

NavisXtend Provisioning Server User's Guide

Ascend Communications, Inc.

Product Code: 80023 Revision 01 December 1997



Copyright © 1997 Ascend Communications, Inc. All Rights Reserved.

This document contains information that is the property of Ascend Communications, Inc. This document may not be copied, reproduced, reduced to any electronic medium or machine readable form, or otherwise duplicated, and the information herein may not be used, disseminated or otherwise disclosed, except with the prior written consent of Ascend Communications, Inc.



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 CONDI-TIONS 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

Software License



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

About This Guide

What You Need to Know	xx
Documentation Reading Path	xx
How to Use This Guide	xxii
What's New in This Guide	xxiii
Related Documents	xxv
Conventions	xxvi
Terminology	xxvii

1 Overview

NavisXtend Provisioning Server	1-1
Application Toolkit	1-4
Synchronous and Asynchronous Functions	1-5
Functions That Take an Argument List	1-8
Function Names	1-9
Toolkit Functionality	1-9
Session Control Functions	1-10
Operational Functions	1-10
Select Loop Processing Functions	1-11
Utility Functions	1-12



Managed Objects	1-14
Object Types	1-14
Containment Hierarchy	1-17
Naming Conventions for Objects	1-20
Descriptions of Object Types	1-22
CVT_Aps	1-22
CVT_AssignedSvcSecScn	1-23
CVT_Card	1-23
CVT_Channel	1-23
CVT_Circuit	1-24
CVT_Customer	1-25
CVT_LPort	1-25
CVT_NetCac	1-26
CVT_Network	1-26
CVT_PerformanceMonitor	1-27
CVT_PFdl	1-27
CVT_PMPCkt	1-27
CVT_PMPCktRoot	1-27
CVT_PMPSpvcLeaf	1-28
CVT_PMPSpvcRoot	1-28
CVT_PPort	1-28
CVT_ServiceName	1-28
CVT_SmdsAddressPrefix	1-29
CVT_SmdsAlienGroupAddress	1-29
CVT_SmdsAlienIndividualAddress	1-29
CVT_SmdsCountryCode	1-30
CVT_SmdsGroupScreen	1-30
CVT_SmdsIndividualScreen	1-30
CVT_SmdsLocalIndividualAddress	1-31
CVT_SmdsNetwideGroupAddress	1-31
CVT_SmdsSSIIndividualAddress	1-31
CVT_SmdsSwitchGroupAddress	1-31
CVT_Spvc	1-32
CVT_SvcAddress	1-32
CVT_SvcConfig	1-33
CVT_SvcCUG	1-33
CVT_SvcCUGMbr	1-34
CVT_SvcCUGMbrRule	1-34



CVT SycNodePrefix 1-34
CVT SvcPrefix 1-34
CVT SvcSecScn 1-36
CVT_SvcSecScnActParam 1-36
CVT SvcUserPart 1-36
CVT Switch
CVT TrafficDesc 1-36
CVT TrafficShaper 1-30
CVT_VDN 1.27
Valid Object Types for Operational Eulerions
Object Attributes 1 40
Dit Mask 141
DIL MIASK
SVC Addressing
String Conversion
E.164native
AESA Addresses 1-44
Example 1 1-44
Example 2 1-45
Example 3 1-45
Example 4 1-45
DefaultRoute1-46
UserPart 1-46
Class B Addressing 1-46
General API Usage 1-47
C Program 1-47
C++ Program 1-48

2 Installation and Administration

Prerequisites	2-2
Provisioning Server Requirements	2-2
Provisioning Client Requirements	2-3
Network Requirements	2-3
Installation Instructions	2-4
Installing the Provisioning Software in a Single-System Configuration	2-4
Installing the Provisioning Software in a Two-System Configuration	2-7
Post-Installation Tasks	2-8
Testing the Server	2-8
Setting Environment Variables	2-9



Testing the CLI	2-10
Recompiling an Existing Provisioning Client	2-10
Installed Files	2-10
Programming Files	2-11
Setting Environment Variables	2-13
Configuring the CLI	2-13
Identifying the Provisioning Server to the CLI	2-14
Specifying Modification Type	2-14
Specifying Security Settings	2-15
Controlling SNMP Parameters	2-15
Configuring the Provisioning Client	2-16
Enabling a Client Trace File	2-16
Controlling SNMP Parameters	2-16
Configuring the Provisioning Server	2-17
Identifying the Provisioning Server Port	2-18
Identifying the MIB Agent Port	2-18
Specifying the Core File Location	2-18
Enabling Server Trace Files	2-19
Controlling SNMP Parameters	2-20
Controlling Context Timeout	2-20
Controlling MIB Cache	2-21
Controlling Object Locking	2-21
Specifying Community Strings	2-22
Controlling SMDS Addresses	2-23
Implementing the Security Feature	2-24
Stopping and Restarting the Provisioning Server	2-24
Stopping and Restarting the CLI	2-25
Troubleshooting Problems	2-25
Problem: Requests Frequently Time Out	2-25
Symptoms	2-25
Possible Causes and Solutions	2-25
Problem: Object Is Locked by Others	2-26
Symptoms	2-26
Possible Causes and Solutions	2-27
Technical Support	2-28
Information Checklist	2-28
Un-installation Instuctions	2-30
Writing a Provisioning Application	2-31



Upgrading an	Existing	Application	2	-3	3
--------------	----------	-------------	---	----	---

3 Using the CLI

Using the CLI	3-2
CLI Usage Overview	3-3
Syntax	3-3
cvadd	3-6
Purpose	3-6
Command Syntax	3-6
Parameters	3-6
Notes	3-6
Examples	3-7
cvaddmember	3-8
Purpose	3-8
Command Syntax	3-8
Parameters	3-8
Notes	3-8
Example	3-9
cvdelete	. 3-10
Purpose	. 3-10
Command Syntax	. 3-10
Parameters	. 3-10
Notes	. 3-10
Example	. 3-11
cvdeletemember	. 3-12
Purpose	. 3-12
Command Syntax	. 3-12
Parameters	. 3-12
Notes	. 3-12
Example	. 3-13
cvget	. 3-14
Purpose	. 3-14
Command Syntax	. 3-14
Parameters	. 3-14
Notes	. 3-14
Examples	. 3-14
cvhelp	. 3-16
Purpose	. 3-16



Command Syntax	3-16
Parameters	3-16
Notes	3-16
Examples	3-16
cvlistallcontained	3-18
Purpose	3-18
Command Syntax	3-18
Parameters	3-18
Notes	3-18
Example	3-20
cvlistcontained	3-22
Purpose	3-22
Command Syntax	3-22
Parameters	3-22
Notes	3-23
Example	3-26
cvmodify	3-27
Purpose	3-27
Command Syntax	3-27
Parameters	3-27
Notes	3-27
Example	3-28
CLI Examples	3-29
Sample CLI Format	3-29
CVT_APS	3-29
CVT_AssignedSvcSecScn	3-30
CVT_Card	3-30
CVT_Channel	3-30
CVT_Circuit	3-30
ServiceName Endpoints	3-30
LPort Endpoints	3-31
CVT_Customer	3-31
CVT_LPort	3-31
CVT_NetCac	3-32
CVT_PerformanceMonitor	3-32
CVT_PFdl	3-32
CVT_PMPCkt	3-33
CVT_PMPCktRoot	3-33



CVT_PMPSpvcLeaf	3-33
CVT_PMPSpvcRoot	3-33
CVT_PPort	3-34
CVT_ServiceName	3-34
CVT_SmdsAddressPrefix	3-34
CVT SmdsAlienGroupAddress	3-34
CVT SmdsAlienIndividualAddress	3-34
CVT SmdsCountryCode	3-34
CVT SmdsGroupScreen	3-34
CVT SmdsIndividualScreen	3-35
CVT_SmdsLocalIndividualAddress	3-35
CVT_SmdsNetwideGroupAddress	3-35
CVT_SmdsSwitchGroupAddress	3-35
CVT_Spvc	3-35
CVT_SvcAddress	3-35
CVT_SvcConfig	3-36
CVT_SvcCUG	3-36
CVT SvcCUGMbr	3-36
CVT SvcCUGMbrRule	3-36
CVT SvcNodePrefix	3-36
CVT SvcPrefix	3-37
CVT SvcSecScn	3-37
CVT SvcSecScnActParam	3-37
CVT SvcUserPart	3-37
CVT Switch	3-37
CVT TrafficDesc	3-37
CVT TrafficShaper	3-38
CVT_VPN	3-38

4 Using the SNMP MIB

About the Enterprise-specific MIB	4-1
Community Strings	4-2
MIB Structure	4-3
Segmented Information in Multiple Tables	4-4
Row Aliasing	4-16
Column Access Specifiers	4-17
Additional Table Entries	4-17
RowStatus Attribute	4-17



ModifyType Attribute	3
Using the MIB	3
Using the SNMP Commands	3
Command Error Table)
MIB Cache and Database Locking)
Row Creation)
Row Modification	1
get-next Operations	3
Specifying the Object Identifier	3
Example 1: get Command	1
Example 2: get-next Command	5
Example 3: set Command to Create an ATM LPort	7
Example 4: set command to Modify an ATM LPort)
Example 5: set Command to Delete an ATM LPort	2
Example 6: set Command to Create an ATM Circuit	1
Example 7: set Command to Modify an ATM Circuit	7
Example 8: set Command to Delete an ATM Circuit 4-40)
Example 9: set Command to Create a VPN Indexed by Name 4-41	1
Example 10: set Command to Create a ServiceName Indexed by Name . 4-44	1
Example 11: set command to Modify a ServiceName Indexed by Name 4-46	5
Example 12: set command to Delete a ServiceName Indexed by Name 4-48	3

A Containment Hierarchy

Containment Tables	A-1
Network A	A- 2
Switch	A-2
STDX 3000/6000 Switch A	A- 2
B-STDX 8000/9000 Switch A	A-3
CBX 500 Switch	A- 4
Card/PPort A	A-5
6-port V.35 Card/PPort	A-5
1-port 24-channel T1 Card/PPort	A-5
1-port 30-channel E1 Card/PPort	A-6
6-port Universal I/O Card/Pport	A-6
8-port Low Speed UIO Card/Pport	A- 7
18-port Low Speed UIO Card/Pport	A-7
8-port Uio Card/Pport A	A-8
4-port 24-channel T1 Card/PPort	A-9



4-port 24-channel PRI T1 Card/PPort	A-10
4-port 30-channel E1 Card/PPort	A-10
2-port HSSI Card/PPort	A-11
10-port DSX-1 Card/PPort	A-12
1-port ATM UNI DS3 Card/PPort	A-13
1-port ATM IWU OC3 Card/PPort	A-13
1-port ATM CS/DS3 Card/PPort	A-14
1-port ATM CS/E3 Card/PPort	A-14
4-port Unchannelized T1 Card/PPort	A-15
12-port E1 Card/PPort	A-16
4-port Unchannelized E1 Card/PPort	A-17
4-port 24-channel DSX Card/PPort	A-18
1-port ATM UNI E3 Card/PPort	A-18
4-port 32-channel PRI E1 Card/PPort	A-19
1-port 28-channel DS3 Card/PPort	A-20
8-port DS3 Card/PPort	A-21
8-port E3 Card/PPort	A-21
4-port OC-3c/STM-1 Card/PPort	A-21
8-port T1 Card/PPort	A-22
8-port E1 Card/PPort	A-22
1-port OC-12c/STM-4 Card/PPort	A-23
+ · · · · · · · · · · · · · · · · · · ·	

Index



List of Figures

Figure 1-1.	Components in the NavisXtend Provisioning Server System 1-3
Figure 1-2.	Application Toolkit Organization 1-5
Figure 1-3.	Flow Between Client and Server for a Synchronous Function 1-6
Figure 1-4.	Flow Between Client and Server for an Asynchronous Function 1-7
Figure 1-5.	Containment Hierarchy for Managed Objects 1-19
Figure 1-6.	Representation of the DS1 Channel Bit Mask 1-41
Figure 4-1.	Creating an ATM LPort 4-28
Figure 4-2.	Modifying an ATM LPort 4-31
Figure 4-3.	Deleting an ATM LPort Using its VPI/VCI Pair 4-33
Figure 4-4.	Deleting an ATM LPort Using its Interface Number 4-34
Figure 4-5.	Creating an ATM Circuit 4-36
Figure 4-6.	Modifying an ATM Circuit Using its Circuit Number 4-39
Figure 4-7.	Deleting an ATM Circuit Using its Circuit Number 4-41
Figure 4-8.	Creating a VPN Indexed by Name 4-43
Figure 4-9.	Creating a ServiceName Indexed by Name 4-45
Figure 4-10.	Modifying a ServiceName Indexed by Name 4-47
Figure 4-11.	Deleting a ServiceName Indexed by Name 4-48



List of Tables

Table 1-1.	Naming Conventions for Toolkit Functions	1-9
Table 1-2.	Object Types Supported by the Provisioning Server	1-15
Table 1-3.	Naming Conventions for Object ID	1-20
Table 1-4.	Valid Object Types for Operational Functions	1-38
Table 1-5.	Calculated nBits Values	1-43
Table 2-1.	Programming Files for Client Development	2-11
Table 3-1.	Valid Parent and Child Object Types	3-19
Table 3-2.	Valid Parent and Child Object Types	3-23
Table 4-1.	Information Required for Creating Specific LPorts	4-4
Table 4-2.	Information Required for Creating Specific Circuits	4-6
Table 4-3.	Groups and Table Indexes of the Provisioning Server MIB	4-10
Table A-1.	Children of the Network Object	A-2
Table A-2.	Children of the STDX 3000/6000 Switch	A-2
Table A-3.	Children of the B-STDX 8000/9000 Switch	A-3
Table A-4.	Children of the CBX 500 Switch	A-4
Table A-5.	Children of the 6-port V.35 Card/PPort	A-5
Table A-6.	Children of the 1-port 24-channel T1 Card/PPort	A-5
Table A-7.	Children of the 1-port 30-channel E1 Card/PPort	A-6
Table A-8.	Children of the 6-port Universal I/O Card/PPort	A-6
Table A-9.	Children of the 8-port Low Speed Universal I/O Card/PPort	A-7
Table A-10.	Children of the 18-port Low Speed Universal I/O Card/PPort	A-7
Table A-11.	Children of the 8-port Universal I/O Card/PPort	A-8
Table A-12.	Children of the 4-port 24-channel T1 Card/PPort	A-9
Table A-13.	Children of the 4-port 24-channel PRI T1 Card/PPort	A-10
Table A-14.	Children of the 4-port 30-channel E1 Card/PPort	A-10
Table A-15.	Children of the 2-port HSSI Card/PPort	A-11
Table A-16.	Children of the 10-port DSX-1 Card/PPort	A-12
Table A-17.	Children of the 1-port ATM UNI DS3 Card/PPort	A-13
Table A-18.	Children of the 1-port ATM IWU OC3 Card/PPort	A-13
Table A-19.	Children of the 1-port ATM CS/DS3 Card/PPort	A-14
Table A-20.	Children of the 1-port ATM CS/E3 Card/PPort	A-14
Table A-21.	Children of the 4-port Unchannelized T1 Card/PPort	A-15
Table A-22.	Children of the 12-port E1 Card/PPort	A-16
Table A-23.	Children of the 4-port Unchannelized E1 Card/PPort	A-17
Table A-24.	Children of the 4-port 24-channel DSX Card/PPort	A-18
Table A-25.	Children of the 1-port ATM UNI E3 Card/PPort	A-18



Table A-26.	Children of the 4-port 32-channel PRI E1 Card/PPort	A-19
Table A-27.	Children of the 1-port 28-channel DS3 Card/PPort	A-20
Table A-28.	Children of the 8-port DS3 Card/PPort	A-21
Table A-29.	Children of the 8-port E3 Card/PPort	A-21
Table A-30.	Children of the 4-port OC-3c/STM-1 Card/PPort	A-21
Table A-31.	Children of the 8-port T1 Card/PPort	A-22
Table A-32.	Children of the 8-port E1 Card/PPort	A-22
Table A-33.	Children of the 1-port OC-12c/STM-4 Card/PPort	A-23





About This Guide

The *NavisXtend Provisioning Server User's Guide* describes how to use the NavisXtend Provisioning Server Application Toolkit to develop a *provisioning client* — an application that runs on a workstation in an Ascend network and interacts with a Provisioning Server. You use the client to query and configure switch nodes, cards, physical ports, logical ports, circuits, and other objects. The Application Toolkit includes a special series of libraries and header files that support client development.

In addition, the *NavisXtend Provisioning Server User's Guide* describes how to use the Toolkit Command Line Interface (CLI) to develop a *provisioning script* — a set of shell commands used for either interactive or batch provisioning of network objects.

The *NavisXtend Provisioning Server User's Guide* also describes how to use the enterprise-specific MIB, which provides SNMP access to the Provisioning Server.

What You Need to Know



What You Need to Know

This guide assumes that you have a working knowledge of network management and provisioning operations. This guide assumes that you have installed the Ascend switch hardware.

To develop a provisioning client, you need to be familiar with programming in C or C++ in a UNIX environment. Programming experience is *not* required if you plan to use the Command Line Interface only.

To use the SNMP MIB, you need to be familiar with the SNMP protocol, operations supported by the protocol, and MIB structure in general.

Documentation Reading Path

Before you read this guide, read the *Software Release Notice (SRN) for NavisXtend Provisioning Server* that accompanies the software. The SRN will alert you to any documentation updates or special conditions that you should be aware of.

The complete document set for the NavisXtend Provisioning Server includes the following manuals:



If you are using the NavisXtend Provisioning Server Application Toolkit for the first time, read the entire *NavisXtend Provisioning Server User's Guide*, which describes the interface, features, and typical applications for the NavisXtend Provisioning Server Application Toolkit. It explains, in step-by-step format, what is involved in developing a provisioning client and a provisioning script. It also describes how to use the SNMP MIB.





Once you are ready to begin developing a provisioning client, use this guide for detailed information on the NavisXtend Provisioning Server Application Programming Interface (API).

Use this guide for detailed information on the various object types supported by the NavisXtend Provisioning Server and their associated attributes.

If you are using the SNMP MIB to access the NavisXtend Provisioning Server, use this guide for detailed information on the MIB.



How to Use This Guide

The following table summarizes the information contained in this guide:

Read	To Learn About	
Chapter 1	General aspects of the NavisXtend Provisioning Server and the client and how they interact with other components on the network. This chapter describes the interface, features, and typical applications of the NavisXtend Provisioning Server Application Toolkit.	
Chapter 2	How to install and administer the various components of the Provisioning Server system. This chapter also describes the steps required to develop a provisioning application.	
Chapter 3	How to use the CLI.	
Chapter 4	How to use the SNMP MIB.	
Appendix A	The containment hierarchy (the parent-child relation) used to build object IDs to name objects in the network.	



What's New in This Guide

The following table lists the new product features in this release.

Provisioning Server 2.0 New Features	Description
Supports: CBX 500 Switch	Configure elements that are part of CBX 500 switches.
Supports: 8-port ATM E3 Card 8-port ATM DS3 Card 8-port ATM E1 Card 8-port ATM T1 Card 4-port ATM OC3 Card 1-port OC-12c Card	Configure elements of these cards for CBX 500 switches.
Supports: ATM IWU Card ATM CS Card 1-port ATM CS/E3 Card 12-port E1 Card	Configure elements of these cards for B-STDX 9000 switches.
Supports these LPorts and their associated attributes: ATM Direct Cell Trunk LPort ATM OPTimum Frame Trunk LPort ATM OPTimum Cell Trunk LPort SMDS OPTimum Trunk LPort Frame OPTimum Trunk LPort Direct Line Trunk LPort ATM NNI LPort	Configure elements of these LPorts.

What's New in This Guide



Provisioning Server 2.0 New Features	Description
Supports these new objects and their associated attributes:	Configure elements of these objects.
Network	
Customer Network Connection Admission Control (applies to B-STDX and CBX LPorts only) Service Name SMDS Country Code SVC CUG SVC CUG Member Rule SVC CUG Member Rule SVC CUG Member SVC Security Screen Traffic Descriptor Virtual Private Network	
Switch	
SVC Node Prefix	
PPort	
Automatic Protection Switching Performance Monitor Extended Super Frame Data Link Traffic Shaper	
LPort	
Assigned SVC Security Screen PMP Circuit Leaf Endpoint PMP Circuit Root Endpoint PMP Spvc Leaf Endpoint PMP Spvc Root Endpoint Soft PVC Endpoint SVC Security Screen Action Param SVC User Part SVC Address SVC Prefix SVC Config	



Provisioning Server 2.0 New Features	Description
Supports these CBX 500 features:	Configure elements related to these
Switch	features.
Switch Billing configuration	
LPort	
Direct Versus Virtual UNI CLLM	
Circuit	
Multihop PVC management	
Miscellaneous	
Flow Control Processors	
Enterprise-specific MIB	Access the Provisioning Server using an SNMP MIB.

Related Documents

This section lists the related Ascend documentation that may be useful to reference.

- Network Management Station Installation Guide (Product code: #80014)
- Network Configuration Guide for B-STDX/STDX (Product code: #80017)
- Sybase 11 SQL Server Upgrade Guide (Product code: #80040)
- Upgrading to Solaris 2.5.1 and HP OpenView 4.11 (Product code: #80045)
- Network Configuration Guide for CBX 500 (Product code: #80049)



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	Program source code.	unsigned long
Courier Bold	User input on a separate line and screen or system output.	eject cdrom
		Please wait
Helvetica	Structure names or other source code in body text.	CvObjectId structure
Boldface	Function name, CLI command, UNIX command, or user input in body text.	CvCreateNetworkId cvaddmember select Type cd install and
Italics	Variable used by a function or command.	UserArg argument
	Book titles, new terms, and emphasized text.	NavisXtend Provisioning Server User's Guide
<key name=""></key>	A keyboard entry.	<return></return>
Black border surrounding text	Notes and warnings.	See the following examples.



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





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

Terminology

The product name for the Provisioning Server product has changed from the NMS Provisioning Server to the *NavisXtend Provisioning Server*. The *NavisXtend Provisioning Server* is referred to in text using any of the following terms:

- NavisXtend Provisioning Server
- Provisioning Server
- server

The *NavisXtend Provisioning Server Application Toolkit* is referred to in text using any of the following terms:

- NavisXtend Provisioning Server Application Toolkit
- Application Toolkit
- toolkit

The *NavisXtend Provisioning client* is referred to in text using any of the following terms:

- NavisXtend Provisioning client
- Provisioning client
- client
- application

The product name for CascadeView has changed to *NavisCore*. The old and new product names are used interchangeably in the software and in the manuals.



Overview

This chapter describes what you need to know before developing an NavisXtend Provisioning client or a provisioning script. It describes the features of the Application Programming Interface (API) and presents some basic procedures that show how to perform tasks with the API.

NavisXtend Provisioning Server

The Ascend NavisXtend Provisioning Server is based on a client-server network management architecture:

NavisXtend Provisioning Client — The client is an application responsible for generating requests to provision Ascend network components. Much of the provisioning functionality of CascadeView is available: the client can query and configure Frame Relay, ATM, ATM Network Interworking, and SMDS objects including switch nodes, cards, physical ports, logical ports, circuits, and so on.

NavisXtend Provisioning Server — The server is a UNIX process that responds to requests from NavisXtend Provisioning clients and updates the Ascend switches and the CascadeView database.

NavisXtend Provisioning Server



The NavisXtend Provisioning client runs on a workstation and interacts with an NavisXtend Provisioning Server. The NavisXtend Provisioning Server responds to client requests to manage Ascend switches and updates the CascadeView database. While there can be multiple instances of CascadeView, the NavisXtend Provisioning client, and the NavisXtend Provisioning Server running on the network, any client typically interacts with only one Provisioning Server at a time. Each Provisioning Server can manage all the Ascend switches and update the shared CascadeView database.

Because the Provisioning Server shares the same Sybase database with other CascadeView processes, the server should reside in the same TCP/IP subnetwork as CascadeView and the Sybase database. As the Provisioning Server makes changes to the Ascend network, it maintains consistency with CascadeView/UX. The Provisioning Server uses a locking mechanism so that NavisXtend Provisioning clients and other CascadeView processes that share the same database cannot update the same object at the same time.

Figure 1-1 shows the relationship among the NavisXtend Provisioning client, the NavisXtend Provisioning Server, and other components on the network.





User-supplied NavisXtend Provisioning Clients

Figure 1-1. Components in the NavisXtend Provisioning Server System



The NavisXtend Provisioning Server product includes the following software:

NavisXtend Provisioning Server — Installed and maintained on a UNIX workstation on a TCP/IP network.

NavisXtend Provisioning Server Application Toolkit — Installed and used by the application developer to create an NavisXtend Provisioning client or script that submits requests to the Provisioning Server.

The next section describes the NavisXtend Provisioning Server Application Toolkit.

Application Toolkit

The Application Toolkit provides the following components:

API — Used by an application developer to write a new Provisioning client or to integrate a client into an existing provisioning system.

For the convenience of the programmer, the API functions are available in various versions. For example, the Application Toolkit provides both a C and C++ interface for each API function. And, the toolkit provides both a synchronous and asynchronous version of each function that performs provisioning operations.

For details on how to use the API to develop a Provisioning client, refer to Chapter 2 in this guide and to the *NavisXtend Provisioning Server Programmer's Reference*.

Command Line Interface (CLI) — Used to build a provisioning script. The CLI is a set of command-line programs that hide the code details of the API. Users can issue these commands from any UNIX shell to provision network objects in either interactive or batch mode.

For details on how to use the CLI to develop a provisioning script, refer to Chapter 3.

Enterprise-specific MIB — Used to access switches in the network via SNMP commands. The MIB supports all the attributes and functionality of the API and provides access via SNMP get, set, and get-next operations.

For details on how to use the MIB to develop a provisioning script, refer to Chapter 4 in this guide and to the *NavisXtend Provisioning Server Enterprise MIB Definitions*.

Figure 1-2 illustrates how the Application Toolkit is organized.





Figure 1-2. Application Toolkit Organization

Synchronous and Asynchronous Functions

The Application Toolkit provides two communication methods for each API function that performs provisioning operations. You can issue either:

- A *synchronous* function and wait for a response to your request. The application waits for the request to complete before continuing with other processing. This is also known as a *blocking* request, because each function blocks to completion.
- An *asynchronous* function and perform other operations while your request is being processed.

Application Toolkit



Figure 1-3 shows the flow between the application and the Provisioning Server for a synchronous request. The application does not regain control of the program until the response returns from the server.



Figure 1-3. Flow Between Client and Server for a Synchronous Function

With an asynchronous function, the application continues with other work while waiting for the response. The application supplies a callback handler to the API. The function invokes this callback handler function to deliver the response from the server.



Figure 1-4 shows the flow between the application and the Provisioning Server for an asynchronous request.



Figure 1-4. Flow Between Client and Server for an Asynchronous Function

During an asynchronous request, the following steps occur:

- Step 1 The application code issues an asynchronous function.
- Step 2 The application sends the request to the Provisioning Server.

Application Toolkit



Step 3	The application immediately returns to the select loop. If the application calls select directly, it needs to know what file descriptor is being used for communication with the server. The application can issue CvGetSelectInfo to obtain the information needed to pass to select .	
Step 4	The Provisioning Server processes the request.	
Step 5	The Provisioning Server sends the response to select.	
Step 6	When select notifies the application of pending messages from the server, the application issues CvProcessEvents , which goes into the API library to receive and process the response.	
Step 7	In turn, the API passes the response to the client response handler.	
Step 8	The client code continues.	

Synchronous functions involve less coding but may be less efficient, because the process waits while they are processed.

Asynchronous functions allow your application to continue processing rather than wait for completion. However, asynchronous functions require additional coding. The application programmer must provide the callback handler and must make sure that the application invokes the API appropriately when the response returns from the server. Specifically, the application must be built around the UNIX **select** system call and must ensure that all processing is done in between calls to **select**.

Functions That Take an Argument List

Most of the C functions that perform provisioning operations on network components take one or more attributes. The attributes are specified in an argument list. The Application Toolkit provides two options for specifying an argument list. Before you issue a function that takes an argument list, you can either:

- Issue a single function (**CvArgsMakeVals** or **CvArgsMakeIds**) that takes a variable number of arguments and builds the required data structure.
- Issue a series of utility functions that create (**CvArgsMake**) and fill in (**CvArgsSet***AttrType*) the required data structure.



In C++, multiple constructors for the argument object (**CvClient::Args**) handle variable argument lists.

Function Names

The name of each function varies, depending on the version of the function. Table 1-1 shows the different names for the same function that adds an object to the database.

 Table 1-1.
 Naming Conventions for Toolkit Functions

Version	Function Name
Asynchronous C function	CvaAddObject
Synchronous C function	CvsAddObject
C++ function	CvClient::addObject
CLI command	cvadd

Toolkit Functionality

The toolkit functions are divided into the following groups:

- Session Control functions open and close sockets and control session settings.
- *Operational* functions perform provisioning operations on network components.
- *Select Loop Processing* functions support loop processing of the **select** system call.
- *Utility* functions build argument lists, handle initialization, and manage storage.

The following sections describe the toolkit functions by group.



Session Control Functions

Session control functions open and close sockets and control session settings.



The CLI uses environment variables to specify session control settings (for details, refer to "Setting Environment Variables" on page 2-13).

The session control functions are:

Connect, open — Establishes a session with the Provisioning Server.

Close — Terminates a session with the Provisioning Server.

SetModifyType — Specifies whether updates are made to the network component and the database, or to the database only.

Operational Functions

Operational functions perform provisioning operations on network components. Most operational functions of the API have a CLI command counterpart.

The operational functions and CLI commands are:

AddObject (Object ID, Attributes) — Creates an object in the database and (optionally) in the switch.

AddMember (Object ID, Object ID) — Adds an address to a screen or netwide group address.

ModifyObject (Object ID, Attributes) — Modifies specific attributes of an object in the database and (optionally) in the switch.

DeleteObject (Object ID) — Deletes an object from the database and (optionally) from the switch.

DeleteMember (Object ID, Object ID) — Deletes an address from a screen or netwide group address.

GetObject (Object ID, Attributes) — Retrieves the values of specific attributes of an object.


ListContainedObjects (Object ID, type, Attributes) — Queries the database for a list of objects of the given type that are contained by a specified parent.

ListAllContainedObjects (Object ID) — Queries the database for a list of objects of any type that are contained by a specified parent.

The following operational functions are used by the API only:

NextObject — Retrieves the next object in a list of objects.

GetErrorMsg — Returns an error message.

GetResponseArgs — Returns the argument list returned by a synchronous function.

The operational functions are supported for most target object types, with a few restrictions. For example, you cannot specify a switch when you issue an Add or Delete command, because the Provisioning Server does not support adding or deleting switches.

Select Loop Processing Functions

Select Loop Processing functions support loop processing of the **select** system call. The functions that support **select** loop processing are:

Callback Handler — Is the prototype for a function supplied by the client. The API calls this callback handler to deliver a response to an asynchronous request.

GetSelectInfo — Obtains information needed to pass to a select system call.

ProcessEvents — Processes activity on file descriptors to receive responses from an asynchronous request.

Timeout — Determines if an outstanding asynchronous request timed out.



Utility Functions

Utility functions build argument lists, handle initialization, and manage storage. The utility functions are:

ArgsMake — Creates an argument list.

ArgsMakeVals, ArgsMakeIds — Creates and adds arguments to an argument list.

ArgsFree — Deletes a pointer to an argument list created either explicitly or implicitly by another function.

ArgsSet*AttrType* — A series of functions that add or modify an argument in an argument list.

ArgsGetAttrType — A series of functions that read values out of an argument list.

ArgsCount — Retrieves the number of arguments in an argument list.

ArgsIdAt — Retrieves a specified argument ID in an argument list.

ArgsTypeAt, **ArgsValueAt** — Retrieves the type or value of a specified argument in an argument list.

ArgsErrorIndex — Indicates if any argument in an argument list has an error status.

ArgsExist — Determines if a specified argument exists in an argument list.

ArgsGetStatus, **ArgsStatusAt** — Returns the error status code of a specified argument in an argument list.

ArgsEqual — Compares two argument lists for equality.

ArgsCombine — Adds two argument lists, resulting in a new argument list that combines both sets of arguments. When an argument exists in both lists, the value from the second list is used.

ArgsRemove — Subtracts two argument lists, resulting in a new argument list that contains the arguments from the first list that were not in the second list.

ArgsSelect — Returns the intersection of two argument lists, resulting in a new argument list that contains the arguments that existed in both lists. Each argument uses the value from the first list.

ArgsToString — Converts an argument list to a printable string format.

Toolkit Functionality



StringFree — Deletes a pointer to a string returned by ArgsToString and ObjectIdToString.

AddArgumentByName — Adds an argument to an argument list by specifying a textual argument name.

AddArgumentByNameValue — Adds an argument to an argument list by specifying a textual argument name and value.

ArgsPrint — Prints the text description of an argument list to a file.

CreateObjectTypeId, setObjectType — A series of functions that create an object identifier of a specified type.

CreateNetworkIdFromString, **CreateSwitchIdFromString** — Creates an object identifier of a specified type, based on a text description.

ObjectIdToString — Converts an object identifier to text descriptions.

GetObjectTypeValue, getObjectType — A series of functions that read a specified value out of an object identifier.

ObjectIdToPrint — Prints the text description of an object identifier to a file.

GetArgumentName — Converts an argument ID to a printable string format.

GetEnumName — Converts an enumerated value to a printable string format.

GetObjectTypeName — Converts an object type to a printable string format.

ParseObjectId — Converts text descriptions of an object to an object identifier.

ParseObjectType — Converts a text description to an object type.

Managed Objects



Managed Objects

Managed objects are network components managed on the network. Each managed object is represented by its *object identifier* (object ID), which is expressed as a concatenated, ordered list of type-value identifiers, each separated by periods. To specify an object ID, you first specify the object's parent (if any), including the parent type and value. Then, you specify the child type and value.

For example, an object ID for a PPort would be expressed as:

switch.100.101.102.103.card.6.pport.4

The object is identified by identifying its parent in the containment hierarchy (switch.100.101.102.103.card.6.), and then identifying the object relative to that parent (pport.4).

Note that the numbering of an object is *not* a globally unique ID; rather, it is relative to the parent object. Thus, this PPort is expressed as the fourth PPort of the card:

... pport.4

In C, an object is represented as a data structure that is manipulated using utility functions. In C++, an object is represented by a class that is manipulated using member functions. For the CLI, an object is represented by string representation.

Object Types

Table 1-2 list the object types supported by the Provisioning Server. These object types are defined in the file CvObjectType.H.



The names of several objects differ from the names used in CascadeView.



Object Types Supported by the Provisioning Server
(

Object Name	Enumerated Object Type (API)	Object Type (CLI)
Automatic Protection Switching	CVT_Aps	Aps
Assigned SVC Security Screen	CVT_AssignedSvcSecScn	AssignedSvcSecScn
Card	CVT_Card	Card
Channel	CVT_Channel	Channel
Circuit	CVT_Circuit	Circuit
Customer	CVT_Customer	Customer
Logical Port	CVT_LPort	Lport
Network Connection Admission Control	CVT_NetCac	NetCac
Network	CVT_Network	Network
Performance Monitor	CVT_PerformanceMonitor	РМ
Extended Super Frame Data Link	CVT_PFdl	Fdl
PMP Circuit Leaf Endpoint	CVT_PMPCkt	PMPCktLeaf
PMP Circuit Root Endpoint	CVT_PMPCktRoot	PMPCktRoot
PMP SPVC Leaf Endpoint	CVT_PMPSpvcLeaf	PMPSpvcLeaf
PMP SPVC Root Endpoint	CVT_PMPSpvcRoot	PMPSpvcRoot
Physical Port	CVT_PPort	Pport
Service Name	CVT_ServiceName	ServiceName
SMDS Address Prefix	CVT_SmdsAddressPrefix	AddressPrefix



 Table 1-2.
 Object Types Supported by the Provisioning Server (Continued)

Object Name	Enumerated Object Type (API)	Object Type (CLI)
SMDS Alien Group Address	CVT_SmdsAlienGroup Address	AlienGroupAddress
SMDS Alien Individual Address	CVT_SmdsAlienIndividual Address	AlienIndividualAddress
SMDS Country Code	CVT_SmdsCountryCode	CountryCode
SMDS Group Screen	CVT_SmdsGroupScreen	GroupScreen
SMDS Individual Screen	CVT_SmdsIndividual Screen	IndividualScreen
SMDS Local Individual Address	CVT_SmdsLocalIndividual Address	LocalIndividualAddress
SMDS Netwide Group Address	CVT_SmdsNetwideGroup Address	NetwideGroupAddress
SMDS Switch Group Address	CVT_SmdsSwitchGroup Address	SwitchGroupAddress
Soft PVC Circuit	CVT_Spvc	Spvc
SVC Address	CVT_SvcAddress	SvcAddress
SVC Config	CVT_SvcConfig	SvcConfig
SVC Close User Group	CVT_SvcCUG	SvcCUG
SVC Close User Group Member	CVT_SvcCUGMbr	SvcCUGMbr
SVC Close User Group Member Rule	CVT_SvcCUGMbrRule	SvcCUGMbrRule
SVC Node Prefix	CVT_SvcNodePrefix	SvcNodePrefix
SVC Prefix	CVT_SvcPrefix	SvcPrefix



Object Name	Enumerated Object Type (API)	Object Type (CLI)
SVC Security Screen	CVT_SvcSecScn	SvcSecScn
SVC Security Screen Action Parameter	CVT_SvcSecScnActParam	SvcSecScnActParam
SVC User Part	CVT_SvcUserPart	SvcUserPart
Switch	CVT_Switch	Switch
Traffic Descriptor	CVT_TrafficDesc	TrafficDesc
Traffic Shaper	CVT_TrafficShaper	TS
Virtual Private Network	CVT_VPN	Vpn

 Table 1-2.
 Object Types Supported by the Provisioning Server (Continued)

Containment Hierarchy

Figure 1-5 shows the containment hierarchy (the parent-child relation) for building object IDs to name objects in the network. Refer to Appendix A for containment tables that indicate the parent-child relation of the various objects.

Keep in mind that network ID is required only when you name an object directly below network in the containment hierarchy. You can omit the network ID for switch and objects lower in the hierarchy.

Containment Hierarchy









Figure 1-5. Containment Hierarchy for Managed Objects



Naming Conventions for Objects

Table 1-3 lists the rules for naming object type-value identifiers.

Table 1-3. Naming Conventions for Object ID

Object Type	How Identified
Aps NetCac PerformanceMonitor PFdl SMDS group screen SMDS individual screen SvcConfig SvcSecScnActParam	The object is unique to its parent and requires no identifying value. Identify the object by the type name and the parent. For example, an SMDS individual screen is expressed as: <i>switch.100.101.102.103.card.6.pport.4.lport.2.individualscreen</i>
AssignedSvcSecScn Customer SvcCUG SvcCUGMbr SvcCUGMbrRule SvcSecScn ServiceName TrafficDesc VPN	By a string name. For SvcCUGMbr, the object is identified by two names: the CUG name and the member name.
Card LPort PMPSpvcLeaf PPort TrafficShaper	By relative number. For example, the fourth PPort on a card is identified as: <i>switch.100.101.102.103.card.6.pport.4</i> The first PMPSpvcLeaf on an LPort is identified as: <i>switch.100.101.102.103.card.6.pport.4.lport.1.PMPspvcleaf.1</i> Use the relative numbering scheme to identify LPorts; do not use the LPort Interface Number displayed in the CascadeView screens. ATM Transport for FR NNI LPorts are identified with VPI and VCI numbers. ATM Virtual UNI LPorts are identified with the VPI start value.



Table 1-3. Naming Conventions for Object ID (Continued)

Object Type	How Identified
Channel	By a number in the range of 1 - 28. The channel object applies only to the channelized DS3 card. For example, a Frame Relay circuit on a channelized DS3 card is identified as: <i>switch.100.101.102.103.card.6.pport.4.channel.25.lport.1.dlci.55</i>
Circuit	By the number(s) of its first endpoint. An endpoint can be an LPort or a Service Name; the object ID representation differs accordingly.
	In the case of LPorts, either endpoint can be a Frame Relay or an ATM endpoint. For Frame Relay endpoints, use the DLCI number. For ATM endpoints, use both the VPI and VCI values. For ATM Network Interworking for Frame Relay NNI endpoints, include the VPI, VCI, and DLCI numbers. For example, a Frame Relay endpoint is expressed as: <i>switch.100.101.102.103.card.6.pport.4.lport.2.dlci.55</i>
	An ATM endpoint is expressed as: switch.100.101.102.103.card.6.pport.4.lport.2.vpi.8.vci.65
	In the case of ServiceName, the endpoint is identified by the network number, the name of the ServiceName binding, and the VPI/VCI pair or DLCI number (depending on endpoint type).
	For example, a ServiceName endpoint is represented as either:
	network.154.188.0.0.ServiceName.xxx.vpi.14.vci.128
	network.154.188.0.0.ServiceName.xxx.dlci.55
	where xxx is the name of the ServiceName binding.
Network	By an IP address with the last 2 bytes set to 0 (Class B addresses) or the last 1 byte set to 0 (Class C addresses). This object type is used to specify a root when you issue a command to list objects contained by a specific parent.
PMPCkt	By the VPI and VCI values of its endpoint.
PMPCktRoot PMPSpvcRoot Spvc	A PMPCktLeaf endpoint is expressed as: switch.100.101.102.103.card.6.pport.4.lport.2.PMPcktleaf.vpi.8. vci.65



Table 1-3. Naming Conventions for Object ID (Continued)

Object Type	How Identified
SMDS address prefix	By an E.164 address string (3 to 6 characters).
SMDS country code	By an E.164 address string (up to 4 characters).
SMDS alien group address SMDS alien individual address SMDS local individual address SMDS netwide group address SMDS switch group address	By an E.164 address string (10 to 16 characters).
SvcAddress SvcNodePrefix SvcPrefix SvcUserPart	By a string that conforms to the convention used to specify addresses. For more information, refer to "SVC Addressing" on page 1-42.
Switch	By an IP address or by a string name.

Descriptions of Object Types

The following sections describe the object types, including the kinds of management operations that you can perform on an object and any operating restrictions. The objects types are listed alphabetically.

CVT_Aps

Automatic Protection Switching (APS) protects SONET media from line outages. Currently, APS support is provided for 1-port OC-12c/STM-4 cards on CBX 500 switches. When the attribute CVA_PPortRedundancy is set to Aps1+1, the protection port forms a pair with the existing PPort on the card.

You can only modify APS objects; the Provisioning Server does not support adding or deleting them. Depending on whether the PPort is the working PPort or the protection line, you can configure a list of PPort and APS parameters. The attribute CVA_ApsApsCommand is supported for the APS PPort pair for sending external switch requests.

CVT_AssignedSvcSecScn

AssignedSvcSecScn specifies the association between an LPort and an SVC security screen. When you add or delete an object of this type, you are actually adding or removing screens from the parent LPort. This object exists only on ATM UNI/NNI LPorts configured on a CBX 500 switch. A limit of 16 screens can be added to one LPort.

CVT_Card

The CascadeView database automatically populates each switch with cards of type "empty". To add a card, use the Modify command to change the card's type from "empty" to a specified type. Specify the appropriate card type using the attribute CVA_CardDefinedType.

The attributes CVA_CardUioDefinedXface, CVA_CardDsx1DefinedXface, and CVA_CardE1DefinedXface provide subtypes for the UIO, Dsx1, and E1 card types, respectively. If you modify a card to one of these card types and do not specify the appropriate subtype, the card defaults to uioXfaceTypeV35, dsx1XfaceTypeRj48, or e1XfaceTypeCoaxPair75Ohm, respectively.

To delete a card, use the Modify command to change the card's type to "empty".

If you modify a card to a type that does not match the actual card type, the Provisioning Server does *not* inform you about the type mismatch.

CVT_Channel

The channel object applies only to the channelized DS3 card on the B-STDX switch. Requests that specify a channel and are sent to an object other than the channelized DS3 card return an error.

A channel is identified by a number in the range of 1 - 28. For example, a Frame Relay circuit on a channelized DS3 card is represented as:

switch.100.101.102.103.card.6.pport.4.channel.25.lport.1.dlci.55

Once a channelized DS3 card has been configured, CascadeView automatically populates the card with all necessary channels. You can only modify channels; the Provisioning Server does not support adding or deleting channels.



CVT_Circuit

A circuit is identified by its first endpoint. An endpoint can be an LPort or a Service Name; the object ID representation differs accordingly.

In the case of LPorts, either endpoint can be a Frame Relay or an ATM endpoint:

- For Frame Relay endpoints, use the DLCI endpoint.
- For ATM endpoints, include both the VPI and VCI values.
- For ATM Network Interworking for Frame Relay NNI endpoints, include the VPI, VCI, and DLCI numbers.
- For ATM Virtual UNI endpoints, use the start VPI value.

Specify a circuit's second endpoint with the attribute CVA_CircuitEndpoint2.

For example, the Frame Relay endpoint that connects switch 100.101.102.103, card 6, PPort 4, LPort 2, DLCI 55 with ATM endpoint 154.188.162.44, card 3, PPort 5, LPort 11, VPI 8, VCI 65 is represented as either:

switch.100.101.102.103.card.6.pport.4.lport.2.dlci.55

switch.154.188.162.44.card.3.pport.5.lport.11.vpi.8.vci.65

The Provisioning Server supports VPI values of 0-15 for ATM circuit endpoints.

If an endpoint of a circuit is defined on a channelized DS3 card, the circuit is identified by the channel ID. For example, the Frame Relay endpoint that connects switch 100.101.102.103, card 6, PPort 4, channel 25, LPort 1, DLCI 55 with ATM endpoint 128.129.130.131, card 1, PPort 2, LPort 3, VPI 14, VCI 128 is represented as either:

switch.100.101.102.103.card.6.pport.4.channel.25.lport.1.dlci.55

switch.128.129.130.131.card.1.pport.2.lport.3.vpi.14.vci.128

When you add a circuit of type VPC, you do not provide the VCI part of the endpoint. For example, the endpoint for a VPC circuit is represented as:

switch.100.101.102.103.card.6.pport.4.channel.25.lport.1.vpi.14

In this case, the second endpoint must also be an ATM endpoint, as VPC circuits only support ATM endpoints.



In the case of ServiceName, the endpoint is identified by the network number, the name of the ServiceName binding, and the VPI/VCI pair or DLCI number (depending on endpoint type). Specify the second endpoint with the attribute CVA_CircuitEndpoint2.

For example, a ServiceName endpoint is represented as either:

network.154.188.0.0.ServiceName.xxx.vpi.14.vci.128

network.154.188.0.0.ServiceName.xxx.dlci.55

where xxx is the name of the ServiceName binding.

CVT_Customer

Customer objects are associated with VPN objects (Virtual Private Networks). Each customer object contains information that identifies both the customer and the VPN with which the customer is associated. The attribute CVA_CustomerVpnName associates the Customer object with a particular VPN.

The attribute CVA_LPortCustomerName specifies a customer name to which the LPort belongs.

CVT_LPort

LPorts have different subtypes: Frame Relay, SMDS, ATM, and Other. When you issue a command, specify only those attributes that are appropriate for the particular LPort's subtype. Refer to the *NavisXtend Provisioning Server Object Attribute Definitions* for the attributes that pertain to each object type and subtype.

For SMDS objects, you can set the CVA_LPortSsiLPort attribute to the object ID of an SMDS SSI DTE LPort. Or, to de-multiplex the LPort, you set the attribute to an object ID of type CVT_Null (use -nullObject in the CLI).

On a channelized DS3 card, an LPort is a child of a channel. Thus, when you issue an Add command, you must specify the channel parent.

Descriptions of Object Types



For ATM Virtual UNI LPorts on the CBX 500 switch, first create a feeder LPort with ATM UNI type. Since the LPort number of a virtual UNI LPort is generated automatically from the combination of its VPI start number and the Interface Number (which is also generated automatically), you can use the VPI start number to identify the LPort. For example, a Virtual UNI LPort with the start VPI number set to 1 is represented as:

switch.128.129.130.131.card.1.pport.2.startvpi.1

ATM Network Interworking for Frame Relay NNI LPorts require a different object identifier. This LPort type is identified by VPI/VCI pair.

For example, an ATM Network Interworking for Frame Relay NNI LPort with VPI 1 and VCI 32 is represented as:

switch.100.101.102.103.card.6.pport.4.vpi.1.vci.32

CVT_NetCac

Network Connection Admission Control (NetCAC) allows you to compute the bandwidth allocation for any virtual circuit. The NetCac object exists under the Network object. This object is supported only for CBX 500 and B-STDX 9000 switches. If the attribute CVA_NetCacCacType is set to Cascade, then only Cell Loss Ratio and Cell Delay Variation parameters are configurable. For Customized CAC configuration, you must supply Port Scale Factors and SCR Limit Scale Factors.

In the case of Customized CAC, you must supply the three SCR Limit Scale Factors of Upper Limit, Scale Factor, and Maximum MBS values together. You can supply a maximum of 10 sets. No default values apply and no upper boundary checks are performed for any of these scale factor values.

CVT_Network

Use this object type to specify a root when you issue a command to list objects contained by a specific parent.



CVT_PerformanceMonitor

This object supports the CURRENT (15-minute) and the ONE-DAY threshold parameters for the 8-port T1/E1, 8-port DS3/E3, 1-port OC12, and 4-port OC3 cards on CBX 500 switches. Default values are set at the time of card configuration, which you can then modify.

CVT_PFdl

Once a card has been configured, CascadeView automatically populates the card with all necessary Physical Ports. In the case of the ATM-T1 card on the CBX 500 switch, the CascadeView database automatically populates the Extended Super Frame object.

You can only modify the Extended Super Frame object; the Provisioning Server does not support adding or deleting it.

CVT_PMPCkt

A Point-to-MultiPoint (PMP) circuit consists of one endpoint acting as the Root and the other endpoints acting as Leaves. Use this object to add PMP Leaves. A PMP Leaf can be added, modified, and deleted. This object type applies only to CBX 500 switches. Since only ATM endpoints are supported, only VPI and VCI values are supported. To add a leaf using the CLI or the API, you must specify the Root object as one of the attributes. To add a leaf using the MIB, you specify the Root as an index.

CVT_PMPCktRoot

A Point-to-MultiPoint (PMP) circuit consists of one endpoint acting as the Root and the other endpoints acting as Leaves. PMP Root can be added or deleted. This object type applies only to CBX 500 switches. Since only ATM endpoints are supported, only VPI and VCI values are supported. To add a PMP Circuit using the CLI or the API, add the Root and the Leaves separately. Root attributes are Create-Only attributes.



CVT_PMPSpvcLeaf

Use this object type to add the Point-to-MultiPoint (PMP) SPVC Leaf.

To add a leaf using the CLI or the API, you must specify the Root object as one of the attributes. To add a leaf using the MIB, you specify the Root as an index.

You must specify the correct instance number when you perform an add, get, modify, or delete operation. To retrieve the correct instance number from the database, use the attribute CVA_PMPSpvcRootNextAvailableLeafNo.

CVT_PMPSpvcRoot

This object type is similar to CVT_Spvc, but is used to add a Point-to-MultiPoint (PMP) SPVC root. This object type applies only to CBX 500 switches. You can add, modify, and delete this object type.

When you add the Root, the first leaf is automatically added. To modify the first leaf, use the object type CVT_PMPSpvcLeaf.

CVT_PPort

Once a card has been configured, CascadeView automatically populates the card with all necessary Physical Ports. You can only modify PPorts; the Provisioning Server does not support adding or deleting PPorts.

CVT_ServiceName

ServiceName binding support allows you to identify a primary port (UNI/NNI) with a name so that a circuit can identify its service endpoint by this name instead of by the LPort name. The primary LPort can be a Frame Relay or an ATM UNI/NNI LPort. To associate a backup binding with the primary service name binding, associate a switch port to act as a backup LPort.

When creating a service name binding, specify only the primary LPort. This primary binding cannot be modified.



To set up or modify a backup binding, modify the ServiceName object by specifying the backup LPort. The attribute CVA_ServiceNameActiveBinding indicates the current status of binding. To revert from backup binding to primary binding, set the attribute CVA_ServiceNameActiveBinding to Primary in the modify request.

CVT_SmdsAddressPrefix

An SMDS address prefix is created on a switch to indicate that the switch handles all E.164 addresses that begin with that prefix. You must create an address prefix before you can create an SMDS local individual address that uses that prefix.

You can create an address prefix at any time. You can delete an address prefix only if it is not referenced by any SMDS local individual address. No attributes apply to address prefixes.

CVT_SmdsAlienGroupAddress

Objects of this type are used only as members of a group screen. Use this object type to add a group address to an SMDS group screen, if the group address does not exist on the switch as a switch group address (refer to "CVT_SmdsSwitchGroupAddress" on page 1-31). In this case, you must first create an SMDS alien group address before you can add the group address to the SMDS group screen. To do so, issue an Add command with no arguments. Then, issue the Add Member command to add the address to the group screen.

CVT_SmdsAlienIndividualAddress

Objects of this type are used only as members of an individual screen. Use this object type to add an address to an SMDS individual screen, if the address does not exist on the switch as a local individual address (refer to

"CVT_SmdsLocalIndividualAddress" on page 1-31). In this case, you must first create an SMDS alien individual address before you can add the address to the SMDS individual screen. To do so, issue an Add command with no arguments. Then, issue the Add Member command to add the address to the individual screen.





The Add command can fail or cause problems in the switch if the alien individual address uses a prefix that is assigned to the switch. By definition, an alien individual address should use a prefix not currently defined anywhere in the network.

CVT_SmdsCountryCode

Specify this object on the network level using up to 3 digits in E.164 format. To use a country code in an SMDS local individual address, include a dash (-) between the country code and the prefix (for example: 1-9789521111). If you omit the country code, the server uses the default country code specified in the environment variable CV_DFLT_SMDS_CC (for details, refer to "Setting Environment Variables" on page 2-13).

CVT_SmdsGroupScreen

There is only one group screen per SMDS LPort, and you must explicitly create it. Create the group screen before you add addresses to it. Once you create the group screen, you can add switch group addresses or alien group addresses as members. When you create a netwide group address, a switch group address is created automatically.

You can use an SMDS screen to either allow or disallow specific addresses for an LPort. Use the attribute CVA_GroupScreenOperation to do so.

CVT_SmdsIndividualScreen

There is only one individual screen per SMDS LPort, and you must explicitly create it. Create an individual screen before you add addresses to it. Once you create the individual screen, you can add addresses or alien addresses as members.

An SMDS screen can be used to allow or disallow specific addresses for the LPort. Use the CVA_IndividualScreenOperation attribute to do so.



CVT_SmdsLocalIndividualAddress

In CascadeView, an SMDS local individual address is known as an individual address. You create this object on an LPort to associate that address with the LPort. The address must use a prefix that has already been created on that switch. You can also use an existing country code specified in the network (refer to "CVT_SmdsCountryCode" on page 1-30) as part of the local individual address. You cannot delete an LPort until you have deleted all its local individual addresses.

CVT_SmdsNetwideGroupAddress

A netwide group address is a collection of SMDS switch group address objects. Normally, you create and manage group addresses through these objects and not through SMDS switch group address objects. You must create a netwide group address in the appropriate subnetwork before you can add members (individual addresses) to it. You must also use the Delete Member command to remove all of the netwide group address members before you can delete the netwide group address itself.

CVT_SmdsSSIIndividualAddress

This object is obsolete, but is maintained for compatibility with previous versions of the Provisioning Server. The SniDxi LPort does not have to subscribe to an address from an SSI LPort's address pool even if the LPort is multiplexed to an SSI LPort. You can perform any SMDS configuration without creating the SSI individual address pool.

CVT_SmdsSwitchGroupAddress

An SMDS switch group address does not appear in CascadeView. This object represents a group address that is local to a switch. A switch group address with a given address should exist on a switch only if the equivalent netwide group address has members on that switch; it lists the addresses of that switch that are members of the equivalent netwide group address. The only time you should need to reference a group address directly is to add one to a group screen.



You should not have to create or delete objects of this type; they are created and deleted automatically during the management of SMDS netwide group addresses. However, since you cannot delete a netwide group address if it contains any group address members, in rare cases, you may need to delete a group address manually. You cannot delete a group address until it no longer contains individual address members.

CVT_Spvc

Soft PVC circuits are identified by an endpoint at one end and the SVC Address at the other end. The other endpoint may not necessarily exist in the same network. You can add, modify, and delete this object. This object type applies only to CBX 500 switches. Since only ATM endpoints are supported, only VPI and VCI values are supported. Specify the SVC Address using attributes.

CVT_SvcAddress

SvcAddress provides an interface to set up full ATM addresses (20 octets) on an LPort. This address is associated with the following LPort types located on CBX 500 switches:

- atmUniDte
- atmUniDce
- atmNni

Frame Relay addresses are associated with the following LPort types located on B-STDX 8000 and 9000 switches:

- frUniDte
- frUniDce
- frNni

There can be zero or more SvcAddresses configured per LPort.

An ATM Address can be one of the following format types:

• E.164native



- AESA addresses:
 - E.164AESA
 - DCCAESA
 - ICDAESA
 - CustomAESA

For AESA addresses, if the address prefix is 39 characters, the Provisioning Server appends a zero to the address to make it 20 octets.

Frame Relay switched virtual circuits can have only E.164native addresses.

For information on the convention used to specify SVC addresses, refer to "SVC Addressing" on page 1-42.

CVT_SvcConfig

Use this object to configure an LPort for switched virtual circuits. For Frame SVCs, only the attribute CVA_SvcConfigHDTimer attribute is valid.

Only one SvcConfig object is associated with an LPort. An LPort is created with a default SvcConfig; you can only modify this object to change an SVC configuration. The SvcConfig is deleted when its LPort is deleted.

CVT_SvcCUG

Use this object to configure SVC Closed User Groups. You can create this object under the Network object. Each SVC Closed User Group can contain up to 128 members. You cannot perform a database-only modification on this object, since the modification has to be distributed throughout the network.



CVT_SvcCUGMbr

Use this object to create the association of an SvcCUG object and an SvcCUGMbrRule. Deleting this type of object disassociates the specified SvcCUG with the SvcCUGMbrRule. Adding, modifying, and deleting an SvcCUGMbr object requires network distribution; you cannot perform a database-only modification on this object.

CVT_SvcCUGMbrRule

During creation of this object, a distribution list is created by matching its rule to ATM SVC prefixes, addresses, and user parts configured on nodes in the network. Adding, modifying, and deleting an SvcCUGMbrRule object requires network distribution; you cannot perform a database-only modification on this object.

CVT_SvcNodePrefix

This object provides an interface to add node prefixes on a switch. The switch imposes no constraints on the node prefixes except to enforce the maximum length (which is the same as a full address length — 20 octets). For AESA addresses, if the address prefix is an odd number of characters, the Provisioning Server stuffs the last 4 bits of the last octet with zeros, thereby appending a zero to the address prefix.

The valid address format types for ATM SVCs are the same as described for the CVT_SvcAddress object (refer to "CVT_SvcAddress" on page 1-32). For Frame Relay SVCs, E.164native format is the only valid address format. There can be zero or more SvcNodePrefixes configured per switch.

For information on the convention used to specify SVC addresses, refer to "SVC Addressing" on page 1-42.

CVT_SvcPrefix

This object provides an interface to set up prefixes on an LPort. Prefix is one of the three classes of static addressing used for switched virtual circuits. The other classes are full ATM Address and Node Prefix.



Prefix is associated with the following LPort types located on CBX 500 switches:

- atmUniDte
- atmUniDce
- atmNni

Frame Relay SVC prefixes are associated with the following LPort types located on B-STDX 8000 and 9000 switches:

- frUniDte
- frUniDce
- frNni

There can be zero or more SvcPrefixes configured per LPort.

An ATM Address prefix can be of the following format types:

- E.164native
- AESA addresses:
 - E.164AESA
 - DCCAESA
 - ICDAESA
 - CustomAESA
- DefaultRoute: Use this format to configure the port with an ATM address length of zero bits. This enables this port to receive messages that could not be routed to other ports because of ATM address mismatch.

For Frame Relay SVCs, E.164native and DefaultRoute are valid formats.

For AESA addresses, if the address prefix is an odd number of characters, the Provisioning Server stuffs the last 4 bits of the last octet with zeros, thereby appending a zero to the address prefix.

For information on the convention used to specify SVC addresses, refer to "SVC Addressing" on page 1-42.



CVT_SvcSecScn

This object exists on the network level. Its name is used by the AssignedSvcSecScn object to apply screening to an LPort. Modify and delete operations require network wide distribution; you cannot perform database-only modification on this object.

CVT_SvcSecScnActParam

This object exists under the LPort object and is populated/deleted automatically with LPort creation. One instance exists for each ATM UNI LPort on a CBX 500 switch.

CVT_SvcUserPart

Use this object to set up the user part on a DTE LPort on a CBX 500 switch. It is used for dynamic address registration at a UNI. The user part address length is 7 octets, of which End System Identifier (ESI) represents 6 octets and the selector represents 1 octet. The user part represents a partial SVC address associated with ATM DTE LPorts on the node. The rest of the address is the network prefix, which is supplied by the network side of the UNI. To obtain an ATM address for a terminal on the user side of a Private UNI, append values for the user part to network prefix(es) for that UNI.

There can be zero or more UserParts configured on an LPort.

For information on the convention used to specify SVC addresses, refer to "SVC Addressing" on page 1-42.

CVT_Switch

The Provisioning Server does not support adding or deleting switches; you must use CascadeView to do so. For an existing switch, you can read or modify any switch attribute except for the CVA_SwitchName attribute, which is Read-Only.

CVT_TrafficDesc

The Provisioning Server provides support for maintaining a pool of ATM traffic descriptors. The traffic descriptors are required for setting up forward and backward traffic descriptors for Soft PVCs. Each traffic descriptor is identified by a name. An ID is automatically associated with each name.

Descriptions of Object Types



Depending on the Quality of Service (QoS) class you select and the Type of Service associated with it, you need to provide the PCR, SCR, and MBS values. Do so using the attributes CVA_TrafficDescParam1, CVA_TrafficDescParam2, and CVA_TrafficDescParam3. Only add and delete operations are supported for ATM traffic descriptors.

CVT_TrafficShaper

Traffic shaper objects are located under PPort in the containment hierarchy. Only PPorts on the following cards can have traffic shaper objects:

- 1-port ATM IWU OC3 card
- 1-Port ATM CS/DS3 card
- 1-port ATM CS/E3 card

A traffic shaper object is not an independent object. It represents a group of traffic shaper attributes under a particular PPort type. All the traffic shaper attributes are essentially the attributes for the belonging PPort. The traffic shaper attributes are treated as a separate object to provide a clear user interface for the attributes.

Once a PPort is created through CascadeView, CascadeView automatically populates the PPort with traffic shaper attributes. You can only modify traffic shaper attributes; the Provisioning Server does not support creating them.

When you issue a ListContained command on a traffic shaper object, the only valid parent object type is PPort.

CVT_VPN

Creation and deletion of VPN objects occur at the network level. Use the attribute CVA_LPortVPNName to specify the VPN to which an LPort belongs. Setting CVA_LPortVPNName to Public makes that LPort a normal public LPort. Similarly, use the attribute CVA_CircuitVPNName to specify the VPN to which a circuit belongs. Both end points of a circuit may not belong to different VPNs.



Valid Object Types for Operational Functions

Table 1-4 lists the object types you can use when you issue the operational functions and commands of the API and CLI.

Object Type	Add Object	Add Member	Delete Object	Delete Member	Get	List All	List	Modify
APS					1		1	1
AssignedSvcSecScn	1		1		1	1	1	1
Card					1	1	1	1
Channel					1	1	1	1
Circuit	1		1		1		1	1
Customer	1		1		1	1	1	1
LPort	1		1		1	1	1	1
NetCac					1		1	1
Performance Monitor					1		1	1
PFdl					1		1	1
PMP Circuit Leaf Endpt	1		1		1		1	1
PMP Circuit Root Endpt	1		1		1	1	1	
PMP SPVC Leaf Endpt	1		1		1		1	1
PMP SPVC Root Endpt	1		1		1	1	1	1
PPort					1	1	1	1
Service Name	1		1		1	1	1	1
SMDS Address Prefix	1		1		1	1	1	1
SMDS Alien Group Address	1		1		1	1	1	1

Table 1-4. Valid Object Types for Operational Functions

Valid Object Types for Operational Functions



Table 1-4. Valid Object Types for Operational Functions (Continued)

Object Type	Add Object	Add Member	Delete Object	Delete Member	Get	List All	List	Modify
SMDS Alien Individual Address	1		1		>	1	1	1
SMDS Country Code	1		1		1	1	1	1
SMDS Group Screen	1	1	1	1	~	1	1	1
SMDS Individual Screen	1	1	1	1	1	1	1	1
SMDS Local Individual Address	1		1		1	1	1	1
SMDS Netwide Group Address	1	1	1	1	1	1	1	1
SMDS Switch Group Address			1		1	1	1	1
Soft PVC Circuit	1		1		1		1	1
SVC Address	1		1		1		1	1
SVC Config					1		1	1
SVC CUG	1		1		1		1	1
SVC CUG Member	1		1		1		1	1
SVC CUG Mbr Rule	1		1		1		1	1
SVC Node Prefix	1		1		1		1	1
SVC Prefix	1		1		1		1	1
SVC Security Screen	1		1		1	1	1	1
SVC SecScnActParam					1	1	1	1
SVC UserPart	1		1		1		1	
Switch					1	1	1	1

Traffic Shaper

Vitual Private Network



Table 1-4.	Valid Object Types for Operational Functions (Continued)							
Object Type	Add Object	Add Member	Delete Object	Delete Member	Get	List All	List	Modify
Traffic Descriptor	1		1		1		1	

1

1

1

1

Object Attributes

1

For each of the managed objects supported by the Provisioning Server, there are arguments (attributes) that can be read or configured through the API or CLI. An argument list is represented as follows:

In:	Argument List represented as:
С	An opaque pointer that is manipulated using utility functions.
C++	A class that is manipulated using member functions.
CLI	String representations of the attributes.

The NavisXtend Provisioning Server Object Attribute Definitions lists the various object types supported by the Provisioning Server and their associated attributes. Refer to that guide to determine which attributes apply to which object types.

Bit Mask



Bit Mask

The 1-port, 24-channel Fractional T1 card and the 4-port, 24-channel Fractional DSX card use a 24-bit bit mask to represent the channels. Bits 0 through 23 represent channels 1 through 24, respectively. To allocate a channel for an LPort, set the appropriate bit to 1. To set the attribute CVA_LPortFractionalDs0s or CVA_PPortAllocatedChannels, create a numeric value where the appropriate bits are set.

For example, to allocate channels 6 and 7, the bit mask would be expressed as:

000000000000000001100000

The decimal number for this bit mask is 96. The hexadecimal number for this bit mask is 0x60.

The 4-port, 30-channel Fractional E1 card uses a 31-bit length bit mask to represent the channels. Bits 0 through 30 representing channels 1 through 31, respectively. To allocate a channel for an LPort, set the appropriate bit to 1. You can set any bit except bit 15, which represents channel 16 (reserved for internal use).

With the CLI, you can input hexadecimal numbers. You must prefix each hexadecimal number with 0x. The CLI output is only in decimal numbers.

The channelized DS3 card uses a 28-bit mask to represent the DS1 channels. 24 DS0 channels are associated with each DS1 channel. When an LPort is created on a channel, the corresponding channel's bit in the DS1 channel bit mask is set to 1. Figure 1-6 shows the corresponding DS1 channel marked with an **X**. All 24 associated DS0s on the selected DS1 channel are automatically set by default, as shown in the figure.



Figure 1-6. Representation of the DS1 Channel Bit Mask



SVC Addressing

SVC addresses are represented as strings, using the following convention:

<Address-Format-Type-ID>-<Address-Prefix>-<nBits>

Address-Format-Type-ID is the number that represents the format of the SVC address. Valid values are as follows:

- **1** E.164native
- 2 DCCAESA
- 3 ICDAESA
- **4** E.164AESA
- 5 CustomAESA
- 6 DefaultRoute
- 7 UserPart

Address-Prefix is the complete SVC address prefix. For AESA addresses, if the address prefix is an odd number of characters, the Provisioning Server stuffs the last 4 bits of the last octet with zeros, thereby appending a zero to the address prefix.

nBits is the number of bits. This field is optional for all address formats except for DefaultRoute. For DefaultRoute, you must specify the number of bits as zero.

For other address formats, if you omit this field, the Provisioning Server calculates the value and appends it to the address prefix. The algorithms used to calculate nBits values are presented in Table 1-5.



Address Format Type	Address-Format-Type-ID	Calculated nBits Value
E.164native	1	string length of address prefix * 8
DCCAESA	2	(integral-part-of((string length of address prefix $+ 1)/2$)*8)
ICDAESA	3	(integral-part-of((string length of address prefix $+ 1)/2$)*8)
E.164AESA	4	(integral-part-of((string length of address prefix $+ 1)/2$)*8)
CustomAESA	5	(integral-part-of((string length of address prefix $+ 1)/2$)*8)
UserPart	7	56

Table 1-5.Calculated nBits Values

String Conversion

In the following cases, the Provisioning Server performs an address string conversion:

- When you omit the nBits field, the Provisioning Server calculates and appends an nBits value to the prefix address.
- When an address prefix of an AESA address is an odd number of characters, the Provisioning Server appends a zero to the prefix address.

A converted string is equivalent to the original string. Either address string can be used in an operation.

Each of the address formats is described in the following sections.

E.164native

Specify an E.164native address as a numeric string of 1 - 15 characters.

For example:

1-12345



where 1 specifies the address format type E.164native, and 12345 represents the address prefix. Since no nBits value is specified, the Provisioning Server calculates a value and appends it to the address. The string is converted to:

1-12345-40

AESA Addresses

Specify an AESA address as a hexadecimal string. The first two characters of the address prefix represent the AFI value. The address prefix must be in the range of 2 - 40 characters. The number of bits must be in the following range:

(integral-part-of((string length of address - 1)/2)*8)< nBits \leq (integral-part-of((string length of address + 1)/2))*8

The minimum value for nBits is 8.

In the case of CustomAESA format, the AFI value can be any two hexadecimal characters.

If the address prefix is an odd number of characters, the Provisioning Server stuffs the last 4 bits of the last octet with zeros, thereby appending a zero to the address prefix.

Standard AFI values are:

39 — DCCAESA

45 — E.164AESA

47 — ICDAESA

Example 1

2-39

where 2 specifies the address format type DCCAESA, and 39 represents the address prefix (consisting of the AFI value only).

Since no nBits value is specified, the Provisioning Server calculates a value and appends it to the address. The string is converted to:

2-39-8



Example 2

2-391234567890abcde

where 2 specifies the address format type DCCAESA, and 39123456789abcde represents the address prefix. Since no nBits value is specified, the Provisioning Server calculates a value and appends it to the address. And, because the address prefix is an odd number of characters, the Provisioning Server appends a zero to the address prefix. The string is converted to:

2-391234567890abcde0-72

Example 3

2-391234567890abcde-70

where 2 specifies the address format type DCCAESA, 391234567890abcde represents the address prefix, and 70 represents the nBits value. In this example, the valid range for nBits is:

 $(integral-part-of((17 - 1)/2)*8) < nBits \le (integral-part-of((17 + 1)/2))*8$

 $64 < nBits \le 72$

Because the address prefix is an odd number of characters, the Provisioning Server appends a zero to the address prefix. The string is converted to:

2-391234567890abcde0-70

Example 4

5-ff1234-23

where 5 specifies the address format type CustomAESA, ff1234 represents the address prefix (with AFI value ff), and 23 represents the nBits value. In this example, the valid range for nBits is:

 $(integral-part-of((6 - 1)/2)*8) < nBits \le (integral-part-of((6 + 1)/2))*8$

 $16 < nBits \le 24$

DefaultRoute

Specify a Default Route address as the address prefix 00 and 0 bits. For example:

6-00-0

where 6 specifies the address format type DefaultRoute, 00 represents the address prefix, and 0 represents the number of bits.

UserPart

Specify a User Part address as a hexadecimal string of 14 characters.

The value for nBits is 56.

For example:

7-1234567890abcd

where 7 specifies the address format type UserPart, and 1234567890abcd represents the address prefix. Since no nBits value is specified, the Provisioning Server calculates a value and appends it to the address:

7-1234567890abcd-56

Class B Addressing

The Provisioning Server treats all IP Addresses as Class B addresses. The server interprets all addresses as follows:

- First 2 bytes of an IP address are used as the network ID.
- Second 2 bytes of an IP address are used as the switch ID.

The Provisioning Server uses the third byte of the address as the Class B subnet number. For example, the server interprets the following network address:

128.100.111.0

as network address 128.100.0.0 and subnet number 111.


General API Usage



General API Usage

This section provides the basic procedures for performing operations with the Provisioning Server API.

The API operates by establishing a session to the Provisioning Server. The session maintains internal context between the client and the server: it opens a socket and an associated file descriptor. More than one session can be open at a time. You should close a session before the program terminates.

C Program

To use most of the C commands, a client program must follow the following general steps:

- 1. Issue CvConnect to establish a session with the Provisioning Server.
- 2. Identify the object to be operated on. To do so, issue **CvCreate***ObjectType***Id** to fill in the CvObjectId structure.
- 3. Identify necessary arguments (object attributes) and set values, if needed. To do so, either:
 - Issue a single function (**CvArgsMakeVals** or **CvArgsMakeIds**) that takes a variable number of arguments and builds the required data structure.
 - Issue a series of utility functions that create (**CvArgsMake**) and fill in (**CvArgsSet***AttrType*) the required data structure.
- 4. Issue an operational function on the object.
- 5. Use **select** loop processing functions to receive and process the response.
- 6. Once the request has been processed, issue **CvArgsFree** to free the memory used by the argument list.
- 7. When the application exits, issue **CvClose** to terminate the session with the Provisioning Server.

General API Usage



C++ Program

To use most of the C++ commands, a client program must follow these general steps:

- 1. Establish a session with the Provisioning Server. To do so, create a **CvClient** class and issue the **CvClient::open** function to pass **CvClient** arguments that provide session context.
- 2. Identify the object to be operated on. To do so, create and set values in a **CvClient::ObjectId** object.
- 3. Identify necessary arguments (object attributes) and set values, if needed. To do so, create and set values in a **CvClient::Args** object.
- 4. Issue an operational function on the object.
- 5. Use **select** loop processing functions to receive and process the response.
- 6. When the application exits, terminate the session with the Provisioning Server. To do so, either:
 - Issue **CvClient::close**. This function does not delete the **CvClient** class object, but does terminate the session with the Provisioning server.
 - Use the **CvClient** destructor.



Installation and Administration

This chapter describes hardware and software requirements and how to perform a new installation of the Provisioning Server and the Application Toolkit. It also describes the steps required for the following administrative tasks:

- Setting environment variables to configure the various components of the Provisioning Server system
- Stopping and restarting the Provisioning Server and the CLI
- Troubleshooting problems with the Provisioning Server
- Developing a provisioning application



Prerequisites

This section describes the hardware and software required by the NavisXtend Provisioning Server. This product requires two workstations, as follows:

- Provisioning Server
- Provisioning client

Provisioning Server Requirements

The minimum workstation requirement for the Provisioning Server is a SunSparc 5 workstation or equivalent with the following minimum hardware:

- 70 MB disk
- CD-ROM drive

The CPU and RAM requirements for the Provisioning Server depend on the number of clients that will issue requests to the server. Typically, CPU or RAM requirements are less than those required for a CascadeView installation. For details, refer to the *CascadeView/UX Network Management Station Installation Guide*.

The Provisioning Server installation utilizes the CascadeView-specific installation procedures. Thus, CascadeView/UX, Version 2.4.3 or 2.5 must be installed on the workstation to at least the point where the Sybase database is installed and configured. If the software is already installed, you do *not* need to reinstall it. For instructions, refer to the *CascadeView/UX Network Management Station Installation Guide*.

Before you install the Provisioning Server software, verify that the following software programs are installed:

Sun Microsystems SunSoft[™] Solaris[®] 2.4 or 2.5.1 operating environment.

Sybase Open ServerTM, Release 4.9.2 or 11 — The relational database software program for storing database information and providing backup and recovery of database files. This software must be installed on the network and the Provisioning Server must be a client of that database.



CascadeView/UX, Version 2.4.3 or 2.5 — The Provisioning Server installation utilizes the CascadeView-specific installation procedures. Thus, at a minimum, this software must be installed to the point where the Sybase database is installed and configured.

Provisioning Client Requirements

The minimum workstation requirement for the Provisioning client is any SunSparc workstation or equivalent.

Before you install the Application Toolkit software, verify that the following software programs are already installed on the workstation:

Sun Microsystems SunSoft Solaris 2.4 or 2.5.1 operating environment.

SPARCWorksTM compiler version 4.0 — The compiler required to compile a C or C++ program. This software is required only if you plan to write a C or C++ program; it is *not* required if you plan to use the CLI only.

SMIv2 MIB compiler — An SMIv2-compliant compiler required to compile the Provisioning Server MIB. This software is required only if you plan to use the Provisioning Server MIB; it is *not* required if you plan to write a C or C++ program or use the CLI only.

Network Requirements

The Provisioning Server must be configured in a TCP/IP network and must have access to the Ascend switches.

The Provisioning client must have access to the Provisioning Server over a local-area or wide-area network.



Installation Instructions

This section describes how to install the Provisioning Server and the Provisioning Server Application Toolkit.

For instructions on upgrading your Provisioning Server software from a previous version, refer to the *Software Release Notice for NavisXtend Provisioning Server*.

For any updates to this installation procedure, see:

- Software Release Notice for NavisXtend Provisioning Server
- On-line version of this manual on CaseView

Installing the Provisioning Software in a Single-System Configuration

This section describes how to install the Provisioning Server software on the same workstation as the SYBASE database and CascadeView/UX. The procedure requires that you already have CascadeView/UX installed and that the database contains information on the switches that you wish to access through the server.



The installation script prompts you for the Sybase DSQUERY name, CascadeView Sybase database name, and Sybase administrator name and password. Determine these values before you begin the installation.

To install the Provisioning Server and the Application Toolkit, perform the following steps:

- 1. Log on as the root user and enter the root password.
- 2. Insert the Provisioning Server media into the media drive.
- 3. Enter the following command to start the installation script:

[media device]/install_NAVISeps

where [media device] is the name of the machine media device (for example, /cdrom/cdrom0).



The pkgadd menu appears, listing the NAVISeps package.

```
The following packages are available:

1 NAVISeps NavisXtend Provisioning Server

(sparc) 02.00.01.00

Select package(s) you wish to process (or `all' to process all

packages). (default: all) [?,??,q]:
```



If the installation utility detects another instance of the Provisioning Server on your system, it prompts whether you want to remove that instance. If you answer **yes**, it removes the instance and performs a fresh install. If you answer **no**, the install script quits.

4. Select the NavisXtend Provisioning Server package.

The installation utility prompts you to select the components you want to install on the machine.

C)	Install NAVISXtend Client		
s)	Install NAVISXtend Server		
b)	Install both NAVISXtend Client and Server		
q)	Exit this install		
Selection:			

5. Specify which components you want to install on the machine. You can install the Provisioning client (which includes the CLI, the Provisioning Server Application Toolkit client libraries, and the client include files), the Provisioning Server, or both. The Provisioning Server and client occupy approximately 50 MBytes of disk space.

Keep in mind that if you choose to install only the Provisioning Server on a machine, the CLI binaries and associated links will not be present on that server machine.

If you choose to install only the client on a machine, skip to Step 19.

6. When prompted, specify whether CascadeView is installed on the machine.

If you answer yes, the installation utility prompts you to enter the base directory where CascadeView is installed.

7. Enter the path to the directory where CascadeView is installed.



Installation Instructions



- 8. If the installation utility detects configuration files on your system, it prompts whether you want to use these existing files for the installation (instead of having to enter configuration values). If you answer **yes**, the utility will create symbolic links to the configuration files.
- 9. The installation utility prompts whether the file start-server.sh was saved from a previous installation and asks whether you want to re-use the file for this installation. If you answer **yes**, the utility prompts you for the path to the file.
- 10. If the installation utility detects MIB files in /opt/CascadeView/snmp_mibs, it prompts whether you want to create symbolic links to the MIB files in the /opt/ProvServ/snmp_mibs directory. And, it prompts whether you want to create a symbolic link for the Provisioning Server MIB file (provserv.mib) to in /opt/CascadeView/snmp_mibs.
- 11. Indicate your choices to these prompts.
- 12. When prompted, enter the Sybase DSQUERY name.
- 13. When prompted, enter the Sybase Database name for the CascadeView database.
- 14. When prompted, enter the Sybase system administrator user name for the CascadeView database.
- 15. When prompted, enter the system administrator password.
- 16. At the verification prompt, re-enter the system administrator password.

The installation utility displays the values you input and allows you to change them.

- 17. Make any necessary changes.
- 18. When prompted, specify whether you want the installation utility to save copies of the configuration files and start-server.sh at de-install time. If you answer yes, the utility prompts you for the path where you want to save the file.
- 19. The installation utility displays the confirmation message:

Install NAVISXtend 02.00.01.00? (y) [y,n,?,q]

20. Enter y to continue.



The installation utility prompts you to enter the package base directory.

Enter path to package base directory [?,q]:

21. Enter the path to the directory where you want the package installed.

The installation utility performs various verification functions and displays the message:

This package contains scripts which will be executed with super-user permission during the process of installing this package.



Do you want to continue with the installation of this package $[{\tt y},{\tt n},?]$

22. Enter y to continue.

The installation utility completes the installation and displays the message:

Installation of <NAVISeps> was successful.

The installation of the Provisioning Server is complete. Before you run the server, perform the post-installation tasks described in "Post-Installation Tasks" on page 2-8.

Installing the Provisioning Software in a Two-System Configuration

This section describes how to install the Provisioning Server software on a separate host from CascadeView/UX and SYBASE. For information on installing a two-system configuration, refer to the *CascadeView/UX Network Management Station Installation Guide*.

- 1. In CascadeView, add an NMS entry to each switch the Provisioning Server will provision. Specify the IP address of the host on which the Provisioning Server will reside.
- 2. In CascadeView, add an NMS path, specifying the IP address of the host on which the Provisioning Server will reside.
- 3. On the host on which the Provisioning Server will reside, log in as the root user and enter the root password.

- 4. Create the /opt/sybase directory.
- 5. On the CascadeView host, copy the file /opt/sybase/interfaces to the /opt/sybase directory on the host on which the Provisioning Server will reside.
- 6. Install the Provisioning Server. Follow the instructions in "Installing the Provisioning Software in a Single-System Configuration" on page 2-4.

The installation of the Provisioning Server in a two-system configuration is complete.

Before you run the server or the CLI, perform the post-installation tasks described in the next section.

Post-Installation Tasks

This section describes post-installation steps you need to perform on the Provisioning Server, the CLI, and the Provisioning client.

Testing the Server

During server installation, the init program (/etc/inittab) was modified to cause the system to automatically resart the server process whenever the system reboots. To start the server manually for testing, issue the following command:

/sbin/init Q <Return>

This command causes the init program to read the file /etc/inittab.

Test the server to make sure that it is running and is accessible. To do so:

- 1. Log on as a user other than root.
- 2. Issue a CLI command for an existing switch in the CascadeView database:

/opt/ProvServ/bin/cvget switch.nn.nn.nn -Location<Return>

where nn.nn.nn is the decimal IP address of the switch. If the Provisioning Server is operating, the **cvget** command prints the location of the switch you specified. Verify that the returned location is valid for that switch.

For instructions on how to troubleshoot problems with the server, refer to "Troubleshooting Problems" on page 2-25.

Setting Environment Variables

There are several environment variables you can set to configure the Provisioning Server. Specifically, you can:

- Specify the server's local port
- Specify the server's core file location
- Enable server trace files
- Control certain SNMP parameters

For instructions on how to set these environment variables, refer to "Configuring the Provisioning Server" on page 2-17.

If the CLI and the Provisioning Server are located on the same host, the CLI can locate the Provisioning Server by default. If the CLI and the Provisioning Server are remote from one another, you need to identify the location and port number of the Provisioning Server. To do so, set the following environment variables in the user's shell start-up script (such as .cshrc, .login, or .profile):

CV_CLI_SERVER_HOST — Set this variable to the IP address of the remote Provisioning Server. Specify the address in either numeric format (such as 152.148.50.2) or in text format (such as provserv.xyz.com).

CV_CLI_SERVER_PORT — Set this variable to the port number of the remote Provisioning Server.

There are other environment variables you can set to configure the CLI. Specifically, you can:

- Specify whether updates are made to the network component and the database, or to the database only
- Specify security settings
- Control certain SNMP parameters

For instructions on how to set these environment variables, refer to "Configuring the CLI" on page 2-13.



Testing the CLI

Test the CLI to make sure that it is running and can access the Provisioning Server. To do so:

- 1. Log on as a user other than root.
- 2. Issue a CLI command for an existing switch in the CascadeView database:

/opt/ProvServ/bin/cvget switch.nn.nn.nn -Location<Return>

where nn.nn.nn is the decimal IP address of the switch. If the CLI is operating and can access the Provisioning Server, the **cvget** command prints the location of the switch you specified. Verify that the returned location is valid for that switch.

For instructions on how to troubleshoot problems with the CLI, refer to "Troubleshooting Problems" on page 2-25.

Recompiling an Existing Provisioning Client

If you have a Provisioning application that was built with a previous version of the Provisioning Server Application Toolkit and you want to use the new features of the Provisioning Server API, you need to make the necessary code changes for the new functions and attributes, and recompile and relink your program with the new API.

If you *do not* want to use the new features of the Provisioning Server API, no code changes are necessary. You need only to recompile and relink your program with the current version of the API include files and libraries.

Installed Files

Once you install the toolkit, the CLI commands and the files you need to write a program with the API are present on the workstation hard disk:

Command line interface and binary file — Contained in the file /opt/ProvServ/bin/cli, as well as various links contained in /opt/ProvServ/bin.

Client libraries — Contained in the directory /opt/ProvServ/lib.

Client include files — Contained in the directory /opt/ProvServ/include.



Sample code — Contained in the directory /opt/ProvServ/src. The C sample code is in the file provSample.c; the C++ sample code is in the files ProvSample.C and ProvSample2.C.

Programming Files

Table 2-1 lists the files (located in the directory /opt/ProvServ/include) necessary for development of an NavisXtend Provisioning client program.

File	Description
ProvClient.h	Header file for the C APIs; contains definitions and function prototypes. If you are programming in C, include this file in your source code.
CvClient.H	Header file for the C++ APIs; contains definitions and function prototypes. If you are programming in C++, include this file in your source code.
CvDefs.H	Contains some definitions that are common to both the C and C++ APIs. The ProvClient.h and CvClient.H files contain a #include statement that incorporates CvDefs.H. Therefore, as long as you include either ProvClient.h or CvClient.H, you do not need to explicitly include CvDefs.H in your source code.
CvObjectType.H	Defines the enumerated object types used by both the C and C++ APIs. The ProvClient.h and CvClient.H files contain a #include statement that incorporates CvObjectType.H. Therefore, as long as you include either ProvClient.h or CvClient.H, you do not need to explicitly include CvObjectType.H in your source code.

 Table 2-1.
 Programming Files for Client Development



File	Description
CvArgld.H	Defines all argument IDs used by both the C and C++ APIs. The ProvClient.h and CvClient.H files contain a #include statement that incorporates CvArgld.H. Therefore, as long as you include either ProvClient.h or CvClient.H, you do not need to explicitly include CvArgld.H in your source code.
CvParamValues.H	Defines the values for each of the enumerated attributes used by both the C and C++ APIs. Include this file in your source code.
CvObjectId.H	Defines the CvObjectId structure used by the C API to identify objects. ProvClient.h contains a #include statement that incorporates CvObjectId.H. Therefore, as long as you include the ProvClient.h file, you do not need to explicitly include CvArgId.H in your source code.
CvUSL.H	Defines simple wrapper classes for various unsigned long data types used by the C++ APIs. CvClient.H contains a #include statement that incorporates CvUSL.H. Therefore, as long as you include CvClient.H, you do not need to explicitly include CvUSL.H in your source code.
CvE164Address.H	Defines a helper class used in the C++ APIs. CvClient.H contains a #include statement that incorporates CvE164Address.H. Therefore, as long as you include CvClient.H, you do not need to explicitly include CvE164Address.H in your source code.
CvSVCAddress.H	Defines a helper class used in the C++ APIs. CvClient.H contains a #include statement that incorporates CvSVCAddress.H. Therefore, as long as you include CvClient.H, you do not need to explicitly include CvSVCAddress.H in your source code.

Table 2-1. Programming Files for Client Development (Continued)



File	Description			
CvErrors.H and CvErrors.h	Define the errors that can be returned by the APIs as well as errors implemented by CascadeView. You do <i>not</i> need to include either of these files in your source code.			

Table 2-1. Programming Files for Client Development (Continued)

Setting Environment Variables

This section describes how to set environment values to configure the behavior of the CLI, the Provisioning client, and the Provisioning Server. To configure the Provisioning Server, add the environment variables to the start-up script that launches the server. To configure the Provisioning client or the CLI, add the environment variables to the user's shell start-up script, such as .cshrc, .login, or .profile.

Configuring the CLI

There are several environment variables you can use to configure the CLI. Specifically, environment variables perform the following:

- Identify the Provisioning Server to which the CLI sends requests.
- Specify whether updates are made to the network component and the database, or to the database only.
- Specify security settings.
- Control certain SNMP parameters.

The best way to set the environment variables is to add them to the user's shell start-up script (such as .cshrc, .login, or .profile)



Identifying the Provisioning Server to the CLI

If the CLI and the Provisioning Server are running on the same host, the CLI can locate the Provisioning Server by default. If the CLI and the Provisioning Server are remote from one another, you need to identify the location and port number of the Provisioning Server. To do so, set the following environment variables:

CV_CLI_SERVER_HOST — Set this variable to the IP address or hostname of the remote Provisioning Server. Specify the address in numeric format (for example, **152.148.50.2**). Specify the hostname in text format (for example, **provserv.xyz.com**). If you do not set this variable, the CLI uses the local host by default.

CV_CLI_SERVER_PORT — Set this variable to the port number of the remote Provisioning Server. If you do not set this variable, the CLI uses port 4001 by default.

Specifying Modification Type

You can specify whether updates are made to the network components and the database, or to the database only. Set the following environment variable:

CV_CLI_MOD_TYPE — Set this variable to the number that represents the update method, as follows:

1 — Sends updates to both the network component and the database. If the network component updates successfully, the database is updated.

4 — Sends updates to the database only.

5 — Sends updates to the database only and sets a flag in the database indicating that the object is out of synchronization with the network component.

If you do not set this variable, the CLI sends updates to both the network component and the database by default.



Specifying Security Settings

By default, the Provisioning Server accepts requests from the CLI without requiring authorization. You can implement a security feature that authenticates user logins. The feature is intended to prevent users from accidentally modifying the database; it is *not* intended to prevent intentional misuse by users. To implement the security feature, you must specify environment variables for both the CLI and the Provisioning Server. To do so for the CLI, set the following environment variables:

CV_CLI_USE_LOGINS — Set this variable to any value (including a null value) to turn on the security feature. If you do not set this variable, the Provisioning Server accepts requests from the CLI without requiring authorization. If you set this variable, you must also set the following variables:

CV_CLI_USERNAME — Set this variable to the username character string required by CascadeView (for example, operator).

CV_CLI_PASSWORD — Set this variable to the password character string required by CascadeView.

The username and password character strings are sent over the network as nonencrypted text.

To fully implement the security feature, you must also specify security settings on the server side. For instructions, refer to "Implementing the Security Feature" on page 2-24.

Controlling SNMP Parameters

You can specify how certain SNMP parameters are controlled. To do so, set the following environment variables:

CV_SNMP_REQUEST_TIMEOUT — Set this variable to the amount of time (in 0.001 second increments) that the CLI waits for a response from the server. If you do not set this variable, the CLI uses the value 256 by default.

CV_SNMP_MAX_RETRIES — Set this variable to the number of times that the CLI retries a request that times out. If you do not set this variable, the CLI uses the value 0 by default.



Configuring the Provisioning Client

There are several environment variables you can set to configure the Provisioning client. Specifically, environment variables perform the following functions:

- Enable a client trace file
- Control certain SNMP parameters

The best way to set the environment variables is to add them to the user's shell start-up script (such as .cshrc, .login, or .profile).

Enabling a Client Trace File

You can specify that the client create a trace file. Such a file can be useful for debugging your Provisioning client. It is recommended that you enable the trace file until the Provisioning Server and your Provisioning client are running in a production environment. To enable a client trace file, set the following environment variable:

CV_CLIENT_TRACE_FILE — Set this variable to the pathname of the file to contain the trace output (for example, /tmp/ctrace.log). If you do not set this variable, no trace file is created.

Once you enable a client trace file, each session of the client is recorded in the file. Output is continuously appended to the file. If you are not debugging your Provisioning client, it is recommended that you periodically delete the file.

Controlling SNMP Parameters

You can specify how certain SNMP parameters are controlled. To do so, set the following environment variables:

CV_SNMP_REQUEST_TIMEOUT — Set this variable to the amount of time (in 0.001 second increments) that the client waits for a response from the server. If you do not set this variable, the client uses the value 256 by default.

CV_SNMP_MAX_RETRIES — Set this variable to the number of times that the client retries a request that times out. If you do not set this variable, the client uses the value 0 by default. For non-MIB clients, you can set this variable to any value. For MIB clients, you can set this variable to any value only if you set the variable CV_SNMP_DISCARD_RETRY to 1.



Configuring the Provisioning Server

There are several environment variables you can set to configure the Provisioning Server. Specifically, environment variables perform the following functions:

- Specify the server's local port and the MIB agent's port
- Specify the server's core file location
- Enable server trace files
- Control certain SNMP parameters
- Control ListContained context timeout
- Control MIB cache
- Control object locking
- Specify SNMP community strings
- Specify how the server formats SMDS addresses
- Implement security feature

The Provisioning Server reads its configuration settings from two files also used by CascadeView. These shared configuration files are as follows:

- /opt/ProvServ/etc/cvdb.cfg
- /opt/ProvServ/etc/cascadeview.cfg

Rather than modifying Provisioning Server's environment variables directly in these files (which would also affect CascadeView), you can enable them in the server's start-up script (/opt/ProvServ/bin/start-server.sh). Look for the invocation of cvdb.cfg and cascadeview.cfg in start-server.sh and make the necessary modifications after that point in the file.

SEFE

Identifying the Provisioning Server Port

The Provisioning Server uses a command line argument to identify which port to listen for API and CLI requests. To specify this command line argument, set the following environment variable in start-server.sh:

CV_PSRV_ARGS — Set this variable to the command **-lport** and the port number. Enclose the command in quotation marks, for example:

```
"-lport 4002"
```

If you do not set this variable, the Provisioning Server uses port 4001 by default.

Identifying the MIB Agent Port

The Provisioning Server implements an SNMP agent as a separate entity to service MIB interface requests. The server uses an environment variable to identify which port to listen for SNMP requests. This port is different from the port number used to listen for API and CLI requests.

To specify this MIB agent port, set the following environment variable in start-server.sh:

CV_SNMP_AGENT_PORT — Set this variable to the port number. If you do not set this variable, the Provisioning Server uses port 9090 by default.

Specifying the Core File Location

If the Provisioning Server crashes, it creates a core file. Such a file can be useful for debugging the server. The core file is written to the Provisioning Server's working directory (/tmp by default). You can specify the directory where the Provisioning Server runs and where it writes any core files. You may want to specify a directory other than the default if /tmp gets deleted frequently and you want to ensure a valid core file. To specify a working directory, set the following environment variable in start-server.sh:

CV_WORKING_DIR — Set this variable to the pathname of the directory. If you do not set this variable, the Provisioning Server writes its core file to the /tmp directory by default.



Enabling Server Trace Files

By default, the Provisioning Server creates three trace files, two of which are enabled by environment variables specified in the configuration file cascadeview.cfg. Rather than turning these trace files on or off directly in cascadeview.cfg (which would also affect CascadeView), you can enable them in the server's start-up script (/opt/ProvServ/bin/start-server.sh). Look for the invocation of cascadeview.cfg in start-server.sh and make the necessary modifications after that point in the file.

These files can be useful for troubleshooting and diagnosing problems. It is recommended that you enable the trace files until the Provisioning Server is running in a production environment. To enable the trace files, set the following environment variables:

CV_TRACE_ENABLE — Set this variable to 1 to enable the application-level trace output for the server. If you set this variable, you must also set the **CV_TRACEFILE** variable.

CV_TRACEFILE — Set this variable to the pathname of the file to contain the application-level trace output for the server. To avoid conflicts with the CascadeView trace file, the suffix .psrv will be appended to the filename you specify. By default, this trace file is written to the /tmp directory.

CVDB_TRACE_FILE_NAME — Set this variable to the pathname of the file to contain the database trace output for the server. To avoid conflicts with the CascadeView trace file, the suffix .psrv will be appended to the filename you specify. By default, this trace file is written to the /tmp directory.

CV_PSRV_TRACE_FILE — Set this variable to the pathname of the file to contain trace output specific to the Provisioning Server. By default, this trace file is written to the /tmp directory and is called strace.log.

Once you enable a trace file, specific activity is recorded in the file. Output is continuously appended to the file. It is recommended that you periodically delete the trace files.

If you are troubleshooting a problem, it can be useful to know what kinds of transactions occur between the Provisioning Server and the Provisioning client. For this reason, you should enable the client trace file as well. For instructions, refer to "Enabling a Client Trace File" on page 2-16.



Controlling SNMP Parameters

You can specify how certain SNMP parameters are controlled. To do so, set the following environment variables in start-server.sh:

CV_SNMP_REQUEST_TIMEOUT — Set this variable to the amount of time (in 0.001 second increments) that the server waits for a response from the switch. If you do not set this variable, the server uses the value 256 by default.

CV_SNMP_MAX_RETRIES — Set this variable to the number of times that the server retries a request that times out. If you do not set this variable, the server uses the value 5 by default. It is recommended that you keep this setting as the default.

CV_SNMP_DISCARD_RETRY — Set this variable to 1 to enable the server to discard multiple SNMP request retries from a MIB client. If you set this variable to 1, the server checks the Request ID, the IP address, and the port number of every SNMP request. If these values match those of a request that the server is currently processing, the server ignores the retry request. If you do not set this variable, the server uses the value 1 by default. It is recommended that you keep this setting as the default, unless your SNMP client generates SNMP PDUs without unique Request IDs or port numbers.

Controlling Context Timeout

The Provisioning Server maintains context for outstanding ListContained requests. The server allows 500 ListContained requests to be outstanding. Any ListContained request for which a NextObject request has not been issued within a configurable time period is subject to deletion to make room for a new ListContained request to be processed. To configure this time period, set the following environment variable in start-server.sh:

CV_PSRV_CONTEXT_TIMEOUT — Set this variable to the amount of time (in minutes) that the server waits for a response to a ListContained request. Any request for which a NextObject request has not been issued in that time period is subject to deletion. If you do not set this variable, the server uses the value 10 by default. It is recommended that you keep this setting as the default.



Controlling MIB Cache

The Provisioning Server implements a MIB cache that stores data in memory for a fixed time period. The server uses this cache to optimize performance of **get-next** requests and to store data to be committed to the database during transactions involving multiple PDUs. Each table row stored in cache has a timestamp. The server uses an environment variable to purge older data by row.

To configure this purge time period, set the following environment variables in start-server.sh:

CV_SNMP_ROWENTRY_TIMEOUT — Set this variable to the amount of time (in seconds) that the server stores a particular row of data in cache during a **get-next** request. Based on this variable, the server flushes out entries in MIB cache that result from a **get-next** operation. Thus, the server uses this variable to optimize performance of **get-next** requests. The minimum value of this variable is 60, the maximum value is 1800. These values apply, even if you set a value lower than the minimum or greater than the maximum. If you do not set this variable, the server uses the timeout value 900 by default.

CV_SNMP_CMDERROR_CACHE_TIMEOUT — Set this variable to the amount of time (in seconds) that the server stores Command Error Table messages in cache. The Command Error Table contains error messages generated by the SNMP agent. Any error message older than this timeout value is subject to deletion. If you do not set this variable, the server uses the timeout value 3600 by default. To save all errors generated during creation and modification transactions, set this variable to a value greater than the value of the CV_SNMP_LOCK_TIMEOUT variable.

Controlling Object Locking

The Provisioning Server uses an object locking scheme for MIB objects in the database that differs from the locking behavior of the Provisioning Server API, CLI, or CascadeView. For these interfaces, the steps associated with locking are transparent to the user. When an object is created or modified, its parent object becomes locked. The user specifies all the information needed to create or modify the object in one PDU. Once the request completes, the parent becomes unlocked.

Setting Environment Variables



By contrast, in the case of the MIB, the information needed to create or modify an object may not be available in one PDU. As a result, the locks in the database must be held for a longer time. Thus, the steps associated with locking are not transparent to the user.

If the user initiates a transaction to create an object, the parent object becomes locked, preventing other users from modifying it. If the user initiates a transaction to modify an object, the object itself becomes locked, preventing other users from modifying it. To configure the time period that objects are locked, set the following environment variable in start-server.sh:

CV_SNMP_LOCK_TIMEOUT — Set this variable to the amount of time (in seconds) that the server:

- Locks a parent object when a child object is being created
- Locks an object that is being modified

The maximum value of this variable is 1800. This maximum value applies, even if you set a greater value. If you do not set this variable, the server uses the timeout value 900 by default.

Specifying Community Strings

The Provisioning Server supports two community names, one for Read-Only operations and one for Read-Write operations. The community name provides a mechanism for authentication and access-control at the SNMP agent.

The community strings are defined using the following environment variables in start-server.sh:

CV_READONLY_COMMUNITY_STRING — Set this variable to the community string to be used when making a Read-Only SNMP request. If you do not set this variable, the server uses the value 'public' by default.

CV_READWRITE_COMMUNITY_STRING — Set this variable to the community string to be used when making a Read-Write SNMP request. If you do not set this variable, the server uses the value 'ascend' by default.

If these environment variables are defined in the script start-server.sh, the specified strings take precedence. If they are not set in the script or if the server shell environment does not define the variables, the server assumes the default values.

Setting Environment Variables



If the community name is not valid when you make an SNMP request, the request is rejected and the SNMP agent returns a genError. You can access the Command Error Table in the MIB to see if the source of the problem is an invalid community name. Specify the Read-Only community name when you access the table, as that community name is used for validation purposes.

When you make an snmp_get request, specify either the Read-Only or the Read-Write community name. If you use a different community name and you encounter an error, the error is not propagated to the Command Error Table.

Controlling SMDS Addresses

You can specify how the Provisioning Server formats SMDS addresses. To do so, set the following environment variables in start-server.sh:

CV_DFLT_SMDS_CC — Set this variable to specify the default country code for SMDS addresses. The server will prepend this default country code to a given address that does not specify the country code. When using multiple country codes, you must specify the country code for addresses that do not use the default code. You have to create a country code before you can specify it as a default.

CV_DFLT_CC_PRT_ENABLE — Set this variable to control the format of individual addresses in responses to List operations. When this variable is set to 1, the default country code part of an address is returned in the List response. For other operations (AddObject, DeleteObject, Get, Modify), the server returns the address in the same format used by the client.

CV_SMDS_MASK_SIZE — Set this variable to specify the character length of the address prefix in SMDS addresses. The server interprets characters preceding a dash (-) as the country code part of the address, the next n characters (specified by this variable) as the prefix part of the address, and the remainder as the address part. For example, if this variable is set to 6, the server interprets the address 1-9789521111 as follows:





Implementing the Security Feature

By default, the Provisioning Server accepts requests from Provisioning client and CLI users without requiring authorization. You can implement a security feature that authenticates user logins. The feature is intended to prevent users from accidentally modifying the database; it is *not* intended to prevent intentional misuse by users. To implement the security feature, you must specify environment variables for both the CLI and the Provisioning Server. To do so for the server, set the following environment variable in start-server.sh:

CV_PSRV_USE_LOGINS — Set this variable to any value (including a null value) to turn on the security feature. If you do not set this variable, the Provisioning Server accepts requests from clients without requiring user authorization.

Once you set this variable, any clients sending requests to the server must send a user ID and password for authorization. For a Provisioning client, this is accomplished when the client establishes the session with the Provisioning Server. For the CLI, the security settings are specified through environment variables. For instructions, refer to "Specifying Security Settings" on page 2-15.

Stopping and Restarting the Provisioning Server

To stop and restart the Provisioning Server running on a workstation:

- 1. On the host that runs the server, log on as the root user and enter the root password.
- 2. Determine the process ID of the Provisioning Server, using the following command:

/bin/ps -ef | grep provserv <Return>

The process ID is the second item in the resulting listing.

3. Kill the current server process, using the following command:

kill [server process id] <Return>

Once the server process is killed, the init program restarts the server.



Stopping and Restarting the CLI

To quit the CLI, press <Ctrl-C>. To restart the CLI, issue a CLI command.

Troubleshooting Problems

This section describes how to troubleshoot problems with the Provisioning Server, the provisioning application, and the CLI.

Problem: Requests Frequently Time Out

Symptoms

Either:

- CLI prints an error message
- API-based application receives an error status

Possible Causes and Solutions

Scenario 1: Error message 4109 ("Request to the server timed out")

- The Provisioning Server may not be running or accessible to the client workstation. Verify that the client can access the server and that the Provisioning Server is running. To do so, follow the procedure in "Testing the CLI" on page 2-10.
- The client's timeout value may be too low. The client timeout value should allow for instances when the server times out and later retries a command to the switch. Since the server's second request may be successful, the client should not timeout while waiting for the server's response. Adjust the client timeout value by setting the client CV_SNMP_REQUEST_TIMEOUT environment variable. Start with the value 3000. If that value does not correct the problem, use the following formula to determine a "worst case" client timeout value:

 $\label{eq:cv_SNMP_REQUEST_TIMEOUT} CV_SNMP_REQUEST_TIMEOUT * CV_SNMP_MAX_RETRIES + n$

Troubleshooting Problems



where *CV_SNMP_REQUEST_TIMEOUT* is the timeout value for the server, *CV_SNMP_MAX_RETRIES* is the retry value for the server and *n* is a factor that allows for client-server round-trip. Start with an *n* value of 300. The result (CV_SNMP_REQUEST_TIMEOUT) is the timeout value to set for the client.

For details on values to use for these environment variables for the client and server, refer to "Setting Environment Variables" on page 2-13.

Scenario 2: Error message 42 ("The SNMP request to the agent timed out"):

The Provisioning Server may take a long time to process a request because it cannot locate the network device specified in the request (such as a switch):

- If the request is intended to modify the switch, verify that the switch is accessible from the server. To do so, remotely log into the Provisioning Server and issue the ping utility to elicit a response from the switch.
- If the request is intended to update the database only, retry the request with the modification type set to update the database only. For a CLI request, set the CV_CLI_MOD_TYPE environment variable to 4 or 5 (refer to "Configuring the CLI" on page 2-13.) For an API request, issue either the C function CvSetModifyType or the C++ member function CvClient::setModifyType, specifying that updates be made to the database only.

Problem: Object Is Locked by Others

Symptoms

Either:

- CLI prints an error message
- API-based application receives an error status

Possible Causes and Solutions

Either a CascadeView user has the object locked or the object appears to be locked when the client retries a request. To determine if the object is locked, change directories to /opt/CascadeView/bin and execute the cv-release-locks.sh shell script. The script lists the objects that are currently locked and who has them locked.



Do not use the cv-release-lock.sh script to release the locks. If you need to release locks, call the Ascend Technical Assistance Center.

Scenario 1: Object Is Locked by CascadeView User

If the cv-release-locks.sh shell script indicates that a CascadeView user has the object locked, either:

- Wait for the user to finish (or request that he or she finish) using the object.
- Call the Ascend Technical Assistance Center.

Scenario 2: Object Appears to Be Locked During Retries

If the cv-release-locks.sh shell script does not indicate that the object is locked, a client timeout may have occurred while the server was still processing the request. Then, when the client automatically retried the request, the object appeared to be locked.

Adjust one of the following environment variables:

- Adjust the client timeout value by setting the client's CV_SNMP_REQUEST_TIMEOUT environment variable to a higher value. To do so, follow the procedure in "Scenario 1: Error message 4109 ("Request to the server timed out")" on page 2-25.
- Adjust the client retry value by setting the CV_SNMP_MAX_RETRIES to 0. For details on this environment variable for the client, refer to "Setting Environment Variables" on page 2-13.



Technical Support

The Ascend Technical Assistance Center (TAC) is available to assist you with any problems encountered while using the NavisXtend Provisioning Server product. To contact the Ascend TAC, call 1-800-DIAL-WAN.

Information Checklist

Before contacting the Ascend TAC, review the following checklist to make sure you have gathered all the information you need:

Software Version Number

Use the UNIX utility pkginfo to obtain information such as version number and install date for the NavisXtend Provisioning Server package:

pkginfo -1 NAVISeps

Note the version number listed in the output.

Problem Report

Collect as much information as possible about the problem:

- For CLI problems, describe what commands caused the problem, what commands preceded the problem, and how did the Provisioning Server respond (such as what error message was returned). If possible, provide the exact text for the commands.
- For API problems, provide the source code that caused the problem. Try to condense the problem to a few lines.
- You can use the API to create a CLI command that recreates the problem. This alternative provides an easy way to recreate a problem without having to provide code. The following code sample illustrates how to use the API to create a CLI command:

```
char *argString = CvArgsToString( args );
char *objString = CvObjectIdToString( objid );
printf( "cvadd %s %s", objString, argString );
CvStringFree( argString );
CvStringFree( objString );
```



Trace Files

Collect any trace files that may exist:

- Server trace files, which you enable using environment variables. By default, these trace files are not produced. The easiest way to turn them on is to edit the start-server.sh script. They are usually written to the /tmp directory with the filenames strace.log or the file suffix .psrv.
- Client trace files, which you enable using environment variables. By default, these trace files are not produced. The easiest way to turn them on is to edit the user's .cshrc file and adding the following line:

```
setenv CV_CLIENT_TRACE_FILE /tmp/ctrace.log
```

This command writes the client trace file ctrace.log to the /tmp directory.

If the resulting trace files are too large, collect the last 5000 lines of each file. If necessary, compress the files using the GZIP program. If you send the compressed files to Ascend by email, UUENCODE the files, if necessary.

For more information on how to enable trace logs, refer to "Enabling Server Trace Files" on page 2-19 and "Enabling a Client Trace File" on page 2-16.

Core Files

If the Provisioning Server crashes, it creates a core file. Collect the core file from the Provisioning Server's working directory, which is either /tmp by default or another directory you specify using an environment variable. For information on how specify the working directory, refer to "Specifying the Core File Location" on page 2-18.



Un-installation Instuctions

If you decide you want to un-install Version 2.0 of the Provisioning Server and Application Toolkit, use the pkgrm utility:

1. To un-install the Provisioning Server components using pkgrm, enter:

```
pkgrm NAVISeps
```

The utility prompts you to verify the un-install:

The following package is currently installed: 1 NAVISeps NavisXtend Provisioning Server (sparc) 02.00.01.00

Do you want to remove this package?

2. To un-install the NavisXtend Provisioning Server package, enter **y**. The un-installation utility displays the message:

Removing installed package instance <NAVISeps>

This package contains scripts which will be executed with super-user permission during the process of removing this package.

Do you want to continue with the removal of this package [y,n,?, q]

3. Enter y to continue.

The un-installation utility performs various verification functions and displays the confirmation message:

Are you sure you want to UNINSTALL the Provisioning Server [y/n]?

4. Enter **y** to continue.

The utility completes the un-installation:

Un-install complete

Removal of <NAVISeps> was successful.

The un-installation of the Provisioning Server components is complete.



Writing a Provisioning Application

To write a Provisioning application, perform the following steps:

- 1. Install the Provisioning Server Application Toolkit, as described in "Installation Instructions" on page 2-4.
- Set the environment variables that control SNMP parameters for the Provisioning client. For instructions, refer to "Configuring the Provisioning Client" on page 2-16.
- 3. Add the following entries to your makefile:

-I/opt/ProvServ/include

-L/opt/ProvServ/lib -lClient

The first line is for all compilations; the second line is for the link step.

- 4. Write the program.
- 5. Compile the program.

Upgrading an Existing Application

If you have a Provisioning application that was built with a previous version of the Provision Server Application Toolkit and you want to use the new features of the Provisioning Server API, you need to make the necessary code changes for the new functions and attributes, and recompile and relink your program with the new API.

If you *do not* want to use the new features of the Provisioning Server API, no code changes are necessary. You need only to recompile and relink your program with the current version of the API include files and libraries.



Using the CLI

This chapter describes how to use the Command Line Interface (CLI) to build a provisioning script instead of a C or C++ program.

To understand the Provisioning Server object hierarchy, first read Chapter 1.

Using the CLI



Using the CLI

The Application Toolkit provides a Command Line Interface (CLI) for users to build a provisioning script instead of a C or C++ program. The CLI is a set of command-line programs that you can issue from any UNIX shell to provision network objects in interactive or batch mode.

There is a CLI command for each operational function of the API. Each command uses a string representation to specify objects and attributes.

cvadd (Object ID, Attributes) — Creates an object in the database and (optionally) in the switch.

cvaddmember (Object ID, Object ID) — Adds a member (usually an address) to an object list.

cvmodify (Object ID, Attributes) - Modifies specific attributes of an object.

cvdelete (Object ID) — Deletes an object from the database and (optionally) from the switch.

cvdeletemember (Object ID, Object ID) — Deletes a member (usually an address) from an object list.

cvget (Object ID, Attributes) — Retrieves specific attribute values from the database.

cvlistcontained (Object ID, type, Attributes) — Retrieves a list of configuration attributes for objects of the given type contained by the specified parent.

cvlistallcontained (Object ID, Attributes) — Retrieves a list of configuration attributes for all objects contained by the specified parent.

The commands are supported for most target object types, with a few restrictions. For example, you cannot specify a switch when you issue an Add or Delete command, as the Provisioning Server does not support adding or deleting switches.

The cvhelp command provides usage help for the CLI.

For a list of the object types you can use when you issue the operational functions of the CLI, refer to Table 1-4 on page 1-38.

There are several environment variables you can use to configure the behavior of the CLI. For details, refer to "Setting Environment Variables" on page 2-13.



CLI Usage Overview

Most of the CLI commands use the following syntax:

command object-name {-attribute-name value}

Syntax

command	The name of the command. If the command is in your path, you can enter just the command name, such as cvadd , cvdelete , cdmodify , cvget , cvaddmember , or cvdeletemember . Otherwise, you must prefix the command name with the path /opt/ProvServ/bin/.
	To use standard input to specify arguments, issue a dash (-) after the command name.
object-name	The object ID. To specify an object ID, you first specify the object's parent (if any), including the parent type and value. Then, you specify the child type and value. For rules on specifying object IDs for various types of objects, refer to "Managed Objects" on page 1-14.
-attribute-name	The attribute ID appropriate to the object ID. Specify the attribute name preceded by a dash (-). Use the attribute ID symbols listed in the <i>NavisXtend</i> <i>Provisioning Server Object Attribute Definitions</i> , but omit the CVA_ <i>ObjectType</i> prefix. For example, specify location as: -Location.


value The valu data typ integer,

The value of the attribute ID. The value requires a data type appropriate for the argument, such as integer, string, and so on. For data types, use the data types listed in the *NavisXtend Provisioning Server Object Attribute Definitions*. Note the following rules for values:

- For integers, specify the integer value.
- For strings, enclose the string in quotes if it contains special characters, such as a blank character. String values cannot begin with a hyphen.
- For enumerated types, specify the text value that represents the integer value. In most cases, the CLI uses an abbreviated text value.
- For Object ID, specify the Object ID that identifies the object in the containment hierarchy (refer to "Managed Objects" on page 1-14).

-attribute-name and *value* are optional parameters. You can specify up to 40 attribute-value pairs.

Before the CLI issues a command to the Provisioning Server, it checks the command for correct syntax. The server checks the input parameters for validity and reports errors back to the client.

To maximize CLI efficiency, *do not* set all possible attributes in a request. Specify only attributes that are mandatory.

In some cases, a CLI command line may become too long for the shell to handle. This can happen most often when adding LPorts. The restriction is most likely to happen when using the sh or csh shells. It occurs only in certain circumstances when using ksh. To work around this buffer restriction, separate the CLI command into multiple lines. At the end of each line, insert the backslash character (\) immediately followed by the <Return> key. This instructs the shell that the next line is part of the same command.



For example:

cvadd switch.1.1.1.1.card.9.pport.1.lport.2 \<Return> -serviceType smds -smdsType SsiDte \<Return> -bandwidth 64000

The sections that follow present the CLI commands in alphabetical order. Usage examples are provided with each command. Use these examples as guidelines for syntax and usage. The exact attributes required by a particular command vary, depending on the type of LPort and Card specified. For additional examples, refer to "CLI Examples" on page 3-29.



cvadd

Purpose

Creates an object in the database and (optionally) in the switch. The attributes specified by the command are used to initialize the object.

Command Syntax

cvadd object-name {-attribute-name value } ...

Parameters

object-name specifies the object to be added. The object is specified by its object ID, based on the containment hierarchy (for information, refer to "Managed Objects" on page 1-14).

-attribute-name value specifies an attribute and its value to be added to the object. The attribute is specified by its argument name. The value uses a data type appropriate for the argument. You can specify up to 40 attribute-value pairs.

Specify only those attributes and values appropriate for the object type. You can specify any attribute except one with either the Read-Only or Create-Only access restriction.

Notes

For a list of object types that you can add with this command, refer to Table 1-4 on page 1-38.

To create a card or PPort, use the Modify command (**cvmodify**). The CascadeView database automatically populates each switch with cards of type "empty". Use the Modify command to change the card's type from "empty" to the specified type. Likewise, once a card has been configured, CascadeView automatically populates the card with all necessary Physical Ports. Use the Modify command to change the PPort specifications. In the case of the channelized DS3 card, once the card has been configured, CascadeView automatically populates the card with all necessary channels. Use the Modify command to change the card with all necessary channels. Use the Modify command to change the card with all necessary channels.



If **cvadd** is successful, it prints the command name followed by the arguments that the Provisioning Server returns. You can use this output to verify that the arguments are the same as those specified in the original request. Any attribute that is missing a valid value is a required attribute that you omitted.

Examples

The following cvadd command creates an LPort:

```
/opt/ProvServ/bin/cvadd - Switch.1.1.1.2.card.4.pport.3.lport.1
-Name lport1 -SmdsType SsiDte -ServiceType Smds -Bandwidth 64000
-ErrorPerMinThreshold 0 -AdminStatus Up -ErrorCheckFlag Off
-HeartBPFlag On -SmdsPduViolTcaFlag Disable -HeartBPInterval 1
-HeartBPNAThresh 1
```

If successful, the command returns the following text:

```
/opt/ProvServ/bin/cvadd Switch.1.1.1.2.card.4.pport.3.lport.1
-Name lport1 -SmdsType SsiDte -ServiceType Smds -Bandwidth
64000 -ErrorPerMinThreshold 0 -AdminStatus Up -ErrorCheckFlag
Off -HeartBPFlag On -SmdsPduViolTcaFlag Disable -HeartBPInterval
1 -HeartBPNAThresh 1
```

The following **cvadd** command creates a circuit connecting a Frame Relay LPort to a PPPto1490 LPort:

```
/opt/ProvServ/bin/cvadd -
Switch.1.1.1.1.card.5.pport.5.lport.5.dlci.22 -Endpoint2
Switch.1.1.1.2.card.5.pport.5.lport.5.dlci.23 -GracefulDiscard
Enabled -AdminStatus Up -Priority Low -RerouteBalance Disabled
```

In this example:

- The first endpoint (Switch.1.1.1.1.card.5.pport.5.lport.5) is a Frame Relay LPort.
- The second endpoint (Switch.1.1.1.2.card.5.pport.5.lport.5) is a PPPto1490 LPort.



cvaddmember

Purpose

Adds an address to a screen or netwide group address. Upon completion of the command, the address represented by the second object parameter is added to the object specified by the first object parameter.

Command Syntax

cvaddmember object-name object-name

Parameters

object-name specifies the objects. Each object is specified by its object ID, based on the containment hierarchy (for information, refer to "Managed Objects" on page 1-14). The first *object-name* specifies the screen or netwide group address; the second *object-name* specifies the address to become a member of the screen or netwide group address.

Notes

For a list of object types that you can add with this command, refer to Table 1-4 on page 1-38.

When you specify the object CVT_SmdsGroupScreen as the container object, the member to be added must be either a CVT_SmdsAlienGroupAddress or a CVT_SmdsSwitchGroupAddress.

When you specify the object CVT_SmdsIndividualScreen as the container object, the member to be added must be either a CVT_SmdsLocalIndividualAddress or a CVT_SmdsAlienIndividualAddress.

When you specify the object CVT_SmdsNetwideGroupAddress as the container object, the member to be added must be a CVT_SmdsLocalIndividualAddress.



If **cvaddmember** is successful, it prints the command name followed by the arguments that the Provisioning Server returns. You can use this output to verify that the arguments are the same as those specified in the original request.

Example

The following **cvaddmember** command adds an SMDS local individual address to an SMDS netwide group address:

```
/opt/ProvServ/bin/cvaddmember -
Network.1.1.1.0.NetwideGroupAddress.1234567899
Switch.1.1.1.1.card.3.pport.4.lport.1.LocalIndividualAddress.12345
67890
```

If successful, the command returns the following text:

```
/opt/ProvServ/cvaddmember Network.1.1.1.0.NetwideGroup
Address.123456789 9 Switch.1.1.1.1.card.3.pport.4.lport.1.
LocalIndividualAddress.1234567890
```



cvdelete

Purpose

Deletes an object from the database and (optionally) from the switch.

Command Syntax

cvdelete object-name

Parameters

object-name specifies the object to be deleted. The object is specified by its object ID, based on the containment hierarchy (for information, refer to "Managed Objects" on page 1-14).

Notes

For a list of object types that you can delete with this command, refer to Table 1-4 on page 1-38.

You only need to delete an SMDS switch group address if the database shows an SMDS switch group address that should not exist.

To remove a card, use the Modify command (**cvmodify**) to change the card's type to "empty".

Some objects cannot be deleted until the objects they contain have been deleted. For example, you cannot delete an LPort until you delete all of its circuits and addresses.

If **cvdelete** is successful, it prints the command name followed by the arguments that the Provisioning Server returns. You can use this output to verify that the arguments are the same as those specified in the original request.

cvdelete



Example

The following **cvdelete** command deletes a circuit:

```
/opt/ProvServ/bin/cvdelete -
Switch.1.1.1.1.card.4.pport.1.lport.1.Dlci.16
```

If successful, the command returns the following text:

```
/opt/ProvServ/bin/cvdelete Switch.1.1.1.1.card.4.pport.1.lport.1.d
lci.16
```



cvdeletemember

Purpose

Deletes an address from a screen or netwide group address. Upon completion of the command, the address represented by the second object parameter is removed from the object specified by the first object parameter.

Command Syntax

cvadeletemember object-name object-name

Parameters

object-name specifies the objects. Each object is specified by its object ID, based on the containment hierarchy (for information, refer to "Managed Objects" on page 1-14). The first *object-name* specifies the screen or netwide group address; the second *object-name* specifies the address to be removed from the screen or netwide group address.

Notes

For a list of object types that you can delete with this command, refer to Table 1-4 on page 1-38.

When you specify the object CVT_SmdsGroupScreen as the container object, the member to be removed must be either a CVT_SmdsAlienGroupAddress or a CVT_SmdsSwitchGroupAddress.

When you specify the object CVT_SmdsIndividualScreen as the container object, the member to be removed must be either a CVT_SmdsLocalIndividualAddress or a CVT_SmdsAlienIndividualAddress.

When you specify the object CVT_SmdsNetwideGroupAddress as the container object, the member to be removed must be a CVT_SmdsLocalIndividualAddress.



If **cvdeletemember** is successful, it prints the command name followed by the arguments that the Provisioning Server returns. You can use this output to verify that the arguments are the same as those specified in the original request.

Example

The following **cvdeletemember** command removes an SMDS alien group address from an SMDS group screen:

```
/opt/ProvServ/bin/cvdeletemember -
Switch.1.1.1.1.card.3.pport.4.lport.1.GroupScreen
Switch.1.1.1.1.AlienGroupAddress.0009998887
```

If successful, the command returns the following text:

```
/opt/ProvServ/cvdeletemember Switch.1.1.1.1.card.3.pport.4.lport.1
.GroupScreen Switch.1.1.1.AlienGroupAddress .0009998887
```



cvget

Purpose

Retrieves the values of specific attributes from the database.

Command Syntax

cvget object-name {-attribute-name } ...

Parameters

object-name specifies the object whose attributes are to be retrieved. The object is specified by its object ID, based on the containment hierarchy (for information, refer to "Managed Objects" on page 1-14).

-attribute-name specifies an attribute to be retrieved. The attribute is specified by its argument name. Specify only attribute names with no values. You can specify up to 40 attributes.

Specify only those attributes appropriate for the object type.

Notes

For a list of object types that you can use with this command, refer to Table 1-4 on page 1-38.

If **cvget** is successful, it prints the command name followed by the arguments that the Provisioning Server returns. You can use this output to verify that the arguments are the same as those specified in the original request.

Examples

The following **cvget** command retrieves the type and administrative status of a card:

```
/opt/ProvServ/bin/cvget -
Switch.1.1.1.1.card.4 -DefinedType -AdminStatus
```



If successful, the command returns the following text:

```
/opt/ProvServ/bin/cvget Switch.1.1.1.1.card.4 -DefinedType
1PortAtmDs3Uni
-AdminStatus Up
```

The following **cvget** command retrieves the location of a switch:

/opt/ProvServ/bin/cvget Switch.152.148.50.2 -Location

If successful, the command returns the following text:

```
/opt/ProvServ/bin/cvget Switch.152.148.50.2
```

-Location "XYZ Corporation"



Purpose

cvhelp

Provides usage help for the CLI.

Command Syntax

cvhelp { object-type -attribute-name }

Parameters

object-type specifies the object type (such as LPort, circuit, etc.) for which you want to print supported attributes or enumerated attribute values.

-attribute-name specifies an enumerated attribute for which you want to print supported enumerated values printed.

Notes

Issue **cvhelp** without arguments to print a command usage statement for each of the CLI commands.

Issue **cvhelp** with the *object-type* argument to print the attribute IDs and attribute types (such as INTEGER, STRING, and so on) that are supported for the specified object.

Issue **cvhelp** with the *object-type* and *-attribute-name* arguments to print the enumerated attribute values that are supported for the specified attribute and object.

Examples

The following **cvhelp** command prints a usage statement for each of the CLI commands:

```
/opt/ProvServ/bin/cvhelp
```



The following **cvhelp** command prints a list of all attributes supported for cards:

/opt/ProvServ/bin/cvhelp card

The following **cvhelp** command prints a list of all enumerated attribute values supported for the enumerated attribute CVA_LPortSmdsType belonging to the object LPort:

```
/opt/ProvServ/bin/cvhelp lport -smdstype
```



cvlistallcontained

Purpose

Queries the database for a list of objects of any type that are immediate children of a specified object.

Command Syntax

cvlistallcontained object-name

Parameters

object-name specifies the parent object that represents the immediate parent of the contained objects (such as a PPort that is a parent of multiple LPorts). The object is specified by its object ID, based on the containment hierarchy (for information, refer to "Managed Objects" on page 1-14).

Notes

You can issue the function on either:

Network level — The function retrieves a list of objects on a network, including all subnets. To issue the function on a network level, specify an IP address, such as 128.100.0.0.

Subnet level — The function retrieves a list of objects on a particular subnet. To issue the function on a subnet level, specify an IP address with a subnet number, such as 128.100.111.0.

Table 3-1 lists the valid parent and child object types you can specify with this command.



Parent Object Type	Child Object Types
Network	Switch VPN Customer SmdsCountryCode SmdsNetwideGroupAddress SvcSecScn SvcCUG SvcCUG SvcCUGMbrRule TrafficDesc NetCac ServiceName
Customer	LPort Circuit
VPN	LPort Circuit
SvcCUGMbrRule	SvcCUGMbr
SvcCUG	SvcCUGMbr
Switch	Card SvcNodePrefix SmdsSwitchGroupAddress SmdsAddressPrefix SmdsAlienIndividualAddress SmdsAlienGroupAddress
Card	PPort
PPort	LPort Channel PFdl (8-port ATM T1 card only) PerformanceMonitor Aps (1-port OC-12c/STM-4 card only)
Channel (valid for channelized DS3 card only)	LPort

Table 3-1. Valid Parent and Child Object Types



Parent Object Type	Child Object Types
LPort	Circuit SvcPrefix SvcAddress SvcUserPart SvcConfig SmdsLocalIndividualAddress SmdsIndividualScreen SmdsGroupScreen PMPCktRoot AssignedSvcSecScn SvcSecScnActParam Spvc PMPSpvcRoot
SmdsIndividualScreen	SmdsLocalIndividualAddress SmdsAlienIndividualAddress
SmdsGroupScreen	SmdsSwitchGroupAddress SmdsAlienGroupAddress
SmdsNetwideGroupAddress	SmdsSwitchGroupAddress SmdsLocalIndividualAddress
PMPCktRoot	PMPCktLeaf
PMPSpvcRoot	PMPSpvcLeaf
ServiceName	Circuit

Table 3-1. Valid Parent and Child Object Types (Continued)

If **cvlistallcontained** is successful, it prints the command name followed by the child objects that the Provisioning Server returns. It prints out one line for each child object. The command does not print any attributes for the listed objects.

Example

The following **cvlistallcontained** command lists all immediate children of a switch:

```
cvlistallcontained switch.1.1.1.1
```

cvlistallcontained



If successful, the command returns the following text:

cvlistallcontained	Switch.1.1.1.1.card.1
cvlistallcontained	Switch.1.1.1.card.2
cvlistallcontained	Switch.1.1.1.card.3
cvlistallcontained	Switch.1.1.1.card.4
cvlistallcontained	Switch.1.1.1.card.5
cvlistallcontained	Switch.1.1.1.card.6
cvlistallcontained	Switch.1.1.1.card.7
cvlistallcontained	Switch.1.1.1.card.8
cvlistallcontained	Switch.1.1.1.1.SwitchGroupAddress.8889998889
cvlistallcontained	Switch.1.1.1.1.AddressPrefix.123456
cvlistallcontained	Switch.1.1.1.1.AddressPrefix.222333
cvlistallcontained	Switch.1.1.1.1.AddressPrefix.890890
cvlistallcontained	Switch.1.1.1.1.AddressPrefix.999000
cvlistallcontained	Switch.1.1.1.1.AlienIndividualAddress.8889998887
cvlistallcontained	Switch.1.1.1.1.AlienGroupAddress.0009998887



cvlistcontained

Purpose

Queries the database for a list of objects of a specified type that are children of a specified object. The children can be positioned anywhere in the containment hierarchy of the root object.

Command Syntax

cvlistcontained object-name object-type {-attribute-name } ...

Parameters

object-name specifies the parent object. The parent object can be the immediate parent of the contained objects (such as a PPort that is a parent of multiple LPorts). Or, the parent object can be positioned higher in the containment hierarchy (such as a switch that is a parent of multiple LPorts). The object is specified by its object ID, based on the containment hierarchy (for information, refer to "Managed Objects" on page 1-14).

object-type specifies the enumerated value that specifies the type of the objects to be retrieved.

-attribute-name specifies an attribute to be retrieved for the object. The attribute is specified by its argument name. Specify only attribute names with no values. You can specify up to 40 attributes.

Specify only those attributes appropriate for the object type.

If you want all attributes to be retrieved, omit the *-attribute-name* argument. The command returns all readable attributes for the child objects.



Notes

You can issue the function on either:

Network level — The function retrieves a list of objects on a network, including all subnets. To issue the function on a network level, specify an IP address, such as 128.100.0.0.

Subnet level — The function retrieves a list of objects on a particular subnet. To issue the function on a subnet level, specify an IP address with a subnet number, such as 128.100.111.0.

Table 3-2 lists the valid parent and child object types you can specify with this command.

Parent Object Type	Child Object Types
Network	Circuit Customer NetCac PMPCktLeaf PMPCktRoot PMPSpvcRoot SmdsAddressPrefix SmdsCountryCode SmdsLocalIndividualAddress SmdsNetwideGroupAddress SvcCUG SvcCUG SvcCUGMbrRule SvcSecScn Switch TrafficDesc ServiceName VPN
Customer	LPort Circuit
VPN	LPort Circuit
SvcCUGMbrRule	SvcCUGMbr
SvcCUG	SvcCUGMbr

Table 3-2. Valid Parent and Child Object Types

cvlistcontained



Table 3-2. Valid Parent and Child Object Types (Continued)

Parent Object Type	Child Object Types
Switch	Aps Card Channel Circuit LPort Performance Monitor PFdl PMPSpvcRoot PPort Spvc SmdsAddressPrefix SmdsAlienGroupAddress SmdsAlienIndividualAddress SmdsIndividualScreen SmdsLocalIndividualAddress SmdsLocalIndividualAddress SwcAddress SvcAddress SvcConfig SvcNodePrefix SvcPrefix SvcUserPart
Card	Aps Channel Circuit LPort Performance Monitor PFdl PMPSpvcRoot PPort Spvc SmdsAlienGroupAddress SmdsAlienIndividualAddress SmdsGroupScreen SmdsIndividualScreen SmdsLocalIndividualAddress SmdsSwitchGroupAddress SvcConfig SvcNodePrefix SvcUserPart

cvlistcontained



Table 3-2.Valid Parent and Child Object Types (Continued)

Parent Object Type	Child Object Types
PPort	Aps Channel Circuit LPort PerformanceMonitor PFdl (8-port ATM T1 card only)
Channel (valid for channelized DS3 card only)	LPort
LPort	AssignedSvcSecScn Circuit PMPCktRoot PMPSpvcRoot SmdsGroupScreen SmdsIndividualScreen SmdsLocalIndividualAddress Spvc SvcAddress SvcConfig SvcCrefix SvcUserPart SvcSecScnActParam
SmdsIndividualScreen	SmdsLocalIndividualAddress SmdsAlienIndividualAddress
SmdsGroupScreen	SmdsSwitchGroupAddress SmdsAlienGroupAddress
SmdsNetwideGroupAddress	SmdsSwitchGroupAddress SmdsLocalIndividualAddress
PMPCktRoot	PMPCktLeaf
PMPSpvcRoot	PMPSpvcLeaf

If **cvlistcontained** is successful, it prints the command name followed by the child objects that the Provisioning Server returns. It prints one line for each child object and includes attributes and values.



Example

The following **cvlistcontained** command lists all LPorts on a given switch by their Names:

cvlistcontained switch.1.1.1.2 lport -Name

If successful, the command returns the following text:

cvlistcontained Switch.1.1.1.2.card.4.pport.2.lport.1 -Name lport1

cvlistcontained Switch.1.1.1.2.card.4.pport.1.lport.1 -Name lport2

cvlistcontained Switch.1.1.1.2.card.4.pport.3.lport.1 -Name lport3



cvmodify

Purpose

Modifies specific attributes of an object in the database and (optionally) in the switch.

Command Syntax

cvmodify object-name {-attribute-name value } ...

Parameters

object-name specifies the object to be modified. The object is specified by its object ID, based on the containment hierarchy (for information, refer to "Managed Objects" on page 1-14).

-attribute-name value specifies an attribute and its value to be modified. The attribute is specified by its argument name. The value uses a data type appropriate for the argument. You can specify up to 40 attribute-value pairs.

Specify only those attributes and values appropriate for the object type. You can specify any attribute except those with either the Read-Only or Create-Only access restriction.

Notes

For a list of object types that you can use with this command, refer to Table 1-4 on page 1-38.

You can use this command to create a card or PPort. The CascadeView database automatically populates each switch with cards of type "empty". Use **cvmodify** to change the card's type from "empty" to the specified type. Likewise, once a card has been configured, CascadeView automatically populates the card with all necessary Physical Ports. Use **cvmodify** to change the PPort specifications. In the case of the channelized DS3 card, once the card has been configured, CascadeView automatically populates the card with all necessary populates the card with all necessary channels. Use **cvmodify** to change the channel specifications.

cvmodify



You can use this command to remove a card. Use **cvmodify** to change the card's type to "empty".

If **cvmodify** is successful, it prints the command name followed by the arguments that the Provisioning Server returns. You can use this output to verify that the arguments are the same as those specified in the original request.

Example

The following **cvmodify** command creates a card by specifying its type and administrative status:

```
/opt/ProvServ/bin/cvmodify Switch.1.1.1.1.card.4 -DefinedType
1PortAtmDs3Uni -AdminStatus Up
```

If successful, the command returns the following text:

```
/opt/ProvServ/bin/cvmodify Switch.1.1.1.1.card.4
-DefinedType 1PortAtmDs3Uni
-AdminStatus Up
```



CLI Examples

This section provides usage examples of each of the managed objects. Use these examples as guidelines for syntax and usage.

Sample CLI Format

Conventions used in the samples are as follows:

<ip_address> — Represents an IP address, such as 130.2.20.1.

<network_no> — Represents a network number.

<id>— Represents any numeric number representation, such as card number, PPort number, LPort number, channel number, DLCI number, VPI number, VCI number, PMPSpvcLeaf number, country code number, and so on. There is no relationship among the values for these numbers.

<name> — Represents a name string, such as the customer name string, Traffic Descriptor name string, VPN name string, and so on. Enclose the string in quotes if it contains special characters, such as a blank character.

<svc_string> — Represents an address string that conforms to the convention for SVC addresses.

<svccug_string> — Represents an SVC CUG string.

<rule_string> — Represents an SVC CUG member rule string.

{-Attribute value}* — Represents the applicable attribute-value pair. An asterisk (*) indicates that you can specify multiple attribute-value pairs.

CVT_APS

There is no identifier for APS.

cvlistcontained switch.<ip_address>.card.<id>.pport.<id> aps

cvget switch.<ip_address>.card.<id>.pport.<id>.aps {-Attribute value}*

cvmodify switch.<ip_address>.card.<id>.pport.<id>.aps {-Attribute value}*



CVT_AssignedSvcSecScn

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id> assignedsvcsecscn

CVT_Card

cvlistcontained switch.<ip_address> card

cvmodify switch.<ip_address>.card.<id> {-Attribute value}*

CVT_Channel

cvget switch.<ip_address>.card.<id>.pport.<id>.channel.<id> {-Attribute value}*

CVT_Circuit

Circuits are always identified by their endpoints. An endpoint can be an LPort or a ServiceName; the object ID representation differs accordingly.

In the case of LPorts, endpoints are represented differently according to different service types for the containing LPort. For Frame Relay, an endpoint is identified by DLCI number; for ATM, an endpoint is identified by the VPI, VCI pair. Specify the first endpoint as the main object identifier in the CLI command. Specify the second endpoint as an mandatory attribute to the first endpoint (using "-Endpoint2").

In the case of ServiceName, the endpoint is identified by the network number, the name of the ServiceName binding, and the VPI/VCI pair or DLCI number (depending on endpoint type). As with LPorts, the second endpoint is represented as an mandatory attribute to the first endpoint using "-Endpoint2."

ServiceName Endpoints

cvadd network.<network_no>.servicename.<name>.vpi.<id>.vci.<id> {-Attribute value}*

cvadd network.<network_no>.servicename.<name>.dlci.<id> {-Attribute value}*



LPort Endpoints

ATM - ATM circuit:

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.vpi.<id>.vci.<id> {-Attribute value}* switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.vpi.<id>.vci.<id>

ATM - Frame Relay circuit:

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.vpi.<id>.vci.<id> {-Attribute value}* switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.dci<id>

Frame Relay - Frame Relay circuit:

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.dlci.<id> {-Attribute value}* switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.dlci<id>

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id> circuit

CVT_Customer

cvlistcontained network.<ip_address> customer

cvget network.<ip_address>.customer.<name> {-Attribute value}*

CVT_LPort

Normal LPort type:

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id> {-Attribute value}*

cvmodify switch.<ip_address>.card.<id>.pport.<id>.lport.<id> {-Attribute value}*

cvget switch.<ip_address>.card.<id>.pport.<id>.lport.<id> {-Attribute value}*

ATM Virtual UNI LPort type:

The LPort number is generated automatically from the Start VPI number and the LPort interface number. Therefore, during creation, you do not need to provide an LPort number. To retrieve information for the LPort, you must specify its LPort number. To obtain this number, use **cvlistcontained**.

cvadd switch.<ip_address>.card.<id>.pport.<id>.startvpi.<id> {-Attribute value}*

cvmodify switch.<ip_address>.card.<id>.pport.<id>.startvpi.<id> {-Attribute value}*

cvget switch.<ip_address>.card.<id>.pport.<id>..startvpi.<id> {-Attribute value}*



cvlistcontained switch.<ip_address>.card.<id>.pport.<id> lport

ATM Network Interworking for Frame Relay NNI LPort type:

The LPort number is identified by VPI/VCI pair.

cvadd switch.<ip_address>.card.<id>.pport.<id>.vpi.<id>.vci.<id> {-Attribute value}*
cvmodify switch.<ip_address>.card.<id>.pport.<id>.vpi.<id>.vci.<id> {-Attribute value}*
cvget switch.<ip_address>.card.<id>.pport.<id>.vpi.<id>.vci.<id> {-Attribute value}*
cvget switch.<ip_address>.card.<id>.pport.<id>.vpi.<id>.vci.<id> {-Attribute value}*
cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.pport.<id>.vpi.<id>.vci.<id> {-Attribute value}*
cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.pport.<id>.vpi.<id>.vci.<id> {-Attribute value}*
cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.pport.<id>.pport</id>

CVT_NetCac

There is no identifier for NetCac. cvlistcontained network.<ip_address> netcac cvmodify network.<ip_address>.netcac {-Attribute value}*

CVT_PerformanceMonitor

There is no identifier for PerformanceMonitor. cvlistcontained switch.<ip_address>.card.<id>.pport.<id> pm cvget switch.<ip_address>.card.<id>.pport.<id>.pm {-Attribute value}* cvmodify switch.<ip_address>.card.<id>.pport.<id>.pm {-Attribute value}*

CVT_PFdl

There is no identifier for PFdl. cvlistcontained switch.<ip_address>.card.<id>.pport.<id> fdl

cvget switch.<ip_address>.card.<id>.pport.<id>.fdl {-Attribute value}*

cvmodify switch.<ip_address>.card.<id>.pport.<id>.fdl {-Attribute value}*

CLI Examples



CVT_PMPCkt

A PMP circuit leaf can be added only when a PMPCktRoot exists. The circuit type of a leaf must be the same as that of the root. For example, if the PMPCktRoot has been created without specifying the VCI value, all the leaves to be added to that particular root should not have their VCI value specified.

```
cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.pmpcktleaf.vpi.<id>.vci.<id>
{-Attribute value}*
```

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.pmpcktroot.vpi.<id>.vci.<id>pmpcktleaf

CVT_PMPCktRoot

For VCC circuit type:

cvadd switch.<ip_address>.card.<id>.pport.<id>.pport.<id>.pmpcktroot.vpi.<id>.vci.<id> {-Attribute value}*

For VPC circuit type:

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.pmpcktroot.vpi.<id> {-Attribute value}*

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id> pmpcktroot

CVT_PMPSpvcLeaf

A PMPSpvc circuit leaf can be added only when a PMPSpvcRoot exists. You must specify the correct instance number when you perform a **cvadd**, **cvget**, **cvmodify**, or **cvdelete**. To retrieve the correct instance number from the database, use the attribute CVA_PMPSpvcRootNextAvailableLeafNo.

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.pmpspvcleaf.<id> {-Attribute value}*

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.pmpspvcroot.vpi.<id>.vci.<id>pmpspvcleaf

CVT_PMPSpvcRoot

For VCC circuit type:

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.pmpspvcroot.vpi.<id>.vci.<id>
{-Attribute value}*



For VPC circuit type:

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.pmpspvcroot.vpi.<id> {-Attribute value}* cvlistcontained switch.<ip_address>.card.<id> pmpspvcroot

CVT_PPort

cvlistcontained switch.<ip_address>.card.<id> pport cvget switch.<ip_address>.card.<id>.pport.<id> {-Attribute value}*

CVT_ServiceName

cvadd network.<ip_address>.servicename.<name> {-Attribute value}*
cvmodify network.<ip_address>.servicename.<name> {-Attribute value}*
cvlistcontained network.<ip_address> servicename

CVT_SmdsAddressPrefix

cvlistcontained switch.<ip_address> addressprefix

CVT_SmdsAlienGroupAddress

cvlistcontained switch.<ip_address> aliengroupaddress

CVT_SmdsAlienIndividualAddress

cvlistcontained switch.<ip_address> alienindividualaddress

CVT_SmdsCountryCode

cvadd network.<ip_address>.countrycode.<id> {-Attribute value}

CVT_SmdsGroupScreen

cvlistcontained switch.<ip_address>.card.<id>.pport.<id> groupscreen



CVT_SmdsIndividualScreen

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id> individualscreen

CVT_SmdsLocalIndividualAddress

 $cvlist contained \ switch. <\!ip_address>.card. <\!id>.pport. <\!id>.lport. <\!id> local individual address>$

CVT_SmdsNetwideGroupAddress

cvlistcontained network.<ip_address> netwidegroupaddress

CVT_SmdsSwitchGroupAddress

cvlistcontained switch. <ip_address> switchgroupaddress

CVT_Spvc

For VCC circuit type:

 $cvadd\ switch.<\!\!ip_address>.card.<\!\!id>.pport.<\!\!id>.lport.<\!\!id>.spvc.vpi.<\!\!id>.vci.<\!\!id>\ \{-Attribute\ value\}^*$

For VPC circuit type:

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.spvc.vpi.<id> {-Attribute value}*

cvlistcontained switch.<ip_address>.card.<id> spvc

CVT_SvcAddress

An SvcAddress is represented by a string that conforms to the convention used to specify SVC addresses. The format of the **cvadd**, **cvmodify**, **cvget**, or **cvdelete** command depends on the format of the SVC address.

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>. svcaddress.<svc_string>
{-Attribute value}*

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id> svcaddress



CVT_SvcConfig

cvmodify switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.svcconfig {-Attribute value}*

CVT_SvcCUG

No attributes are needed for addition. cvadd network.<ip_address>.svccug.<svccug_string> cvlistcontained network. <ip_address> svccug

CVT_SvcCUGMbr

cvadd network.<ip_address>.svccug_string>.svccugmbr.<rule_string> {-Attribute value}* cvlistcontained network.<ip_address>.svccug_string> svccugmbr

CVT_SvcCUGMbrRule

cvadd network.<ip_address>.svccugmbrrule.<rule_string> {-Attribute value}*

cvlistcontained network.<ip_address>.svccugmbrrule

CVT_SvcNodePrefix

An SvcNodePrefix is represented by a string that conforms to the convention used to specify SVC addresses. The format of the **cvadd**, **cvmodify**, **cvget**, or **cvdelete** command depends on the format of the SVC address.

cvadd switch.<ip_address>.svcnodeprefix.<svc_string> {-Attribute value}*

cvlistcontained switch.<ip_address> svcnodeprefix

CLI Examples



CVT_SvcPrefix

An SvcPrefix is represented by a string that conforms to the convention used to specify SVC addresses. The format of the **cvadd**, **cvmodify**, **cvget**, or **cvdelete** command depends on the format of the SVC address.

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>.svcprefix.<svc_string> {-Attribute value}*

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id> svcprefix

CVT_SvcSecScn

cvlistcontained network.<ip_address> svcsecscn

CVT_SvcSecScnActParam

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id> svcsecscnactparam

CVT_SvcUserPart

An SvcUserPart is represented by a string that conforms to the convention used to specify SVC addresses. The format of the **cvadd**, **cvget**, or **cvdelete** command depends on the format of the SVC address.

cvadd switch.<ip_address>.card.<id>.pport.<id>.lport.<id>. svcuserpart.<svc_string>

cvlistcontained switch.<ip_address>.card.<id>.pport.<id>.lport.<id> svcuserpart

CVT_Switch

cvlistcontained network.<ip_address> switch

CVT_TrafficDesc

cvadd network.<ip_address>.trafficdesc.<name> {-Attribute value}*

cvlistcontained network.<ip_address> trafficdesc



CVT_TrafficShaper

cvlistcontained switch.<ip_address>.card.<id>.pport.<id> ts

 $cvget\ switch.{<}ip_address{>}.card.{<}id{>}.pport.{<}id{>}.ts.{<}id{>}\ \{-Attribute\ value\}*$

CVT_VPN

cvadd network.<ip_address>.vpn.<name>



Using the SNMP MIB

This chapter describes how to use the SNMP MIB to access the Provisioning Server. To understand the Provisioning Server object hierarchy, first read Chapter 1.

About the Enterprise-specific MIB

The enterprise-specific MIB interface provides SNMP access to the Provisioning Server. Use the Provisioning Server MIB to provision via SNMP instead of using a C or C++ program or the CLI.



The Provisioning Server MIB is different than the Ascend Enterprise MIB.

The Provisioning Server MIB supports SNMPv1 protocol. The following SNMP operations are supported:

- get
- get-next
- set (used for creating, modifying, and deleting)




The Provisioning Server does not generate or process SNMP traps.

The Provisioning Server MIB is defined according to Structure of Management Information version 2 (SMIv2). You can view the MIB with an SMIv2-compliant MIB browser.

To compile the MIB, use an SMIv2-compliant compiler.

The MIB is defined in the file provserv.mib, which is installed in the directory /opt/ProvServ/snmp_mibs.

If you install the Provisioning Server on a separate machine from CascadeView, and you want to use the HP OpenView MIB browser to view the MIB, perform the following steps:

- 1. File transfer provserv.mib from the Provisioning Server machine to the directory opt/CascadeView/snmp_mibs on the CascadeView machine.
- 2. Load the MIB file from the CascadeView machine.

For a listing of the variables in the Provisioning Server MIB, refer to the *NavisXtend Provisioning Server Enterprise MIB Definitions*.

Community Strings

The Provisioning Server implements an SNMP agent as a separate entity within the server to service MIB interface requests. The community name provides a mechanism for authentication and access-control at the agent. The Provisioning Server supports two community names, one for Read-Only operations and another for Read-Write operations.

The community strings are defined using the environment variables CV_READONLY_COMMUNITY_STRING (default value 'public') and CV_READWRITE_COMMUNITY_STRING (default value 'ascend'). If the environment variables are defined in the script start-server.sh, the specified strings take precedence. If they are not set in the script or if the server shell environment does not define the variables, the server assumes the default values.



If the community name is not valid when you issue an snmp_set request, the request is rejected and the SNMP agent returns a genError. You can access the Command Error Table in the MIB to see if the source of the problem is an invalid community name. Specify the Read-Only community name when you access the table, as that community name is used for validation purposes.

When you make an snmp_get request, specify either the Read-Only or the Read-Write community name. If you use a different community name and you encounter an error, the error is not propagated to the Command Error Table.

For details on setting environment variables to configure how the Provisioning Server handles MIB requests, refer to "Setting Environment Variables" on page 2-13.

MIB Structure

The Provisioning Server MIB defines objects that a client can configure or read. The MIB is organized into logical groups by object (node, card, LPort, PPort, and so on). Each group contains table entries that map to the attributes of the API.

The various groups of the MIB are placed under the Provisioning Server object identifier (OID):

1.3.6.1.4.1.277.9.1

where the last term in the OID represents the version number of the MIB.

Each group can have one or more tables and/or scalar objects. The tables are two-dimensional, with each column representing an attribute and each row representing an object instance on a switch. Because a column contains rows for all possible object instances, many of which may not actually use that attribute, a table can contain *holes*. Holes are non-applicable elements of the matrix. For example, in the PPort table, the column that contains the attribute pportChannelIsInUse is sparse because it contains values only for PPort instances present on channelized cards.

Row instances in a table are uniquely identified by Index information. The Index represents the information you need to provide when issuing a command on a particular object. For example, to configure an LPort, you need to specify the IP address of the switch that contains the LPort and the ifIndex.



Segmented Information in Multiple Tables

LPorts are always identified by specifying the IP address of the switch that contains the LPort and the LPort's ifIndex. Because LPorts are complex objects, additional information (such as LPortId, DLCI number, or VPI/VCI pair) is required to obtain an ifIndex.

The MIB uses Translation Tables to convert the information required for a specific LPort type into the ifIndex value. The Translation Table provides a unique key to access a specific row in the Configuration Table (the table that contains the configuration attributes of the LPort). Table 4-1 lists the information required to create each type of LPort and which specific Translation Table to use.

LPort Type	Information Required	Table to Use
ATM Direct Trunk	Switch	lportIdIndexTransTable
ATM UNI DCE/DTE	Card	lportIdChannelIndexTransTable (for LPorts on
Direct Line Trunk	PPort	the channelized DS3 card)
Encapsulation FRAD	LPort Id	
FR NNI		
FR UNI DCE/DTE		
PPP to1490 Encapsulation		
SMDS DXI/SNI DTE/DCE		
SMDS OPT Trunk		
SMDS SSI DTE		
FR OPT PVC Trunk	Switch	dlciIndexTransTable
	Card	dlciChannelIndexTransTable (for LPorts on
	PPort	the channelized DS3 card)
	DLCI	

Table 4-1. Information Required for Creating Specific LPorts

ATM OPT Frame Trunk



Information **LPort** Type Required Table to Use Virtual UNI DCE/DTE Switch vpiStartIndexTransTable Card PPort VPI start number ATM OPT Cell Trunk Switch vpiVciIndexTransTable Card vpiVciChannelIndexTransTable (for LPorts on the channelized DS3 card) PPort VPI (1-15)/ VCI 0 ATM Network Interworking for Switch vpiVciIndexTransTable FR NNI Card vpiVciChannelIndexTransTable (for LPorts on

 Table 4-1.
 Information Required for Creating Specific LPorts (Continued)

Circuits are segmented into several categories of tables, based on technology. You access various tables to configure circuit endpoints and configure the cross-connections between endpoints. Table 4-2 lists the information required to create each type of circuit endpoint and which specific endpoint table to use. The type of endpoint table to use depends on the type of services offered on the card on which the endpoint is created. Note that the ServiceName can be used with either or both endpoints.

the channelized DS3 card)

Once the endpoints are created, use the CircuitCrossConnectTable.

PPort VPI (0-15)/ VCI (32-255)



Circuit Type	Card Type	Information Required	Table to Use
FR-FR	All cards on B-STDX 9000 and STDX 6000 except: 1-port ATM IWU OC3 1-port ATM-CS/DS3	{switchIdIndex lportIfIndex DlciIdIndex}	frCircuitEndpointTable
FR-FR (with either endpoint using ServiceName)	All cards on B-STDX 9000 and STDX 6000 except: 1-port ATM IWU OC3 1-port ATM-CS/DS3	<pre>{switchIdIndex lportIfIndex DlciIdIndex } (for non-ServiceName based endpoint) {networkIdIndex networkServiceName Index dlciIdIndex } (for ServiceName based endpoint)</pre>	frCircuitEndpointTable (for non-ServiceName endpoint) frCircuitServiceNameEndpoint Table (for ServiceName based endpoint)
ATM-ATM	All cards on CBX 500 and 1-port ATM IWU OC3 and 1-port ATM CS/DS3 card on B-STDX 9000	{switchIdIndex lportIfIndex vpiIndex vciIndex}	atmCircuitEndpointTable

Table 4-2. Information Required for Creating Specific Circuits



Information **Circuit Type Card Type** Required Table to Use ATM-ATM All cards on CBX 500 {switchIdIndex atmCircuitEndpointTable (for (with either and 1-port ATM IWU non-ServiceName endpoint) lportIfIndex OC3 and 1-port ATM endpoint using atmCircuitServiceNameEndpoint vpiIndex ServiceName) CS/DS3 card on Table for ServiceName based B-STDX 9000 vciIndex} (for endpoint) non-ServiceName based endpoint) {networkIdIndex networkServiceName Index vpiIndex vciIndex} (for ServiceName based endpoint) ATM-ATM One endpoint on any For category A without For category A or card that is one of the category B without ServiceName use following (category A): ServiceNames use: atmCircuitEndpointTable. All cards on CBX 500 {switchIdIndex For category A with ServiceNames use 1-port ATM IWU OC3 lportIfIndex atmCircuitServiceNameEndpoint card and 1- port ATM vpiIndex Table. CS/DS3 card on the vciIndex} **B-STDX 9000** For category B without ServiceName use For category A or B The other endpoint on with ServiceNames use: interworkingCircuitEndpointTable. any card that is one of the following (category {networkIdIndex For category B with ServiceNames B): use networkServiceName interworkingCircuitServiceName All cards on STDX Index EndpointTable. 6000 and all cards on vpiIndex B-STDX 9000 except vciIndex} 1-port ATM IWU OC3 card and 1-port ATM CS/DS3 card

Table 4-2. Information Required for Creating Specific Circuits (Continued)



Table 4-2. Information Required for Creating Specific Circuits (Continued)

Circuit Type	Card Type	Information Required	Table to Use
ATM-ATM	Both endpoints exist on cards that belong to category B as explained above.	For category B without ServiceNames use: {switchIdIndex lportIfIndex vpiIndex vciIndex} For category B with ServiceNames use: {networkIdIndex networkServiceName Index vpiIndex vciIndex}	For category B without ServiceName use interworkingCircuitEndpointTable. For category B with ServiceNames use interworkingCircuitServiceName EndpointTable.
FR-ATM Interworking	Does not depend on the card type of the endpoints.	For the FR endpoint: {switchIdIndex lportIfIndex dlciIdIndex} For the ATM endpoint: {switchIdIndex lportIfIndex vpiIndex vciIndex}	For the ATM endpoint use interworkingCircuitEndpointTable For the FR endpoint use frCircuitEndpointTable



 Table 4-2.
 Information Required for Creating Specific Circuits (Continued)

Circuit Type	Card Type	Information Required	Table to Use
FR-ATM Interworking with either endpoint using ServiceName	Does not depend on the card type of the endpoints.	For the ServiceName based FR endpoint: {networkIdIndex networkServiceName Index, dlciIdIndex} For the ServiceName based ATMendpoint: {networkIdIndex networkServiceName Index vpiIndex vciIndex}	For the FR endpoint use frCircuitServiceNameEndpoint Table For the ATM endpoint use interworkingCircuitServiceName EndpointTable

By segmenting information into separate tables based on technology or specific features, the MIB improves performance of **get-next** operations because it minimizes holes in matrices.

Table 4-3 lists the main groups of the MIB and the indexing scheme for each group.



Table 4-3. Groups and Table Indexes of the Provisioning Server MIB

Group Name Group Value	Description	Sub-group or Table	Table Index
server psMibRev2.1	Contains operational information about the Provisioning Server. Also contains the Command Error Table, which supplies information about the errors encountered during snmp_set operations.	cmdErrorTable	{cmdErrorOrigIpAddr, cmdErrorRequestId, cmdErrorPortId}
network	Contains information	networkTable	{networkNetId}
psMibRev2.2	about the network or	networkCACTable	{networkIdIndex}
	sub-network.	networkCUGTable	{networkIdIndex, networkCUGNameIndex}
		networkCUGMemberRule Table	{networkIdIndex, networkCUGMemberRule NameIndex}
		networkCUGMemberTable	{networkIdIndex, networkCUGNameIndex, networkCUGMemberRule NameIndex}
		networkSvcSecurityScreen Table	{networkIdIndex, networkSvcSecurityScreen NameIndex}
		networkServiceNameTable	{networkIdIndex, networkServiceNameIndex}
		vpnTable	{networkIdIndex, vpnName}
		trafficDescriptorPoolTable	{networkIdIndex, trafficDescriptorName}
		customerTable	{networkIdIndex, customerName}



Table 4-3.Groups and Table Indexes of the Provisioning Server MIB
(Continued)

Group Name Group Value	Description	Sub-group or Table	Table Index
node psMibRev2.3	Contains information pertaining to the managed switch node.	switchTable switchAtmBillingTable	{switchIdIndex} {switchIdIndex}
card psMibRev2.4	Contains information pertaining to card objects on the managed switch node.	cardTable	{switchIdIndex, slotIdIndex}
pport psMibRev2.5	Contains information pertaining to PPort objects on the card.	pportTable pportTrafficShaperTable pportApsTable ds1pmThresholdTable ds3pmThresholdTable sonetpmThresholdTable pportPFdlTable	<pre>{switchIdIndex, slotIdIndex, pportIdIndex {switchIdIndex, slotIdIndex, pportIdIndex, pportTrafficShaperIndex} {switchIdIndex, slotIdIndex, pportIdIndex} {switchIdIndex, slotIdIndex, pportIdIndex} {switchIdIndex, slotIdIndex, pportIdIndex} {switchIdIndex, slotIdIndex, pportIdIndex} {switchIdIndex, slotIdIndex, pportIdIndex}</pre>
lport psMibRev2.6	Contains two categories of tables: Translation and Configuration.	lportTranslation lportConfiguration	



Table 4-3.Groups and Table Indexes of the Provisioning Server MIB
(Continued)

Group Name Group Value	Description	Sub-group or Table	Table Index	
lportTranslation lport.1	Contains tables that translate service- specific index	vpiVciIndexTransTable	{switchIdIndex, slotIdIndex, pportIdIndex, vpiIndex, vciIndex}	
	information of an LPort to a more generic index (like an <i>ifIndex</i>). Used mostly during LPort creation.	information of an LPort to a more	lportIdIndexTransTable	{switchIdIndex,slotIdIndex, pportIdIndex,lportIdIndex}
		vpiVciChannelIndexTrans Table	{switchIdIndex, slotIdIndex, pportIdIndex, ds1ChannelIdIndex, vpiIndex, vciIndex}	
		lportIdChannelIndexTrans Table	{switchIdIndex, slotIdIndex, pportIdIndex, ds1ChannelIdIndex, lportIdIndex}	
		dlciIndexTransTable	{switchIdIndex,slotIdIndex, pportIdIndex,dlciIdIndex}	
		dlciChannelIndexTrans Table	{switchIdIndex, slotIdIndex, pportIdIndex, ds1ChannelIdIndex, dlciIdIndex}	
		vpiStartIndexTransTable	{switchIdIndex, slotIdIndex, pportIdIndex, vpiStartIndex}	



Table 4-3.Groups and Table Indexes of the Provisioning Server MIB
(Continued)

Group Name Group Value	Description	Sub-group or Table	Table Index
lportConfiguration	Contains configuration	lportAdminTable	{switchIdIndex, lportIfIndex}
ipon.2	information pertaining to the	lportFrTable	{switchIdIndex, lportIfIndex}
	LPorts created on PPorts. Tables are	lportAtmTable	{switchIdIndex, lportIfIndex}
	based on service type (FR, ATM, SMDS,	lportAtmFcpTable	{switchIdIndex, lportIfIndex}
	etc.). Within a service category, further breakdown is based on specific functionality (For example: lportAtmFcpTable is specific to FlowControlProcessor related configuration for an ATM LPort).	lportAtmNtmTable	{switchIdIndex, lportIfIndex}
		lportAtmBillingTable	{switchIdIndex, lportIfIndex}
		lportSmdsTable	{switchIdIndex, lportIfIndex}
		svcAtmConfigTable	{switchIdIndex, lportIfIndex}
		lportSvcSecurityScreen ActionTable	{switchIdIndex, lportIfIndex}
		lportSvcSecurityScreen Table	{switchIdIndex, lportIfIndex, networkSvcSecurityScreen NameIndex}
circuit	Contains two categories of tables:	circuitTranslation	
psivilukev2.7	Translation and Configuration.	circunconinguration	



Table 4-3.Groups and Table Indexes of the Provisioning Server MIB
(Continued)

Group Name Group Value	Description	Sub-group or Table	Table Index
circuitConfiguration circuit.2	Contains tables used in the configuration of circuit endpoints and	interworkingCircuit EndpointTable	{switchIdIndex, lportIfIndex, vpiIndex, vciIndex}
	the circuit cross-connections between the two ordepoints	atmCircuitEndpointTable	{switchIdIndex, lportIfIndex, vpiIndex, vciIndex}
	enupoints.	frCircuitEndpointTable	{switchIdIndex, lportIfIndex, dlciIndex}
		circuitCrossConnectTable	circuitIndex
		circuitPmpRootTable	{switchIdIndex, lportIfIndex, vpiIndex, vciIndex}
		circuitPmpLeafTable	{switchIdIndex, lportIfIndex, vpiIndex, vciIndex, circuitLeafSwitchIdIndex, circuitLeafLportIfIndex, circuitLeafVpiIndex, circuitLeafVciIndex}
		atmCircuitBillingTable	circuitIndex
		atmCircuitNdcTable	circuitIndex
		circuitSpvcTable	{switchIdIndex, lportIfIndex, vpiIndex, vciIndex}
		circuitSpvcPmpRootTable	{switchIdIndex, lportIfIndex, vpiIndex, vciIndex}
	circuitSpvcPmpLeafTable interworkingCircuitService	{circuitSpvcPmpLeafIndex, switchIdIndex, lportIfIndex, vpiIndex, vciIndex}	
		NameEndpointTable	{networkIdIndex, networkServiceNameIndex, vpiIndex, vciIndex}



Table 4-3.Groups and Table Indexes of the Provisioning Server MIB
(Continued)

Group Name Group Value	Description	Sub-group or Table	Table Index
circuitConfiguration circuit.2		atmCircuitServiceName EndpointTable	{networkIdIndex, networkServiceNameIndex, vpiIndex, vciIndex}
(Continuea)		frCircuitServiceName EndpointTable	{networkIdIndex, networkServiceNameIndex, dlciIdIndex}
svc psMibRev2.9	Contains groups associated with SVC address configuration.	svcaddress svcmgt	
svcaddress	Contains tables associated with SVC	svcNodePrefixTable	{switchIdIndex, svcNodePrefixIndex}
	addressing at various levels node, port, etc.	svcPortPrefixTable	{switchIdIndex, lportIfIndex, svcPortPrefixIndex}
		svcAddrTable	{switchIdIndex, lportIfIndex, svcAddrAddressIndex}
		svcAtmDteUserPartTable	{switchIdIndex, lportIfIndex, svcAtmDteUserPartIndex}
svcmgmt svc.2	Contains tables related to SVC configuration (but not solely associated with SVC addresses).	svcConfigTable	{switchIdIndex, lportIfIndex}
channel psMibRev2.13	Contains information related to DS1 Channels associated with DS3 channelized ports.	ds1ChannelTable	{switchIdIndex, slotIdIndex pportIdIndex ds1ChannelIdIndex}



Row Aliasing

For objects in the MIB that have attributes dispersed in several tables, some attributes are common to multiple tables. In the tables of the LPort and circuit groups, the following attributes are common attributes:

- RowStatus (refer to "RowStatus Attribute" on page 4-17)
- ModifyType (refer to "ModifyType Attribute" on page 4-18)
- lportIfIndex (for tables in the LPort group only)
- CircuitNumber (for tables in the circuit group only)

The tables containing common attributes are considered linked. Thus, an operation on a common attribute in one linked table affects the common attribute in the other linked tables. For example, for a Frame Relay UNI DCE LPort, when the RowStatus attribute is modified in one table (such as the lportIdIndexTransTable), the value of that attribute is updated in other linked tables (such as the lportAdminTable and lportFrTable). For a Frame Relay to Frame Relay circuit, when the RowStatus attribute is modified in the frCircuitEndpointTable, the value of the attribute is updated in the linked circuitCrossConnectTable.

This feature, known as *row aliasing*, gives the user the flexibility to set an attribute in only one table rather than set it in all related tables. Using row aliasing, the Provisioning Server reflects the same value for a common attribute for the same row across linked tables. Row aliasing assures that the status of a row and its common attributes are always the same irrespective of the table.

The lportIfIndex attribute is an attribute that is not directly set by the user. It is generated when the user sets the RowStatus attribute to the createAndWait state in a translation table (such as the lportIdIndexTransTable). Once lportIfIndex is generated, the attribute is updated in the linked LPort tables lportIdIndexTransTable, lportAdminTable, and lportFrTable.

The user does not directly set the CircuitNumber attribute. It is generated when the user generates circuit endpoints. Once CircuitNumber is generated, the attribute is updated in the linked circuit tables frCircuitEndpointTable and circuitCrossConnectTable.



Column Access Specifiers

Access specifiers for a table column are specified as Read-Only, Read-Write, or Not-Accessible. Because SNMP does not support the category Create-Only, attributes with this restriction are defined as Read-Write. These attributes are usually mandatory attributes that you provide when creating an object. Refer to the *NavisXtend Provisioning Server Object Attribute Definitions* for attributes that are Create-Only.

For most tables, the index attributes are specified as Not-Accessible. Instead of accessing these index columns directly, you use the Translation Tables to convert required information into index attributes.

Additional Table Entries

Most table entries have the attributes RowStatus and ModifyType. These attributes are used in **set** operations.

RowStatus Attribute

The RowStatus attribute specifies the state of the table entry at a given time. Valid values are as follows:

active (1) — Entry is active, such as when it has been created and definitions have been made to it.

notInService (2) — Entry is not in service, such as when modifications are being made to it.

notReady (3) — Entry is under creation.

createAndWait (5) — Entry is being created, and is waiting for definitions to be made to it. When you set the RowStatus attribute to 5, it gets set to 3.

destroy (6) — Entry has been removed.

You must include the RowStatus attribute when you:

- Create an object
- Modify an object by specifying the attribute modifications in multiple PDUs
- Destroy an object

Using the MIB



ModifyType Attribute

The ModifyType attribute specifies the update method, as follows:

1 sends updates to both the network component and the database. The database is updated only if the network component updates successfully.

4 sends updates to the database only.

5 sends updates to the database only and sets a flag indicating that the object is out of synchronization in the database.

By default, updates are made to both the component and the database.

You must include the ModifyType attribute when you want an update to be made to the database only. The setting applies only to the current request. Subsequent requests revert to the default setting.

Using the MIB

This section describes how to use the MIB to list, create, and modify a given component on the network.

Using the SNMP Commands

The Provisioning Server supports the following SNMP commands:

get — Reads a single attribute of a row in a table.

get-next — Walks the MIB (similar to performing a ListContained command in the API or CLI). The command is based on a lexicographical ordering of the complete OID for various row instances. Thus, the command walks a table by reading all row values of the first column before starting the second column.

set — Creates a new object, or modifies or deletes an existing object.



Command Error Table

The Command Error Table supplies information about any errors you encounter during snmp_set operations to create or modify objects. This information is useful for troubleshooting problems.

Entries in the table contain the following information:

- IP address of the host machine where the MIB client's request originated.
- The request ID of the request sent to the server.
- UDP port number of the client.
- Error code encountered when the server executed the command.
- Error message string.
- The OID of the attribute (column) that is in error. If several attributes are in error, only the first one is reported.
- The timestamp at which the error occurred.

If several MIB clients use the same host, it can be difficult to distinguish the various entries in the table based on IP address. To determine uniqueness, use the request ID, UDP port number, or the timestamp of the entry.

The Provisioning Server purges entries in this table based on the value of environment variable CV_SNMP_CMDERROR_CACHE_TIMEOUT. The default setting of this variable is 6000 seconds. For more information on this environment variable, refer to "Controlling MIB Cache" on page 2-21.

MIB Cache and Database Locking

The Provisioning Server implements a MIB cache that stores data in memory for a fixed time period. The server uses the cache to optimize performance of **get-next** requests and to store data to be committed to the database during transactions involving multiple PDUs. The caching behavior varies depending on which operation is being performed. For details, refer to "Row Creation", "Row Modification", and "get-next Operations" later in this section.

Using the MIB



The object locking behavior for MIB objects in the database differs from the locking behavior of the Provisioning Server API, CLI, or CascadeView. For these interfaces, the steps associated with locking are transparent to the user. When an object is created or modified, its parent object gets locked. The user specifies all the information needed to create or modify the object in one request. Once the request is complete, the parent gets unlocked.

By contrast, in the case of the MIB, the information needed to create or modify an object may not be available in one PDU. As a result, the locks in the database must be held for a longer time. Thus, the steps associated with locking are not transparent to the user.



If the API, CLI, or CascadeView makes modifications to an object in the database at the same time that the object is present in MIB cache during a **get** or **get-next** request, the MIB values in cache become stale.

Row Creation

When an object is created, a new row is created in the database. During a successful row creation, you perform the following steps:

1. Initiate the transaction by setting the RowStatus attribute to the createAndWait state.

The parent object gets locked.

- 2. Issue one or more snmp_set requests to assign values to other attributes of the row. The attribute values are stored in MIB cache.
- 3. Complete the transaction by setting the RowStatus attribute to the active state.

When no errors are encountered, the changes are committed to the switch and to the database, the row is flushed from MIB cache, and the lock is released.

If an error is encountered, the row remains in cache and the lock remains in effect. You can correct the error by modifying the contents of the cache (by returning to step 2). Once you have corrected the error and set the RowStatus to the active state, the row creation is completed, the row is flushed from MIB cache, and the lock is released. Note that it can take several iterations before all the errors are corrected.



If a user initiates but does not complete a transaction to create an object, the partially-created row remains in MIB cache for the amount of time specified by the environment variable CV_SNMP_LOCK_TIMEOUT. And, the parent object remains locked for the time specified by the CV_SNMP_LOCK_TIMEOUT value, preventing other users from accessing the parent object. Thus, users should make sure to complete all transactions. Once the CV_SNMP_LOCK_TIMEOUT timer expires, the partially-created row is flushed out of cache and the lock is removed.

For more information on this environment variable, refer to "Controlling Object Locking" on page 2-21.

Row Modification

When an object is modified, a row is modified in the database. Before modifying an object, perform an snmp_get request on the RowStatus attribute to check if another user is currently accessing the entry. If the entry is in use, retry your request later.

Row modification can be performed with or without modifying the RowStatus attribute.

PDU Modification without Modifying RowStatus

Simple modifications do not require you to set the RowStatus to the notInService state; the RowStatus remains Active through the transaction. When you do not have to modify the RowStatus attribute, the locking and unlocking of the object becomes transparent.

If you want to modify only a few attributes, you can issue one PDU containing the appropriate values for the varbinds. Or, you can issue a PDU multiple times.



To maximize MIB efficiency, you should specify all varbinds in one PDU whenever possible.



With complex modifications involving a number of attributes, you can issue multiple PDUs containing the appropriate values for the varbinds. However, to maximize MIB efficiency, you should specify all varbinds in one PDU whenever possible.

Because of attribute dependencies, you should first set the RowStatus to the notInService state before making the modifications.

During a complex modification, you perform the following steps:

- 1. Issue an snmp_get request on the RowStatus attribute to make sure that no other user is currently accessing the object.
- 2. Initiate the transaction by setting the RowStatus attribute to the notInService state. The object gets locked.
- 3. Issue one or more snmp_set requests to assign values to other attributes of this row.

The attribute values are stored in MIB cache.

4. Complete the transaction by setting the RowStatus attribute to the active state.

When no errors are encountered, the changes are committed to the switch and to the database, and the lock is released.

If an error is encountered during modification, the lock remains in effect. You can correct the error by modifying the contents of the cache (by returning to step 3). Once you have corrected the error and set the RowStatus to the active state, the row creation is completed and the lock is released.

If a user initiates (but does not complete) a transaction to modify an object, the partially-modified row remains in MIB cache for the amount of time specified by the environment variable CV_SNMP_LOCK_TIMEOUT. And, the object remains locked for the time specified by the CV_SNMP_LOCK_TIMEOUT value, preventing other users from accessing the object. Thus, users should make sure to complete all transactions. Once the CV_SNMP_LOCK_TIMEOUT timer expires, the partially-created row is flushed out of cache and the lock is removed.



get-next Operations

You can perform a **get-next** request starting at any location in the MIB (including the top of the MIB), at any group of the MIB, any column of a table, or a specific column of an instance.

get-next requests are performance-intensive operations. The Provisioning Server uses MIB cache to cache objects (rows), thus optimizing performance of **get-next** requests. When the objects are initially loaded into cache from the database, the response to a **get-next** request may be slow. However, once the caching is complete, the response becomes significantly faster.

Be aware that using **get-next** operations on tables with many entries in the database may take some time to retrieve. These operations can significantly affect performance of the server. Although a **get-next** operation will not block other requests, it can slow the response to the other requests.

The Provisioning Server purges entries in MIB cache resulting from a **get-next** operation based on the value of environment variable CV_SNMP_ROWENTRY_TIMEOUT. The default setting of this variable is 900 seconds. For more information on this environment variable, refer to "Controlling MIB Cache" on page 2-21.

Specifying the Object Identifier

When you want to access a specific variable from a MIB group, you enter an OID that uses the following format:

{Provisioning Server OID}.{Group}.{Sub-group}.{Table}.{Entry}.{Column}.{Index}

Complex objects, such as LPorts and circuits, require a sub-group; simple objects do not.

The Provisioning Server OID is:

1.3.6.1.4.1.277.9.1

where the last term in the OID represents the version number of the MIB.



Example 1: get Command

To find out what type of card is located in a particular slot of a switch, use the following steps to determine the OID of the command you want to issue:

1. Determine the group value by locating the Card Group in the beginning of the MIB document. The following line indicates that the group value is **4**:

card OBJECT IDENTIFIER ::= { psMibRev2 4 }

Cards are simple objects that do not require a sub-group name.

2. Determine the Table value by locating the Table index, cardTable. The line ::= { card 1 } indicates that the Table value is 1:

```
cardTable OBJECT-TYPE
SYNTAX SEQUENCE OF CardEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
     "Table representing information about all cards in the network"
::= { card 1}
```

Determine the Entry value by locating the Entry index, cardEntry. The line
 ::= { cardTable 1 } indicates that the Entry value is 1:

```
cardEntry OBJECT-TYPE
SYNTAX CardEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
     "Entry representing information about one card"
INDEX { switchIdIndex, slotIdIndex }
::= { cardTable 1 }
```

- 4. Determine the Column value for the MIB variable you want to access. To retrieve a card's type, you need to access the variable cardDefinedType. The line ::= { cardEntry 1 } indicates that the Column value is 1.
- 5. Determine the Index items by locating them in the cardEntry variable you located in step 3. The line **INDEX** { switchIdIndex, slotIdIndex } indicates the index items you need to provide to complete this command.

Using the MIB



The **switchIdIndex** represents the IP address of the switch. The **slotIdIndex** represents the slot where the card is located. If the switch that contains the card has IP address 152.148.10.19 and the card for which you are requesting information is in slot 8, then the index is 152.148.10.19.8.

6. Enter the following command to retrieve the card type for the card (this example uses MIT SNMP Tools command syntax):

```
snmpget -h <server-machine-name> -p<server-port> -c<community-name>
1.3.6.1.4.1.277.9.1.4.1.1.1152.148.10.19.8
```

```
where {Provisioning Server OID = 1.3.6.1.4.1.277.9.1}.{Group = 4}.
{Table = 1}.{Entry = 1}.{Column = 1}.{Index = 152.148.10.19.8}
```

The system responds by displaying the command as the full MIB tree index, 1.3.6.1.4.1.277.9.1.4.1.1.1.152.148.10.19.8, and retrieves an integer that represents the type of the card. Refer to the cardDefinedType variable to interpret this integer.

Example 2: get-next Command

To retrieve the Admin status for all LPorts on a switch, use the following steps to determine the OID of the command you want to issue:

1. Determine the group value by locating the LPort Group in the beginning of the MIB document. The following line indicates that the group value is **6**:

```
lport OBJECT IDENTIFIER ::= { psMibRev2 6 }
```

2. Admin status is a configuration attribute. To determine the Sub-group value, locate the LPortConfiguration table in the beginning of the MIB document. The following line indicates that the Sub-group value is **2**.

```
lportConfiguration OBJECT IDENTIFIER ::= { lport 2 }
```

3. Determine the Table value by locating the Table index, **lportAdminTable**. The line **::=** { **lportConfiguration 1** } indicates that the Table value is 1:

lportAdminTable OBJECT-TYPE

```
SYNTAX SEQUENCE OF LportAdminEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION "List of logical port common attribute entries."
::= { lportConfiguration 1 }
```



4. Determine the Entry value by locating the Entry index, **lportAdminEntry**. The line **::=** { **lportAdminTable 1**} indicates that the Entry value is 1:

```
lportAdminEntry OBJECT-TYPE
```

```
SYNTAX LportAdminEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
     "Logical Port Configuration Entry"
INDEX { switchIdIndex, lportIfIndex }
::= { lportAdminTable 1 }
```

 Determine the Column value for the MIB variable you want to access. To retrieve the Admin status, you need to access the variable lportAdminAdminStatus. The line ::= { lportAdminEntry 19 } indicates that the Column value is 19:

```
lportAdminAdminStatus OBJECT-TYPE
SYNTAX INTEGER {
    up(1),
    down(2),
    testing(3)
  }
MAX-ACCESS read-write
STATUS current
DESCRIPTION "LPort Administrative Status. This attribute is mandatory
for lport creation"
  ::= { lportAdminEntry 19}
```

6. Enter the following command to request Admin status of all LPort instances in the table (this example uses MIT SNMP Tools command syntax):

```
snmpget -h <server-machine-name> -p<server-port> -c<community-name>
1.3.6.1.4.1.277.9.1.6.2.1.1.19
```

```
where {Provisioning Server OID = 1.3.6.1.4.1.277.9.1}.{Group = 6}.
{Sub-group = 2}.{Table = 1}.{Entry = 1}.{Column = 19}
```

Index items are omitted, because the request is for Admin status of *all* LPort instances in the table.

For each LPort instance in the network, the system responds by displaying the command as the full MIB tree index, 1.3.6.1.4.1.277.9.1.6.2.1.1.19, and retrieves an integer that represents the Admin status of the LPort. If the value is 1, the Admin Status of an LPort is up; if the value is 2, the Admin Status is down.



The following examples illustrate how to use the Provisioning Server MIB to create, modify, and delete objects. Several examples involve ATM objects. You would use a similar approach to manage Frame Relay objects, except that you access different tables in the MIB. For example, to manage a Frame Relay LPort, you use the appropriate LPort Translation Table, the lportAdminTable, and the lportFrTable (refer to Table 4-3 on page 4-10).

Example 3: set Command to Create an ATM LPort

To create an ATM LPort, you use the lportIdIndexTransTable to map between the card, PPort, and LPort ID and the LPort interface number. You must specify the LPort ID and request an interface number for it.

To create an ATM LPort, use the following steps:

1. Issue an snmp_set request to obtain an LPort interface number based on the LPort ID. Set the lportIdIndexTransRowStatus to the createAndWait state, specifying the switchIdIndex 1.1.1.1, slotIdIndex 7, pportIdIndex 8, and lportIdIndex 1.

The SNMP agent processes the request and returns a successful snmp_setResponse.

2. Issue an snmp_get request to obtain the interface number (lportIfIndex) that will be used to create a new entry in the lportAdminTable and the lportAtmTable. Issue the request on the lportIdIndexTransIfIndex, specifying the switchIdIndex, slotIdIndex, pportIdIndex, and lportIdIndex values.

The SNMP agent processes the request and returns an snmp_getResponse with the lportIfIndex 7.

3. Issue a series of snmp_set requests that assign values to the attributes of the LPort in both the lportAdminTable and the lportAtmTable.

The SNMP agent processes the requests by storing the values in MIB cache. Then, the agent returns a successful snmp_setResponse.

4. Issue an snmp_set request to commit the new entry. Set the lportAdminRowStatus to the active state, specifying the switchIdIndex 1.1.1.1 and the lportIfIndex 7. This command automatically sets the lportIdIndexTransRowStatus to the active state.

The SNMP agent processes the request by committing the new entry to the switch and to the CascadeView database.



On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse to the MIB client.

Figure 4-1 shows the request-response message flow between the MIB client, the SNMP agent, and the database when adding the LPort.



Figure 4-1. Creating an ATM LPort



Example 4: set command to Modify an ATM LPort

You can modify an LPort using either of the following methods:

- Specifying the interface number of the LPort
- Specifying the LPort's VPI/VCI pair

Before modiying any attribute, perform an snmp_get request on the RowStatus attribute to check if another user is currently accessing the entry. If the entry is in use, retry your request later.

Before modifying the LPort attributes, set the lportIdIndexTransRowStatus to the notInService state. You can skip this step if you specify the attribute modifications in a single PDU.

To modify attributes of an ATM LPort for which you do not know the interface number, use the following steps:

1. Issue an snmp_set request to set the LPort to the notInService state, based on the LPort ID. Set the lportIdIndexTransRowStatus to the notInService state, specifying the switchIdIndex 1.1.1.1, slotIdIndex 7, pportIdIndex 8, and lportIdIndex 1.

The SNMP agent processes the request and returns a successful snmp_setResponse.

2. Issue an snmp_get request to obtain the interface number (lportIfIndex) that will be used to modify the entry in the lportAdminTable and the lportAtmTable. Issue the request on the lportIdIndexTransIfIndex, specifying the switchIdIndex, slotIdIndex, pportIdIndex, and lportIdIndex values.

The SNMP agent processes the request and returns an snmp_getResponse with the lportIfIndex 7.

3. Issue a series of snmp_set requests that modify values of the attributes of the LPort in both the lportAdminTable and the lportAtmTable.

The SNMP agent processes the requests by storing the values in MIB cache. Then, the agent returns a successful snmp_setResponse.

Using the MIB



4. Issue an snmp_set request to commit the modified entry. Set the lportAdminRowStatus to the active state, specifying the switchIdIndex 1.1.1.1 and the lportIfIndex 7. This command automatically sets the lportIdIndexTransRowStatus to the active state.

The SNMP agent processes the request by committing the modified entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse.

Figure 4-2 shows the request-response message flow between the MIB client, the SNMP agent, and the database when modifying attributes of the LPort.





Figure 4-2. Modifying an ATM LPort



Example 5: set Command to Delete an ATM LPort

You can delete an LPort using either of the following methods:

- Specifying the interface number of the LPort
- Specifying the LPort ID

Before deleting an object, perform an snmp_get request on the RowStatus attribute to check if another user is currently accessing the object. If the object is in use, retry your request later.

To delete an ATM LPort for which you do not know the interface number, use the following step:

1. Issue an snmp_set request to delete an LPort based on the LPort ID. Set the lportIdIndexTransRowStatus to the destroy state, specifying the switchIdIndex 1.1.1.1, slotIdIndex 7, pportIdIndex 8, and lportIdIndex 1. (As an alternative, you could set the lportAdminRowStatus to the destroy state, as these attributes are linked by aliasing.)

The SNMP agent processes the request by committing the modified entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse.

Figure 4-3 shows the request-response message flow between the MIB client, the SNMP agent, and the database when deleting the LPort.

MIB Client



SNMP Agent



Figure 4-3. Deleting an ATM LPort Using its VPI/VCI Pair

To delete an ATM LPort for which you know the interface number, use the following step:

1. Issue an snmp_set request to delete an LPort based on the LPort's interface number. Set the lportAdminRowStatus to the destroy state, specifying the switchIdIndex 1.1.1.1 and the lportIfIndex 7.

The SNMP agent processes the request by committing the modified entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse.

Figure 4-4 shows the request-response message flow between the MIB client, the SNMP agent, and the database when deleting the LPort.





Figure 4-4. Deleting an ATM LPort Using its Interface Number

Example 6: set Command to Create an ATM Circuit

To create an ATM circuit, you define the two circuit endpoints using the atmCircuitEndpointTable and establish their interconnection using the circuitCrossConnectTable (refer to Table 4-2).

To create an ATM circuit, use the following steps:

 Issue an snmp_set request to define the two circuit endpoints and establish their interconnection. Set the atmCircuitEndpointRowStatus to the createAndWait state, specifying both endpoint 1 (switchIdIndex 1.1.1.1, lportIfIndex 10, vpiIdIndex 8, and vciIdIndex 34) and endpoint 2 (switchIdIndex 2.2.2.2, lportIfIdIndex 4, vpiIdIndex 4, and vciIdIndex 54) in a single PDU.

The SNMP agent processes the request and returns a successful snmp_setResponse.

Using the MIB



2. Issue a series of snmp_set requests that assign values to the attributes of the circuit endpoints in the atmCircuitEndpointTable.

The SNMP agent processes the requests by storing the values in MIB cache. Then, the agent returns a successful snmp_setResponse.

3. Issue an snmp_get request to obtain the circuit number that will be used to create a new entry in the circuitCrossConnectTable. Specify the switchIdIndex, lportIfIndex, vpiIdIndex, and vciIdIndex values for one of the endpoints (the circuit number is the same for both endpoints).

The SNMP agent processes the request and returns an snmp_getResponse with the atmCircuitEndpointCircuitNumber 10.

4. Issue a series of snmp_set requests that assign values to the attributes of the circuit interconnection in the circuitCrossConnectTable.

The SNMP agent processes the requests by storing the values in MIB cache. Then, the agent returns a successful snmp_setResponse.

5. Issue an snmp_set request to commit the new entry. Set the circuitCrossConnectRowStatus to the active state, specifying the atmCircuitEndpointCircuitNumber 10. This command automatically sets the atmCircuitEndpointRowStatus of the two endpoints to the active state.

The SNMP agent processes the request by committing the new entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse.

Figure 4-5 shows the request-response message flow between the MIB client, the SNMP agent, and the database when adding the ATM circuit.





Figure 4-5. Creating an ATM Circuit



Example 7: set Command to Modify an ATM Circuit

You can modify a circuit using either of the following methods:

- Specifying the circuit number
- Specifying the circuit's endpoints

Before performing a modification on any attribute, perform an snmp_get request on the RowStatus attribute to check if another user is currently accessing the entry. If the entry is in use, retry your request later.

Before modifying the circuit attributes, set the circuitCrossConnectRowStatus to the notInService state. You can skip this step if you specify the attribute modifications in a single PDU.

To modify attributes of a circuit, use the following steps:

1. Issue an snmp_get request to obtain the circuit number in the appropriate Circuit Endpoint Table. Specify the switch IP address, lportIfIndex, vpiIdIndex, and vciIdIndex values for one of the endpoints (the circuit number is the same for both endpoints).

If you know the circuit number, skip to step 2.

2. Issue an snmp_set request to set the circuit to the notInService state, based on the circuit number. Set the circuitCrossConnectRowStatus to the notInService state, specifying the circuit number 10.

The SNMP agent processes the request and returns a successful snmp_setResponse.

3. Issue a series of snmp_set requests that modify values of the attributes of the circuit. Modifications are made to the circuitCrossConnectTable and the atmCircuitEndpointTable.

The SNMP agent processes the requests by storing the values in MIB cache. Then, the agent returns a successful snmp_setResponse.

4. Issue an snmp_set request to commit the modified entry. Set the circuitCrossConnectRowStatus to the active state, specifying the circuit number 10. This command automatically sets the atmCircuitEndpointRowStatus to the active state.
Using the MIB



The SNMP agent processes the request by committing the modified entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse.

Figure 4-6 shows the request-response message flow between the MIB client, the SNMP agent, and the database when modifying attributes of the circuit.

Using the MIB





Figure 4-6. Modifying an ATM Circuit Using its Circuit Number



Example 8: set Command to Delete an ATM Circuit

You can delete a circuit using either of the following methods:

- Specifying the circuit number
- Specifying the circuit's endpoints

Before deleting an object, perform an snmp_get request on the RowStatus attribute to check if another user is currently accessing the object. If the object is in use, retry your request later.

To delete an ATM circuit for which you know the circuit number, use the following step:

1. Issue an snmp_set request to delete a circuit based on the circuit number. Set the circuitCrossConnectRowStatus to the destroy state, specifying the circuit number 10.

The SNMP agent processes the request by committing the modified entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse.

Figure 4-7 shows the request-response message flow between the MIB client, the SNMP agent, and the database when deleting the circuit.





Figure 4-7. Deleting an ATM Circuit Using its Circuit Number

Example 9: set Command to Create a VPN Indexed by Name

To specify a string value when you create an object, you specify the length of the string and the ASCII representation of each of the characters in the string.

To create a VPN indexed by name, use the following steps:

1. Issue an snmp_set request to set the VPN name "abc". Set the vpnRowStatus to the createAndWait state, specifying the networkIdIndex 100.100.0.0, the length of the name (3 characters), and the ASCII values of each letter in the name (97, 98, and 99, respectively).

The SNMP agent processes the request and returns a successful snmp_setResponse.

2. Issue a series of snmp_set requests that assign values to the attributes of the VPN in the vpnTable.

The SNMP agent processes the requests by storing the values in MIB cache. Then, the agent returns a successful snmp_setResponse.

Using the MIB



3. Issue an snmp_set request to commit the new entry. Set the vpnRowStatus to the active state, specifying the networkIdIndex 100.100.0.0, the length of the name, and the ASCII values of each letter in the name.

The SNMP agent processes the request by committing the new entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse to the MIB client.

Figure 4-8 shows the request-response message flow between the MIB client, the SNMP agent, and the database when adding the VPN.





Figure 4-8. Creating a VPN Indexed by Name



Example 10: set Command to Create a ServiceName Indexed by Name

To specify a string value when you create an object, you specify the length of the string and the ASCII representation of each of the characters in the string.

When you create a ServiceName, the first PDU should contain the networkServiceNameRowStatus as the first varbind and the networkServiceNamePrimaryLPort as the second varbind.

To create a ServiceName indexed by the name "abc", use the following steps:

1. Issue an snmp_set request to define the primary ServiceName binding. In a single PDU, set the networkServiceNameRowStatus to the createAndWait state and set the networkServiceNamePrimaryLPort to the objectId (lportAdminIfIndex) of the LPort. Specify the networkIdIndex 100.100.0.0, the length of the name (3 characters), and the ASCII values of each letter in the name (97, 98, and 99, respectively).

The SNMP agent processes the request and returns a successful snmp_setResponse.

2. Issue a series of snmp_set requests that assign values to the attributes of the ServiceName in the networkServiceNameTable.



Do not set the networkServiceNameBackupLPort attribute in an add request. Otherwise, an error will be reported when the new entry is committed to the database.

The SNMP agent processes the requests by storing the values in MIB cache. Then, the agent returns a successful snmp_setResponse.

3. Issue an snmp_set request to commit the new entry. Set the networkServiceNameRowStatus to the active state, specifying the networkIdIndex 100.100.0.0, the length of the name, and the ASCII values of each letter in the name.

The SNMP agent processes the request by committing the new entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse to the MIB client.



Figure 4-9 shows the request-response message flow between the MIB client, the SNMP agent, and the database when adding the ServiceName binding.



Figure 4-9. Creating a ServiceName Indexed by Name

Using the MIB



Example 11: set command to Modify a ServiceName Indexed by Name

Before performing a modification on any attribute, perform an snmp_get request on the RowStatus attribute to check if another user is currently accessing the entry. If the entry is in use, retry your request later.

Before modifying the ServiceName attributes, set the networkServiceNameRowStatus to the notInService state. You can skip this step if you specify the attribute modifications in a single PDU.

To modify attributes of a ServiceName indexed by the name "abc", use the following steps:

1. Issue an snmp_set request to set the ServiceName "abc" to the notInService state. Set the networkServiceNameRowStatus to the notInService state, specifying the networkIdIndex 100.100.0.0, the length of the name (3 characters), and the ASCII values of each letter in the name (97, 98, and 99, respectively).

The SNMP agent processes the request and returns a successful snmp_setResponse.

2. Issue a series of snmp_set requests that assign values to the attributes of the ServiceName in the networkServiceNameTable.

The SNMP agent processes the requests by storing the values in MIB cache. Then, the agent returns a successful snmp_setResponse.

3. Issue an snmp_set request to commit the new entry. Set the networkServiceNameRowStatus to the active state, specifying the networkIdIndex 100.100.0.0, the length of the name, and the ASCII values of each letter in the name.

The SNMP agent processes the request by committing the new entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse to the MIB client.

Figure 4-10 shows the request-response message flow between the MIB client, the SNMP agent, and the database when modifying the ServiceName binding.



MIB Client

SNMP Agent





Using the MIB



Example 12: set command to Delete a ServiceName Indexed by Name

Before deleting an object, perform an snmp_get request on the RowStatus attribute to check if another user is currently accessing the object. If the object is in use, retry your request later.

To delete a ServiceName indexed by the name "abc", use the following steps:

1. Issue an snmp_set request to set the ServiceName "abc" to the destroy state. Set the networkServiceNameRowStatus to the destroy state, specifying the networkIdIndex 100.100.0.0, the length of the name (3 characters), and the ASCII values of each letter in the name (97, 98, and 99, respectively).

The SNMP agent processes the request by committing the modified entry to the switch and to the CascadeView database.

On receipt of noError messages from the switch and database, the SNMP agent returns a successful snmp_setResponse.

Figure 4-11 shows the request-response message flow between the MIB client, the SNMP agent, and the database when deleting the ServiceName.

MIB Client

SNMP Agent



Figure 4-11. Deleting a ServiceName Indexed by Name



Containment Hierarchy

This appendix lists the containment hierarchy (the parent-child relation) used to build object IDs to name objects in the network.

Containment Tables

Each child object is defined as follows:

Child Type — The object contained by another object positioned higher in the containment hierarchy.

Cardinality — The number of occurrences of the child object within the switch.



Network

Table A-1. Children of the Network Object

Child Type	Cardinality
ATM Traffic Descriptors	MANY
Connection Admission Control (CAC)	ONE
Service Name	MANY
Customers	MANY
SMDS Country Code	MANY
SMDS Netwide Group Address	MANY
SVC CUG Member Rule	MANY
SVC CUGs	MANY
SVC Security Screen	MANY
Switch	MANY
Virtual Private Networks	MANY

Switch

STDX 3000/6000 Switch

Table A-2. Children of the STDX 3000/6000 Switch

Child Type	Cardinality
6-port V.35 Card	MANY
1-port 24-channel T1 Card	MANY
1-port 30-channel E1 Card	MANY



Table A-2. Children of the STDX 3000/6000 Switch (Continued)

Child Type	Cardinality
6-port Universal I/O Card	MANY
8-port Low Speed UIO Card	MANY
18-port Low Speed UIO Card	MANY

B-STDX 8000/9000 Switch

Table A-3. Children of the B-STDX 8000/9000 Switch

Child Type	Cardinality
СР	ONE
8-port Universal I/O Card	MANY
4-port 24-channel T1 Card	MANY
4-port 24-channel PRI T1 Card	MANY
4-port 30-channel E1 Card	MANY
2-port HSSI Card	MANY
10-port DSX-1 Card	MANY
1-port ATM UNI DS3 Card	MANY
1-port ATM IWU OC3 Card	MANY
1-port ATM CS/E3 Card	MANY
1-port ATM CS/DS3 Card	MANY
4-port Unchannelized T1 Card	MANY
12-port E1 Card	MANY
4-port Unchannelized E1 Card	MANY



Table A-3. Children of the B-STDX 8000/9000 Switch (Continued)

Child Type	Cardinality
4-port 24-channel DSX Card	MANY
1-port ATM UNI E3 Card	MANY
4-port 32-channel PRI E1 Card	MANY
1-port 28-channel DS3 Card	MANY
SMDS group address	MANY
SMDS address prefix	MANY
SMDS alien individual address	MANY
SMDS alien group address	MANY
SvcNodePrefix	MANY

CBX 500 Switch

Table A-4.Children of the CBX 500 Switch

Child Type	Cardinality
SP20	ONE
SP10	ONE
8-port DS3 Card	MANY
8-port E3 Card	MANY
4-port OC-3c/STM-1 Card	MANY
8-port T1 Card	MANY
8-port E1 Card	MANY
1-port OC-12c/STM-4 Card	MANY



Table A-4. Children of the CBX 500 Switch (Continued)

Child Type	Cardinality
SvcNodePrefix	MANY

Card/PPort

6-port V.35 Card/PPort

Table A-5. Children of the 6-port V.35 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY
Direct Line Trunk LPort	MANY
Encapsulation FRAD LPort	MANY
PPP-to-1490 Translation LPort	MANY

1-port 24-channel T1 Card/PPort

Table A-6. Children of the 1-port 24-channel T1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY



Table A-6. Children of the 1-port 24-channel T1 Card/PPort (Continued)

Child Type	Cardinality
Direct Line Trunk LPort	MANY
Encapsulation FRAD LPort	MANY
PPP-to-1490 Translation LPort	MANY

1-port 30-channel E1 Card/PPort

Table A-7. Children of the 1-port 30-channel E1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
Direct Line Trunk LPort	MANY
Encapsulation FRAD LPort	MANY
PPP-to-1490 Translation LPort	MANY

6-port Universal I/O Card/Pport

Table A-8. Children of the 6-port Universal I/O Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY



Table A-8. Children of the 6-port Universal I/O Card/PPort (Continued)

Child Type	Cardinality
Direct Line Trunk LPort	MANY
Encapsulation FRAD LPort	MANY
PPP-to-1490 Translation LPort	MANY

8-port Low Speed UIO Card/Pport

Table A-9. Children of the 8-port Low Speed Universal I/O Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY
Direct Line Trunk LPort	MANY
Encapsulation FRAD LPort	MANY
PPP-to-1490 Translation LPort	MANY

18-port Low Speed UIO Card/Pport

Table A-10. Children of the 18-port Low Speed Universal I/O Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY



Table A-10. Children of the 18-port Low Speed Universal I/O Card/PPort

Child Type	Cardinality
FR OPTimum PVC Trunk LPort	MANY
Direct Line Trunk LPort	MANY
Encapsulation FRAD LPort	MANY
PPP-to-1490 Translation LPort	MANY

8-port Uio Card/Pport

Table A-11. Children of the 8-port Universal I/O Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
SMDS OPTimum Trunk LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY
Others Direct Line Trunk LPort	MANY



Table A-11. Children of the 8-port Universal I/O Card/PPort (Continued)

Child Type	Cardinality
Otbers Encapsulation FRAD Lport	MANY
Others PPP-to-1490 Translation LPort	MANY

4-port 24-channel T1 Card/PPort

Table A-12. Children of the 4-port 24-channel T1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others PPP-to-1490 Translation LPort	MANY



4-port 24-channel PRI T1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI LPort	MANY
FR NNI LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others ISDN PRI D - Channel LPort	MANY
Others PPP-to-1490 Translation LPort	MANY

Table A-13. Children of the 4-port 24-channel PRI T1 Card/PPort

4-port 30-channel E1 Card/PPort

Table A-14. Children of the 4-port 30-channel E1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI LPort	MANY
FR NNI LPort	MANY
SMDS DXI/SNI DCE LPort	MANY





Table A-14. Children of the 4-port 30-channel E1 Card/PPort (Continued)

Child Type	Cardinality
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others PPP-to-1490 Translation LPort	MANY

2-port HSSI Card/PPort

Table A-15. Children of the 2-port HSSI Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
SMDS OPTimum Trunk LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY



Table A-15. Children of the 2-port HSSI Card/PPort (Continued)

Child Type	Cardinality
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others PPP-to-1490 Translation LPort	MANY

10-port DSX-1 Card/PPort

Table A-16. Children of the 10-port DSX-1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
SMDS OPTimum Trunk LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Optimum Frame Trunk LPort	MANY



Table A-16. Children of the 10-port DSX-1 Card/PPort (Continued)

Child Type	Cardinality
ATM Network Internetworking for FR NNI LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others PPP-to-1490 Translation LPort	MANY

1-port ATM UNI DS3 Card/PPort

Table A-17. Children of the 1-port ATM UNI DS3 Card/PPort

Child Type	Cardinality
ATM UNI DTE LPort	MANY
ATM Direct Trunk LPort	MANY
ATM Optimum Cell Trunk LPort	MANY
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY

1-port ATM IWU OC3 Card/PPort

Table A-18. Children of the 1-port ATM IWU OC3 Card/PPort

Child Type	Cardinality
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Direct Trunk LPort	MANY



Table A-18. Children of the 1-port ATM IWU OC3 Card/PPort (Continued)

Child Type	Cardinality
ATM Optimum Cell Trunk LPort	MANY
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY

1-port ATM CS/DS3 Card/PPort

Table A-19. Children of the 1-port ATM CS/DS3 Card/PPort

Child Type	Cardinality
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Direct Trunk LPort	MANY
ATM Optimum Cell Trunk LPort	MANY
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY

1-port ATM CS/E3 Card/PPort

Table A-20. Children of the 1-port ATM CS/E3 Card/PPort

Child Type	Cardinality
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Direct Trunk LPort	MANY



Table A-20. Children of the 1-port ATM CS/E3 Card/PPort (Continued)

Child Type	Cardinality
ATM Optimum Cell Trunk LPort	MANY
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY

4-port Unchannelized T1 Card/PPort

Table A-21. Children of the 4-port Unchannelized T1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
SMDS OPTimum Trunk LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY
Others Direct Line Trunk LPort	MANY



Table A-21. Children of the 4-port Unchannelized T1 Card/PPort (Continued)

Child Type	Cardinality
Otbers Encapsulation FRAD LPort	MANY
Others PPP-to-1490 Translation LPort	MANY

12-port E1 Card/PPort

Table A-22. Children of the 12-port E1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
SMDS OPTimum Trunk LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others PPP-to-1490 Translation LPort	MANY



4-port Unchannelized E1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
SMDS OPTimum Trunk LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others PPP-to-1490 Translation LPort	MANY

Table A-23. Children of the 4-port Unchannelized E1 Card/PPort





4-port 24-channel DSX Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others PPP-to-1490 Translation LPort	MANY

Table A-24. Children of the 4-port 24-channel DSX Card/PPort

1-port ATM UNI E3 Card/PPort

Table A-25.	Children of the 1-port ATM UNI E3 Card/PPort
-------------	--

Child Type	Cardinality
ATM UNI DTE LPort	MANY
ATM Direct Trunk LPort	MANY
ATM Optimum Cell Trunk LPort	MANY



Table A-25. Children of the 1-port ATM UNI E3 Card/PPort (Continued)

Child Type	Cardinality
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY

4-port 32-channel PRI E1 Card/PPort

Table A-26. Children of the 4-port 32-channel PRI E1 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
SMDS DXI/SNI DCE LPort	MANY
SMDS DXI/SNI DTE LPort	MANY
SMDS SSI DTE LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others ISDN PRI D - Channel LPort	MANY
Others PPP-to-1490 Translation LPort	MANY

1-port 28-channel DS3 Card/PPort

Child Type	Cardinality
FR UNI DCE LPort	MANY
FR UNI DTE LPort	MANY
FR NNI LPort	MANY
FR OPTimum PVC Trunk LPort	MANY
SMDS DXI/SNI DCE Lport	MANY
SMDS DXI/SNI DTE Lport	MANY
SMDS SSI DTE LPort	MANY
SMDS OPTimum Trunk LPort	MANY
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM Optimum Frame Trunk LPort	MANY
ATM Network Internetworking for FR NNI LPort	MANY
Others Direct Line Trunk LPort	MANY
Otbers Encapsulation FRAD LPort	MANY
Others PPP-to-1490 Translation LPort	MANY

Table A-27. Children of the 1-port 28-channel DS3 Card/PPort



8-port DS3 Card/PPort

Table A-28. Children of the 8-port DS3 Card/PPort

Child Type	Cardinality
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM NNI LPort	MANY
ATM Direct Trunk LPort	MANY
ATM Optimum Cell Trunk LPort	MANY

8-port E3 Card/PPort

Table A-29. Children of the 8-port E3 Card/PPort

Child Type	Cardinality
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM NNI LPort	MANY
ATM Direct Trunk LPort	MANY
ATM Optimum Cell Trunk LPort	MANY

4-port OC-3c/STM-1 Card/PPort

Table A-30. Children of the 4-port OC-3c/STM-1 Card/PPort

Child Type	Cardinality
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY



Table A-30. Children of the 4-port OC-3c/STM-1 Card/PPort (Continued)

Child Type	Cardinality
ATM NNI LPort	MANY
ATM Direct Trunk LPort	MANY
ATM Optimum Cell Trunk LPort	MANY

8-port T1 Card/PPort

Table A-31. Children of the 8-port T1 Card/PPort

Child Type	Cardinality
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM NNI LPort	MANY
ATM Direct Trunk LPort	MANY
ATM Optimum Cell Trunk LPort	MANY

8-port E1 Card/PPort

Table A-32. Children of the 8-port E1 Card/PPort

Child Type	Cardinality
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM NNI LPort	MANY
ATM Direct Trunk LPort	MANY
ATM Optimum Cell Trunk LPort	MANY



1-port OC-12c/STM-4 Card/PPort

Table A-33. Children of the 1-port OC-12c/STM-4 Card/PPort

Child Type	Cardinality
ATM UNI DCE LPort	MANY
ATM UNI DTE LPort	MANY
ATM NNI LPort	MANY
ATM Direct Trunk LPort	MANY
ATM Optimum Cell Trunk LPort	MANY



INDEX



Index

A

API usage recompiling an existing application 2-31 writing a C program 1-47, 2-31 writing a C++ program 1-48, 2-31 **Application Toolkit** installation instructions 2-4 to 2-10 installed files 2-10 to 2-13 overview 1-4 to 1-5 post-installation tasks 2-8 to 2-10 recompiling an existing application 2-31 un-installation instructions 2-30 upgrading an existing application 2-31 writing a provisioning application 2-31 Aps object ID 1-20 operations and limitations 1-22 Argument list for the CLI 1-40 in C 1-40 in C++ 1-40 methods for specifying variable arguments 1-8, 1-47 AssignedSvcSecScn object ID 1-20 operations and limitations 1-23 Asynchronous functions 1-5 to 1-8 Asynchronous Transfer Mode. See ATM ATM Network Interworking for Frame Relay NNI object ID 1-21 ATM Transport for FR NNI LPorts object ID 1 - 20ATM Virtual UNI LPorts object ID 1-20

Attribute list. *See* Argument list Attributes, how represented for the CLI 3-3 Automatic Protection Switching. *See* Aps

C C

argument list 1-40 interface for API functions 1-4 writing a program 1-47, 2-31 C++ argument list 1-40 interface for API functions 1-4 writing a program 1-48, 2-31 Cache, used to store MIB data in memory 2-21, 4-19 to 4-23 Card object ID 1-20 operations and limitations 1-23 Channel object ID 1-21 operations and limitations 1-23 Circuit DLCI for Frame Relay circuits 1-21 object ID 1-21 object ID for ATM Network Interworking for Frame Relay NNI 1-21 operations and limitations 1-24 VCI for ATM circuits 1-21 VPI for ATM circuits 1-21 Class B addressing 1-46 CLI argument list 1-40

INDEX

commands 3-6 to 3-28 controlling SNMP parameters 2-15 defined 1-4, 3-2 enclosing strings in quotes 3-4, 3-29 examples 3-29 to 3-38 identifying the Provisioning Server to the CLI 2-14 installed files 2-10 specifying abbreviated attribute IDs 3-3 specifying abbreviated enumerated attribute values 3-4 specifying modification type 2-14 specifying security settings 2-15 stopping and restarting 2-25 testing the CLI 2-10 troubleshooting problems 2-25 to 2-29 usage 3-2 to 3-5 writing a provisioning script 2-31 Client include files 2-10 to 2-13 Client libraries 2-10 Column access specifiers in MIB tables 4-17 Command Error Table 2-21, 2-23, 4-3, 4-19 Command Line Interface. See CLI Community name, for authentication and access-control 2-22, 4-2 Configuration variables. See Environment variables Containment hierarchy 1-17 to 1-19, A-1 to A-23 Core file, specifying location 2-18 Customer object ID 1-20 operations and limitations 1-25 CvArgId.H header file 2-12 CvClient.H header file 2-11 CvDefs.H header file 2-11 CvE164Address.H header file 2-12 CvErrors.H header file 2-13

CvErrors.h header file 2-13 CvObjectId.H header file 2-12 CvObjectType.H header file 2-11 CvParamValues.H header file 2-12 CvSVCAddress.H header file 2-12 CvUSL.H header file 2-12

D

Data link connection identifier. *See* DLCI Database locking, for MIB objects 4-20 to 4-23 DLCI, for Frame Relay circuits 1-21

Е

Environment variables configuring the CLI 2-13 to 2-15 configuring the MIB 2-21 to 2-23 configuring the Provisioning client 2-16 configuring the Provisioning Server 2-17 to 2-24 Extended Super Frame. *See* Pfdl

F

Files installed with Provisioning Server and Application Toolkit 2-10 to 2-13 FR NNI LPort object ID 1-20 Functions asynchronous 1-5 to 1-8 naming conventions 1-9 operational functions 1-9, 1-10 to 1-11, 1-47, 1-48 select loop processing functions 1-9, 1-11, 1-47, 1-48session control functions 1-9, 1-10, 1-47, 1-48synchronous 1-5 to 1-6 utility functions 1-9, 1-12 to 1-13, 1-47, 1-48


Header files 2-10 to 2-13

I

Include files for client 2-10 to 2-13 Installation instructions 2-4 to 2-10 Installed files for CLI 2-10 for MIB 4-2 for Provisioning Server and Application Toolkit 2-10 to 2-13

L

Libraries for client 2-10 Locked database 1-2, 2-21, 4-20 to 4-23 Logical port. *See* LPort LPort object ID 1-20 object ID for ATM Transport for FR NNI LPorts 1-20 object ID for ATM Virtual UNI LPorts 1-20 operations and limitations 1-25 start VPI for Virtual UNI LPort 1-26

Μ

MIB

cache 2-21, 4-19 to 4-23 column access specifiers 4-17 Command Error Table 4-3, 4-19 community name 4-2 compiling 4-2 controlling object locking 2-21, 4-20 to 4-23examples 4-24 to 4-48 identifying agent port 2-18 installed file 4-2 ModifyType attribute 4-18



OID for MIB objects 4-3 overview 1-4 row aliasing 4-16 RowStatus attribute 4-17 SNMP commands supported 4-18 specifying an OID 4-23 to 4-48 structure 4-3 to 4-18 various tables of 4-4 to 4-15 viewing 4-2 ModifyType attribute in MIB tables 4-18

N

Naming conventions for functions 1-9 for object IDs 1-20 NavisXtend Provisioning Server Application Toolkit. *See* Application Toolkit. NavisXtend Provisioning Server. *See* Provisioning Server NetCac object ID 1-20 operations and limitations 1-26 Network object ID 1-21 operations and limitations 1-26

0

Object Attributes 1-40 Object ID Aps 1-20 AssignedSvcSecScn 1-20 ATM Network Interworking for Frame Relay NNI 1-21 ATM Transport for FR NNI LPorts 1-20 card 1-20 channel 1-21 circuit 1-21 customer 1-20

NavisXtend Provisioning Server User's Guide

Index-3





TrafficDesc 1-20 TrafficShaper 1-20 VPN 1-20 Object identifier. *See* Object ID Object types, supported 1-14 to 1-37 Operational functions 1-9, 1-10 to 1-11

P

PerformanceMonitor object ID 1-20 operations and limitations 1-27 PFd1 object ID 1-20 operations and limitations Physical port. See PPort **PMPCkt** object ID 1-21 operations and limitations 1-27 **PMPCktRoot** object ID 1-21 operations and limitations 1-27 **PMPSpvcLeaf** object ID 1-20 operations and limitations 1-28 **PMPSpvcRoot** object ID 1-21 operations and limitations 1-28 Point-to-MultiPoint circuit leaf. See PMPCkt Point-to-MultiPoint circuit root. See **PMPCkRoot** Point-to-MultiPoint SPVC leaf. See **PMPSpvcLeaf** Point-to-MultiPoint SPVC root, See **PMPSpvcRoot** Post-installation tasks 2-8 to 2-10 **PPort** object ID 1-20 operations and limitations 1-28

NavisXtend Provisioning Server User's Guide



R

Row aliasing in MIB tables 4-16 RowStatus attribute in MIB tables 4-17

S

Sample code 2-11 Security settings CLI 2-15 **Provisioning Server 2-24** Select loop processing functions 1-9, 1-11 Server port, identifying 2-18 ServiceName object ID 1-20 operations and limitations 1-28 ServiceName endpoint, object ID 1-21 Session control functions 1-9, 1-10 SMDS address prefix object ID 1-22 operations and limitations 1-29 SMDS alien group address object ID 1-22 operations and limitations 1-29 SMDS alien individual address object ID 1-22 operations and limitations 1-29 SMDS country code object ID 1-22 operations and limitations 1-30 SMDS group screen object ID 1-20 operations and limitations 1-30 SMDS individual screen object ID 1-20 operations and limitations 1-30 SMDS local individual address object ID 1-22 operations and limitations 1-31 SMDS netwide group address object ID 1-22 operations and limitations 1-31 SMDS SSI individual address, operations and limitations 1-31

NavisXtend Provisioning Server User's Guide





SvcCUGMbrRule object ID 1-20 operations and limitations 1-34 SvcNodePrefix object ID 1-22 operations and limitations 1-34 **SvcPrefix** object ID 1-22 operations and limitations 1-34 SvcSecScn object ID 1-20 operations and limitations 1-36 SvcSecScnActParam object ID 1-20 operations and limitations 1-36 **SvcUserPart** object ID 1-22 operations and limitations 1-36 Switch object ID 1-22 operations and limitations 1-36 Synchronous functions 1-5 to 1-6

Т

Testing CLI 2-10 Provisioning Server 2-8 Trace file enabling client trace file 2-16 enabling server trace files 2-19 TrafficDesc object ID 1-20 operations and limitations 1-36 TrafficShaper object ID 1-20 operations and limitations 1-37 Troubleshooting problems 2-25 to 2-29

NavisXtend Provisioning Server User's Guide



U

Un-installation instructions 2-30 Utility functions 1-9, 1-12 to 1-13

V

Variable argument list 1-8, 1-47 VCI for ATM circuits 1-21 Virtual Channel Identifier. *See* VCI Virtual Path Identifier. *See* VPI VPI (start) for Virtual UNI LPort 1-26 VPI for ATM circuits 1-21 VPN object ID 1-20 operations and limitations 1-37

W

Writing a provisioning script using CLI 2-31 Writing programs basic steps in C 1-47, 2-31 basic steps in C++ 1-48, 2-31 recompiling existing application 2-31 upgrading an existing application 2-31