PSC needs a unique VPCI. For example, using the Add VPCI Attributes dialog box (Figure 13-4 on page 13-6), configure the following VPCI mapping on a PSA port:

- VPCI 1 means VPI 0 on port 12
- VPCI 2 means VPI 0 on port 13

If the SVC request does not include a VPCI (UNI 3.X signaling), the PSA port performs a routing lookup on the calling party address to determine the appropriate PSC. It matches the calling party address to a logical port, and then uses the VPCI that corresponds to the logical port.

Proxy Signaling Client

To define SVC proxy signaling functions, you must first modify the ILMI/Signaling/OAM attributes for a UNI DCE logical port, and set the Proxy Admin Status to Client. Signaling is disabled for this UNI DCE logical port. For each PSC, you select the switch/logical port combination that represents the controlling PSA.

If you are using Ascend's VNN trunk protocol, clients and agents may reside on different switches; if you are using PNNI, clients and agents must reside on the same switch.

VPCI/SVC Address Association

There is no direct association between VPCIs and SVC addresses. The SVC address can be associated with a VPCI because the address is configured on a logical port that corresponds to the VPCI. For example, if VPCI 1 represents VPI 3 on logical port 4 and logical port 4 is configured for SVC address 5085551212, then the address 5085551212 is implicitly associated with VPCI 1.

With this configuration, an incoming SVC request at the PSA port that specifies VPCI 1 is set up on logical port 4 (proxy on VPCI); an SVC request with no VPCI selected and a calling party address of 5085551212 is also set up on logical port 4 (proxy on calling party).

Configuring the Proxy Signaling Agent

To configure the PSA:

1. Use the instructions beginning on page 3-2 to access the Set All Logical Ports dialog box.
2. Use the instructions beginning on page 3-6 to configure an ATM UNI DCE logical port. Complete the sections "Administrative Attributes" (page 3-9) and "ATM Attributes" (page 3-11).
3. Use the Set Attributes pull-down menu to access "ILMI/Signaling/OAM Attributes" on page 3-15.
4. Set the Signaling Admin Status to Enabled.
5. Select Agent for the Proxy Admin Status. The Proxy Signaling attributes appear as shown in Figure 13-2.

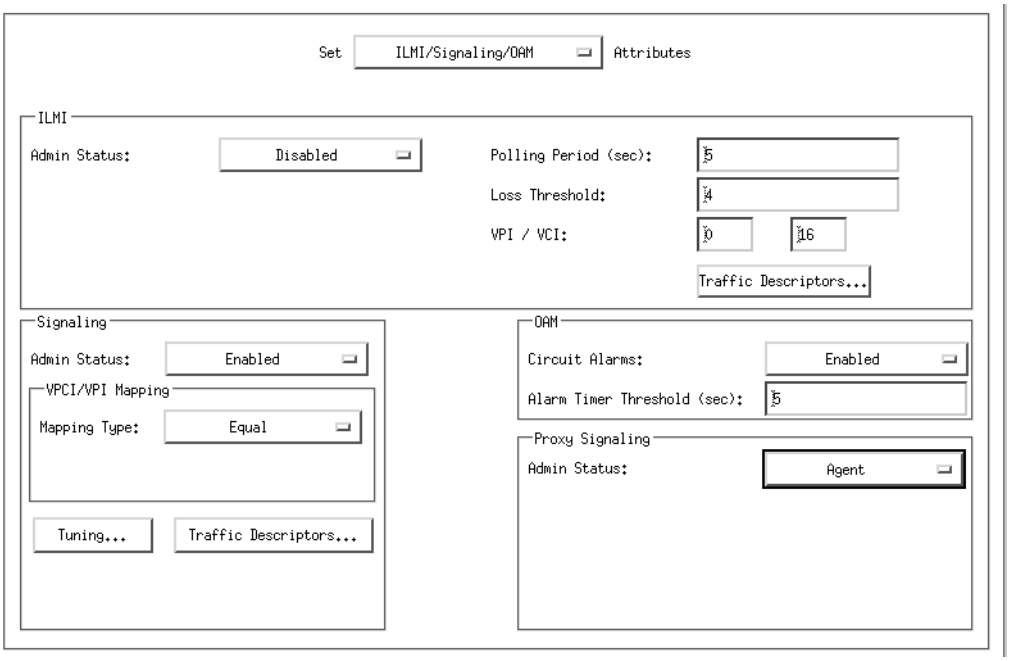

Figure 13-2. Set ILMI/Signaling/OAM Attributes (Agent)

6. Continue with the instructions beginning on page 3-16 to configure the remaining logical port attributes (including SVC Attributes) and return to the Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-2).

Once you have defined these basic logical port attributes, you must modify this logical port to define the VPCI table. Continue with the following section, "Configuring the VPCI Table."

Configuring the VPCI Table

To configure the VPCI table:

1. From the Set All Logical Ports in PPort dialog box (Figure 3-1 on page 3-2), select the PSA logical port and choose Modify.
2. Choose OK to acknowledge the logical port type.
3. Use the Set Attributes pull-down menu to access ILMI/Signaling/OAM Attributes.
4. In the VPCI/VPI Mapping box, select Mapping Type Table.
5. Choose the Show Table button to display the following dialog box.

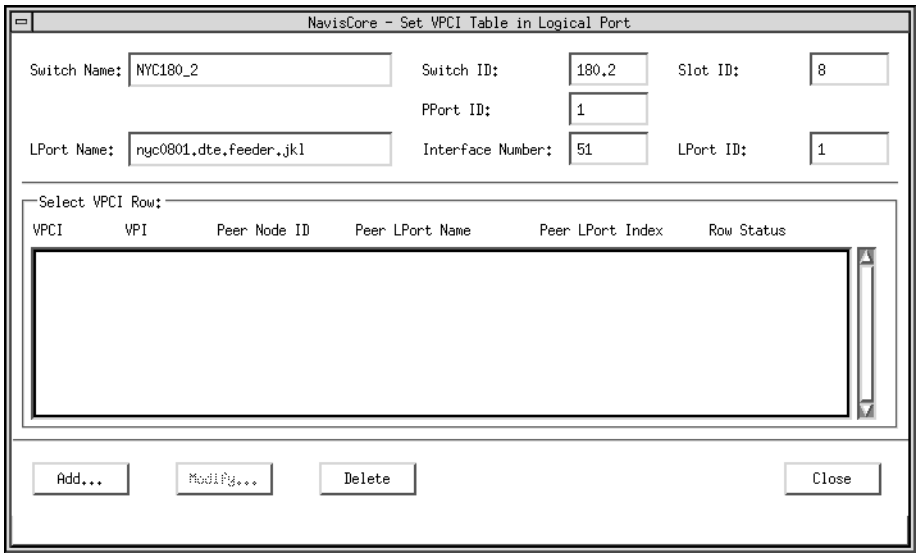

Figure 13-3. Set VPCI Table in Logical Port Dialog Box

The Set VPCI Table in Logical Port dialog box presents a series of screens that enable you to add or delete a row in the VPCI table. Each row in the VPI/VPCI table corresponds to the VPI/VPCI of the switch/logical port combination that represents the proxy signaling client.

6. Choose Add to create a VPI/VPCI combination. The following dialog box appears.

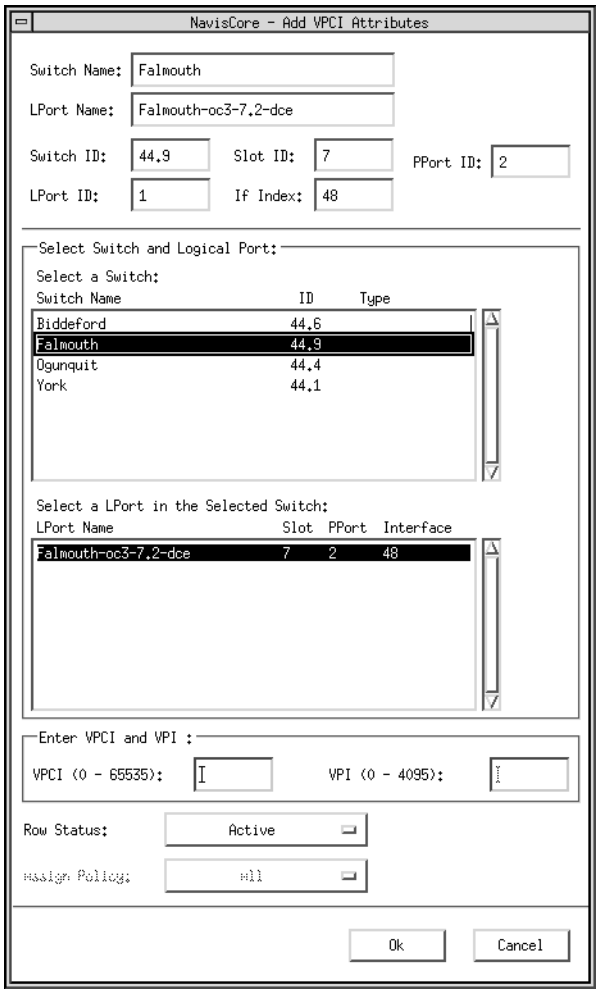

Figure 13-4. Add VPCI Attributes

7. Select the switch and logical port combination that identifies the PSC.

8. Enter values for the VPCI and VPI that identify the PSC.

9. Select Row Status. *Active* assigns this VPCI/VPI combination to the PSC; *Inactive* saves this VPCI/VPI combination in the database for future use.

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

11. Choose OK to complete the PSA table configuration.

Configuring the Proxy Signaling Client

To configure the PSC:

1. Use the instructions beginning on page 3-2 to access the Set All Logical Ports dialog box.
2. Use the instructions beginning on page 3-6 to configure an ATM UNI DCE logical port. Complete the sections "Administrative Attributes" (page 3-9) and "ATM Attributes" (page 3-11).
3. Use the Set Attributes pull-down menu to access the ILMI/Signaling/OAM Attributes (see Figure 13-5).
4. Set the Signaling Admin Status to Disabled (*default*).
5. Select Client for the Proxy Admin Status. The Proxy Signaling attributes appear as shown in Figure 13-5.

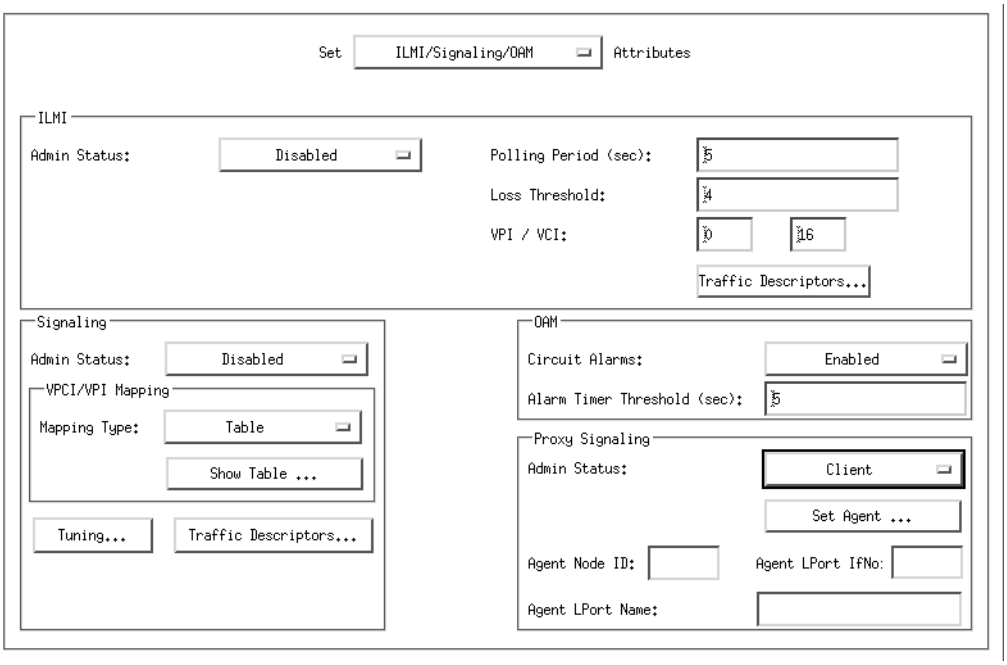

Figure 13-5. Set ILMI/Signaling/OAM Attributes (Client)

6. Choose Set Agent. The following dialog box appears.

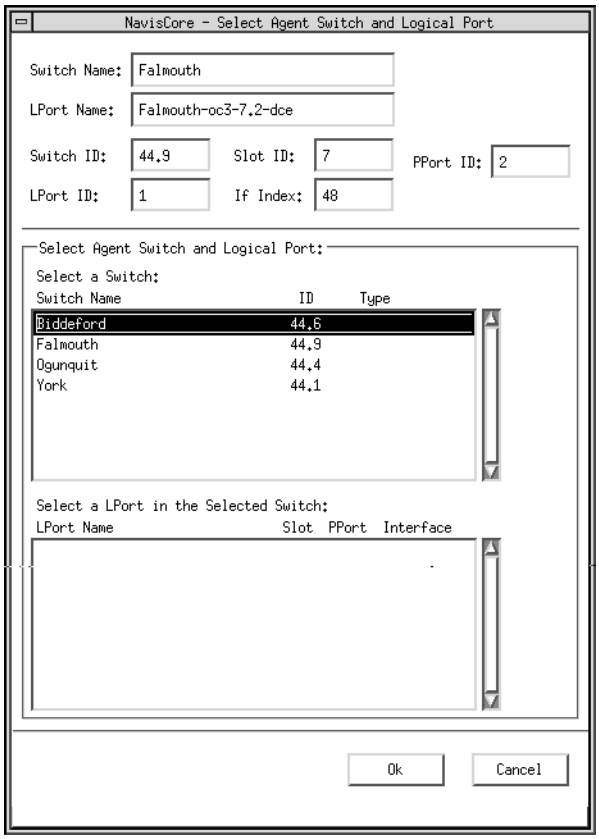

Figure 13-6. Select Agent Switch and Logical Port Dialog Box

7. Select the switch and logical port combination that identifies the PSA.

8. Choose OK to return to the Add Logical Port dialog box.

9. Continue with the instructions beginning on page 3-16 to configure the remaining logical port attributes. You do not need to configure SVC signaling parameters for this logical port.

14

Configuring SPVCs

A permanent virtual circuit (PVC) is established administratively (that is, by network management) rather than on demand (i.e., using signaling across the UNI). A soft PVC (SPVC) is established by the network using signaling. Once the SPVC configuration is in place, the switch at one end of the SPVC initiates the signaling. This release supports up to 4096K SPVCs per card.

The network management system provisions one end of the SPVC with the address identifying the egress interface from the network. The calling end has the responsibility for establishing, releasing, and re-establishing the call.

The B-STDX does not support the SPVC feature. You also cannot configure SPVCs on the 6-port Frame-based DS3 module.

About SPVCs

There are two types of ATM virtual connections: virtual channel connections (VCCs) and virtual path connections (VPCs). These virtual connections are made up of a series of virtual links which form a path between two endpoints. Based on the type of virtual connection you are using (VCC or VPC), you can create either a soft permanent virtual channel connection (SPVCC) or a soft permanent virtual path connection (SPVPC).

When working with SPVCs, you can configure a connection that is point-to-point or point-to-multipoint. In a point-to-multipoint configuration, the CBX/GX endpoint defined as the root can access several terminating endpoints (configured as “leaves”).

When you create an SPVC, you configure one endpoint (known as the *originating endpoint*) as you would a PVC. You select the logical port on which the endpoint will reside, and assign a virtual path identifier/virtual channel identifier (VPI/VCI) value. You configure the other endpoint(s) (*terminating endpoints*) with addresses, as you would an SVC. Optionally, you may also specify the remote VPI/VCI values. The originating endpoint uses signaling to access the terminating endpoints.

Using PVC/PVP Termination

Before you can configure SPVCs, you must first configure the SVC address or prefix you want to assign to the SPVC terminating endpoint. This endpoint may not actually terminate the SPVC. When you configure an SVC port address, you enable or disable PVC/PVP termination. If you disable termination, the egress logical port signals the SPVC on as a regular SVC.

PVC and PVP termination enable you to send calls through the network to a non-SVC endpoint, using an SVC. Table 14-1 on page 14-3 shows the results of using PVC/PVP termination.

As you configure PVC/PVP termination, keep the following points in mind:

- If you enable PVC termination, you can optionally specify a VPI/VCI or allow the SPVC originator or the network to choose a VPI/VCI. The switch terminates the SPVCC on the logical port that is associated with the VPI/VCI, and the traffic then continues on the local PVC segment.
- If you enable PVP termination, you can optionally specify a VPI or allow the SPVC originator or the network to choose a VPI, and the CBX 500 terminates the SPVPC on the associated logical port.
- If you enable both, you must allow the SPVC originator or the network to select the VPI/VCI or VPI.

For more information about configuring PVC/PVP Termination on the SVC, see page 12-31.

Specifying the Target Select Type

The originating endpoint may optionally specify the remote VPI or VPI/VCI for an SPVC. This feature is called the “target select type”. A target select type of “any” means that the appropriate VPI or VPI/VCI has been locally configured at the terminating endpoint or that the network is free to select a VPI or VPI/VCI.

A target select type of “specified” means that the terminating endpoint is obligated to use a specific VPI or VPI/VCI, as determined by the originating endpoint. This information is propagated by signaling. However, use of the “specified” target select type has the following limitations:

- You have Ascend equipment at both the originating and terminating endpoints. As long as this is the case, the connecting portion of the network can contain network equipment from any vendor, using any protocol.
- You only have Ascend equipment at one endpoint, but the SPVC traverses only Ascend Virtual Network Navigator (VNN) or PNNI links. Some LAN-based ATM networks currently support the PNNI protocol.
- If the SPVC must traverse UNI or IISP links and one end of the SPVC is not Ascend equipment, you cannot use the specified target select type.

Table 14-1 summarizes the results of using SPVC target select type in conjunction with PVC/PVP termination.

Table 14-1. SPVC Target Select Type

Originating Endpoint Target Select Type	Terminating Endpoint Termination Type	Behavior at Terminating Endpoint
Any	Any	Network allocates any available VPI or VPI/VCI.
Any	Specified VPI or VPI/VCI	Accept SPVC on a specified VPI or VPI/VCI. The SVC port address is dedicated to terminating this single SPVC.
Specified VPI or VPI/VCI	Any	Accept SPVC on a specified VPI or VPI/VCI; the SVC port address may terminate additional SPVCs.
Specified VPI or VPI/VCI	Specified VPI or VPI/VCI	Accept SPVC if VPI or VPI/VCI match; reject SPVC if they do not. The SVC port address is dedicated to terminating this single SPVC.

About Point-to-Point SPVCs

To access the Set All Point-to-Point SPVCs dialog box, from the Administer menu, select Ascend Parameters ⇒ Set All Soft PVCs ⇒ Point-to-Point. The following dialog box appears.

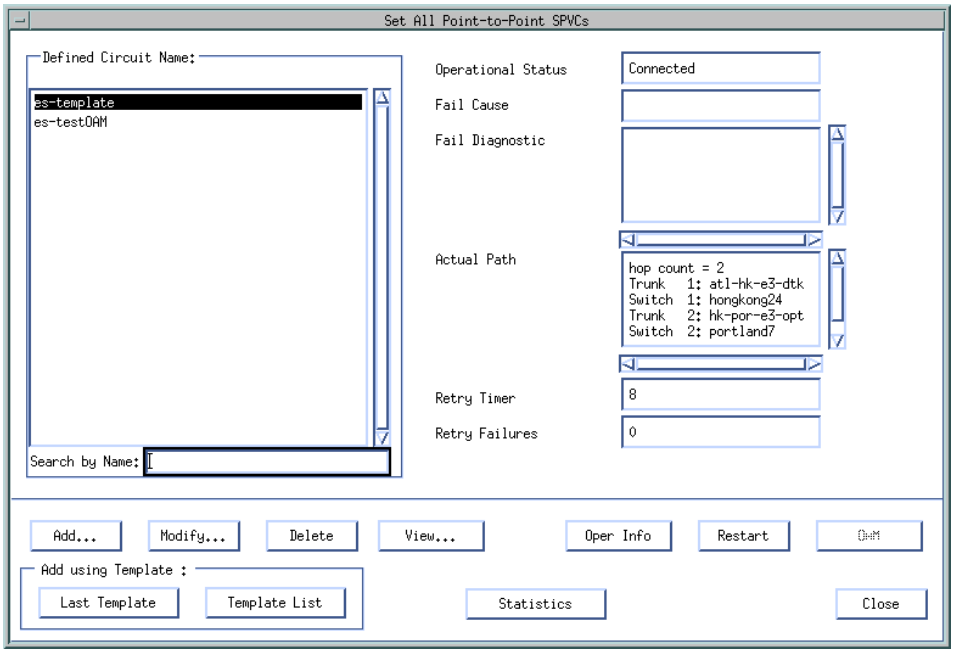

Figure 14-1. Set All Point-to-Point SPVCs

To view a list of configured SPVCs, position the cursor in the Search by Name field and press Return. To use a wildcard search to find a specific SPVC 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 \\

The Set All Point-to-Point SPVCs dialog box displays information about the configured options for the Defined Circuit Name you select. Table 14-2 describes the dialog box status indicators and commands.

Table 14-2. Set All Point-to-Point SPVCs Dialog Box Fields and Buttons

Field/Command	Action/Description
Operational Status	Displays a status message: Establishing, Connected, Failed, or Other.
Fail Cause	Displays the ID of the last release failure.
Fail Diagnostic	Displays an eight-character diagnostic of the last release failure (as applicable).
Actual Path	Displays a character string that represents the actual Ascend VNN path the SPVC used.
Retry Timer	Displays the current value of the retry timer in seconds.
Retry Failures	Displays the number of failed attempts to establish a connection.
Add	Enables you to add an SPVC.
Modify	Enables you to modify an SPVC.
Delete	Enables you to delete an SPVC.
View	Enables you to view the configured attributes for the selected SPVC.
Oper Info	Updates the information in the following Operational Status fields: • Fail Cause ID • Fail Diag • Actual Path • Retry Time • Retry Failure
Restart	Restarts the selected SPVC no matter what condition it is in. If the SPVC is not established, it restarts the process to establish the connection. If it is established, this command clears the existing connection and establishes it again.
OAM	Runs the Operations, Administration, and Management loopback diagnostics for the selected circuit.
Statistics	Displays the summary statistics for the selected SPVC.

Table 14-2. Set All Point-to-Point SPVCs Dialog Box Fields and Buttons

Field/Command	Action/Description
Add using Template	If you have already defined an SPVC configuration and saved it as a template, use this command to define a new circuit using similar parameters. • Choose Last Template to use the last template you used to establish a circuit in this NMS session. • Choose Template List to display a list of templates previously defined for this map.
Close	Exit this dialog box and return to the network map.

Setting the VPI/VCI Values for SPVCs

For each SPVC you configure, you must specify a value from 0 – *nnnn* to represent the VPI for the SPVC. The maximum value is based on the Valid Bits in VPI that is configured for the logical port, as follows:

Maximum value = $2^P - 1$

where *P* is the value in the Valid Bits in VPI field. For example, if you entered 5 in the Valid Bits in VPI field, the maximum value is 31 ($2^5 - 1 = 31$) which would give you up to 32 virtual paths (numbered 0-31).

If you are defining a soft permanent virtual channel connection (SPVCC), you must also specify a value to represent the VCI for an ATM circuit. The maximum value is based on the Valid Bits in VCI value that is configured for the logical port, as follows:

Maximum value = $2^C - 1$

where *C* is the value in the Valid Bits in VCI field. For example, if you entered 6 in the Valid Bits in VCI field, the maximum VCI value you can enter is 63 (which would give you 32 virtual channels, numbered 32 to 63).

These VPI/VCI range restrictions only apply to SPVCCs. You can provision SPVPCs to use the following values:

- For UNI, use the VPI=0-255 range.
- For NNI cell header format, use the VPI=0-4095

For more information on the Valid Bits in VPI/VCI fields, see page 3-14.

Adding an SPVC

To add an SPVC:

1. From the Set All Point-to-Point SPVCs dialog box (Figure 14-1 on page 14-4), choose Add. The following dialog box appears.

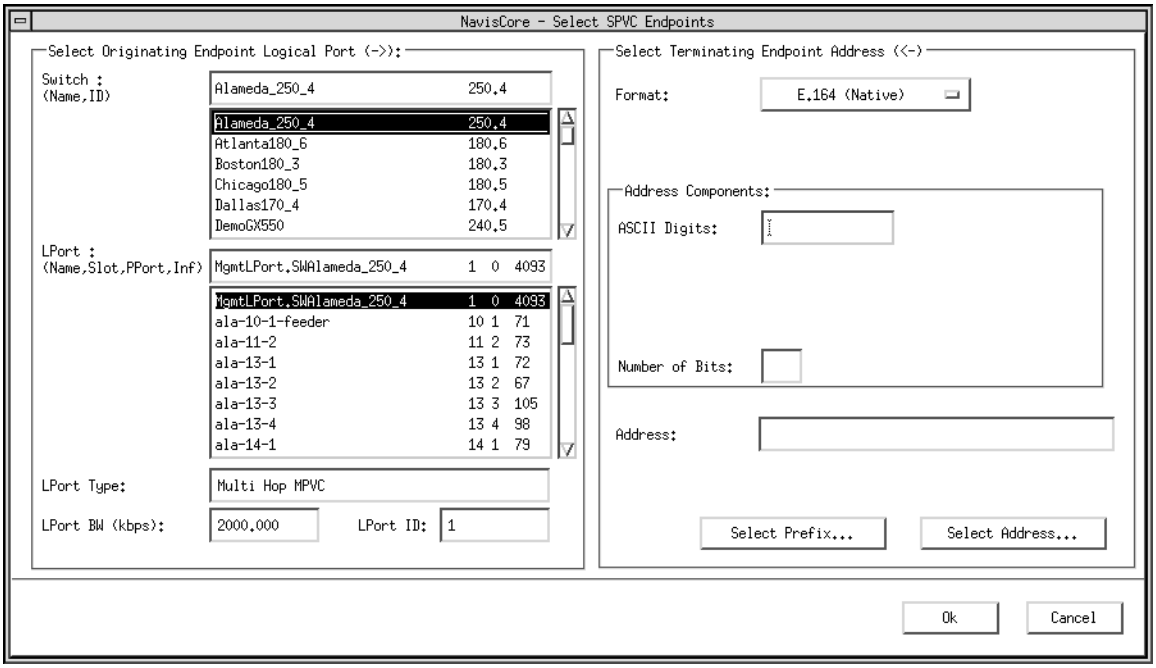

Figure 14-2. Select SPVC Endpoints Dialog Box

2. Use the following steps to configure the originating endpoint logical port.
   a. Select the name of the switch on which the endpoint will reside.
   b. Select the name of the logical port that will be used as the endpoint. The dialog box displays the logical port type, bandwidth, and ID for this endpoint.
3. To complete this configuration:
   - If you know the SVC terminating endpoint address, use Table 14-3 on page 14-8 to select the address format and configure the terminating endpoint address. For more information on AESA address formats, see page 11-2.
   - If you do not know this address, or if you need to configure the terminating endpoint address, see page 14-12 for instructions on using the Select Prefix and Select Address buttons to configure this address.

Table 14-3. Configuring the SPVC Terminating Endpoint Address

Address Format	Address Components
E.164 (Native)	In the ASCII Digits field, enter all of the 1-15 ASCII digits that represent the E.164 address. The value you enter is then converted to the ASCII hex values that represent each digit in the number (this value is displayed in the Address).
DCC and ICD AESA (or anycast)	In the Hex Digits field, enter the data country code (DCC) of the country in which the address is registered, or the international country designator (ICD) that identifies the international organization to which this address applies. DCCs and ICDs consist of 4 hex digits, and occupy two octets. Then enter the appropriate HO-DSP, ESI and SEL values.
E.164 AESA (or anycast)	• In the Hex Digits field, enter the full or partial E.164 AESA address. Since the IDI portion of the address is 8 octets (16 hex digits), but the E.164 address format is a maximum of 15 digits, you must terminate the IDI portion with Fh. For example, 5085551234 should be entered as 000005085551234F. • After typing in the IDI portion of the address, enter the appropriate HO-DSP, ESI, and SEL portions to complete the address.
Custom AESA	• In the AFI Digits field, enter the custom AFI you want to use. • In the Hex Digits field, enter the customized address format you want to use, starting with the HO-DSP, and followed by the ESI and SEL values (in that order). This address must be the full 19 octets (38 hex digits) long, with 12 octets used for the HO-DSP, 6 octets used for the ESI, and 1 octet used for the SEL.

4. Choose OK. The following dialog box appears (Figure 14-3).

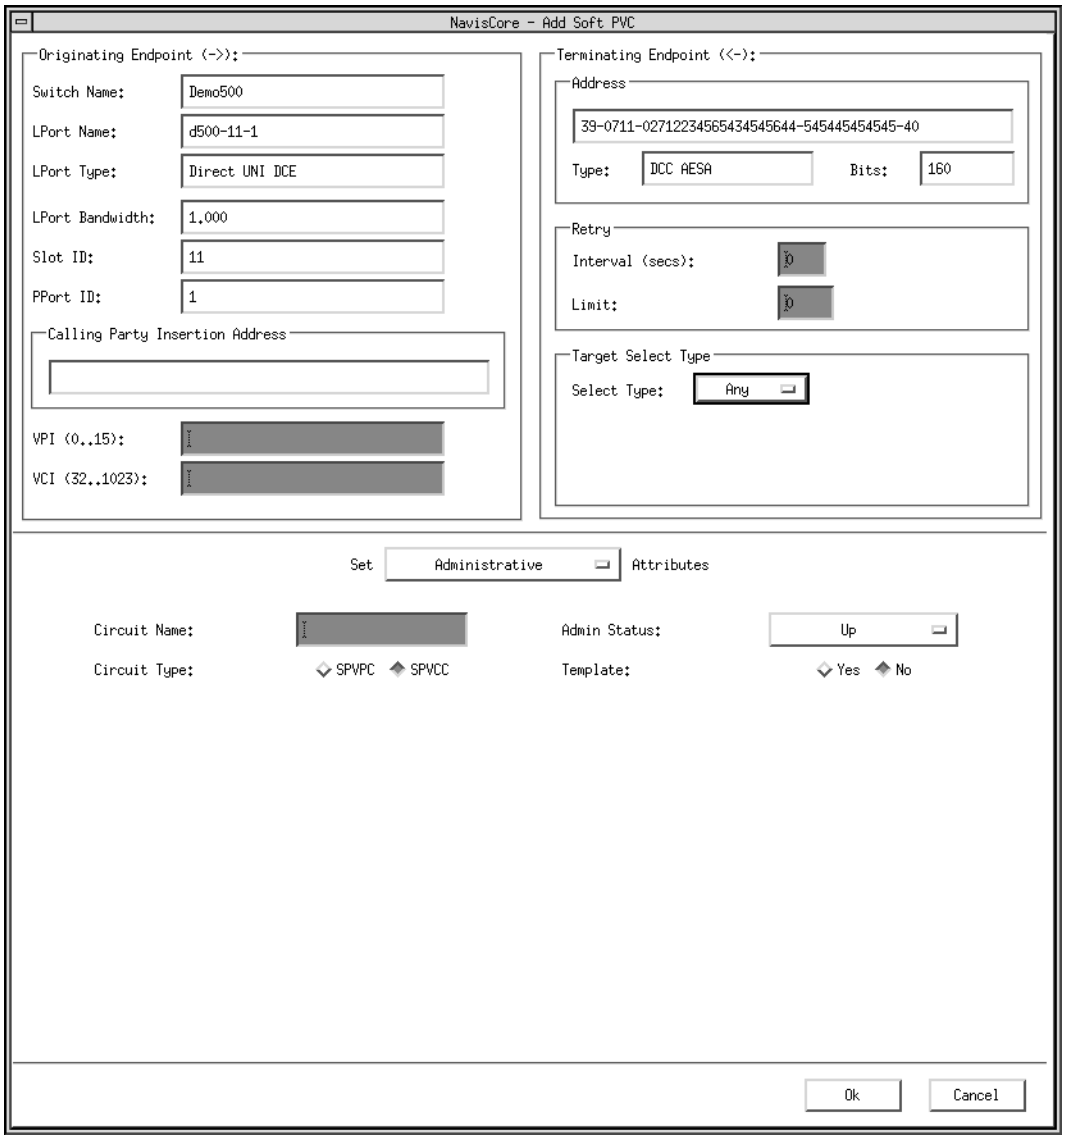

Figure 14-3. Add Soft PVC Dialog Box

5. Use the instructions in Table 14-4 to configure these fields.

Table 14-4. Add Soft PVC Dialog Box Fields

Field	Action/Description
VPI	Enter a value from 0 – *nnnn* to represent the VPI for the SPVC; enter a value from 0 – 4095 if you use the NNI cell header format. The maximum value you can enter is based on the Valid Bits in VPI that is configured for the logical port. See page 14-6 for more information.
VCI (*SPVCCs only*)	If you are configuring a soft permanent virtual channel connection (SPVCC), enter a value to represent the VCI for the SPVCC.

Table 14-4. Add Soft PVC Dialog Box Fields (Continued)

Field	Action/Description
Retry Interval	The originating endpoint makes several attempts to connect to the terminating endpoint. This value indicates the number of seconds the originating endpoint waits before trying to reestablish a connection. Specify this interval in seconds (1 – 3600), or enter 0 (default) for no retries.
Retry Limit	This value indicates the number of times the originating endpoint tries to connect to the terminating endpoint. Specify the number of retries (1 – 65535), or enter 0 (default) for an unlimited number of retries.
Target Select Type	Review the information in "Specifying the Target Select Type" on page 14-3 first to determine your network needs. Then Select one of the following Types: *Any* – Indicates that the terminating endpoint uses any available VPI/VCI value. If you need to specify a VPI/VCI for the terminating endpoint, you must complete the PVC/PVP Termination fields on the Add SVC Port Address dialog box. See page 12-31. *Specified* – The terminating endpoint uses the VPI/VCI address you specify. If this is an SPVPC, enter the VPI; for an SPVCC, enter the VPI and VCI. 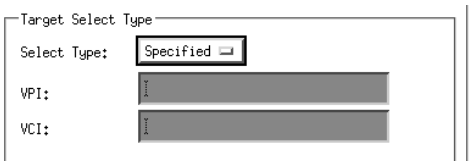
Circuit Name	Enter any unique, alphanumeric name to identify the SPVC. Do not use parentheses and asterisks. This name must be unique to the entire map.
Admin Status	Select Up (the default) to activate the SPVC at switch start-up, or Down if you do not want to activate the SPVC at switch start-up.
Circuit Type	Specify whether the circuit is a soft permanent virtual path connection (SPVPC) or SPVCC (the default). If you select SPVPC, the VCI field is set to 0 and cannot be changed.
Template *(Optional)*	You can save these settings as a template to configure another SPVC with similar options. To create a template, choose Yes in the "Is Template" field.

6. Choose Set [Traffic Type] Attributes to select the traffic descriptors for this SPVC. The following dialog box appears.

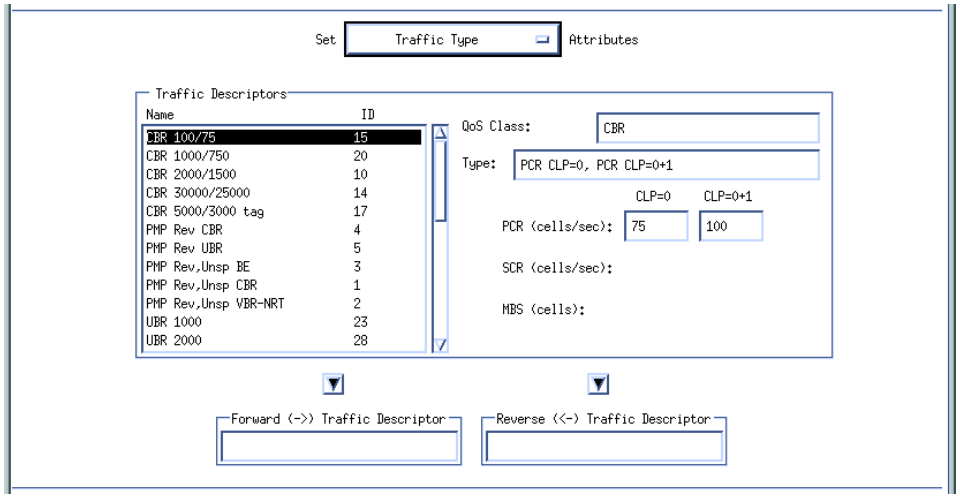

Figure 14-4. Add Soft PVC Dialog Box - [Set Traffic Type]

7. See "Defining Traffic Descriptor Attributes" on page 8-9 for instructions to configure these attributes.

8. Choose OK to create the new SPVC and return to the Set All Point-to-Point SPVCs dialog box (Figure 14-1 on page 14-4).

9. Choose Close to return to the network map.

Configuring the Terminating Endpoint Address

The Select SVC Endpoints dialog box (Figure 14-2 on page 14-7) contains the Select Prefix and Select Address buttons you can use to either select an existing or configure a new SVC terminating endpoint address.

Selecting an SVC Port Prefix

1. To select an SVC Port Prefix, choose Select Prefix. The following dialog box appears.

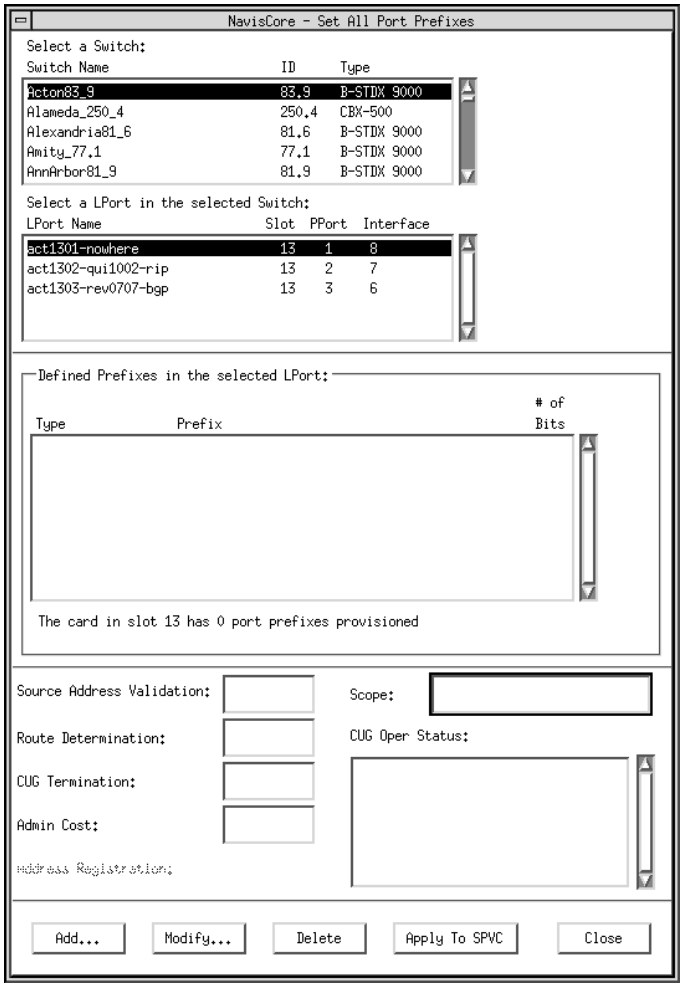

Figure 14-5. Set All Port Prefixes Dialog Box

- To add a new SVC prefix, choose Add and use the instructions beginning on page 12-12.
- To select an existing prefix, complete the following steps.

2. Select the name of the switch where the endpoint resides.
3. Select the name of the logical port that will be used as the endpoint. The dialog box displays the address information for this endpoint.

4. Choose Apply to SVC. The Select SVC Endpoints dialog box (Figure 14-2 on page 14-7) reappears.
5. Enter the remaining digits of the port address. If necessary, use the following procedure, "Selecting an SVC Port Address," to select a preconfigured address or to configure a new one.
6. Continue with Step 4 on page 14-8.

Selecting an SVC Port Address

1. To select an SVC Port address, choose Select Address. The following dialog box appears.

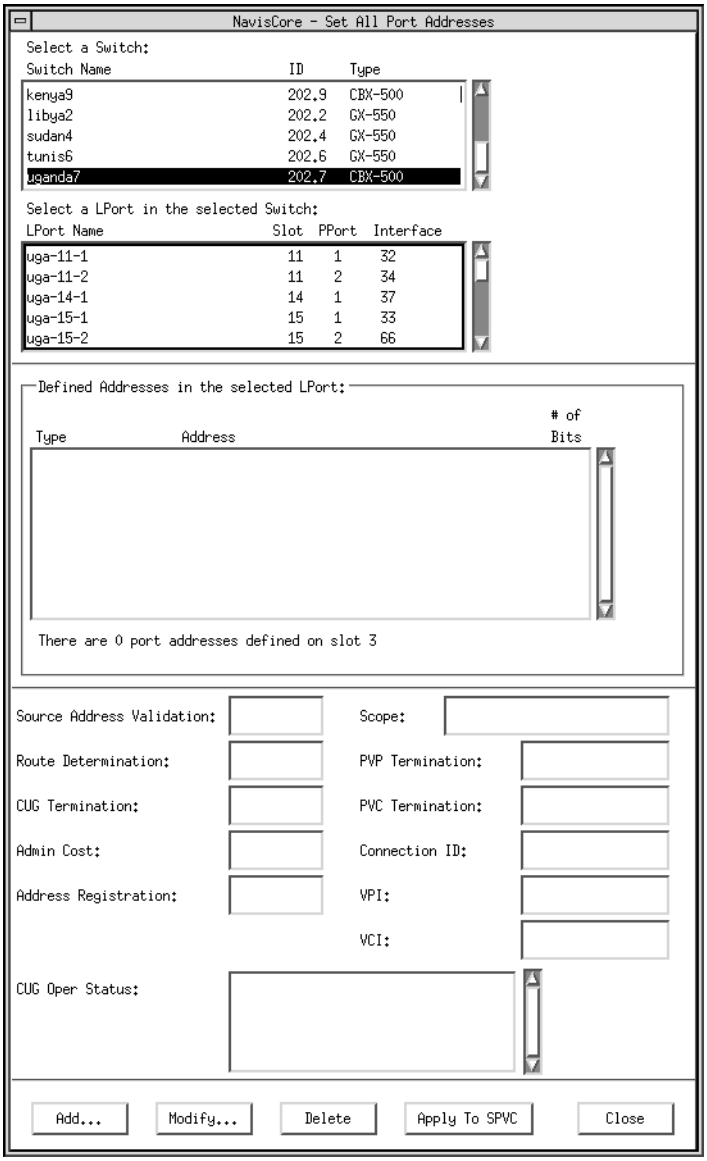

Figure 14-6. Set All Port Addresses

- To add a new SVC address, choose Add and use the instructions beginning on page 12-24.
- To select an existing prefix, complete the following steps.

2. Select the name of the switch where the endpoint resides.
3. Select the name of the logical port that will be used as the endpoint. The dialog box displays the address information for this endpoint.
4. Choose Apply to SVC. The Select SVC Endpoints dialog box (Figure 14-2 on page 14-7) reappears.
5. Continue with Step 4 on page 14-8.

Configuring Point-to-Multipoint SPVCs

To access the Set All Point-to-Multipoint SPVCs dialog box, from the Administer menu, select Ascend Parameters ⇒ Set All Soft PVCs ⇒ Point-to-Multipoint. The following dialog box appears.

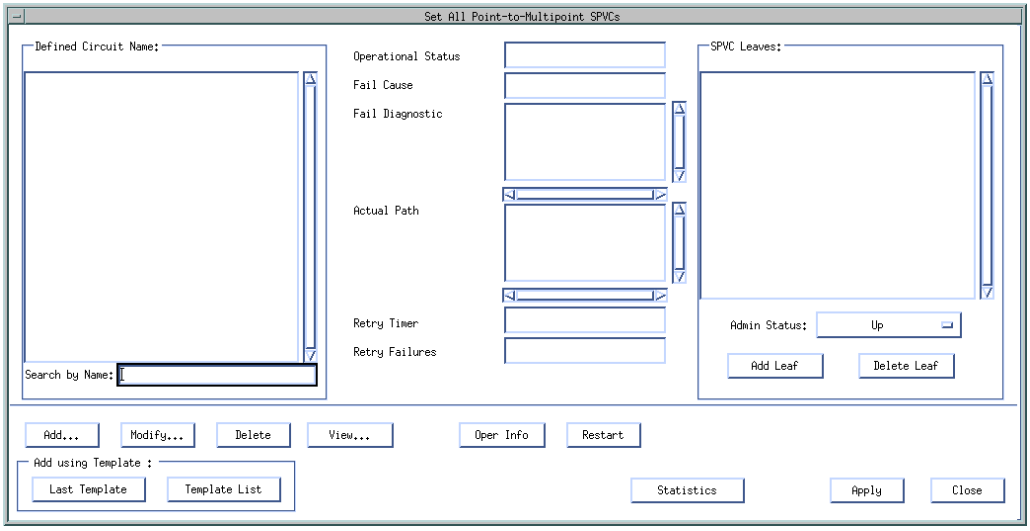

Figure 14-7. Set All Point-to-Multipoint SPVC Dialog Box

To view a list of configured Point-to-Multipoint SPVCs, position the cursor in the Search by Name field and press Return. To use a wild card search to find a specific SPVC name, you can do any of the following:

- 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 \\

The Set All Point-to-Multipoint SPVCs dialog box displays information about the configured options for the Defined Circuit Name you select. Table 14-2 on page 14-5 describes the dialog box status indicators and commands.

Adding a Point-to-Multipoint SPVC

When you configure a Point-to-Multipoint SPVC, you first define an SPVC consisting of a root (originating endpoint) and one leaf (terminating endpoint). This procedure is similar to the one for creating Point-to-Point SPVCs. Once you define the initial root/leaf combination, you can create additional leafs.

Defining the Point-to-Multipoint SPVC Root

1. From the Set All Point-to-Multipoint SPVC dialog box, choose Add. The Select SPVC Endpoints dialog box appears (Figure 14-2 on page 14-7).
2. Configure the originating endpoint logical port (root):
   a. Select the name of the switch on which the endpoint will reside.
   b. Select the name of the logical port that will be used as the endpoint.

      The dialog box displays the logical port type, bandwidth, and ID for this endpoint.
3. If you know the SVC terminating endpoint address, use Table 14-3 on page 14-8 to select the address format and configure the terminating endpoint address. For more information on AESA address formats, see page 11-2.

   If you do not know this address, or if you need to configure the terminating endpoint address, see page 14-12 for instructions on using the Select Prefix and Select Address buttons to configure this address.
4. Choose OK. The Add Soft PVC dialog box appears (Figure 14-3 on page 14-9).
5. Use the instructions in Table 14-5 to configure these fields.

Table 14-5. Add Point-to-Multipoint Dialog Box Fields

Field	Action/Description
VPI	Enter a value from 0 – *nnnn* to represent the VPI for the SPVC. Enter a value from 0 – 4095 if you use the NNI cell header format. The maximum value you can enter is based on the Valid Bits in VPI that is configured for the logical port. (See page 14-6 for more information.)
VCI (*SPVCCs only*)	If you are configuring a SPVCC, enter a value to represent the VCI for the SPVCC.
Retry Interval	The originating endpoint makes several attempts to connect to the terminating endpoint. This value indicates the number of seconds the originating endpoint waits before trying again to establish a connection. Specify this interval in seconds (1 – 3600), or enter 0 (default) for no retries.

Table 14-5. Add Point-to-Multipoint Dialog Box Fields (Continued)

Field	Action/Description
Retry Limit	This value indicates the number of times the originating endpoint tries to connect to the terminating endpoint. Specify the number of retries (1 – 65535), or enter 0 (default) for an unlimited number of retries.
Target Type	Review the information in "Specifying the Target Select Type" on page 14-3 first to determine your network needs. Then Select one of the following Types: *Any* – Indicates that the terminating endpoint uses any available VPI/VCI value. If you need to specify a VPI/VCI for the terminating endpoint, you must complete the PVC/PVP Termination fields on the Add SVC Port Address dialog box. See page 12-31. *Specified* – The terminating endpoint uses the VPI/VCI address you specify. If this is an SPVPC, enter the VPI. For an SPVCC, enter the VPI and VCI. 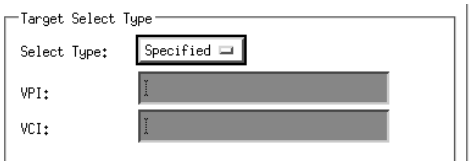
Circuit Name	Enter any unique, alphanumeric name to identify the SPVC. Do not use parentheses and asterisks. This name must be unique to the entire map.
Admin Status	Select Up (the default) to activate the SPVC at switch start-up, or Down if you do not want to activate the SPVC at switch start-up.
Circuit Type	Specify whether the circuit is an SPVPC or SPVCC (the default). If you select SPVPC, the VCI field is set to 0 and cannot be changed.
Template *(Optional)*	You can save these settings as a template to configure another SPVC with similar options. To create a template, choose Yes in the "Is Template" field.

6. See "Defining Traffic Descriptor Attributes" on page 8-9 for instructions on configuring these attributes.
7. Choose OK to create the new SPVC and return to the Set All Point-to-Multipoint SPVCs dialog box.

Defining Additional SPVC Leafs

1. Once you configure the circuit root and initial leaf, choose Add Leaf to configure additional SPVC leaves. The following dialog box appears.

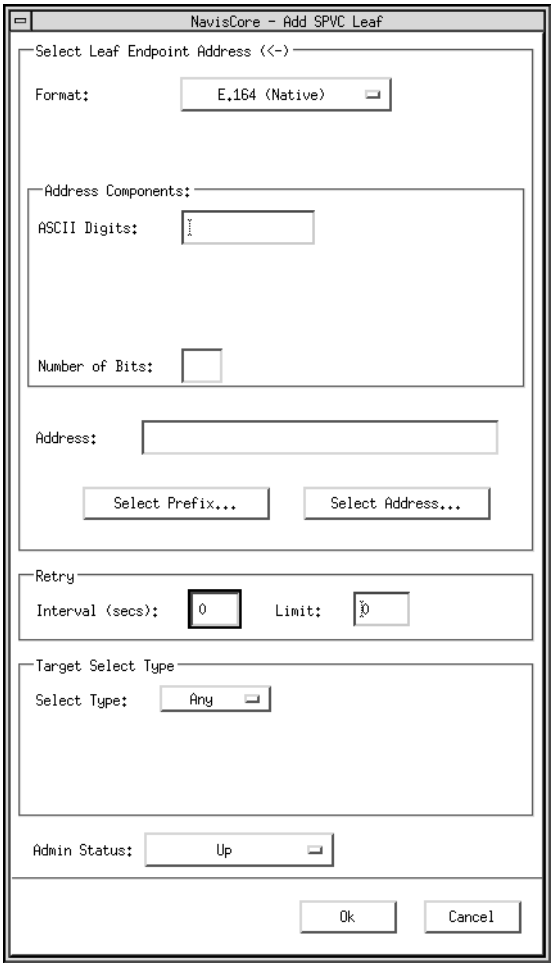

Figure 14-8. Add SPVC Leaf Dialog Box

2. If you know the SVC terminating endpoint address, use Table 14-3 on page 14-8 to select the address format and configure the address components for this leaf. For more information on AESA address formats, see "Address Formats" on page 11-2.

   If you do not know this address, or if you need to configure the terminating endpoint address, see page 14-12 for instructions on using the Select Prefix and Select Address buttons to configure this address.

3. Use the instructions in Table 14-6 to define these fields.

Table 14-6. Add SPVC Leaf Dialog Box Fields

Field	Action/Description
Retry Interval	The originating endpoint makes several attempts to connect to the terminating endpoint. This value indicates the number of seconds the originating endpoint waits before trying again to establish a connection. Specify this interval in seconds (1 – 3600), or enter 0 (default) for no retries.
Retry Limit	This value indicates that the number of times the originating endpoint tries to connect to the terminating endpoint. Specify the number of retries (1 – 65535), or enter 0 (default) for an unlimited number of retries.
Target Select Type	Review the information in "Specifying the Target Select Type" on page 14-3 first to determine your network needs. Then Select one of the following Types: *Any* – Indicates the terminating endpoint uses any available VPI/VCI value. If you need to specify a VPI/VCI for the terminating endpoint, you must complete the PVC/PVP Termination fields on the Add SVC Port Address dialog box. See page 12-31. *Specified* – The terminating endpoint uses the VPI/VCI address you specify. If this is an SPVPC, enter the VPI. For an SPVCC, enter the VPI and VCI. 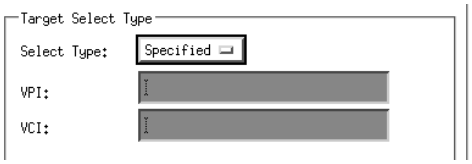

4. Choose OK to return to the Set All Point-to-Multipoint SPVC dialog box (Figure 14-7 on page 14-14).

15

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 (see page 12-23). You configure CUG member rules in either AESA or E.164 address format. When you configure a member rule, you can replace some digits with the * or ? UNIX wildcard characters. If a member rule does not contain a wildcard character, it maps to a specific network user. If the member rule includes a wildcard, 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 the 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 wildcard to replace a single digit. You can only use the * once in a string. Keep in mind that an AESA digit is 4 bits and an E.164 digit is 8 bits.

The following examples show how you can use wildcards 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 AESA), 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 which 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 also 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.

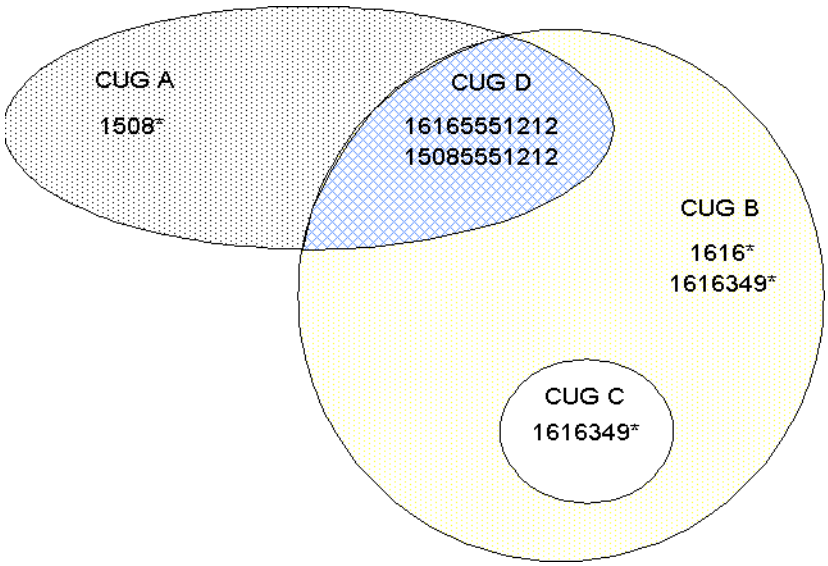

Figure 15-1. Implementing CUGs

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, Table 15-1 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 15-1. Incoming and Outgoing Calls Barred Example

	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 15088881212	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 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 15-1 on page 15-4, Table 15-2 illustrates how a single configured address can match multiple member rules, and can belong to more than one CUG.

Table 15-2. 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 can only 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 which 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 which 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 (see page 12-23).

Step 2. Define the CUG member rules that represent the member addresses and call access (see page 15-7).

Step 3. Define the CUG names (see page 15-10).

Step 4. Associate CUG members to specific CUGs. You can also modify call access attributes for a specific CUG (see page 15-11).

Defining CUG Members

A CUG member is defined by a rule that matches one or more port addresses/prefixes and attributes which 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 (see Figure 15-2 on page 15-8).

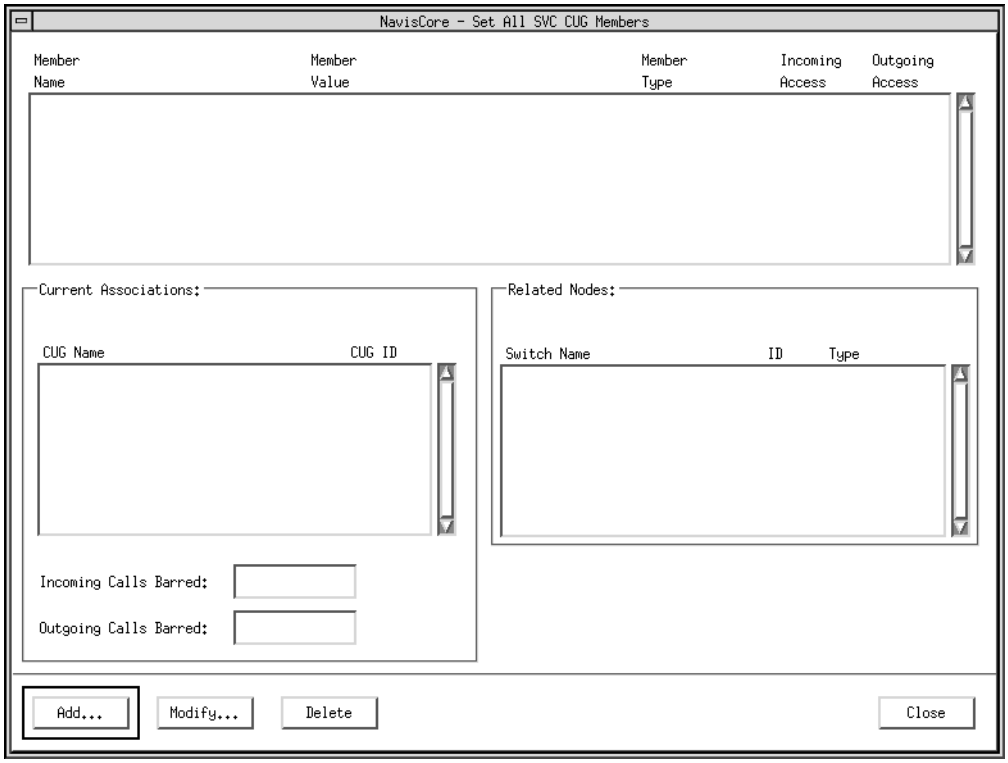

Figure 15-2. Set All SVC CUG Members Dialog Box

This dialog box provides a list of previously defined CUG members and their rules. For each member name you select, it provides the name of the CUG(s) with which the member is associated.

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

2. Choose Add to define a CUG member. The following dialog box appears.

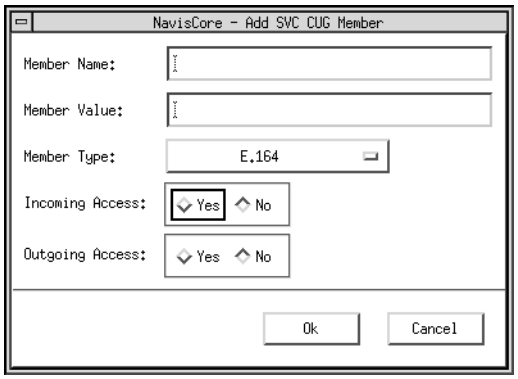

Figure 15-3. Add SVC CUG Member Dialog Box

3. Configure the member attributes as described in Table 15-3.

Table 15-3. Add SVC CUG Member Fields

Field	Action/Description
Member Name	Enter a name (up to 32 characters).
Member Value	Enter the CUG member rule using the guidelines on page 15-2. Do not enter more than 15 characters for an E.164 address or more than 40 characters for an AESA address.
Member Type	Select either AESA 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

Before you assign member rules, you need to 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 following dialog box appears.

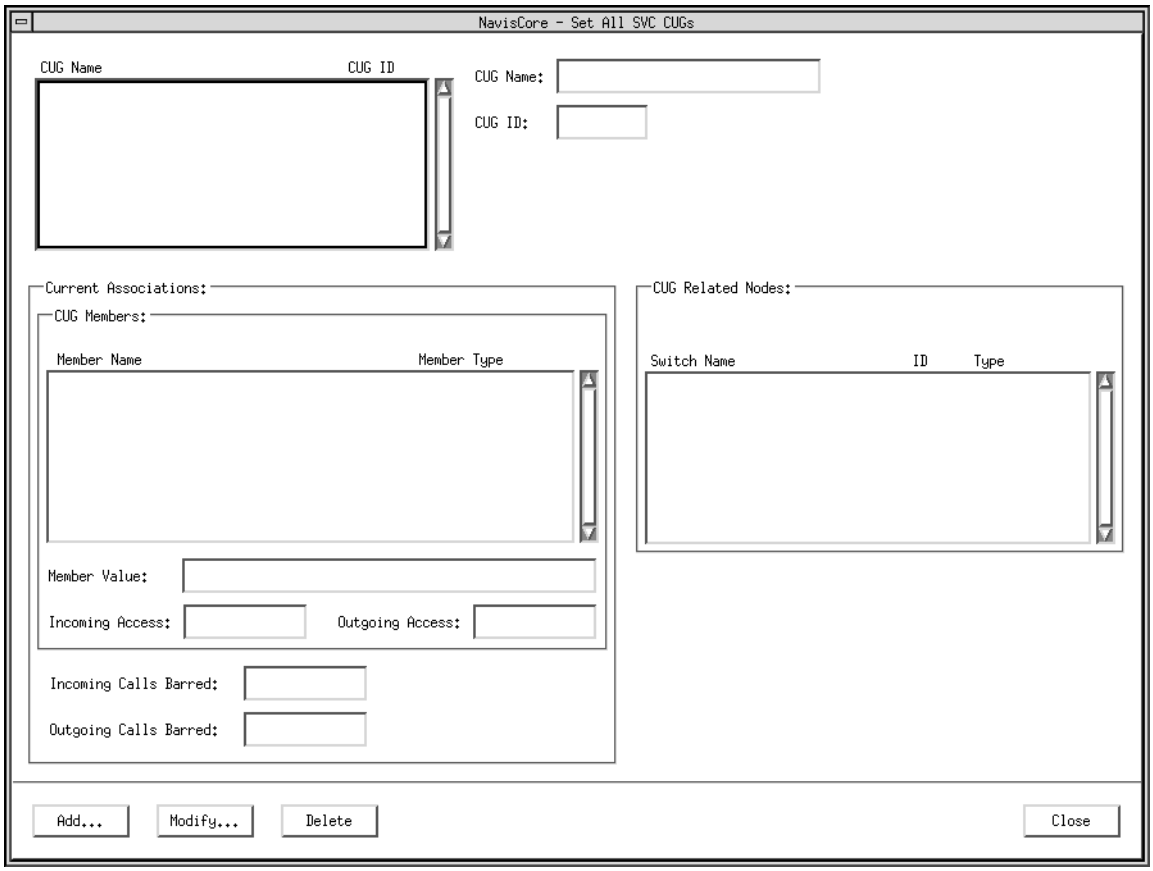

Figure 15-4. Set All SVC CUGs Dialog Box

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.

2. Choose Add to create a new CUG. The following dialog box appears.

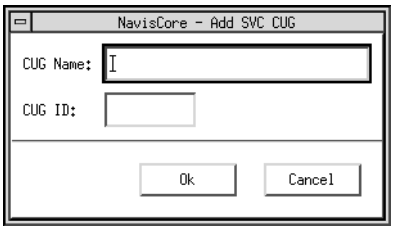

Figure 15-5. Add SVC CUG Dialog Box

3. Enter a CUG name (up to 32 characters). The NMS assigns a CUG ID.
4. 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 15-4 on page 15-10).
2. From the CUG Name list, select the CUG to which you want to add members.
3. Choose Modify. The following dialog box appears (Figure 15-6 on page 15-12).

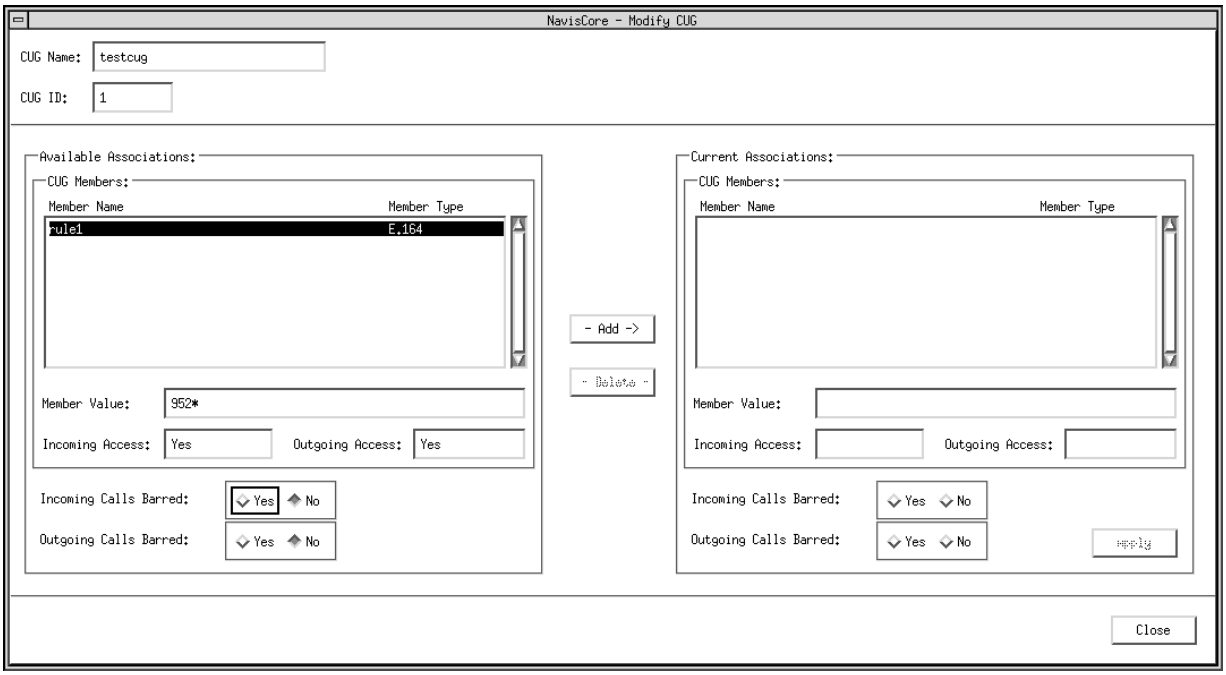

Figure 15-6. Modify CUG Dialog Box

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 the ICB and OCB options:

   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 15-4 on page 15-10).
2. From the CUG Name list, select the CUG that contains this member.
3. Choose Modify. The Modify CUG dialog box appears (Figure 15-6 on page 15-12).
4. Select the CUG member name from the Current Associations list.

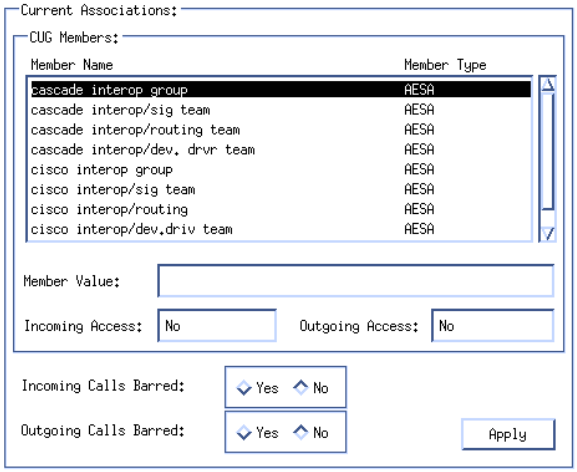

Figure 15-7. Modify CUG Dialog Box

5. Use the instructions on page 15-9 to modify incoming and outgoing call access.
6. Choose Apply. Choose Close to exit this dialog box.

16

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, E.164-AESA, DCC-AESA, ICD-AESA, and custom AESA.

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

Once you develop a set of screens, you can apply them to any ATM 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 16-1 describes the default ingress and egress security screen options. These defaults represent the port screen activation parameters.

Table 16-1. 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 16-1. 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 AESA-style address. You can enter the entire address as a number, or enter a UNIX-style expression using wildcards. When you use a UNIX expression, a single screen can match multiple endpoint addresses. Use the ? wildcard to replace a single digit or the * wildcard to replace one or more digits. You can only use the * once in a string. See “Address Formats” on page 11-2 for more information on 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 16-2 describes some examples using the port security screens.

Table 16-2. 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 (see Figure 16-3 on page 16-11), 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 the 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).

Step 2. Configure SVCs (see Chapter 12).

Step 3. Create a set of security screens (see page 16-7).

Step 4. Define the logical port security screening defaults. If necessary, assign the security screens that provide exceptions to these defaults (see page 16-10).

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 following dialog box appears.

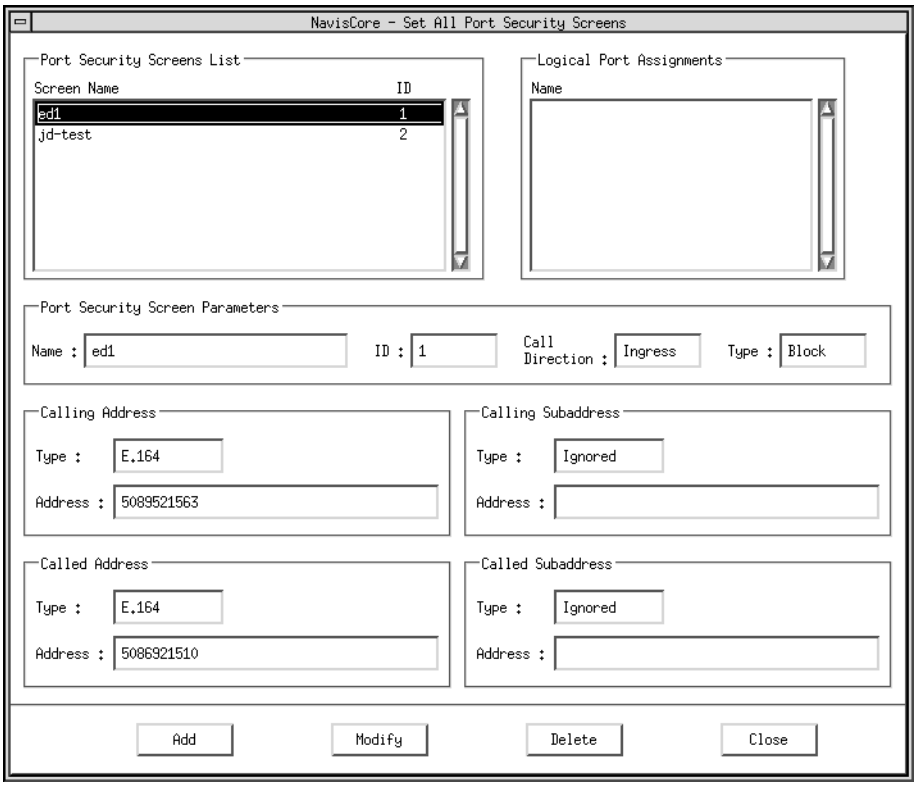

Figure 16-1. Set All Port Security Screens Dialog Box

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 following dialog box appears.

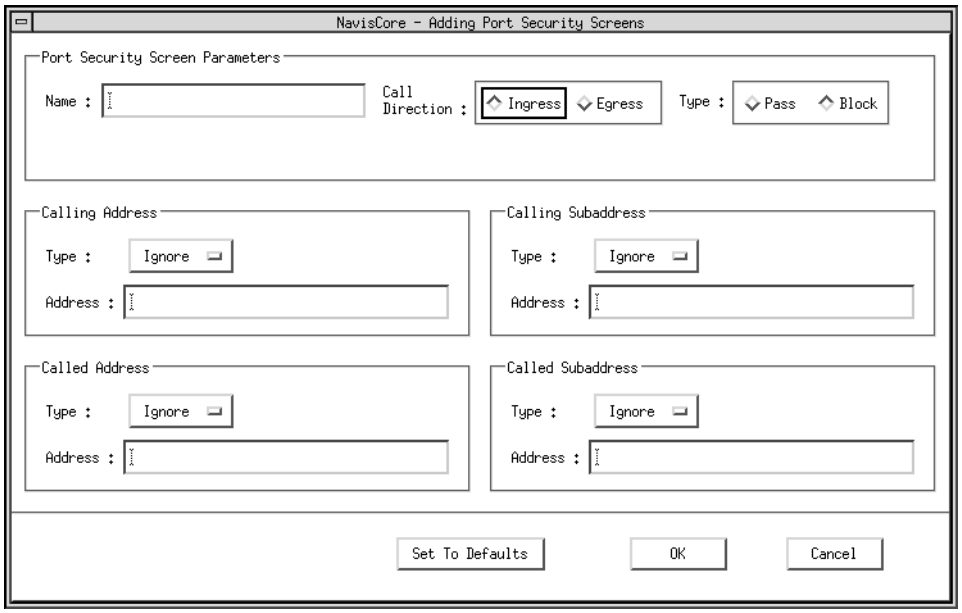

Figure 16-2. Adding Port Security Screens Dialog Box

3. Complete the dialog box fields as described in Table 16-3.

Table 16-3. 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.

Table 16-3. Adding Port Security Screens Fields (Continued)

Field	Action/Description
Calling Address	Configure the Calling Address if this screen is for incoming SVCs. *Type* – Select the address type, either AESA or E.164. Select Ignore (default) if the screen does not use this parameter. *Address* – Enter the address screen using the guidelines on page 16-3. Enter up to 15 characters for an E.164 address; enter up to 40 characters for an AESA address.
Calling Subaddress	Configure the Calling Subaddress for incoming AESA SVCs only. 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 16-3.
Called Address	Configure the Called Address if this screen is for outgoing SVCs. *Type* – Select the address type, either AESA or E.164. Select Ignore (default) if the screen does not use this parameter. *Address* – Enter the address screen using the guidelines on page 16-3. Enter up to 15 characters for an E.164 address; enter up to 40 characters for an AESA address.
Called Subaddress	Configure the Called Subaddress for outgoing AESA SVCs only. 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 16-3.

4. Choose OK to create the new screen.
5. The Adding Port Security Screen dialog box (Figure 16-2 on page 16-8) 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 16-9. 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 16-12.

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 16-13.

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 3-2).
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 following dialog box appears.

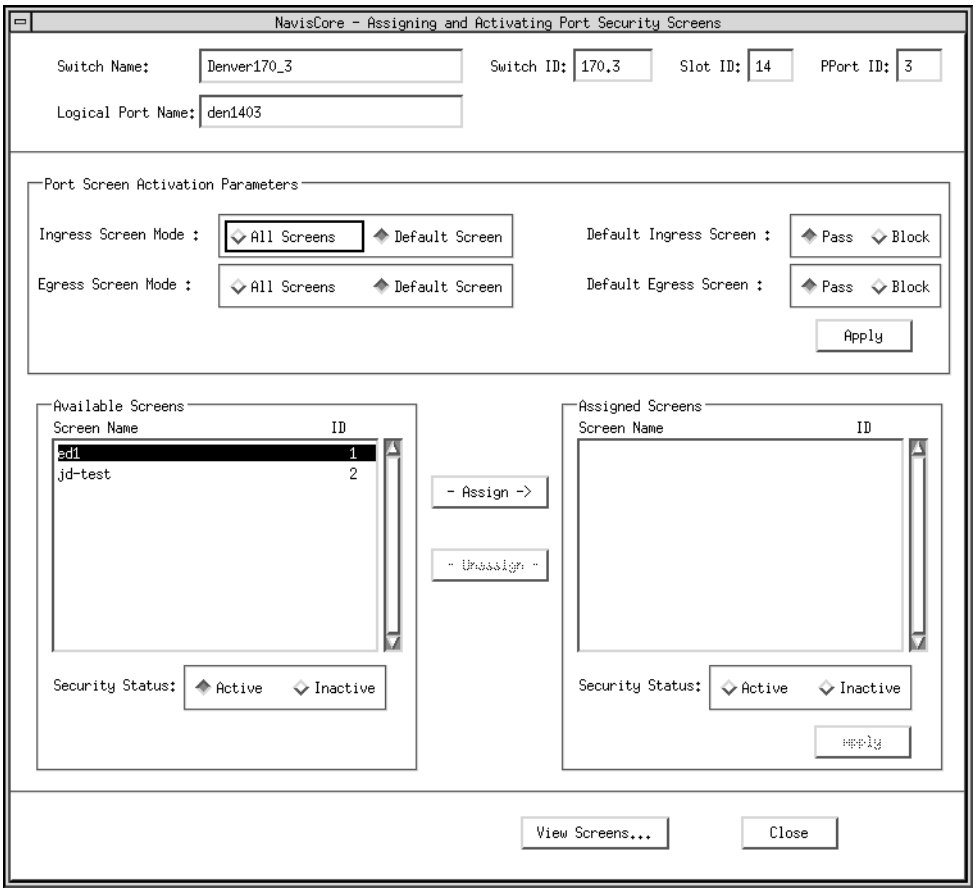

Figure 16-3. Assigning and Activating Port Security Screens

5. See Table 16-1 on page 16-2 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 16-3 on page 16-11) 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 16-12 to access the Assigning and Activating Port Security Screens dialog box (Figure 16-3 on page 16-11).
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 3-2).
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 16-3 on page 16-11).
5. Review the information configured in the Port Screen Activation Parameters group box. See Table 16-1 on page 16-2 if you need information to modify these parameters.

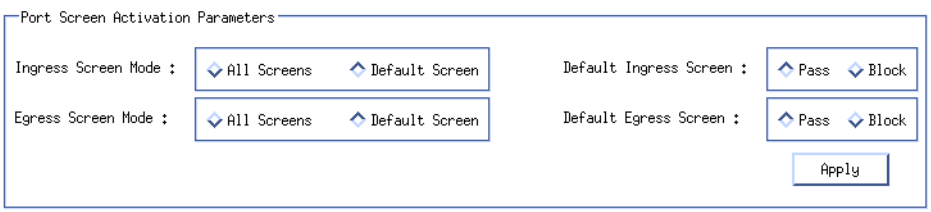

6. Choose Apply to activate the default screen.
7. Choose Close to exit this dialog box.

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 3-2).
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 16-3 on page 16-11).
5. In the Assigned Screens list, select the screen and modify the Security Status as necessary (Active or Inactive).

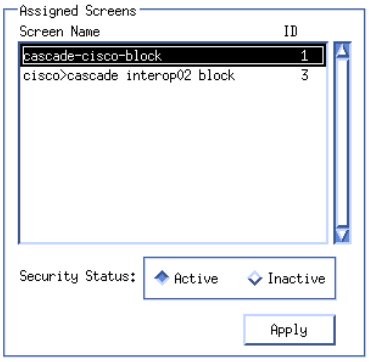

> You can choose the View Screens button to view the parameters configured for the screen you want to modify.

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 3-2).
2. Select the logical port for which you want to view screen assignments.
3. Use the Select:Options menu to select Screen Assignments.
4. Choose View. The following dialog box appears.

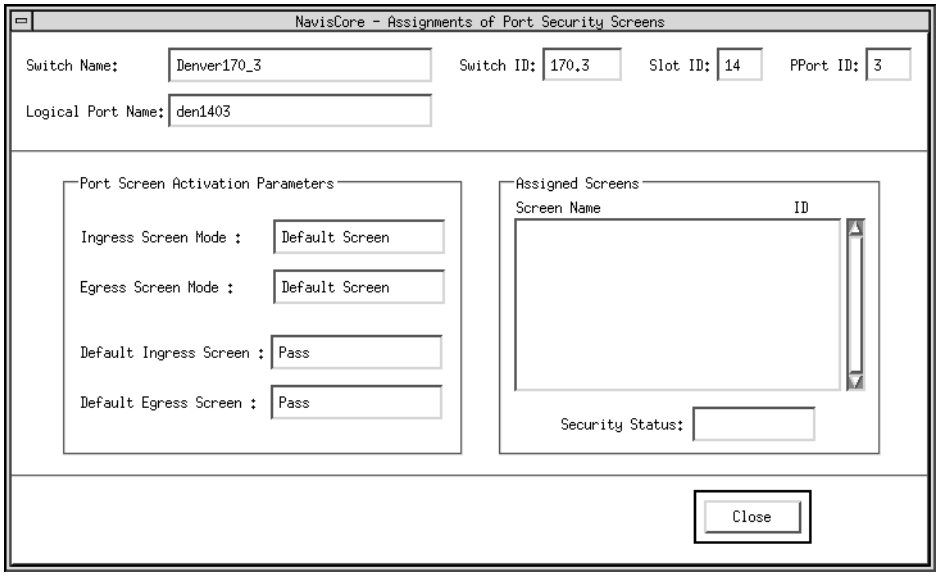

Figure 16-4. Assignments of Port Security Screens Dialog Box

- The *Port Screen Activation Parameters* fields provide the default security screen settings for this logical port. See Table 16-1 on page 16-2 for field descriptions.
- The *Assigned Screens* list provides each screen name assigned to this logical port.

5. Choose Close to exit this dialog box.

17

Configuring PNNI Routing

This chapter describes how to configure the ATM PNNI routing protocol in your Ascend network. Private Network-to-Network Interface (PNNI) is a standard designed by the ATM Forum. This standard defines both an ATM routing protocol and an ATM signaling protocol. Ascend supports PNNI on both the CBX 500 and GX 550 switch platforms. For a detailed explanation of PNNI routing, see the *ATM Forum Technical Committee Private Network-Network Interface Specification Version 1.0* (af-pnni-0055.000), available from the ATM Forum's web site (http://www.atmforum.com).

PNNI Routing Protocol Overview

The PNNI routing protocol provides for dynamic routing configuration and a highly scalable routing scheme. In an ATM network, nodes that support PNNI routing are organized into peer groups. Each peer group is identified by a peer group identifier. A peer group identifier consists of two parameters:

1st byte — peer group level (0 – 104)

bytes 2-14 — peer group identifier

On Ascend switches, the peer group identifier appears by default at the beginning of each peer group member's ATM End System Address (AESA). For example, suppose that five CBX 500 switches are in a peer group, and the ID of that peer group is 39999999. This means that 39999999 would appear at the beginning of the AESA of each of the five switches.

The way in which peer group IDs are configured depends on the vendor implementation.

Peer groups can be organized hierarchically. To accomplish a hierarchical organization of peer groups, each peer group is represented to the next level of the hierarchy by an abstract entity called a logical group node. The lowest-level node in the peer group (called the peer group leader) performs the logical group node functions. Members of the peer group communicate to elect the peer group leader based on leadership priority. The member that has the highest leadership priority is chosen to be the peer group leader and becomes the parent of the peer group, which is now referred to as a child peer group.

For its first release of PNNI routing, Ascend switches do not support PNNI hierarchy, and cannot act as border nodes, peer group leaders, or logical group nodes. Ascend switches can participate in peer group leader election.

As the parent of its child peer group, the logical group node is eligible to join the next highest peer group in the hierarchy, which can be made up of other parents (i.e., logical group nodes) representing other child peer groups and lowest-level nodes. In turn, the members of the next highest peer group in the hierarchy choose a logical group node, which represents them to the third level of the hierarchy — and on up the hierarchical chain, forming a kind of ancestry consisting of children, parents, grandparents, and so on.

PNNI Routing Example

Figure 17-1 shows a simple two-tiered PNNI routing hierarchy, with six lowest-level nodes divided into two child peer groups (PG1 and PG2). The logical group nodes that are the parents of each of the child peer groups form a top-tier peer group (PG3). As a result, each logical group node is a member of two groups — its child group and its peer group.

Peer groups may contain both logical group nodes and lowest-level nodes. For example, in Figure 17-1, a lowest-level node could also be a member of PG3.

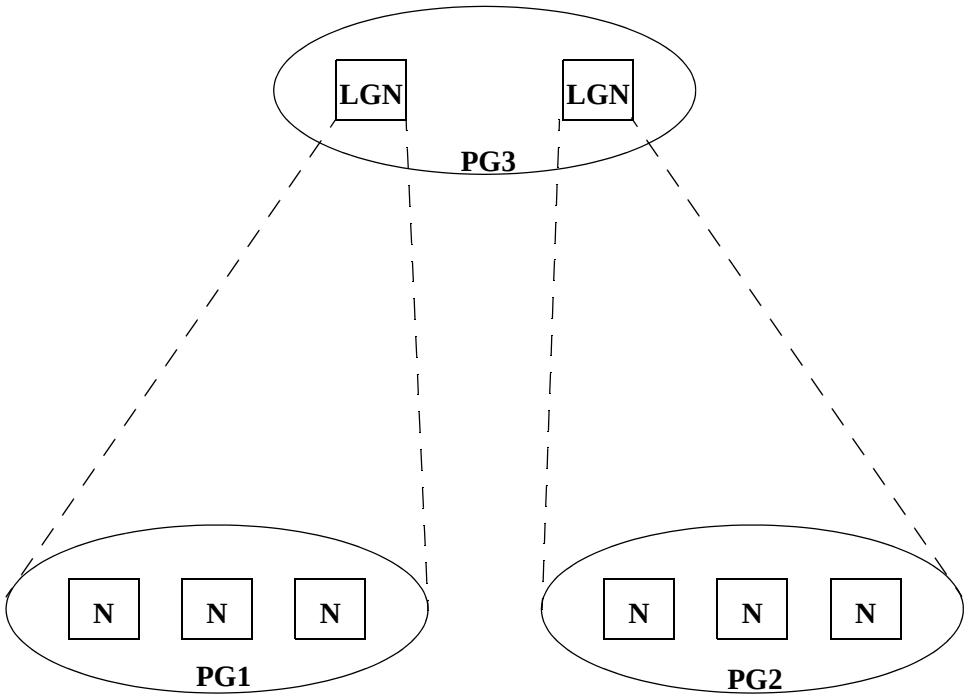

Legend
N = Lowest-level Node
LGN = Logical Group Node
PG = Peer Group

Figure 17-1. Simple PNNI Routing Hierarchy

Neighboring lowest-level nodes within a peer group exchange information to synchronize their topology databases. The topology database contains information on the peer group in which a node resides and information that allows the node to reach destinations in other peer groups. A node receives information about the network beyond the peer group from its peer group leader.

In its role as logical group node, the peer group leader collects routing information from all of the nodes in its child peer group and propagates (distributes) a summarized version of that information to the higher-level peer group. In turn, the peer group leader receives summarized routing information from its peers and distributes that information to the other nodes in its child group.

This automated collection and propagation process eliminates the need for manual configuration and maintenance of routing information on network nodes. In effect, PNNI allows network nodes to automatically learn the topology of the network, and use the topological knowledge they acquire to route data to its correct destination.

Figure 17-2 illustrates the flow of PNNI topology information within peer groups and between peer groups. The neighboring nodes in each peer group exchange topology information to synchronize each other's topology databases. Logical group nodes also propagate information about how to reach their child groups to other logical group nodes.

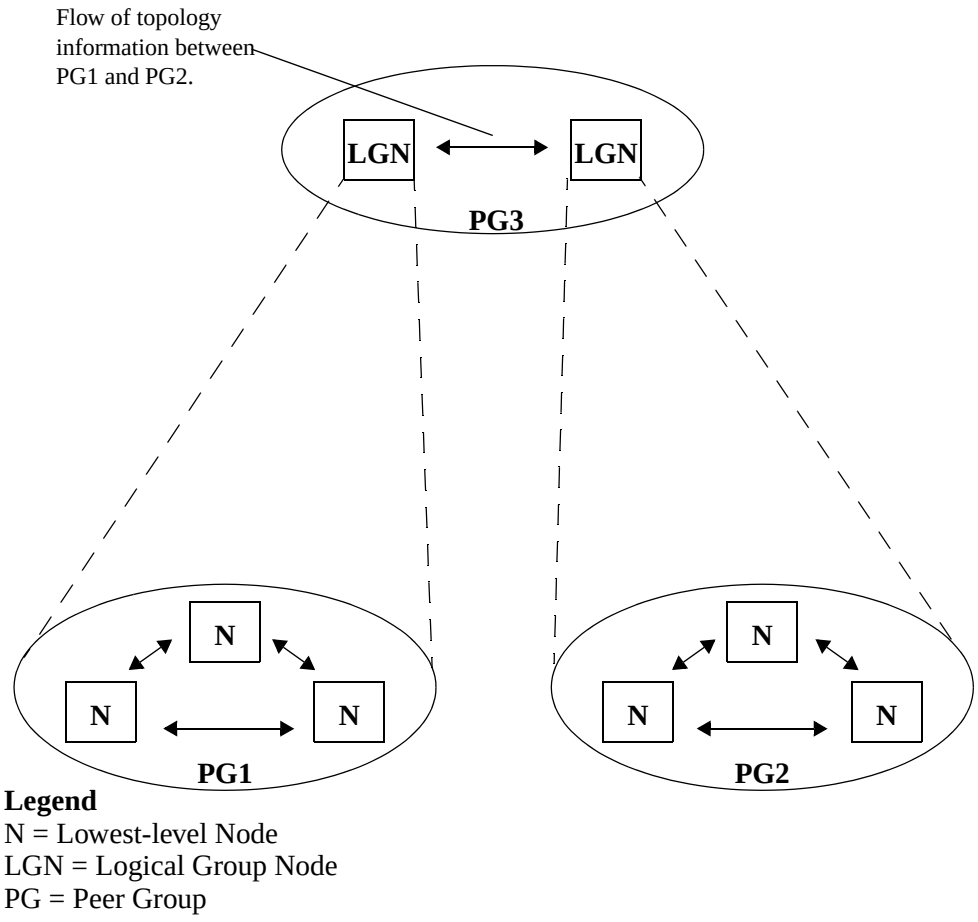

Figure 17-2. Flow of PNNI Topology Information

The following packets carry PNNI control information during exchanges between neighbors:

Hello Packets — Contain information that neighboring nodes exchange to discover and verify each other's identity and to determine the status of the links that connect them.

Database Summary Packets — Contain the identifying information of all PNNI Topology State Elements (PTSEs) in a node's topology database. A PTSE is a collection of PNNI topology information that is sent to all nodes in a peer group. When a node first learns that a neighboring peer node residing in the same peer group exists, it initiates a database exchange process in order to synchronize its topology database with its neighbor. When one neighbor sends a database summary packet to another neighbor, the other neighbor responds with its own database summary packet.

PTSE Request Packets — Contain one or more entries that request PTSEs. When a node examines received database summary packets from neighbors and detects one or more missing PTSEs in its topology database, it builds a PTSE request packet. This packet contains a list of IDs that identify the missing PTSEs. The node sends the PTSE request packet to neighbors, which respond with a PTSP.

PNNI Topology State Packets (PTSPs) — Contain one or more PTSEs. A node sends PTSPs when it:

- Detects that its local topology information has changed, in which case it immediately sends PTSP(s) containing information about the change to its neighbors.
- Receives a PTSP containing new topology information from a neighbor; the node then propagates this information to other neighbors in PTSP packets.
- Responds to PTSE requests during topology database synchronization.

Note that the first two bullets describe the most common reasons for sending PTSPs. The last bullet describes the least common reason.

PTSE Acknowledgment Packets — Contain acknowledgments of receipt of PTSEs. When a node receives PTSEs, the receiving node acknowledges receipt by sending one or more PTSE acknowledgment packets.

In Ascend's PNNI implementation, each CBX 500 switch and GX 550 switch requires a logical port for each of its neighbors. The logical port type is ATM NNI, and the protocol type is PNNI 1.0. For example, each lowest-level node in PG1 in Figure 17-1 would have two ATM NNI logical ports (with protocol type of PNNI) configured – one for each of its neighboring nodes in PG1.

When you configure each logical port, you can assign an administrative weight to each QoS category. This weight allows you to configure the network to favor one path over another path for a given QoS category, when the path constraint for an SVC call is administrative weight. The weights of all the network interfaces along a path are added up, and switches choose the path with the lowest cumulative weight when making routing decisions.

For example, suppose that VBR Real Time traffic has two available paths for reaching a given destination. One path has a weight of 1000 while another path has a weight of 4000. The switch will choose the path with the weight of 1000, if the call requests VBR-RT QoS and administrative weight as a metric.

PNNI Signaling Overview

This section provides a brief overview of PNNI signaling. For a detailed explanation of PNNI signalling, see the *ATM Forum Technical Committee Private Network-Network Interface Specification Version 1.0* (af-pnni-0055.000), available from the ATM Forum's web site (http://www.atmforum.com).

PNNI signaling allows ATM SVC calls to be set up across a private network that supports the PNNI protocol. It is based on a subset of UNI 4.0 signaling. It does not support some UNI 4.0 signalling features such as leaf initiated joint capability or user-to-user supplementary service, but adds new features which support the use of PNNI routing for dynamic call setup, and PNNI crankback for the dynamic re-routing of call setups around failed nodes or links or links with insufficient resources.

PNNI signalling makes use of PNNI routing information. PNNI uses the route calculations derived from the reachability, connectivity, and resource information dynamically maintained by PNNI routing. These routes are calculated as needed from the node's view of the current topology.

Configuring PNNI Routing

Figure 17-3 displays the sequence of steps to follow to configure PNNI routing for your network:

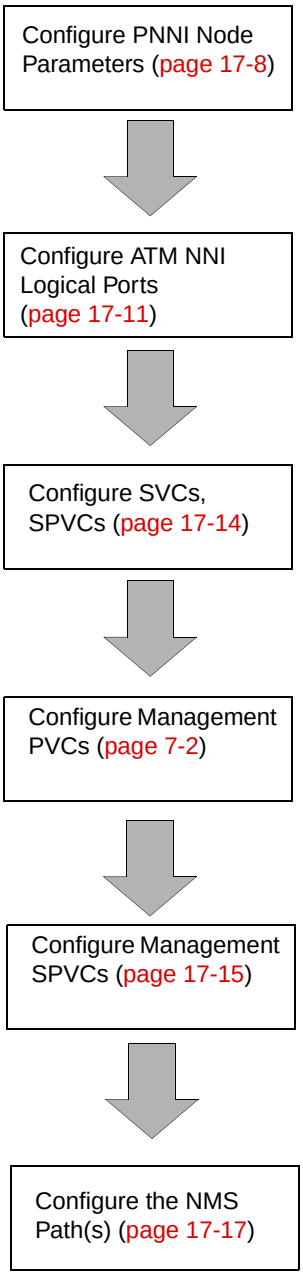

Figure 17-3. Configuring PNNI Routing

Configuring PNNI Node Parameters

To begin using the PNNI routing protocol in your Ascend network, you need to configure the PNNI node parameters for each switch that supports PNNI in the network.

1. Select the appropriate switch icon on the network map.
2. From the Administer menu, select Ascend Parameters ⇒ Set All PNNI Node Parameters. The following dialog box appears.

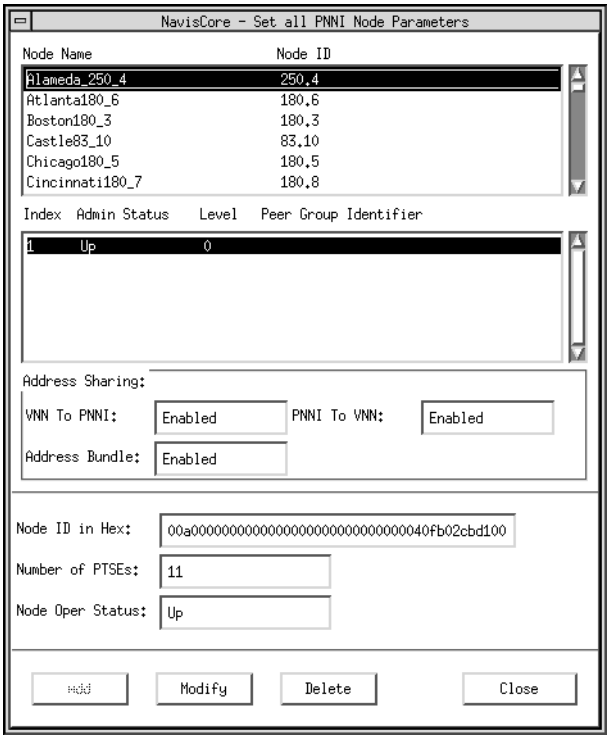

Figure 17-4. Set All PNNI Node Parameters Dialog Box

Once you Add a PNNI node, the Set All PNNI Node Parameters dialog box displays the configured information for that node. Table 17-1 describes the dialog box fields.

Table 17-1. Set All PNNI Node Parameters Fields.

Field	Description
Node Name	Displays the name of the switch.
Node ID	Displays the switch number used as the host assignment in the switch's internal IP address.
Index	Displays the index value that identifies the row entry for the switch in the table that stores the switch parameters. This value is always 1.
Admin Status	Displays the PNNI administrative status (Up or Down). If PNNI Admin Status is Down, then PNNI Node Oper Status is also Down.
Level	Displays the number of significant bits available for forming the PNNI Peer Group Identifier. The value can be from 0 to 104. By determining the number of bits allocated for the peer group identifier, the PNNI Level also determines the level of the switch in the PNNI routing hierarchy. The Peer Group level decreases as you move higher up in the hierarchy. For example, a node that is the grandparent of a peer group two levels lower will have fewer bits reserved for its peer group identifier than its grandchildren. As a result, the grandparent will have a smaller peer group identifier than the peer group identifier of its grandchildren.
Peer Group Identifier	Displays the identifier of the peer group to which the switch belongs. The identifier is determined by the PNNI Level value.
VNN to PNNI	Indicates whether or not internal addresses known to PNNI should be advertised within the VNN domain.
Address Bundle	Indicates whether or not to support address bundling within the PNNI domain on a per-node basis.
PNNI to VNN	Indicates whether or not internal addresses known to VNN should be advertised within the PNNI domain.
Node ID in Hex	Displays the PNNI node ID of the switch. This ID is not configurable. The switch derives this ID by concatenating the PNNI Level, the hexadecimal value 0xA0, the PNNI Peer Group Identifier, and the Media Access Control (MAC) address of the Ethernet interface of the switch. This method of deriving the ID guarantees uniqueness.
Number of PTSEs	Displays the number of PTSEs in the topology database of the switch. A PTSE is a collection of PNNI topology information that is sent to all nodes in a peer group.
Node Oper Status	Displays the PNNI operating status for the switch (Up or Down). If PNNI Admin Status is Down, then PNNI Node Oper Status is also Down.
Add	Adds PNNI node information for the switch you select.
Modify	Modifies an existing PNNI node that you select.
Delete	Deletes an existing PNNI node that you select.

3. From the list box at the top, select the switch for which you want to configure PNNI node parameters.
4. Choose Add. The following dialog box appears.

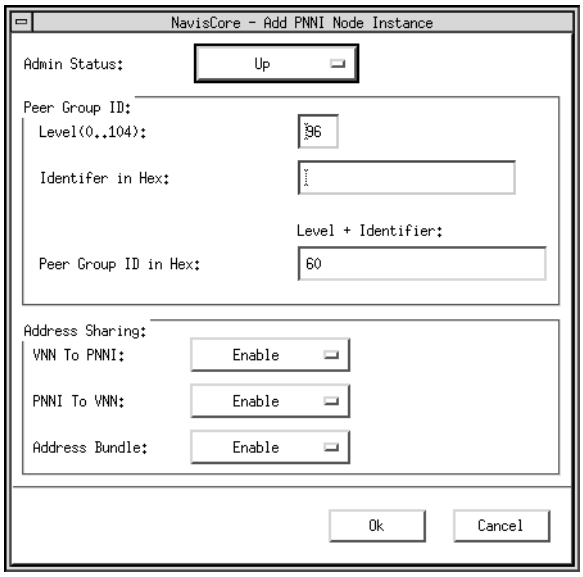

Figure 17-5. Add PNNI Node Instance

5. Complete the dialog box fields as described in Table 17-2.

Table 17-2. PNNI Node Instance Fields.

Field	Action/Description
Admin Status	Set the Admin Status Up (*default*) or Down. If PNNI Admin Status is Down, then PNNI Node Oper Status is also Down.
Level	Sets the number of significant bits available for forming the PNNI Peer Group Identifier. The value can be from 0 to 104. By determining the number of bits allocated for the peer group identifier, the PNNI Level also determines the level of the switch in the PNNI routing hierarchy. As you ascend the hierarchy, the number of bits allocated for peer group identifiers decreases, resulting in smaller peer group identifiers. For example, a node that is the grandparent of a peer group two levels lower will have fewer bits reserved for its peer group identifier than its grandchildren. As a result, the grandparent will have a smaller peer group identifier than the peer group identifier of its grandchildren.
Identifier in Hex	Enter a PNNI peer group identifier. This identifier provides a type of routing prefix.

Table 17-2. PNNI Node Instance Fields. (Continued)

Field	Action/Description
Peer Group ID in Hex	Displays the complete PNNI peer group identifier for the switch. This ID is not configurable. The switch derives this ID by concatenating the following elements: • PNNI Level you enter • Hexadecimal value 0xA0 • PNNI Peer Group Identifier you enter • Media Access Control (MAC) address of the switch's Ethernet interface This method of deriving the ID guarantees uniqueness.
VNN to PNNI	Indicates whether or not internal addresses known to PNNI should be advertised within the VNN domain. Choose enable to advertise an internal address within the VNN domain.
Address Bundle	Indicates whether or not to support address bundling within the PNNI domain on a per-node basis. Choose enable to support address bundling within the PNNI domain.
PNNI to VNN	Indicates whether or not internal addresses known to VNN should be advertised within the PNNI domain. Choose enable to advertise an internal address within the PNNI domain.

6. Choose OK to set the PNNI node parameters and return to the Set All PNNI Node Parameters dialog box. Choose Close to return to the network map.

Configuring an ATM NNI Logical Port

You configure PNNI routing protocol using ATM NNI logical ports. This logical port type uses many of the same attributes as the ATM NNI logical port configured for BICI 1.1. PNNI only requires you to configure a few additional PNNI logical port attributes.

To configure an ATM NNI logical port for PNNI:

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

5. Choose OK. The Add Logical Port dialog box appears (Figure 3-3 on page 3-7).
6. Use the instructions in Table 17-3 to set the logical port attributes.

Table 17-3. Configuring an ATM NNI Logical Port

Use the instructions on	To set the
page 3-9	Administrative Attributes to enter a logical port name
page 3-11	ATM Attributes to select the ATM Protocol, PNNI 1.0
page 3-15	ILMI/Signaling/OAM to enable signaling
page 3-21	Flow Control Processor Attributes (*optional*)
page 3-23	SVC VPI/VCI Range (*optional*)
page 3-31	SVC Attributes (*optional*)
page 3-39	SVC Routing Priorities (*optional*)

7. Select Set PNNI Parameters Attributes to display the following fields.

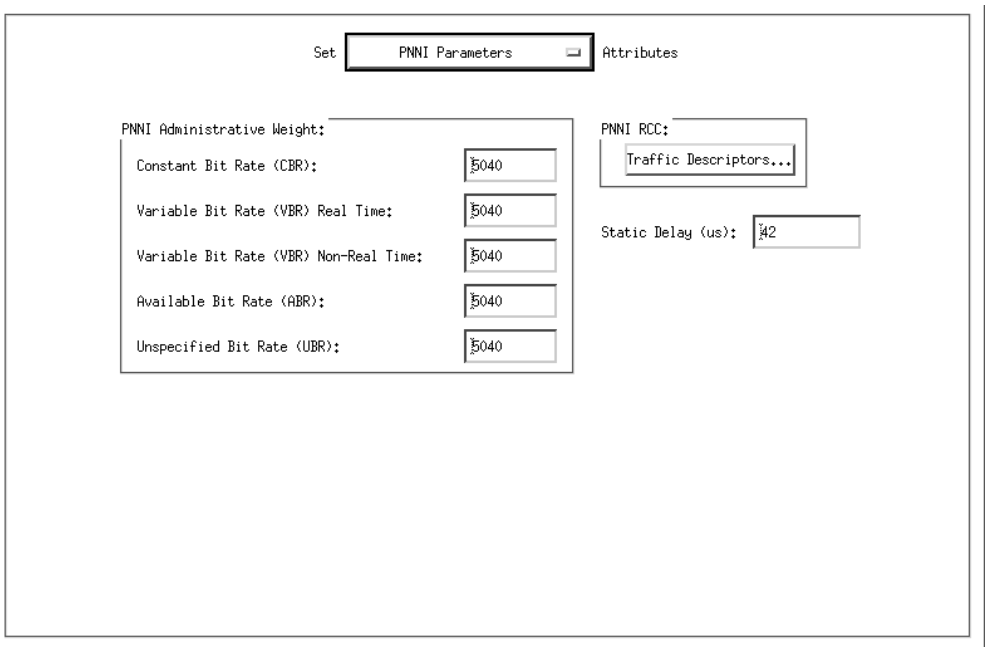

Figure 17-6. Set PNNI Parameters Attributes

PNNI Attributes enable you to configure the PNNI Administrative Weight Status. Assign an administrative weight to each QoS category. This weight allows you to configure the network to favor one path over another path for a given category. The weights of all the network interfaces along a path are added up, and switches choose the path with the lowest cumulative weight when making routing

decisions. For example, suppose that VBR Real Time traffic has two available paths for reaching a given destination. One path has a weight of 1000 while another path has a weight of 4000. The switch will choose the path with the weight of 1000.

8. Use Table 17-4 to complete the logical port PNNI parameters.

Table 17-4. PNNI Administrative Weight

Field	Action/Description
Constant Bit Rate (CBR)	Configures the administrative weight assigned to the CBR QoS category for the network interface associated with the logical port.
Variable Bit Rate (VBR) Real Time	Configures the administrative weight assigned to the VBR Real Time QoS category for the network interface associated with the logical port.
Variable Bit Rate (VBR) Non-Real Time	Configures the administrative weight assigned to the VBR Non-Real Time QoS category for the network interface associated with the logical port.
Available Bit Rate (ABR)	Configures the administrative weight assigned to the ABR QoS category for the network interface associated with the logical port.
Unspecified Bit Rate (UBR)	Configures the administrative weight assigned to the UBR QoS category for the network interface associated with the logical port.
Static Delay	If you enable this option, an SVC originating from this logical port 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 in a value of 500 µsec., the SVC will not be routed over a path whose total end-to-end delay exceeds 500 µsec. 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. The valid range for this field is 0 – 167777214 µsec; the default is 42 µsec. ***Note:*** *The value you enter should reflect your network topology. If an SVC typically traverses high speed trunks, set the delay rate lower. You need to increase the delay if the SVC uses low-speed trunks.*
PNNI RCC Traffic Descriptors	Choose the Traffic Descriptors button to display the dialog box shown on Figure 8-3 on page 8-10. Use Step 3 through Step 5 beginning on page 8-10 to configure these values.

9. Use the instructions on page 3-27 to complete the logical port configuration.
10. Repeat Step 2 through Step 9 beginning on page 17-11 for each ATM NNI logical port you need to configure.

For information on configuring Virtual NNI logical ports, see page 2-6. Virtual logical ports allow you to configure more than one logical port on the same physical port. Each logical port that you configure uses a portion of the total physical port bandwidth.

Configuring SVCs and SPVCs for PNNI

PNNI routing connects your network of switches using SVCs instead of PVCs. Once you have configured your ATM NNI logical ports, you need to configure SVCs using an AESA address format. If necessary, review Chapter 11 for an overview of SVC addressing. To begin configuring SVCs, continue with the instructions in Chapter 12.

In addition to SVCs, PNNI routing can access switches using SPVCs. For information on configuring SPVCs, see Chapter 14.

Configuring a Management PVC

A Management PVC (MPVC) provides the connection from the NMS to the gateway switch, while the remaining switches in your PNNI network are connected using Management SPVCs (MSPVCs). Figure 17-7 illustrates this concept.

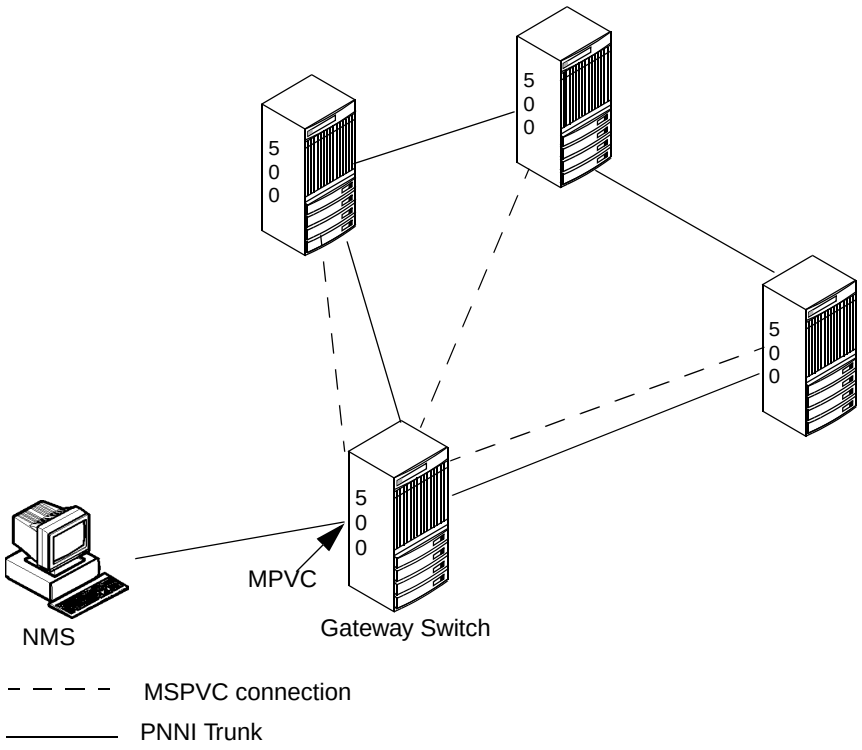

Figure 17-7. Connecting a PNNI Network

Perform the following tasks to make this MPVC connection:

- To define an MPVC, use the instructions beginning on page 7-2 .
- To connect the NMS to the gateway switch in the PNNI network, use the instructions beginning on page 7-4 to define an NMS path for an MPVC.

Configuring Management SPVCs

A management SPVC (MSPVC) connects the switch management port to an SVC terminating address located on an adjacent switch. This management connection is used as the NMS path which enables the NMS to manage the switch.

MSPVCs originate at an internal logical port located on the switch's processor module (either SP or NP, respectively). They terminate at the switch's I/O interface: IOM for a CBX, and BIO for the GX 550. MSPVCs provide a data path that accesses internal network management functions. The MSPVC internal logical port is designated as MgmtLPort.SW[*switch name*]. It uses an interface number (ifnum) of 4093. To form the MSPVC, connect the MgmtLPort. SW[*switch name*] endpoint to any target AESA address configured on an ATM UNI logical port.

To configure the MSPVC:

1. From the Administer menu, select Ascend Parameters ⇒ Set All Soft PVCs ⇒ Point-to-Point. The Set All Point-to-Point SPVCs dialog box appears. (See page 14-4 for a description of the dialog box fields.)
2. Choose Add. The following dialog box appears.

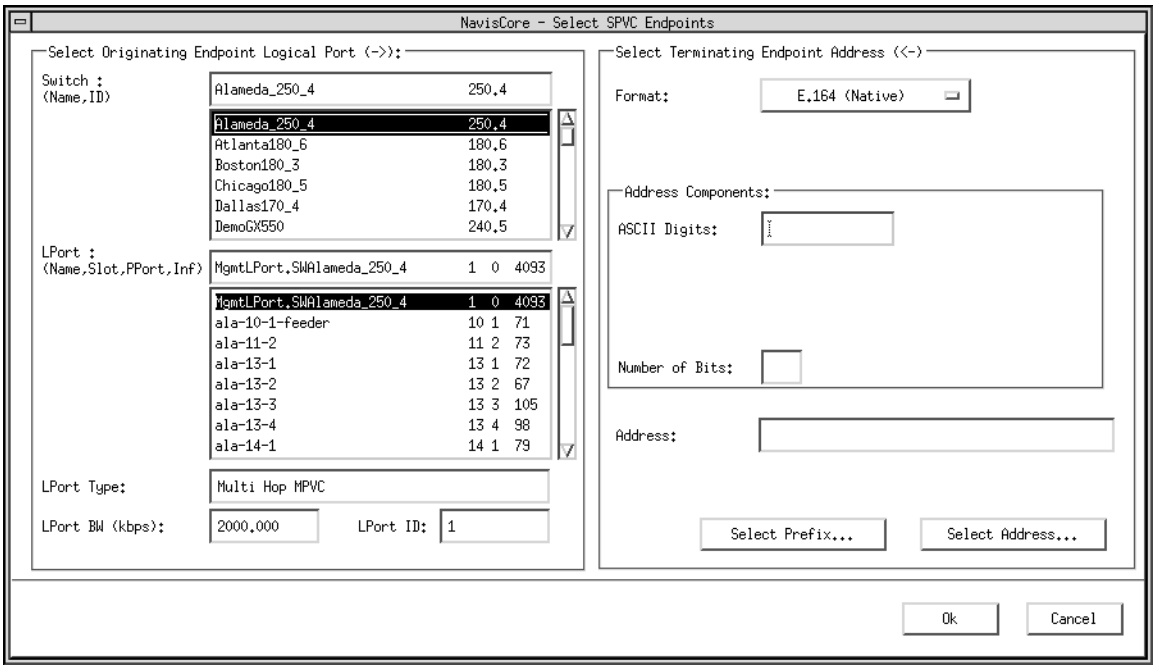

Figure 17-8. Select SPVC Endpoints Dialog Box

3. Use the following steps to configure the originating endpoint logical port.
   a. Select the name of the switch on which the MSPVC endpoint will reside.
   b. Select the MgmtLPort.SW[*switch name*] endpoint.
4. To complete this configuration:
   - If you know the SVC terminating endpoint address, use Table 14-3 on page 14-8 to select the address format and configure the terminating endpoint address. For more information on AESA address formats, see page 11-2.
   - If you do not know this address, or if you need to configure the terminating endpoint address, see page 14-12 for instructions on using the Select Prefix and Select Address buttons to configure this address.
5. Choose OK. The following dialog box appears.

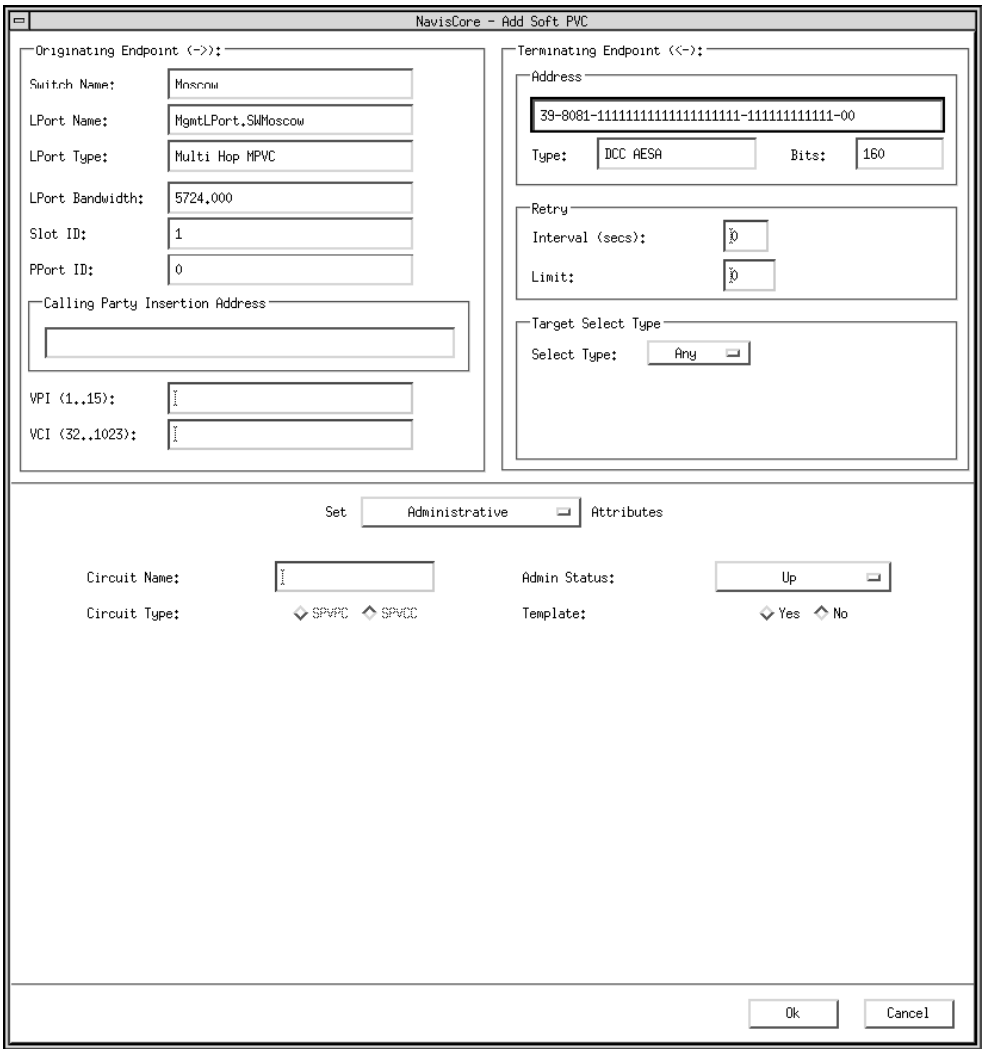

Figure 17-9. Add Soft PVC Dialog Box

6. Use the instructions in Table 14-4 on page 14-9 to configure these fields.

7. Choose Set [Traffic Type] Attributes to select the traffic descriptors for this SPVC. The Set Traffic Type Attributes appear (Figure 14-4 on page 14-11).
8. See "Defining Traffic Descriptor Attributes" on page 8-9 for instructions to configure these attributes.
9. Choose OK to create the new SPVC and return to the Set All Point-to-Point SPVCs dialog box (Figure 14-1 on page 14-4).
10. Choose Close to return to the network map.

Defining the NMS Path

To define an NMS path and complete the MSPVC configuration:

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 following dialog box appears.

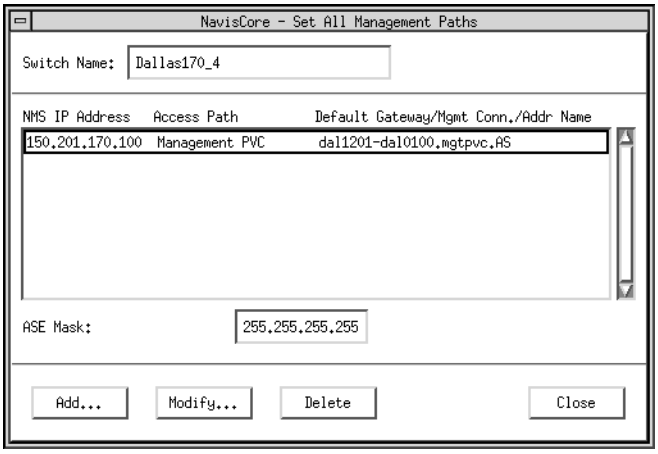

Figure 17-10. Set All Management Paths

3. Choose Add. The following dialog box appears (see Figure 17-11 on page 17-18).

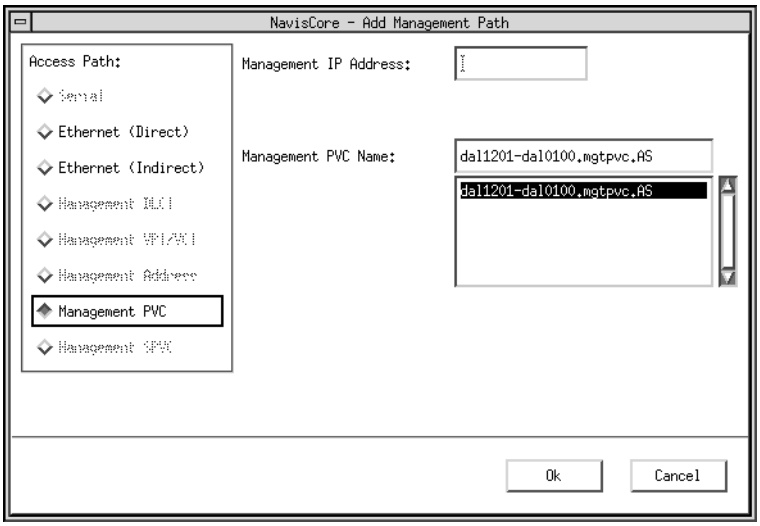

Figure 17-11. Add Management Path

4. Select Management SPVC as the Access Path.
5. Enter the NMS IP Address. This is the IP address of the SPARCstation that will manage this switch.
6. Select the name of the Management SPVC Name you created in the previous section.
7. Choose OK. Choose Close to return to the network map.

A

Adjusting the CAC

This appendix describes how to tune the Ascend Call Master Connection Admission Control (CAC) to achieve a desired cell loss ratio objective across all physical ports in your network. The Ascend CAC is responsible for the bandwidth allocation on all ATM cards on the CBX 500, GX 550, and B-STDX. It is also responsible for bandwidth allocation on all frame cards with the priority frame capability.

When you create a circuit, the CAC function computes a bandwidth allocation for that circuit and updates the bandwidth allocation for the circuit's QoS class. This bandwidth allocation depends on the specified CAC implementation, the circuit's QoS class, and the circuit's specified traffic descriptor. If you try to create a circuit that causes the allocated bandwidth for a given QoS class to exceed the bandwidth available for that class, the circuit will not be created.

The CAC configuration option enables you to choose between three CAC implementations. You can choose the Ascend CAC implementation or configure one of two customized CAC implementations: "customize VBR-NRT and ABR" and "customize VBR-RT, VBR-NRT, and ABR."

The Ascend CAC implementation allows you to control the Quality of Service and bandwidth allocation by specifying cell loss ratio and cell delay variation objectives while the customized CAC implementations allow you to directly control the bandwidth allocation for circuits. The “customize VBR-NRT and ABR” CAC implementation allows you to control the amount of bandwidth that is reserved for VBR Non-Real Time and available bit rate (ABR) circuits, while the “customize VBR-RT, VBR-NRT, and ABR” CAC implementation allows you to control the amount of bandwidth that is reserved for VBR Real Time, VBR-NRT, and ABR circuits. In either implementation, you can control the amount of bandwidth reserved based on either the physical port type or the configurable range of SCR values, or both. With the two customized implementations you can also control circuit establishment based on the configurable range of maximum MBS values.

Quality of service for the VBR-RT and VBR-NRT QoS classes is not guaranteed when you use the “customize VBR-RT,VBR-NRT, and ABR” CAC implementation. Also, the Quality of Service for the VBR-NRT QoS class is not guaranteed when you use the “customize VBR-NRT and ABR” CAC implementation.

When you adjust the CAC function, choose only one of these options. Whether you are tuning the Ascend CAC or configuring a customized CAC, the adjustments you make apply only to the VBR-RT,VBR-NRT, and ABR traffic types.

Before tuning the Ascend CAC or configuring a customized CAC, you should closely monitor your network to achieve a good understanding of the network’s traffic profile. Be conservative when you adjust the CAC to ensure quality of service. After you make adjustments, monitor the network closely to determine the effect of these adjustments, making sure you have not adversely impacted the Quality of Service on the network.

About the Customizable CAC Options

Both of the customizable CAC implementations enable you to directly control the amount of bandwidth reserved for VBR-NRT and ABR circuits. In addition, you can control the amount of bandwidth reserved for VBR-RT circuits if you choose the "customize VBR-RT, VBR-NRT, and ABR" CAC implementation. You control the amount of bandwidth reserved based on either the physical port type, the SCR requirements of the circuit, or both. When you use the customized CAC options, the following formula determines the amount of bandwidth required for a given circuit:

```
Bwidthreq = SCR*F1*F2
```

where F1 is the physical port factor (entered as a percentage), and F2 is the SCR scale factor (entered as a percentage). You can configure only an F1 factor, only an F2 factor, or both factors. If you do not configure one of these factors, then the value of that factor is, by default, 100%.

On frame cards with priority frame capability, the Bwidthreq=SCR.

Customizable CAC Example

A circuit request is made, and the circuit needs to reserve bandwidth based on an SCR of 10,000 cells/sec. You configure the F1 factor for DS3 ports at 150%, the F1 factor for OC3c ports at 80%, and the F2 factor for circuits with an SCR from 8,001-15,000 cells/sec at 80%.

- If the circuit request is made on a DS3 port, then the bandwidth requirements of the circuit will be based on an SCR of 12,000 cells/sec, instead of 10,000 cells/sec (10,000 x 150% x 80% = 12,000).
- If the circuit request is made on an OC3c port, then the circuit bandwidth requirements will be based on an SCR of 6,400 cells/sec (10,000 x 80% x 80%= 6,400).

Configuring the CAC

To configure a customized CAC:

1. On the network map, select the switch for which you need to adjust the CAC.
2. From the Administer menu, select Ascend Parameters ⇒ Set All CAC Parameters. The Set All CAC Parameters dialog box appears.

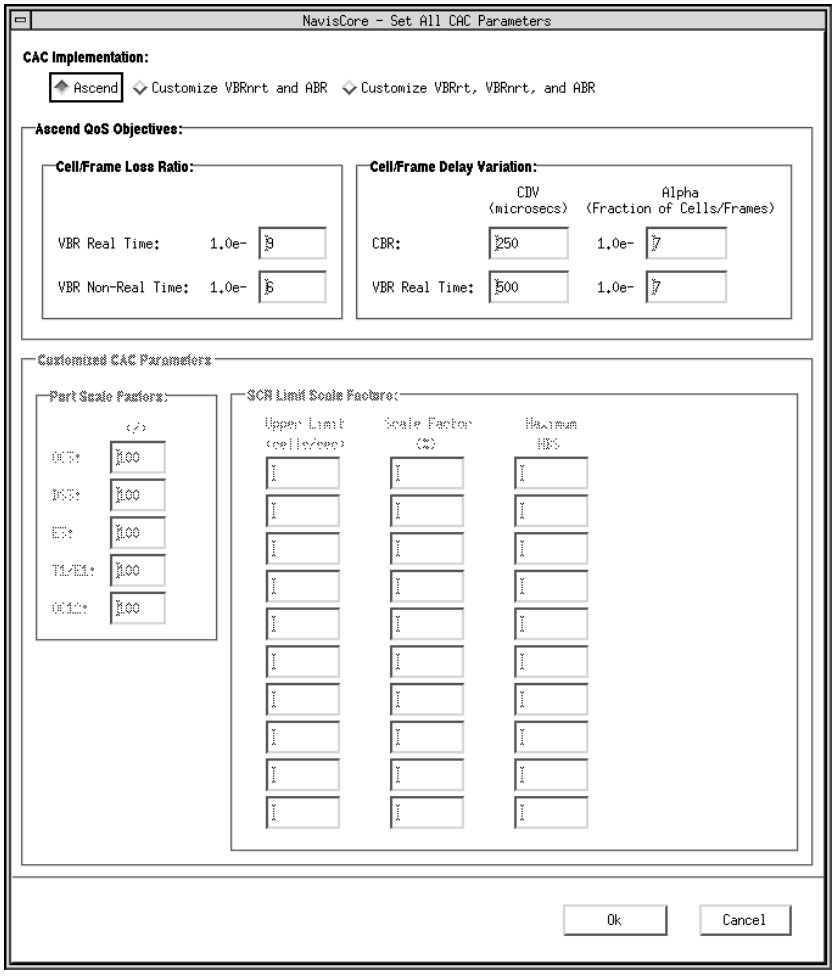

Figure A-1. Set All CAC Parameters Dialog Box

3. Select one of the following CAC Implementations:

   Ascend — Enables you to tune the Cell Loss Ratio and Cell Delay Variation only. Refer to the following section, "Tuning the Ascend CAC" on page A-5, for more information.

   Customize VBR-NRT and ABR — Enables you to tune the Cell Loss Ratio, Cell Delay Variation, and Customized CAC Parameters.

   Customized VBR-RT, VBR-NRT, and ABR — Enables you to tune Customized CAC Parameters only.

Tuning the Ascend CAC

To tune the Ascend CAC, specify the cell loss ratio objectives you want to meet across your network. You can specify a cell loss ratio objective in the range of 10^{-1} to 10^{-12}. For example, an entry of 10^{-5} specifies that circuits will not be created on any physical port on which:

- The cell drop ratio is currently 1 in 100,000 (because 10^{-5} is equal to 1/100,000)

OR

- The creation of the circuit would potentially cause the cell drop ratio to exceed 1 in 100,000

Ascend recommends that you adjust the CAC when you first configure a switch. Adjusting the CAC after several circuits have been created will not automatically change the bandwidth allocation for these circuits and may not guarantee the defined Quality of Service.

To tune the Ascend CAC:

1. On the Modify CAC Parameters dialog box (Figure A-1 on page A-4), select the Ascend CAC implementation.
2. In the Cell Loss Ratio Objectives VBR Real Time and VBR Non-Real Time fields, specify the cell loss ratio objective you want to meet for each of these traffic types. This value is a negative power of ten (1.0e–). For example, if you enter 5, your cell loss ratio objective is a maximum of 1 dropped cell for every 100,000 cells. If the CAC determines that the creation of a circuit on a physical port will cause more than 1 in 100,000 cells to be dropped, then the circuit will not be created on that physical port.

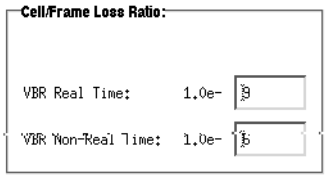

By default, VBR Real Time is set to 9 (1 in 1,000,000,000) and VBR Non-Real Time is set to 6 (1 in 1,000,000).

3. In the Cell Delay Variation CBR and VBR Real Time fields, specify the CDV (in microseconds). This value represents the CDV objective for the CBR and VBR Real Time QoS class. Although this value represents an upper bound on the delay variation for most physical interfaces, the CAC algorithm allows the actual CDV values on slow interfaces (such as T1 cards) to exceed this configured value.

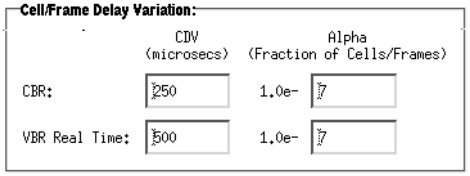

Keep in mind that since both the CDV and CLR calculations are non-linear in nature, the resulting Equivalent Bandwidth for VBR-RT and VBR-NRT circuits may not be the same as it was in previous releases. Since the circuits might end up with a larger Equivalent Bandwidth as a result of the CDV objectives, one or more existing circuits many no longer be admitted because of insufficient bandwidth on a port where they were previously admitted.

4. In the Alpha field, specify the fraction of the CBR (or VBR Real Time) cells that can exceed this CDV objective. This value is a negative power of ten (1.0e–). By default, the Alpha field for each of the CBR and VBR Real Time classes is set to 7 (1 in 10,000,000).
5. When you finish, choose OK to send the values you entered to the selected switch.
6. To send these values to another switch on the network map, select the switch and use Step 2 on page A-4 to access the Modify CAC Parameters dialog box. Choose OK.

Customizing the CAC for the VBR-RT, VBR-NRT, and ABR Classes

To customize the CAC for VBR-RT, VBR-NRT, and ABR:

1. On the Modify CAC Parameters dialog box (Figure A-1 on page A-4), select the customize VBR-RT, VBR-NRT, and ABR CAC implementation.
2. In the Port Scale Factors box, enter a scale factor percentage to use for computing bandwidth requirements on the physical port.

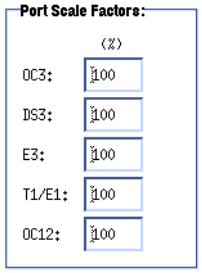

For example, if you enter a value of 125% in the DS3 field, a circuit that would normally reserve bandwidth based on an SCR of 10,000 cells/sec would be allocated bandwidth of 12,500 cells/sec.

3. To customize the CAC based on the SCR and MBS values:

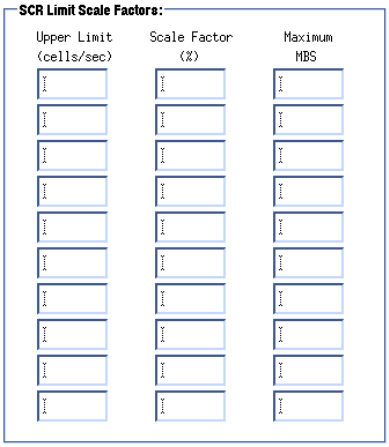

a. In the Upper Limit column, enter the upper limit of the SCR range for which you want to customize the amount of bandwidth reserved. You can specify up to ten upper limits. The following list shows several examples.

Example 1	Example 2	Example 3
10,000	10,000	8,000
20,000	16,000	12,000
35,000	20,000	15,000
—	24,000	20,000
—	28,000	25,000
—	35,000	30,000
—	—	35,000

b. This would give you the following ranges of SCR values:

Range	Example 1	Example 2	Example 3
1	0-10,000	0-10,000	0-8,000
2	10,001-20,000	10,001-16,000	8,001-12,000
3	20,001-35,000	16,001-20,000	12,001-15,000
4	—	20,001-24,000	15,001-20,000
5	—	24,001-28,000	20,001-25,000
6	—	28,001-35,000	25,001-30,000
7	—	—	30,001-35,000

To determine the ranges you should configure, monitor the VBR traffic on your network, then group your VBR circuits into appropriate SCR ranges.

c. In the Scale Factor column, enter a scale factor percentage to use when computing bandwidth requirements for circuits in each of the SCR ranges you defined.

For example, if you enter a value of 125%, a circuit with an SCR of 12,000 cells/sec would be allocated a bandwidth of 15,000 cells/sec (assuming you did not define physical port scale factors).

d. In the Maximum MBS column, enter an MBS value that defines the maximum MBS value allowed for each range of SCR values.

For example, if you enter a maximum MBS value of 256 for the range of SCR values (0-10000), a circuit with an SCR of 7,000 cells/sec and MBS of 300 cells is rejected by the CAC function because its MBS exceeds the specified maximum MBS.

4. When you finish, choose OK to send the values you entered to the selected switch.

5. To send these values to another switch on the network map, select that switch and use Step 2 on page A-4 to access the Modify CAC Parameters dialog box. Choose OK.

Customizing the CAC for the VBR-NRT and ABR Classes

To customize the CAC for VBR-NRT and ABR:

1. On the Modify CAC Parameters dialog box (Figure A-1 on page A-4), select the customize VBR-NRT and ABR CAC implementation.
2. Refer to "Tuning the Ascend CAC" on page A-5 to enter the desired values for the Cell Loss Ratio and Cell Delay Variation objectives.
3. Refer to "Customizing the CAC for the VBR-RT, VBR-NRT, and ABR Classes" on page A-7 to enter the desired values in the Port Scale Factors, SCR ranges, and SCR scale factors and maximum MBS fields.
4. When you finish, choose OK to send these values to the selected switch.

B

ATM Traffic Descriptors

This appendix describes how each traffic descriptor combination affects the cell streams under different traffic conditions. When you create either a PVC or a point-to-multipoint circuit, you select one of several traffic descriptor combinations. The traffic descriptor combination specifies which traffic parameters are used for traffic control. It also determines the number and type of cells that are admitted into a congested queue, and whether or not high-priority cells are tagged as low-priority cells when traffic exceeds the traffic parameter thresholds.

PCR CLP=0, PCR CLP=0+1

You can select this option for constant bit rate (CBR) traffic. Traffic conformance is based on the peak cell rate (PCR) of both the cell loss priority (CLP)=0 and CLP=0+1 cell streams with no Tagging. The cell streams are checked for traffic conformance as follows:

- The switch checks the cell rate of the CLP=0 stream; if the cell rate exceeds the PCR of CLP=0, the switch drops the CLP=0 cells arriving above that rate.
- The switch checks the cell rate of the CLP=0+1 stream; if the cell rate exceeds the PCR of CLP=0+1, the switch drops cells arriving above that rate. Cells are dropped according to a ratio of CLP=0 to CLP=1 cells.

  For example, if the ratio of CLP=0 to CLP=1 cells is 8 to 5, approximately 8 CLP=0 cells are dropped for every 5 CLP=1 cells that are dropped.

Table B-1 illustrates what would happen to CLP=0 and CLP=1 cells in different situations if you select this option. This example assumes you set the PCR for CLP=0 to 50,000 cells/sec and the PCR for CLP=0+1 to 70,000 cells/sec.

All values in the table represent the measured traffic rate at a given point in time.

Table B-1. PCR CLP=0, PCR CLP=0+1

CLP=0 Cells/sec	CLP=1 Cells/sec	Result
45,000	22,000	The switch does not drop any cells because the CLP=0 and CLP=0+1 streams did not exceed the PCR.
50,000	22,000	The switch drops 2,000 cells/sec because the cell transmission rate exceeded the PCR of the CLP=0+1 cell stream. Since the ratio of CLP=0 to CLP=1 cells is 50 to 22, approximately 50 CLP=0 cells are dropped for every 22 CLP=1 cells that are dropped.
55,000	17,000	Since CLP=0 exceeds the PCR, the switch drops 5,000 CLP=0 cells/sec. This leaves 67,000 cells/sec in the CLP=0+1 stream, which is below the PCR of CLP=0+1. Therefore, no additional cells are dropped.
55,000	22,000	Since CLP=0 exceeds the PCR, the switch drops 5,000 CLP=0 cells/sec. This leaves 72,000 cells/sec in the CLP=0+1 stream, which also exceeds the traffic contract. Therefore, 2,000 additional cells/sec are dropped. Since the ratio of CLP=0 to CLP=1 cells is 50 to 22, approximately 50 CLP=0 cells will be dropped for every 22 CLP=1 cells that are dropped.

PCR CLP=0, PCR CLP=0+1, Tagging

You can select this option for CBR traffic. Traffic conformance is based on the PCR of both the CLP=0 and CLP=0+1 cell streams with Tagging enabled. The cell streams are checked for traffic conformance as follows:

- The switch checks the cell rate of the CLP=0 stream; CLP=0 cells arriving above the PCR of CLP=0 are tagged as CLP=1 cells.
- The switch checks the cell rate of the CLP=0+1 stream; if the cell rate exceeds the PCR of CLP=0+1, the switch drops additional cells, based approximately on the ratio of CLP=0 to CLP=1 cells.

  For example, if the ratio of CLP=0 to CLP=1 cells is 8 to 5, approximately 8 CLP=0 cells are dropped for every 5 CLP=1 cells that are dropped.

Table B-2 illustrates what would happen to CLP=0 and CLP=1 cells in different situations if you select this option. This example assumes you set the PCR for CLP=0 to 50,000 cells/sec and the PCR for CLP=0+1 to 70,000 cells/sec.

All values in the table represent the measured traffic rate at a given point in time.

Table B-2. PCR CLP=0, PCR CLP=0+1, Tagging

CLP=0 Cells/sec	CLP=1 Cells/sec	Result
45,000	22,000	The switch does not tag or drop any cells because the CLP=0 and CLP=0+1 streams did not exceed the PCR.
50,000	22,000	The switch drops 2,000 cells/sec because the cell transmission rate exceeded the PCR of the CLP=0+1 cell stream. Since the ratio of CLP=0 to CLP=1 cells is 50 to 22, approximately 50 CLP=0 cells are dropped for every 22 CLP=1 cells that are dropped.
55,000	17,000	Since CLP=0 exceeds the PCR, 5,000 CLP=0 cells/sec are tagged as CLP=1. This still leaves 72,000 cells/sec in the CLP=0+1 stream, which exceeds the PCR of CLP=0+1. Therefore, 2,000 cells/sec are dropped. Since the ratio of CLP=0 to CLP=1 cells is 50 to 22, approximately 50 CLP=0 cells are dropped for every 22 CLP=1 cells that are dropped.
55,000	22,000	Since CLP=0 exceeds the PCR, 5,000 CLP=0 cells/sec are tagged as CLP=1 cells. This still leaves 77,000 cells/sec in the CLP=0+1 stream, which exceeds the PCR of CLP=0+1. Therefore, 7,000 cells/sec are dropped. Since the ratio of CLP=0 to CLP=1 cells is 50 to 27, approximately 50 CLP=0 cells are dropped for every 27 CLP=1 cells that are dropped.

PCR CLP=0+1

You can select this option for CBR and unspecified bit rate (UBR) traffic. Traffic conformance is based only on the PCR of the CLP=0+1 aggregate cell stream with no best effort. If you select this option, when the cell rate of the aggregate cell stream exceeds the specified PCR of CLP=0+1, the switch drops all non-conforming cells, whether they are CLP=0 or CLP=1 cells.

PCR CLP=0+1, Best Effort

You can select this option only for UBR traffic. A “best effort” attempt is made to deliver all traffic, but there is no guarantee the switch will not drop cells due to congestion.

PCR CLP=0+1, SCR CLP=0, MBS CLP=0

You can select this option only for variable bit rate (VBR) traffic. Traffic conformance is based on the PCR of the CLP=0+1 aggregate cell stream, as well as the sustainable cell rate (SCR) and maximum burst size (MBS) of the CLP=0 cell stream with no Tagging. The cell streams are checked for traffic conformance as follows:

- The switch checks the cell rate of the CLP=0+1 stream; the switch drops cells arriving above the PCR. The number of CLP=0 and CLP=1 cells dropped is based approximately on the ratio of CLP=0 to CLP=1 cells.

  For example, if the ratio of CLP=0 to CLP=1 cells is 8 to 5, approximately 8 CLP=0 cells are dropped for every 5 CLP=1 cells that are dropped.

- The switch checks the SCR and the MBS of the CLP=0 stream. If the cell rate exceeds the SCR, cells arriving above the SCR are admitted until the stream exceeds tolerance for such cells. Tolerance is based on the MBS, PCR, and cell delay variation tolerance (CDVT). The switch drops cells that arrive above the SCR once the stream exceeds this tolerance level.

For more information about these traffic conformance parameters, see the *ATM UNI Specification, Version 3.1* or *Bellcore’s GR-1110-CORE Specification*.

Table B-3 illustrates what happens to CLP=0 and CLP=1 cells in different situations if you select this option. This example assumes you set the traffic parameters as follows:

- PCR of CLP=0+1 is 70,000 cells/sec
- SCR of CLP=0 is 40,000
- MBS of CLP=0 is 32

All values in the table represent the measured traffic rate at a given point in time.

Table B-3. PCR CLP=0+1, SCR CLP=0, MBS CLP=0

CLP=0+1 Cells/sec	SCR of CLP=0 Stream	MBS of CLP=0 Stream	Result
68,000	40,000	30	The switch does not drop any cells because the stream does not exceed traffic parameters.
70,000	40,000	60	The switch drops CLP=0 cells from the aggregate cell stream if the burst tolerance is exceeded. The number of cells that are dropped depends on the traffic pattern combination of sustained and burst cells. The larger the burst, the more cells are dropped.
70,000	50,000	30	The switch drops 10,000 CLP=0 cells/sec because CLP=0 exceeds the SCR. It may drop additional cells because the cell burst of 30 cells at PCR, combined with the sustained traffic, may exceed the burst tolerance.
77,000	40,000	60	The switch drops 7,000 cells/sec from the CLP=0+1 stream because the stream exceeds the PCR. The number of CLP=0 and CLP=1 cells dropped depends on the ratio of CLP=0 to CLP=1 cells in the aggregate stream. In addition, the switch will drop some CLP=0 cells if they exceed the burst tolerance.

PCR CLP=0+1, SCR CLP=0, MBS CLP=0, Tagging

You can select this option only for VBR traffic. Traffic conformance is based on the PCR of the CLP=0+1 aggregate cell stream, as well as the SCR and MBS of the CLP=0 cell stream with Tagging enabled. The cell streams are checked for traffic conformance as follows:

- The switch checks the cell rate of the CLP=0+1 stream; the switch drops cells arriving above the PCR of CLP=0+1. The number of CLP=0 and CLP=1 cells dropped is based approximately on the ratio of CLP=0 to CLP=1 cells.

  For example, if the ratio of CLP=0 to CLP=1 cells is 8 to 5, approximately 8 CLP=0 cells are dropped for every 5 CLP=1 cells that are dropped.

- The switch checks the SCR and the MBS of the CLP=0 stream. If the stream exceeds SCR, cells arriving above the SCR are admitted until the stream exceeds tolerance for such cells. Tolerance is based on the MBS, PCR, and CDVT. The switch tags cells that arrive above the SCR once the stream exceeds this tolerance level.

For more information about these traffic conformance parameters, see the *ATM UNI Specification, Version 3.1* or *Bellcore's GR-1110-CORE Specification*.

Table B-4 illustrates what happens to CLP=0 and CLP=1 cells in different situations if you select this option. This example assumes you set the traffic parameters as follows:

- PCR of CLP=0+1 is 70,000 cells/sec
- SCR of CLP=0 is 40,000
- MBS of CLP=0 is 32

All values in the table represent the measured traffic rate at a given point in time.

Table B-4. PCR CLP=0+1, SCR CLP=0, MBS CLP=0, Tagging

CLP=0+1 Cells/sec	SCR of CLP=0 Stream	MBS of CLP=0 Stream	Result
68,000	40,000	30	The switch does not drop or tag any cells because the stream does not exceed traffic parameters.
70,000	40,000	60	CLP=0 cells from the aggregate cell stream are tagged if the burst tolerance is exceeded. The number of cells that are tagged depends on the traffic pattern combination of sustained and burst cells. The larger the burst, the more cells are tagged.
70,000	50,000	30	The switch tags as many as 10,000 CLP=0 cells/sec because CLP=0 exceeds the SCR. It may tag additional cells because the cell burst of 30 cells at PCR, combined with the sustained traffic, may exceed the burst tolerance.
77,000	40,000	60	The switch drops 7,000 cells/sec from the CLP=0+1 stream because CLP=0+1 exceeds the PCR. The number of CLP=0 and CLP=1 cells that are dropped depends on the ratio of CLP=0 to CLP=1 cells in the aggregate stream. In addition, the switch will tag some CLP=0 cells if they exceed the burst tolerance.

PCR CLP=0+1, SCR CLP=0+1, MBS CLP=0+1

You can select this option only for VBR traffic. Traffic conformance is based on the PCR, SCR, and MBS of the CLP=0+1 cell stream with no Tagging. The cell streams are checked for traffic conformance as follows:

- The switch checks the cell rate of the CLP=0+1 stream; the switch drops cells arriving above the PCR of CLP=0+1. The number of CLP=0 and CLP=1 cells that it drops is based approximately on the ratio of CLP=0 to CLP=1 cells.

  For example, if the ratio of CLP=0 to CLP=1 cells is 8 to 5, approximately 8 CLP=0 cells are dropped for every 5 CLP=1 cells that are dropped.

- The switch checks the SCR and the MBS of the CLP=0+1 stream. If the stream exceeds SCR, cells arriving above the SCR are admitted until the stream exceeds tolerance for such cells. Tolerance is based on the MBS, PCR, and CDVT. The switch drops cells that arrive above the SCR once the stream exceeds this tolerance level.

For more information about these traffic conformance parameters, see the *ATM UNI Specification, Version 3.1* or *Bellcore's GR-1110-CORE Specification*.

Table B-5 illustrates what happens to CLP=0 and CLP=1 cells in different situations if you select this option. This example assumes you set the traffic parameters as follows:

- PCR of CLP=0+1 is 70,000 cells/sec
- SCR of CLP=0+1 is 40,000
- MBS of CLP=0+1 is 32

All values in the table represent the measured traffic rate at a given point in time.

Table B-5. PCR CLP=0+1, SCR CLP=0+1, MBS CLP=0+1

CLP=0+1 Cells/sec	SCR of CLP=0+1 Stream	MBS of CLP=0+1 Stream	Result
68,000	40,000	30	The switch does not drop any cells because the streams do not exceed traffic parameters.
70,000	40,000	60	CLP=0+1 cells are dropped from the aggregate cell stream if the burst tolerance is exceeded. The number of cells that are dropped depends on the traffic pattern combination of sustained and burst cells. The larger the burst, the more cells are dropped.

Table B-5. PCR CLP=0+1, SCR CLP=0+1, MBS CLP=0+1 (Continued)

CLP=0+1 Cells/sec	SCR of CLP=0+1 Stream	MBS of CLP=0+1 Stream	Result
70,000	50,000	30	The switch drops 10,000 CLP=0+1 cells/sec because CLP=0+1 exceeds the SCR. It may drop additional cells because the cell burst of 30 cells at PCR, combined with the sustained traffic, may exceed the burst tolerance.
77,000	40,000	60	The CLP=0+1 stream drops 7,000 cells/sec. because CLP=0+1 exceeds the PCR. The number of CLP=0 and CLP=1 cells that the switch drops depends on the ratio of CLP=0 to CLP=1 cells in the aggregate stream. In addition, the switch may drop some CLP=0+1 cells if they exceed the burst tolerance.

C

Allocating Logical Port Bandwidth on CBX 500 Shared SP Threads

CBX 500 chassis slots 3-4, 5-6, 7-8, 9-1, 10-2, 11-12, 13-14, and 15-16 are associated with the SP threads. This means that if you have an IOM installed in slots 3 and 4, you are “sharing” an SP thread. If you have an IOM in slot 9 or 10, you are sharing a thread with the SP itself. In this case, there are no thread limitations; the IOM has the full 599.040 Mbps of bandwidth available.

If two IOMs share the same SP thread, the maximum user cell bandwidth available to the two IOMs is 599.040 Mbps (599040 kbs. or 1412830 cps.). The NMS now enforces this limit such that the combined sum of all logical port bandwidths on the two IOMs cannot exceed 599.040 Mbps. These bandwidth limitations ensure the QoS guarantee even when you install two IOMs on the same SP fabric thread. Even with this thread bandwidth enforcement, you may still oversubscribe the VBR and UBR service classes on some or all of the IOM ports to utilize the statistical multiplexing gains that are an inherent part of running with two IOMs on one SP thread. However, you should carefully plan such oversubscription according to the intended service offerings and network engineering considerations of the different logical ports that share the thread.

The 599.040 Mbps number is derived from the maximum user cell bandwidth supported by the OC12/STM-4 interface. The OC12/STM-4 physical layer bandwidth is 622.080 Mbps, but the maximum user traffic bandwidth that any OC12/STM-4 port can support is 599.040 Mbps. This 599.040 thread limitation is also derived from the maximum user cell bandwidth that the four OC3/STM-1 interfaces can support. OC3/STM-1 physical layer bandwidth is 155.020 Mbps, but the maximum user traffic bandwidth that any OC3/STM-1 port can support is 149.76 Mbps. Refer to “About Logical Port Bandwidth” on page 2-11 for a detailed description of mapping physical port bandwidth to logical port bandwidth.

The 599.040 Mbps bandwidth value is available exclusively for user cell traffic. Management and internal switch control traffic have the potential to use a maximum of 11 Mbps of thread bandwidth, but this value is already factored into the total available thread bandwidth. The total available thread bandwidth starts at 611 Mbps, and once the NMS reserves 11 Mbps for management and control traffic, 599.040 Mbps remains exclusively for user cell traffic. At no time does management or internal control traffic conflict with the 599.040 Mbps of user cell traffic. If user cell traffic exceeds 599.040 Mbps, user traffic may be lost (depending on the QoS class of the user cell traffic) if the following conditions exist:

- It is a lesser priority than the management and internal control traffic
- It exceeds the overall 611 Mbps thread capacity

This NMS enforcement of SP thread bandwidth only applies when the switch has two IOMs installed on the same SP thread. If the switch only has one IOM on a thread, the maximum possible logical port bandwidth for all ports on the IOM is supported by the 599.040 Mbps limit.

Shared SP Thread Example

When a switch has two IOMs installed on an SP thread, you will notice the NMS enforcement of the SP thread bandwidth whenever you attempt to provision two OC3 cards on the same SP fabric thread. As you provision logical ports, the NMS subtracts the assigned bandwidth from the 599.040 Mbps total. After you provision four OC3 logical ports on the first OC3 card using the maximum 149.76 Mbps of bandwidth, there will not be any bandwidth left for the other OC3 card and its logical ports.

Consequently, when you have two cards installed on the same fabric thread, Ascend recommends that you allocate the bandwidth accordingly, across all of the IOM ports. In this example, you would allocate approximately 75 Mbps to each of the eight logical ports. This enables each logical port to support 75 Mbps of CBR traffic, and consequently allows full use of the thread bandwidth.

Even when you use 75 Mbps per logical port, you can still oversubscribe the logical port to overbook the VBR and UBR service classes on the port. For example, by reserving 10% of each logical port’s bandwidth (i.e., 75 Mbps) for UBR traffic, and overbooking the UBR bandwidth, hundreds of UBR circuits can be set up. Since UBR circuits are not policed, these best-effort UBR circuits can potentially utilize the full port bandwidth of each logical port, and consequently the full thread bandwidth. However, at periods when the combined UBR traffic exceeds thread bandwidth, the excess UBR traffic is dropped.

D

Implementing CBX 500 ATM Flow Control

The CBX 500 ATM Flow-Control Processor supports ATM traffic management through binary, hop-by-hop, and closed-loop, flow-control algorithms that shift network congestion to the edge of the network. In addition, the CBX 500 ATM Flow-Control Processor uses several per-virtual circuit (VC) cell/packet queuing and discarding mechanisms for additional network congestion control.

Based on the ATM Forum's *Traffic Management Specification*, Version 4.0, the ATM Flow-Control Processor delivers a fair, deterministic service for bursty ATM traffic, including:

- Dynamically adjusting the allowable cell rate (ACR) in response to Resource Management (RM) cell feedback
- Reducing congestion in the network by adjusting the data rate at which a VC sends cells
- Fair resource allocation based on the minimum cell rate (MCR)
- Per VC-queuing with early packet discard/partial packet discard (EPD/PPD) capability

Supported ATM Service Classes

The ATM Flow-Control Processor supports three ATM service classes:

Enhanced Unspecified Bit Rate (UBR+) Class — The ATM Flow-Control Processor provides a UBR+ service for the UBR Quality of Service (QoS) class by applying:

- Closed-loop flow control
- Dynamic cell rate adjustment
- Minimum Cell Rate (MCR) guarantee

The MCR is set at the minimum rate of the ATM Flow-Control Processor. The ACR is adjusted by the Rate Increase Factor (RIF) and the Rate Decrease Factor (RDF). Both the RIF and the RDF are configurable through NavisCore. See the *NavisCore Physical Interface Configuration Guide* for more information about configuring the RDF and RIF.

Available Bit Rate (ABR) Class — The MCR is configured during circuit admission. The ACR is adjusted by the RDF and RIF. You can configure the RDF and RIF in proportion to the MCR. See the *NavisCore Physical Interface Configuration Guide* for more information about configuring the RDF and RIF.

Variable Bit Rate-Non-Real Time (VBR-nrt) Class — The ATM Flow-Control Processor can manage the VBR-nrt QoS class. VBR-nrt is selectable through NavisCore. The sustainable cell rate (SCR) is configured during circuit admission. The SCR is used in the same way as the MCR during ACR adjustments. The ACR is adjusted by the RIF and RDF. You can configure the RIF and RDF in proportion to the SCR. See the *NavisCore Physical Interface Configuration Guide* for more information about configuring the RDF and RIF.

ATM Flow-Control Processor Architecture

The ATM Flow-Control Processor provides per-VC queuing, and supports the CBX 500 quad-plane buffer architecture. Figure D-1 shows the ATM Flow-Control Processor output buffers relative to the CBX 500 quad-plane output buffers.

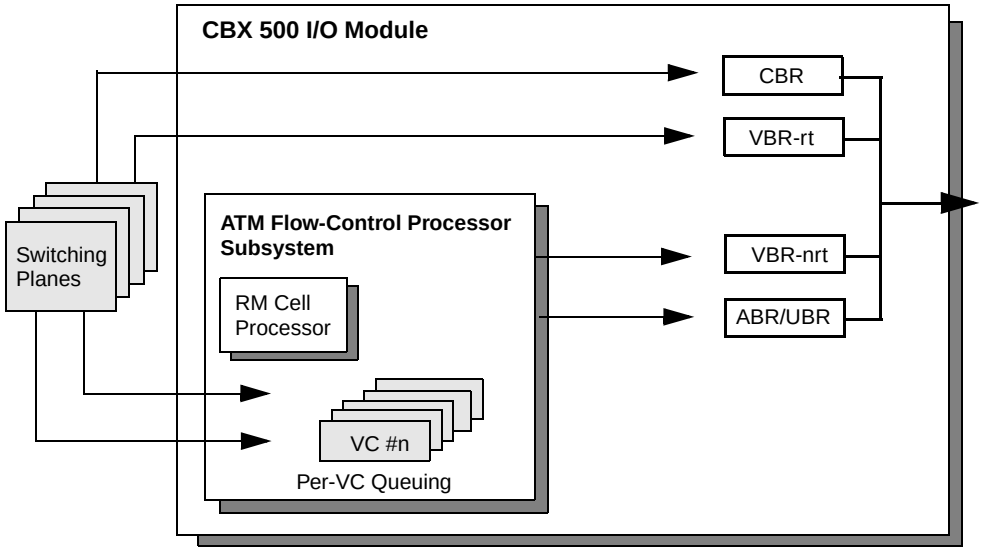

Figure D-1. CBX 500 Queues and the ATM Flow-Control Processor

Cells from the CBX 500 switching fabric are queued at the ATM Flow-Control Processor queues. Note that the ATM Flow-Control Processor only queues non-real time QoS VCs.

Cells are queued and dequeued based on the configured rate for the VC. Each VC is subject to discard mechanisms. Cells entering the output CBX 500 quad-plane queues are scheduled based on the Connection Admission Control (CAC) scheduling algorithm. See Appendix A, “Adjusting the CAC” for more information about the CAC.

Closed-Loop Flow Control

Ascend’s closed-loop flow-control architecture is based on hop-by-hop control loops with binary feedback. The hop-by-hop control loops push congestion at central nodes to switches at the edge of the network, thereby providing more efficient use of network bandwidth. In addition, with less network congestion at central nodes, there is increased network throughput.

Flow Control Mechanisms

The ATM Flow-Control Processor supports three closed-loop, flow-control mechanisms:

Cascade Communications Resource Management (CCRM) Cells — CCRM cells are a subset of the ATM Forum's *ATM Traffic Management Specification*, Version 4.0, ABR RM cells. The Protocol ID field in each RM cell is defined as the CCRM ID, indicating that it is a CCRM cell. The default value for the CCRM ID is 6. You can change the Protocol ID in the event that another switch vendor is using the default value for their proprietary loops. See the *NavisCore Physical Interface Configuration Guide* for information about provisioning CCRM cells.

Backward Congestion Message (BCM) Cells — BCM cells provide a different RM cell mechanism and may also provide interoperability with other manufacturers' ATM switches. The Protocol ID field in each BCM cell is defined as the BCM ID. The default value for the BCM ID is 5. You can change the Protocol ID in the event that another switch vendor is using the default value for their proprietary loops. See the *NavisCore Physical Interface Configuration Guide* for information about provisioning BCM cells.

Because the CBX 500 communicates with either CCRM or BCM cells for hop-by-hop control loops, both CCRM and BCM cells can be configured within a single network, allowing conversion between one closed-loop, flow-control algorithm to another.

Available Bit Rate (ABR) RM Cells — The ATM Flow-Control Processor marks ABR RM cells with binary notification as defined in the ATM Forum's *Traffic Management Specification*, Version 4.0. The Protocol ID for an ABR RM cell is 1. The ATM Flow-Control Processor identifies any RM cell with a Protocol ID of 1 as an ABR RM cell.

RM Cell Generation (General)

You can configure any port on an IOM to generate:

- CCRM cells
- BCM cells
- No RM-type cells

This allows for different closed-loop, flow-control algorithms to be implemented on the same IOM.

Because RM cells are generated in the backward direction, the type of RM cells generated depends on the logical port configuration through which they are transmitted.

In general, RM-type cells can be generated at 30 to 250 millisecond (ms) intervals per VC. The default value for this parameter is 100 ms.

Table D-1 shows an example of the maximum number of circuits you can configure when using a particular RM Cell Interval.

Table D-1. Minimum RM Cell Intervals

RM Cell Interval	Maximum Supported VCs
100 ms	12K
50 ms	6K
30 ms	4K

Figure D-2 shows hop-by-hop, closed-loop flow control between four CBX 500 switches. The flow-control loops are shown as solid lines. The data paths are shown as dotted lines.

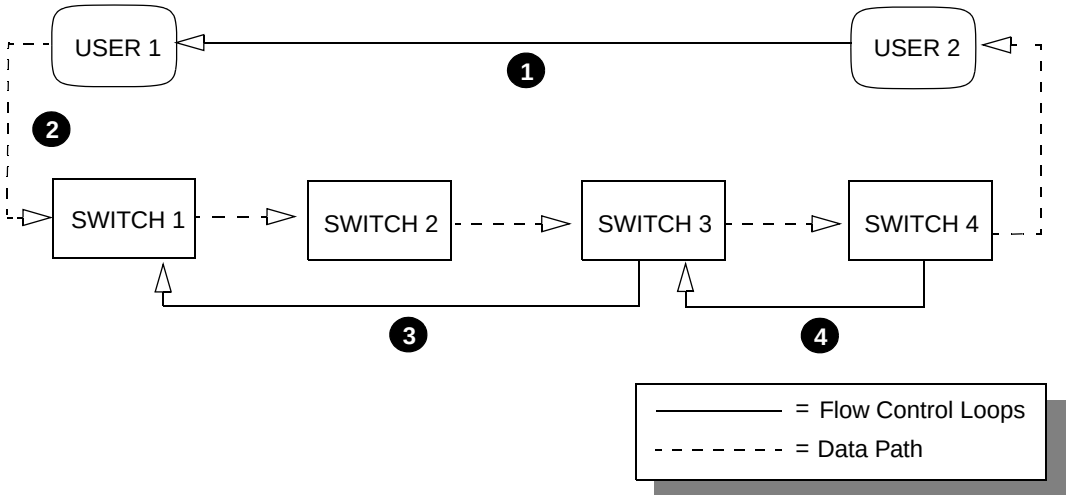

1 End-to-End User Control Loops

End-to-end user flow-control loops are “outer” loops. The switches do not change their cell rates in response to this flow-control loop. Instead, they mark the congestion indication (CI) and no increase (NI) bits based on the local congestion state, as defined in the ATM Forum’s *Traffic Management Specification*, Version 4.0.

2 Different Logical Port Types on the Same I/O Module

The ATM Flow-Control Processor supports different types of flow-control loops on the same I/O module. USER 1 has a user-to-network Interface (UNI) connection. SWITCH 2 has a trunk connection to a different port on the same I/O module in SWITCH 1. Enabling and disabling of loop control is provisioned per port.

3 Switches Without Flow-Control Loops

SWITCH 2 does not generate or terminate flow-control loops to the other switches. SWITCH 2 generates a forward notification of congestion to SWITCH 3 (explicit forward congestion indication - (EFCI) marking can be configured on a CBX 500 switch through NavisCore. When SWITCH 2 marks EFCI in the data cells, SWITCH 3 can be configured to include EFCI notification in the decision of the backward notification to SWITCH 1.

4 Rate Control at the Output Switch

The SWITCH 4 cell rate fills the available bandwidth and is adjusted based on local congestion. The flow-control loop between SWITCH 4 and USER 2 can be configured as either BCM or CCRM termination. If configured as BCM, SWITCH 4 will adjust rates according to the port congestion. If configured as CCRM, SWITCH 4 will perform traffic shaping to the ICR of each VC.

Figure D-2. Closed-Loop Flow Control

CCRM Closed-Loop Flow Control

Ascend's closed-loop, flow-control architecture can use CCRM cells to notify CBX 500 switches of network congestion or availability.

CCRM Closed-Loop Flow Control on a Trunk

Figure D-3 shows an example of CCRM closed-loop flow control between two CBX 500 switches.

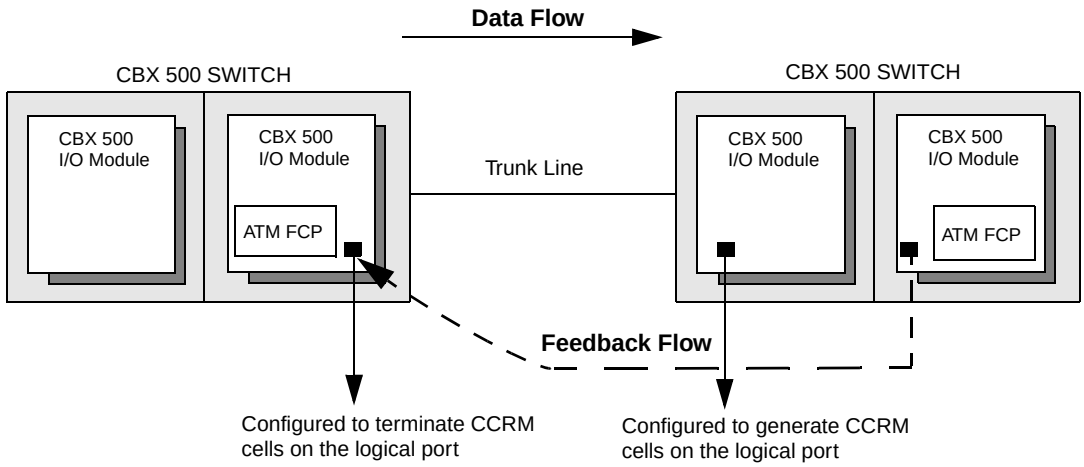

Figure D-3. CCRM Closed-Loop Flow Control

CCRM Closed-Loop Flow Control on a UNI (Traffic Shaping)

For information on CCRM Closed-Loop Flow Control on a UNI, see "Per-VC Traffic Shaping" on page D-13.

CCRM Cell Generation

The following results occur when a CCRM cell is generated:

- The direction (DIR) and backward indicator (BI) bits are set, indicating that this is a switch-generated backward RM cell.
- The congestion indication (CI) and no increase (NI) bits are set according to the current congestion status of the VC.

The destination ATM switch periodically sends backward binary notification through CCRM cells to the source ATM switch, indicating the state of the destination ATM switch's queue for a VC. The binary notification is reflected in the CI and NI bits of the CCRM cell. The CCRM cell indicates a cell rate increase, decrease, or no change. The source ATM switch then responds by adjusting the cell rate accordingly for that VC and terminates the CCRM cell.

BCM Closed-Loop Flow Control

The CBX 500 can also utilize a BCM closed-loop, flow-control algorithm. Unlike CCRM cells, BCM cells only indicate cell rate decreases. BCM cells are sent on periodic intervals only when congestion exists.

During the RM cell generation interval, the allowable cell rate (ACR) for a VC is increased if:

- A BCM cell is not received over the previous RM cell interval.
- The port is not congested.

BCM Closed-Loop Flow Control on a Trunk

Figure D-4 shows an example of BCM closed-loop flow control between two CBX 500 switches.

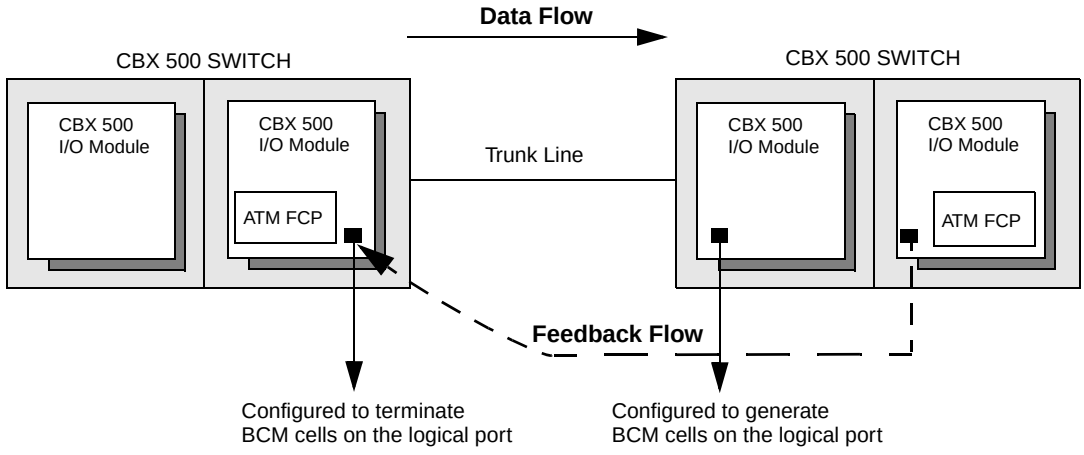

Figure D-4. BCM Closed-Loop Flow Control

BCM Closed-Loop Flow Control on a UNI

You can configure an output UNI logical port to allow ATM Flow-Control Processor-managed VCs going through that logical port to increase their cell rates. This enables the logical port to use all available non-real time bandwidth. This is done by setting the RM termination type on that logical port to BCM, as shown in Figure D-5.

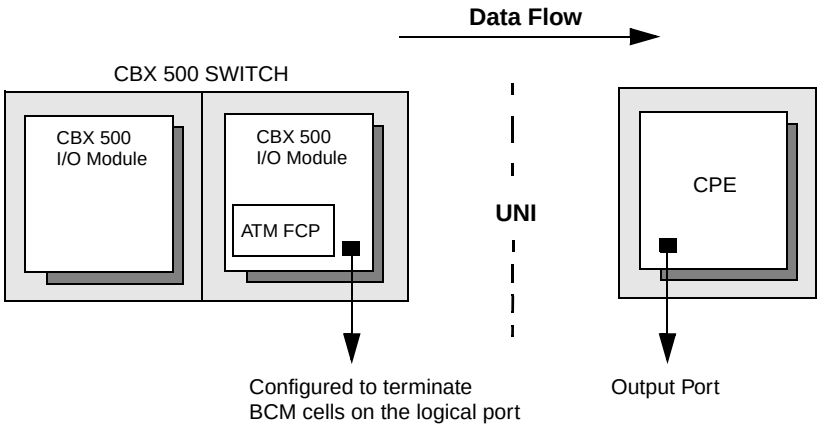

Figure D-5. Output UNI Logical Port RM Termination

Because the logical port does not receive any BCM cells from the UNI, the ACR of the VCs keeps increasing until the logical port becomes congested. The ACR will increase fairly, corresponding to the RIF and PCR values of the VCs. See "Flow Control Processor Attributes" on page 3-21 for information on setting the RM termination type.

Generating BCM Cells

You can configure any port on an I/O module to generate BCM cells. If you select the BCM generation option when configuring the ATM Flow-Control Processor, BCM cells are generated when the port is congested. See Table D-1 on page D-5 for the RM cell intervals and the number of supported VCs. See the *NavisCore Physical Interface Configuration Guide* for information on the BCM generation option.

Terminating CCRM and BCM Cells

When the CBX 500 terminates either a CCRM or BCM cell, the CBX 500 makes a decision on whether or not to increase or decrease the ACR. This decision is based upon one or more of the following:

- The local port congestion state
- The current ACR being above the fair bandwidth for the VC
- The CI and NI state in the CCRM cell
- If no BCM cells were received within the RM generation interval (if the port is configured for BCM termination)
- If BCM cells are received (if the port is configured for BCM termination)

If BCM cells are received, but the port is not configured for BCM termination, the BCM cells are forwarded.

The fair bandwidth for a VC is the proportional allocation of the total bandwidth for managed (non-real time) circuits, based on the MCR of the VC relative to all of the managed VCs. The total, non-real time bandwidth is the total port bandwidth, less the bandwidth allocated to unmanaged (real-time) circuits.

Note that the ATM Flow-Control Processor can increase the ACR well beyond its fair bandwidth. Once other circuits attempt to use that bandwidth, (causing a congestion condition), the ATM Flow-Control Processor will throttle back the ACR towards the fair bandwidth for the circuit until the congestion condition is removed.

ABR RM Closed-Loop Flow Control

ABR RM closed-loop flow control is an additional flow-control loop for switches that generate ABR RM cells. Because the ABR RM flow-control loop is an end-to-end loop, the CBX 500 does not generate or terminate ABR RM cells. Instead, the ATM Flow-Control Processor marks the CI and NI bits in the ABR RM cell based on the local ATM Flow-Control Processor congestion state. The ATM Flow-Control Processor then forwards the ABR RM cells through the network.

Cell Rate Adjustment

When a VC initially becomes active, its ACR is set to its initial cell rate (ICR). The ICR for a VC is determined by its PCR, MCR, and ICR Constant.

ICR Constant

You configure the ICR Constant through NavisCore. The default value is 0. The following formula shows how to calculate the ICR Constant:

$$ICR = MCR + \frac{PCR-MCR}{2^{ICR\ CONSTANT}}$$

See the *NavisCore Physical Interface Configuration Guide* for information on configuring the ICR Constant.

Idle VC Factor

The specified number of RM intervals for a VC to go idle is configurable through NavisCore. This is called the Idle VC Factor. The default value for the Idle VC Factor is 8. See the *NavisCore Physical Interface Configuration Guide* for information about configuring the Idle VC Factor.

If no cells are received for a specified number of RM cell intervals, the VC is marked “idle,” and the ACR is set to the ICR. RM cells are not generated for idle VCs.

Rate Decrease Factor (RDF) and Rate Increase Factor (RIF)

This section provides information on configuring RDF and RIF values. The cell rate of a VC is decreased according to the following formula:

$$ACR = ACR - (RDF \times ACR)$$

Where: $1/32768 \leq RDF \leq 1$

The ACR is lower-bounded by the MCR.

The rate of a VC is increased according to the following formula:

$$ACR = ACR + (RIF \times PCR)$$

Where: $1/32768 \leq RIF \leq 1$

The ACR is upper-bounded by the PCR.

The RDF and the RIF values are configurable through NavisCore. See "Rate Profile Tables" on page D-13.

Table D-2 lists the minimum allocated MCR for ABR and UBR circuits.

Table D-2. Cell Scheduling

Port Bandwidth	Max. Port Cell Rate (cells/sec)	Max. Number of Circuits (connections/port)	Min. Allocated MCR (cells/sec)
OC12	1412830	16K	88
OC3	353207	4K	88
DS3	96000	2K	48
E3	8000	2K	40
DS1	3622	2K	1.8
E1	4528	2K	2.4

Rate Profile Tables

You can load two rate profile tables into the ATM Flow-Control Processor. The ATM Flow-Control Processor uses these tables to determine the Rate Increase Exponent (RIE) and the Rate Decrease Exponent (RDE) for each VC on a port.

Rate Increase Exponent (RIE) — The RIE is a provisionable value that is the negative exponent for the RIF calculation ($RIF=2^{-RIE}$). For example, a RIE of 3 translates to a RIF of 1/8. The RIE must be less than 16.

Rate Decrease Exponent (RDE) — The RDE is a provisionable value that is the negative exponent for the RDF calculation ($RDF=2^{-RDE}$). For example, a value of 3 translates to a RDF of 1/8. The RDE must be less than 16.

You use the RIE and RDE values to compute the RIF and the RDF. Each table consists of 256 entries. See the *NavisCore Physical Interface Configuration Guide* for information on using NavisCore to download these tables.

The RIF and RDF value for any VC is obtained from indexing the corresponding rate profile table with the VC's MCR class. The MCR (SCR for VBR-nrt VCs) of any VC is mapped to one of 256 MCR classes. Note that MCR class 0 is reserved for UBR VCs. See the *NavisCore Physical Interface Configuration Guide* for information about MCR class mappings per I/O module.

Per-VC Traffic Shaping

You can configure the ATM Flow-Control Processor to perform traffic shaping for ATM Flow-Control Processor-managed VCs on a trunk or UNI port by turning off the control loops for these VCs. For any direction of data flow, you can:

- Configure the RM cell generation of all input logical ports the VC passes through to "no loop." See page 3-21 for information.
- Configure all output logical ports the VC passes through to terminate CCRM cells. See page 3-21 for information.

VCs are shaped at their Initial Cell Rate (ICR). See "ICR Constant" on page D-11 for a description of the ICR Constant. Because control loops are disabled, the ACR will stay at the ICR. Note that there is no guarantee of ICR if it is overbooked.

ATM Flow-Control Processor Queues

The ATM Flow-Control Processor provides per-VC queueing. Per-VC queuing provides independent buffer allocation to each VC, thereby isolating congestion on one VC from other VCs. Each per-VC queue has two configurable thresholds:

- Local congestion threshold
- Local discard threshold

The congestion and discard thresholds for a specific VC are obtained by indexing the congestion and discard tables with the MCR class of the VC. The MCR class of the VC is obtained from its MCR. See the *NavisCore Physical Interface Configuration Guide* for information about MCR classes.

MCR class 0 is reserved for UBR VCs.

You can configure the congestion and discard threshold tables through NavisCore. See the *NavisCore Physical Interface Configuration Guide* for information.

In addition to the local thresholds, each port on an IOM is assigned one:

- Global congestion threshold
- Global discard threshold
- Global CLP0+1 threshold

You can configure all of the above thresholds through NavisCore. See "Flow Control Processor Attributes" on page 3-21 for information.

Table D-3. ATM FCP Logical Port Threshold Defaults

Card Type:	T1/E1 DS3/E3	OC3/STM-1	OC12/STM-1
Port Buffers	8K Cells	16K Cells	64K Cells
CLP 0+1	7168	15360	64512
Discard	5120	13312	62464
Congestion	1024	2662	12493

Both local and global thresholds are used for congestion notification and discarding. A VC is considered congested only if its queue is above the local congestion threshold, and the global queue length on the port is above the global congestion threshold.

Similarly, a VC enters a discard state only if the VC queue length is greater than the VC discard threshold, and the port queue length is greater than the global discard threshold.

Figure D-6 shows the five ATM Flow-Control Processor buffer thresholds.

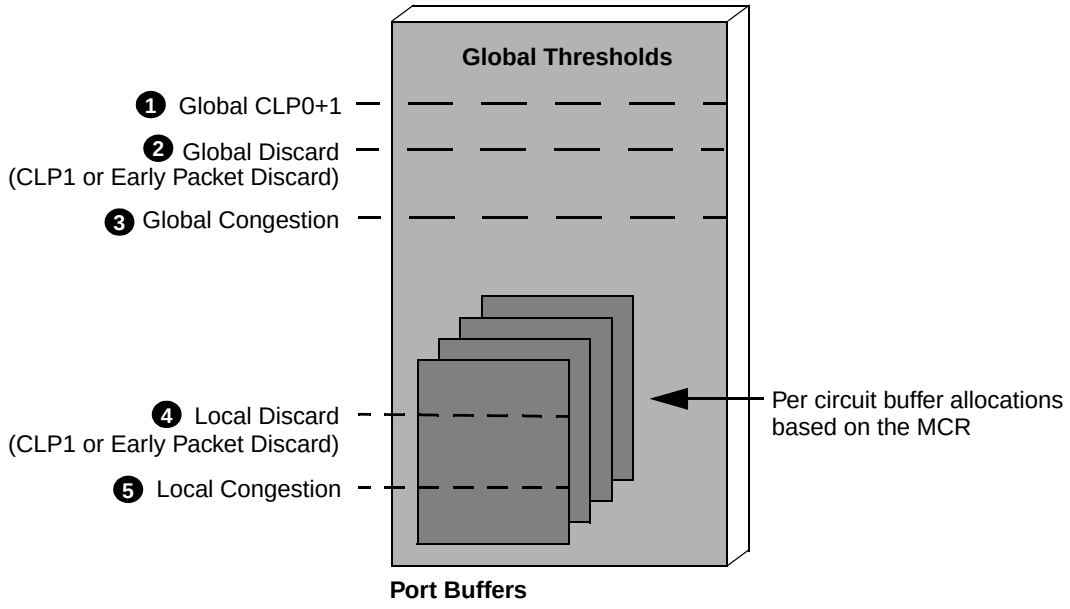

Figure D-6. ATM Flow-Control Processor Buffers

The difference between the early packet discard (EPD) and CLP0+1 threshold allows the VCs to continue to queue cells due to the EPD state. When the EPD threshold is exceeded, cells are queued on the current packet, and the next packet is discarded for the same VC.

The CLP0+1 threshold enables you to reserve buffers before the maximum buffer capacity is reached. Ascend recommends that you reserve a sufficient number of buffers to allow idle circuits to get access to buffers. Idle circuits are those that have temporarily stopped sending traffic.

ATM Flow-Control Discard Mechanisms

The ATM Flow-Control Processor supports three mechanisms for discarding cells:

Early Packet Discard (EPD) — The ATM Flow-Control Processor performs EPD for UBR, ABR, and VBR-nrt VCs. If a cell causes the queue for a VC 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 VC.

Selective Discard (CLP1) — Selective CLP1 discard can be provisioned for UBR, ABR, and VBR-nrt VCs. If the current cell causes the queue for a VC 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.

Partial Packet Discard (PPD) — If the global CLP0+1 threshold for a port is reached, PPD is performed for circuits that are configured for EPD. Unlike EPD, however, all of the remaining cells in the current packet are discarded. Note that the end-of-file (EOF) cell is discarded as well. This results in the loss of the next packet even if the packet is transmitted.

Note that the following conditions exist:
– A circuit is set for EPD and does not send AAL5 PDUs (e.g. AAL0 data)
– A port becomes sufficiently congested (CLP 0+1 threshold is reached)

The PPD results in no further throughput for this circuit. To regain service, you must re-establish the circuit.

Multicast Cells

All Multicast cells are placed into a single queue. There is one queue per I/O module. Multicast cells are discarded when the ATM Flow-Control Processor multicast queue length reaches a certain threshold. You can configure this threshold for each installed I/O module. See the *NavisCore Physical Interface Configuration Guide* for information.

Multicast cells are dequeued at the assigned Multicast cells shaping rate. This rate is configurable using NavisCore. See the *NavisCore Physical Interface Configuration Guide* for information.

E

Priority Routing

Priority Routing enables you to prioritize circuits in your network. This provides the following advantages:

- Higher up time for high-priority circuits
- Optimal paths for high-priority circuits which results in lower delay
- Higher capacity to burst past the guaranteed QoS rates for high-priority circuits

Priority routing introduces convergence time in the network. With priority routing, when a trunk is provisioned or fails, circuits are bumped based on priority. Higher priority circuits converge quickly compared to lower priority circuits. Total convergence time is directly proportional to the number of priorities defined in the network.

The switch treats priority routing, QoS class, and circuit priority as independent elements. Priority routing rules are used for connection setup. QoS class is applied after the connection is set up. Circuit priority rules are applied once QoS class is established. Keep in mind that you must assign a higher priority to real-time QoS classes.

About Routing Priorities

When you configure a logical port's SVC Routing Priority (see page 3-39) or PVC routing priority (see page 6-16), you specify the *bandwidth priority* and *bumping priority*, or level of importance, of each SVC or PVC in the network. The lower the number, the higher the priority.

Bandwidth priority — A value from 0 – 15, where 0 is the highest priority (default), is set for the bandwidth priority. This value is used in route calculations.

Bumping priority — A value from 0 – 7, where 0 is the highest priority (default). Setting this value to zero (0) for non-real time circuits means "keep it up always." This value is ignored for real time circuits. Bumping priority is only used at the time of bumping; this is after route calculations establish circuit priorities based on bandwidth priority values.

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 prioritizing occurs at the cost of the lower-priority circuits.

Priority Routing and Path Cost

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

Priority Routing and Path Cost 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. Multiple circuits within the network are defined with a priority routing of 2,0 and these virtual circuits 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.

Routing Priority Rules

The switch uses the following rules to implement priority routing at the time of circuit provisioning, trunk-failure recovery, and balance rerouting.

Circuit Provisioning

At the time of provisioning and load balance rerouting (following trunk failure), a circuit selects a path ignoring all circuits with lower bandwidth priority. In doing so, a circuit will force lower bandwidth priority circuits from their selected path until available link bandwidth is positive and can accommodate circuit bandwidth needs. The following sequence is used to force circuits from their path:

1. Bandwidth priority order, where lowest bandwidth priority circuits are chosen first. Keep in mind that bandwidth priority values range from 0 to 15, with 15 being the lowest priority.
2. Bumping priority order, where lowest bumping priority circuits are chosen first. Bumping priority values range from 0 to 8, with 8 being the lowest priority.
3. Equivalent bandwidth (EBW) order, where higher EBW circuits are chosen first.
4. Virtual channel identifier (VCI) order.

Trunk-Failure Recovery

PVCs always attempt to reroute themselves when a trunk goes down. The 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 the 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 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, PVCs with lower priority may be forced onto a path that is not optimal. See "Circuit Provisioning" on page E-3 for details about path selection.

F

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. Errors that can occur are presented 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 following conditions are listed:

- Effect on NMS database
- State of both switches, out of sync status
- Effect of performing a PRAM sync
- Special considerations

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 causes both endpoint cards to be 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 F-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 F-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 F-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 F-3. Errors Encountered During Circuit Delete Procedure

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. ***Note:*** *Retry process starts with 1st switch, which has a deleted circuit that results in an error message.*

G

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 G-1. Abbreviations

Abbreviation	Meaning
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)

Table G-1. Abbreviations (Continued)

Abbreviation	Meaning
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
usec	microsecond (abbreviate with lowercase “u” for micro)
sec	second
vs.	versus
#	number; pound
x	by (multi)

Acronyms

The following table lists some of the acronyms used in Ascend's WAN-switching documentation.

Table G-2. Acronyms

Acronym	Meaning
AAL	ATM Adaptation Layer
ABR	available bit rate
ACR	allowable cell rate
AESA	ATM End System Address
AFI	authority and format identifier
APS	automatic protection switching
ARP	Address Resolution Protocol
ASE	Autonomous System External
ASR	Application Specific Route
ATM	Asynchronous Transfer Mode
BCM	backward congestion message
BER	bit error rate
BI	backward indicator
B-ICI	BISDN-Inter-Carrier Interface
BIO	Base Input/Output
CAC	Connection Admission Control
CBR	constant bit rate
CCRM	Cascade Communications Resource Management
CDV	cell delay variation
CDVT	cell delay variation tolerance
CFR	constant frame rate
CI	congestion indication
CIC	carrier identification code

Table G-2. Acronyms (Continued)

Acronym	Meaning
CIR	committed information rate
CLP	cell loss priority
CLR	cell loss ratio
CP	control processor
CPE	customer premise equipment
CRC	cyclic redundancy check
CSR	Customer Specific Route
CSU	channel service unit
CTD	cell transfer delay
CUG	closed user group
DCC	data country code
DCE	data communications equipment
DE	discard eligible
DLCI	data link connection identifier
DNIC	data network identification code
DSL	digital subscriber line
DSU	data service unit
DTE	data terminal equipment
DXI	data exchange interface
EBR	excess burst rate
EBW	equivalent bandwidth
EFCI	explicit forward congestion indication
EPD	early packet discard
ESI	end system identifier
FEAC	far-end alarm and control
FECN	forward explicit congestion notification

Table G-2. Acronyms (Continued)

Acronym	Meaning
FCP	Flow Control Processor
FR	Frame Relay
FTP	File Transfer Protocol
FUNI	Frame-User-to-Network Interface
GUI	graphical user interface
HDLC	High-level Data Link Control
HO-DSP	high-order domain-specific part
IA	incoming access
IARP	Inverse Address Resolution Protocol
ICB	incoming calls barred
ICD	international country designator
ICMP	Internet Control Message Protocol
ICR	initial cell rate
IDI	initial domain identifier
IDP	initial domain part
IE	information element
IFNUM	interface number
IGMP	Internet Group Multicast Protocol
ILMI	Interim Link Management Interface
IOA	input/output adapter
IOM	input/output module
IOP	input/output processor
ISDN	Integrated Services Digital Network
IXC	inter-exchange carrier
KA	keep alive
LAN	local area network

Table G-2. Acronyms (Continued)

Acronym	Meaning
LAP	Link Access Protocol
LATA	Local Access and Transport Area
LTP	Link Trunk Protocol
MAC	Media Access Control
MBS	maximum burst size
MCR	minimum cell rate
MIB	Management Information Base
MPT	multipoint-to-point tunnel
MPVC	management permanent virtual circuit
NDC	network data collection
NHRP	Next Hop Routing Protocol
NI	no increase
NIC	network interface card
NMS	Network Management Station
NNI	Network-to-Network Interface
NP	node processor
NPA	node processor adapter
NPC	network parameter control
NTM	network traffic management
OA	outgoing access
OAM	Operations, Administration, and Maintenance
OCB	outgoing calls barred
OPTimum	Open Packet Trunking
OSPF	Open Shortest Path First
PAD	packet assembler/disassembler
PCR	peak cell rate

Table G-2. Acronyms (Continued)

Acronym	Meaning
PDN	public data network
PDU	protocol data unit
PMP	point-to-multipoint
PNNI	Private Network-to-Network Interface
PPD	partial packet discard
PPP	Point-to-Point Protocol
PRAM	parameter random access memory
PRI	Primary Rate Interface
PSA	proxy siganling agent
PSC	proxy signaling client
PTSP	PNNI Topology State Packet
PTSE	PNNI Topology State Element
PVC	permanent virtual circuit
PVP	permanent virtual path
QoS	Quality of Service
RADIUS	remote authentication dial-in user service
RBOC	Regional Bell Operating Company
RDE	Rate Decrease Exponent
RDF	Rate Decrease Factor
RFC	request for comments
RIE	Rate Increase Exponent
RIF	Rate Increase Factor
RIP	Routing Information Protocol
RM	resource management
SCR	sustainable cell rate
SD	signal degrade

Table G-2. Acronyms (Continued)

Acronym	Meaning
SF	signal fail
SLIP	Serial Line over Internet Protocol
SMDS	Switched Multimegabit Data Service
SNMP	Simple Network Management Protocol
SONET	Synchronous Optical Network
SP	switch processor
SPVC	soft permanent virtual circuit
SPVCC	soft permanent virtual channel connection
SPVPC	soft permanent virtual path connection
SVC	switched virtual circuit
SVCC	switched virtual channel connection
SVPC	switched virtual path connection
TCP	Transmission Control Protocol
TD	traffic descriptor
TDM	Time Division Multiplexing
TM	timing module
TNS	transit network selection
UBR	unspecified bit rate
UDP	User Datagram Protocol
UFR	unspecified frame rate
UIO	universal input/output
UNI	user-to-network interface
UPC	usage parameter control
VBR-RT/ NRT	variable bit rate-real time/non-real time
VC	virtual circuit
VCC	virtual channel connection

Table G-2. Acronyms (Continued)

Acronym	Meaning
VCI	virtual channel identifier
VCL	virtual circuit link
VFR-RT/ NRT	variable frame rate-real time/non-real time
VNN	Virtual Network Navigator™
VP	virtual path
VPC	virtual path connection
VPCI	virtual path connection identifier
VPI	virtual path identifier
VPN	Virtual Private Network
WAN	wide area network

Glossary

A

AAL

See *ATM Adaptation Layer*.

ABR

See *available bit rate*.

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*.

AFI

See *authority and format identifier*.

AIS

See *alarm indication signal*.

alarm

Message notifying an operator or administrator of a network problem.

alarm indication signal

An error or alarm signal transmitted in lieu of the normal signal. This signal maintains transmission continuity to the receiving node indicating that there is a transmission fault located either at the sending node or upstream of the sending node.

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.

asynchronous communications server

A LAN server that enables a network user to dial out of the network into either the public-switched telephone system, or to accessed leased lines for asynchronous communications. This device also is called a dial-in/dial-out server or modem server.

Asynchronous Transfer Mode

A method used for transmitting voice, video, and data over high-speed LAN and WAN networks. See also *cell relay*.

ATM

See *Asynchronous Transfer Mode*.

ATM Adaptation Layer

The standards layer that allows multiple applications to have data converted to and from the ATM cell. A protocol used that translates higher layer services into the size and format of an ATM cell. This layer is divided into four different levels of service:

AAL-1 offers AAL functions in support of constant bit rate, time-dependent traffic such as voice and video.

AAL-2 is still undefined by the International Standards bodies. It is a place holder for variable bit rate video transmission.

AAL-3/4 functions in support of variable bit rate, delay-tolerant data traffic requiring some sequencing and/or error detection support. Originally, two AAL types existed (connection-oriented and connectionless); these have been combined.

AAL-5 functions in support of variable bit rate, delay-tolerant, connection-oriented data traffic requiring minimal sequencing or error detection support.

ATM service interworking feeder

A service that enables Frame Relay network traffic to be fed into an ATM network, enabling a Frame Relay end user to communicate with an ATM end user.

ATM/DXI trunk

See *OPTimum PVC trunk*.

ATM/DXI trunk interface

An ATM circuit used as a trunk between two Frame Relay networks that are built with Ascend switches.

authority and format identifier

This identifier is part of the network level address header.

available bit rate

An ATM layer service category for which the limiting ATM layer transfer characteristics provided by the network may change subsequent to connection establishment. A specified flow control mechanism supports several types of feedback to control the source rate in response to changing ATM layer transfer characteristics. An end-system that adapts its traffic in accordance with the feedback is expected to experience a low cell loss ratio and obtain a fair share of the available bandwidth according to a network specific allocation policy. Cell delay variation is not controlled in this service, although admitted cells are not delayed unnecessarily.

B

backbone

The part of a network that carries the bulk of the network traffic, e.g., over Ethernet cabling, and fiber-optic cabling.

background diagnostics

Programs that run continuously in the background of the NMS to provide current operating status for all active switches. These programs do not interfere with switch operations.

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 transmission capacity of a computer or a communications channel.

bandwidth-on-demand

A WAN feature that enables users to dial up additional bandwidth as their applications demand.

Bc

See *committed burst size*.

Be

See *excess burst*.

BECN

See *Backward Explicit Congestion Notification*.

best-effort packets

Packets delivered to the best of the network's ability, after the requirements for delivering the guaranteed packets are met. See also *guaranteed packets*.

B-ICI

See *B-ISDN Inter-Carrier Interface*.

B-ISDN Inter-Carrier Interface

An ATM Forum-defined specification for the interface between public ATM networks to support user services across multiple public carriers.

bit

A binary unit of measurement, which may be either a one or a zero.

bits per second

The number of bits transmitted every second during a data transfer.

bps

See *bits per second*.

broadband network

A type of network that allows for the transmission of large amounts of information (including voice, data, and video) over long distances using the same cable.

broadcast

A message that is sent to all users currently logged in to the network.

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.

byte

A series of consecutive binary digits that are operated upon as a unit (for example, an eight-bit byte).

C

CAC

See *Connection Admission Control*.

CBR

See *constant bit rate*.

CDVT

See *cell delay variation tolerance*.

cell

Any fixed-length data packet. For example, ATM uses fixed-length, 53-byte cells. See also *cell relay*.

cell delay variation tolerance

Parameter defined by the ATM Forum for ATM traffic management. In CBR transmissions, determines the level of jitter that is tolerable for the data samples taken by the PCR queue. See also *cell loss priority* and *peak cell rate*.

cell loss priority

A field in the ATM cell header that indicates the eligibility of the cell for discard by the network under congested conditions. Cells with CLP = 0 are insured traffic, which is unlikely to be dropped. Cells with CLP = 1 are best-effort traffic, which might be dropped in congested conditions in order to free up resources to handle insured traffic.

cell loss ratio

CLR is a negotiated QoS parameter and acceptable values are network-specific. The objective is to minimize CLR, provided the end-system adapts the traffic to the changing ATM layer transfer characteristics. Cell loss ratio is defined for a connection as lost cells/total transmitted cells. The CLR parameter is the value of CLR that the network agrees to offer as an objective over the lifetime of the connection.

cell relay

A form of packet transmission that uses a fixed-length, 53-byte cell over a packet-switched network; also known as Asynchronous Transfer Mode (ATM).

cell switching

An operational feature of cellular networks that enables callers to move from one location to another without losing the call connection. The cellular system is designed to switch calls to a new cell with no noticeable drop in the conversation. Cell switching is sometimes called “handing off.” While not noticeable in voice communications, the approximate 300 milliseconds this switching requires can be a problem in data transmission.

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.

client

A device that makes use of the services provided by a server.

CLP

See *cell loss priority*.

CLR

See *cell loss ratio*.

cold boot

A reboot enabling the user to restart the switch as if it were powered off, then on. Compare with *warm boot*.

committed burst size

The maximum amount of data, in bits, that the network agrees to transfer 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.

community names

The name given to an SNMP community for purposes of identification. A member has associated access rights: read-only or read/write. The Ascend switch has the following default community names: public (read-only) and cascade (read/write).

concentrator

A repeater or hub that joins communications channels from several different network nodes. Concentrator devices provide bridging, routing, and other management functions.

congestion

The point at which devices in the network are operating at their highest activity rate. Congestion is handled by employing a congestion avoidance mechanism. See also *mild congestion*, *absolute congestion*, and *severe congestion*.

Connection Admission Control

The 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.

Connectionless Network Service

OSI network layer service that does not require a circuit to be established before data is transmitted. CLNS routes messages to their destinations independently of any other messages.

connectivity

The degree to which any given computer or application can cooperate with other network components in a shared-resource network environment.

console commands

SNMP protocol supports three important commands: Get, Set, and Next. Get enables an NMS to query one or more objects or variables in an agent MIB. Set enables an NMS to modify a value of a MIB object or variable and may be used to boot or reboot devices. Next enables an NMS to query agent MIB tables and lists.

constant bit rate

A *Quality of Service* class defined by the ATM Forum for ATM networks. CBR is used for connections that depend on precise clocking to ensure undistorted delivery of bits.

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 I/O Processor modules (IOPs).

CP

See *control processor*.

crankback

A mechanism for partially releasing an in-progress connection setup that has encountered a failure. This mechanism allows PNNI to perform alternate routing.

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. See also *frame check sequence*.

D

data bits

In asynchronous transmission, the bits that actually contain the data being sent. Also called “payload” in some transmission methods.

data communications equipment

Any device that connects a computer or terminal to a communications channel or public network.

data country code

One of two ATM address formats developed by the ATM Forum for use by private networks. Adapted from the subnetwork model of addressing in which the ATM layer is responsible for mapping network layer addresses to ATM addresses. See *international code designator*.

Data Exchange Interface

A specification, described in RFC 1483, that defines how a network device can be used to convert data for interworking between different network services (i.e., Frame Relay to ATM).

Data Link Connection Identifier

A 10-bit address that identifies frame relay PVCs. See also *Local Management Interface*.

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 terminal ready

A hardware signal, defined by the RS-232 standard, exchanged between devices. For example, an RS-232-C circuit that alerts a DCE device that the DTE device is ready to send and receive data.

data transfer rate

The speed at which data is transferred, usually measured in megabits per second (Mbps) or megabytes (MB) per second.

datagram

A message unit that contains source- and destination-address information, as well as the data itself, which is routed through a packet-switched network.

DCC

See *data country code*.

DCE

See *data communications equipment*.

DE

See *discard eligible (DE)*.

dedicated line

A communications circuit used for one specific purpose, and not used by or shared between other users.

dedicated server

A computer on the network that functions only as a server performing specific network tasks.

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 that represents 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.

direct Ethernet

A method used to connect the NMS to the switch 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.

DLCI

See *Data Link Connection Identifier*.

domain

A network community of users sharing the same database information.

DSU

See *data service unit*.

DTE

See *data terminal equipment*.

DXI

See *Data Exchange Interface*.

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. ATM uses E.164 addressing for public network addressing.

EFCI

See *explicit forward congestion indicator*.

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. Compare with *Internet Protocol address*.

Ethernet packet

A variable-length unit of data transmitted on an Ethernet LAN.

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.

explicit forward congestion indicator

EFCI is an indicator in the ATM cell header. A network element in an impending-congested state or a congested state may set EFCI so that the destination end-system can examine this indicator. For example, the end-system may use this indicator to implement a protocol that adaptively lowers the cell rate of the connection during congestion or impending congestion. A network element that is not in a congestion state or an impending congestion state, will not modify the value of this indicator. Impending congestion is the state when a network equipment is operating around its engineered capacity level.

F

fault-tolerant PVCs

A set of backup ports (Permanent Virtual Circuits) on the switch that are 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.

FECN

See *forward explicit congestion notification bit*.

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.

foreground diagnostics

A set of tests used to check for non-fatal errors indicated by background diagnostics or statistics. Foreground tests may also run at start-up to test new equipment functions.

forward explicit congestion notification bit

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 (FRAD) 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 RFC1294 multi-protocol encapsulation

A specification allowing for a single circuit to be established between two devices.

FTP

See *File Transfer Protocol*.

G

gateway

A shared connection between a LAN and a larger system (such as a mainframe computer), or a large packet-switched network whose communication protocols differ.

generic flow control

The field in the ATM cell that controls the flow of traffic across the User-Network Interface (UNI) and into the network. The mechanisms for using this field are still under development.

GFC

See *generic flow control*.

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.

guaranteed packets

Data delivered according to some time constraint with high reliability.

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.

header error control

In ATM, a feature that provides protection against misdelivery of cells due to addressing errors.

HEC

See *header error control*.

heartbeat polling process

An exchange of sequence numbers between the network and a user device to ensure that both are operational and communicating.

HELLO

A routing protocol used principally by NSFnet nodes (nodes in the National Science Foundation Network). Hello allows trusting packet switches to discover minimal delay routes.

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.

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.

hot swappable

A feature that allows the user to add, replace, or remove interface processors in an Ascend switch without interrupting switch operations.

HP OpenView

The UNIX-based network management application used with NavisCore on an NMS to manage an Ascend switch network.

I

ICD

See *international code designator*.

ICMP

See *Internet Control Message Protocol.*

IDP

See *initial domain part.*

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.

initial domain part

The part of a CLNP address that contains an authority and format identifier and a domain identifier.

input/output adapter

A module that connects the various IOP and IOP Plus modules in a switch. IOA configurations vary according to the specific IOP module they support.

input/output processor

A switch module 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.

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.

Interim Local Management Interface

Specifications developed by the ATM Forum for incorporating network-management capabilities into the ATM UNI.

internal clocking

An Ascend switch hardware function that provides the transmit and receive clocks to the user equipment.

internal testing

A hardware diagnostic that performs an internal loopback test on the I/O card and other cards.

international code designator

One of two ATM address formats developed by the ATM Forum for use by private networks. Adapted from the subnetwork model of addressing in which the ATM layer is responsible for mapping network layer addresses to ATM addresses.

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 Control Message Protocol

The IP portion of TCP that provides the functions used for network layer management and control.

Internet Protocol

The TCP/IP session-layer protocol that regulates packet forwarding. See also *Internet Control Message Protocol*.

Internet Protocol address

A 32-bit address assigned to hosts using TCP/IP. The address is written as four octets separated with periods (dotted decimal format), which are made up of a network section, an optional subnet section, and a host section. Compare with *Ethernet address*.

IOA

See *input/output adapter*.

IOP

See *input/output processor*.

IP

See *Internet Protocol*.

IP address

See *Internet Protocol address*.

ISDN

See *Integrated Services Digital Network*.

ITU-T

See *International Telecommunication Union Telecommunication Standard Sector*.

K

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

LAP

See *Link Access Protocol*.

LAP-B

A bit-oriented data-link protocol used to link terminals and computers to packet-switched networks.

Link Access Protocol

The link-level protocol used for communications between DCE and DTE devices.

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 Management Interface

See *Link Management Interface*.

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 an 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.

marginal LED

An amber status indicator on a switch module that indicates a non-fatal system fault (such as low memory).

maximum burst size

Specifies the largest burst of data above the insured rate that will be allowed temporarily on an ATM PVC, but will not be dropped at the edge by the traffic policing function, even if it exceeds the maximum rate. This amount of traffic will be allowed only temporarily; on average, the traffic source needs to be within the maximum rate. Specified in bytes or cells.

maximum rate

Maximum total data throughput allowed on a given virtual circuit, equal to the sum of the insured and uninsured traffic from the traffic source. The uninsured data might be dropped if the network becomes congested. The maximum rate, which cannot exceed the media rate, represents the highest data throughput the virtual circuit will ever deliver, measured in bits or cells per second.

Mbps

Megabits per second.

MBS

See *maximum burst size*.

MCR

See *minimum cell rate*.

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.

multicast

A type of broadcast transmission that sends copies of the message to multiple stations, but not to all possible stations.

minimum cell rate

Parameter defined by the ATM Forum for ATM traffic management. MCR is defined only for ABR transmissions, and specifies the minimum value for allowed cell rate.

multiplexer (mux)

A device that merges several low-speed transmission channels into one high-speed channel at one end of the link. Another mux reverses this process at the opposite end.

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 personal computer, that enables you to communicate with other users on a LAN; also called adapter.

network parameter control

Network parameter control is defined as 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 other already established connections by detecting violations of negotiated parameters and taking appropriate actions. See also *traffic policing*.

Network-to-Network Interface

The standard that defines the interface between ATM switches and Frame Relay switches. In an SMDS network, an NNI is referred to as Inter-Switching System Interface (ISSI).

NIC

See *Network Interface Card*.

NNI

See *Network-to-Network Interface*.

node

Any device such as a personal computer, 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.

noise

Extraneous signals on a transmission channel that degrade the quality or performance of the channel.

NPC

See *network parameter control*.

O

OAM

See *Operation, Administration, and Maintenance*.

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. See also *shortest path routing*.

Operation, Administration, and Maintenance

ATM Forum specification for cells used to monitor virtual circuits. OAM cells provide a virtual circuit-level loopback in which a router responds to the cells, demonstrating that the circuit is up and that the router is operational.

OPTimum PVC trunk

A logical port configuration that optimizes inter-operability 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*.

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 processor

The Ascend switch module that performs the frame format validation, routing, queuing, and protocol conversion for the STDX switch. This module is not hot swappable.

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.

path

The complete location of a directory or file in the file system. See *define path* and *alternate path*.

payload

The portion of a frame that contains the actual data.

payload type

Payload type is a 3-bit field in the ATM cell header that discriminates between a cell carrying management information or one which is carrying user information.

PCR

See *peak cell rate*.

peak cell rate

In ATM transmission, the maximum cell transmission rate. PCR is equivalent to Be for Frame Relay, and is measured in cells per second and converted internally to bits per second. PCR defines the shortest time period between two cells.

PDN

See *public data network*.

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.

polling

An access control method in which one master device, such as the NMS, polls or queries other network devices, requesting them to transmit one at a time.

PPP

See *Point-to-Point Protocol*.

PRAM

See *parameter random access memory*.

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.

proxy service

A management service provided for one or more devices by another. For example, the Ascend SMDS Access Servers/switches are proxy-managed through the SMDS network.

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

random access memory

The main system memory in a computer used for the operating system, applications, and data.

RAM

See *random access memory*.

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.

reboot

To restart the computer and reload the operating system, usually after a crash.

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 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.

RFC1294

A specification documenting multi-protocol access over Frame Relay.

RIP

See *Routing Information Protocol*.

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 Information Protocol

A routing protocol that maintains a list of accessible networks and calculates the lowest hop count from a particular location to a specific network.

routing protocol

A protocol that implements routing using a specific routing algorithm. Routing protocols include IGRP, OSPF, and RIP.

S

SCR

See *sustainable cell rate*.

SEAL

See *Simple and Efficient Adaption Layer*.

Serial Line over Internet Protocol

A protocol that enables point-to-point serial communication over IP using serial lines or telephone connections and modems.

serial management port

A management port on the packet processor card in an Ascend switch.

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 jeopardizes 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 *Open Shortest Path First*.

Simple and Efficient Adaption Layer

In ATM, an extension of the Type 3 AAL. It simplifies the SAR portion of the Adaption layer to pack all 48 bytes of the cell information field with data. This AAL makes ATM look like high-speed Frame Relay. It also assumes that only one message is crossing the UNI at a time. That is, multiple end users at one location cannot interleave messages on the same virtual circuit, but must queue them for sequential transmission.

Simple Network Management Protocol

A standard network management protocol used to manage and monitor nodes and devices on a network.

SLIP

See *Serial Line over Internet Protocol*.

smart hub

A concentrator with certain network management features built into the firmware. This capability enables the user to manage LAN configurations.

SMDS

See *Switched Multimegabit Data Services*.

SNMP

See *Simple Network Management Protocol*.

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.

subnet address

An extension of the Internet addressing scheme that allows a site to use a single Internet address for multiple physical networks.

subnet mask

A 32-bit address mask used in IP to specify a particular subnet. See also *address mask.*

sustainable cell rate

The average cell transmission rate in ATM transmission. Equivalent to CIR for Frame Relay, SCR is measured in cells per second and converted internally to bits per second. Usually, SCR is a fraction of the peak cell rate. Cells are sent at this rate if there is no credit.

SVC

See *switched virtual circuit.*

SVCC

See *switched virtual channel connection.*

SVPC

See *switched virtual path connection.*

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.*

switched virtual channel connection

A switched VCC is one which is established and taken down dynamically through control signaling. A VCC is an ATM connection where switching is performed on the VPI/VCI fields of each cell. See also *virtual channel connection*.

switched virtual path connection

A switched VPC is one which is established and taken down dynamically through control signaling. A VPC is an ATM connection where switching is performed on the VPI field only of each cell. See also *virtual path connection*.

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

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.

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 policing

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 agreed upon flow can be tagged (where the CLP bit is set to 1) and can be discarded en route if congestion develops. Traffic policing is used in ATM, Frame Relay, and other types of networks. Also known as *admission control*, *rate enforcement*, and UPC (usage parameter control).

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.

transceiver

A device that connects a host interface to a LAN. A transceiver transmits and receives data.

Transmission Control Protocol

The Internet standard, transport-level protocol that provides the reliable, full-duplex, stream service on which many application protocols depend.

trap

An unsolicited message generated by an SNMP agent on a network device (e.g., a switch) due to a predefined event occurring or an alarm threshold being exceeded, which triggers an alarm at the NMS.

trunk

The communications circuit between two switches.

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. Backup trunks 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

UBR

See *unspecified bit rate*,

UDP

See *User Datagram Protocol*.

UNI

See *User-to-Network Interface*.

UNI DCE

See *user network interface data communications equipment*.

UNI DTE

See *user network interface data terminal equipment*.

unspecified bit rate

QoS class defined by the ATM Forum for ATM networks. UBR allows any amount of data up to a specified maximum to be sent across the network, but there are no guarantees in terms of cell loss rate and delay. Compare with *ABR*, *CBR*, and *VBR*.

UPC

Usage parameter control. See *traffic policing*.

usage parameter control

See *traffic policing*.

User Datagram Protocol

An unreliable transport-layer protocol from the TCP/IP protocol suite. It simply acts as an interface to various applications through the use of different ports.

user network interface data communications equipment

A device that performs the DCE functions for link management and expects a DTE device (e.g., an Ascend switch) to be attached to it.

user network interface data terminal equipment

A device that performs the DTE functions for link management. The user specifies this option on the NMS to connect to a 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

VBR

See *variable bit rate*.

VC

See *virtual channel*; *virtual circuit*.

VCC

See *virtual channel connection*.

VCI

See *virtual circuit identifier*.

VCL

See *virtual channel link*.

variable bit rate

QoS class defined by the ATM Forum for ATM networks. VBR is subdivided into a real time (RT) class and non-real time (NRT) class. VBR-RT is used for connections in which there is a fixed timing relationship between samples. VBR-NRT is used for connections that have no fixed timing relationship between samples, but that still need a guaranteed QoS. Compare with *ABR*, *CBR*, and UBR.

virtual bandwidth

Channel capacity calculated to allow for oversubscription of channel usage.

virtual channel

A connection between two communicating ATM networks.

virtual channel connection

A concatenation of VCLs that extends between the points where the ATM service users access the ATM layer. The points at which the ATM cell payload is passed to, or received from, the users of the ATM Layer (i.e., a higher layer or ATM-entity) for processing signify the endpoints of a VCC. VCCs are unidirectional.

virtual channel link

A means of unidirectional transport of ATM cells between the point where a VCI value is assigned and the point where that value is translated or removed.

virtual circuit

A logical circuit set up to ensure reliable communication between two network devices. See also *permanent virtual circuit* and *switched virtual circuit*.

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 connection

A concatenation of VPLs between virtual path terminators (VPTs). VPCs are unidirectional.

virtual path identifier

An 8-bit field in the ATM cell header that is used as an addressing identifier to route cell traffic.

virtual path link

A means of unidirectional ATM cell transport between the point where a VPI value is assigned and the point where that value is translated or removed.

VP

See *virtual path*.

VPC

See *virtual path connection*.

VPI

See *virtual path identifier*.

VPL

See *virtual path link*.

W

WAN

See *Wide Area Network*.

warm boot

A reboot performed after the operating system has been running for a period of time. Compare with *cold boot*.

Wide Area Network

A network that usually consists of packet-switching nodes over a large geographical area.

Y

yellow alarm

A T1 alarm that is generated when the interface receives a red alarm signal from the remote end.

Index

Numerics

A

B

C

D

E

F

G

H

I

K

L

M

N

O

P

Q

R

S

T

U

V